



**IEEE Standard for
Information technology—
Telecommunications and information
exchange between systems—
Local and metropolitan area networks—
Specific requirements**

**Part 3: Carrier Sense Multiple Access with
Collision Detection (CSMA/CD) Access Method
and Physical Layer Specifications**

**Amendment 4: Media Access Control Parameters, Physical
Layers, and Management Parameters for 40 Gb/s and
100 Gb/s Operation**

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

IEEE
3 Park Avenue
New York, NY 10016-5997, USA

22 June 2010

IEEE Std 802.3ba™-2010
(Amendment to
IEEE Std 802.3™-2008)

802.3baTM

IEEE Std 802.3ba™-2010

(Amendment to
IEEE Std 802.3™-2008)

**IEEE Standard for
Information technology—
Telecommunications and information
exchange between systems—
Local and metropolitan area networks—
Specific requirements**

**Part 3: Carrier Sense Multiple Access with
Collision Detection (CSMA/CD) Access Method
and Physical Layer Specifications**

**Amendment 4: Media Access Control Parameters, Physical
Layers, and Management Parameters for 40 Gb/s and
100 Gb/s Operation**

**LAN/MAN Standards Committee
of the
IEEE Computer Society**

Approved 17 June 2010
IEEE-SA Standards Board

Approved 5 May 2011
American National Standards Institute

Abstract: This amendment to IEEE Std 802.3-2008 includes changes to IEEE Std 802.3-2008 and adds Clause 80 through Clause 88, Annex 83A through Annex 83C, Annex 85A, and Annex 86A. This amendment includes IEEE 802.3 Media Access Control (MAC) parameters, Physical Layer specifications, and management parameters for the transfer of IEEE 802.3 format frames at 40 Gb/s and 100 Gb/s.

Keywords: 802.3ba, 40 Gb/s Ethernet, 40GBASE-CR4, 40GBASE-KR4, 40GBASE-LR4, 40GBASE-R, 40GBASE-SR4, 100 Gb/s Ethernet, 100GBASE-CR10, 100GBASE-ER4, 100GBASE-LR4, 100GBASE-R, 100GBASE-SR10, AN, Auto-Negotiation, Backplane, BP, CAUI, CGMII, CPPI, FEC, MMF, SMF, Optical Transport Network, OTN, XLAUI, XLGMII, XLPPI

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Copyright © 2010 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 22 June 2010. Printed in the United States of America.

IEEE and 802 are registered trademarks in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

PDF: ISBN 978-0-7381-6322-2 STD96079
Print: ISBN 978-0-7381-6323-9 STDPD96079

Second Printing 3 August 2010

IEEE prohibits discrimination, harassment and bullying. For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information or the soundness of any judgments contained in its standards.

Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied **“AS IS.”**

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, or every ten years for stabilization. When a document is more than five years old and has not been reaffirmed, or more than ten years old and has not been stabilized, 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.

In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon the advice of a competent professional in determining the exercise of reasonable care in any given circumstances.

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 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 societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration. A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered the official position of IEEE or any of its committees and shall not be considered to be, nor be relied upon as, a formal interpretation of the IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position, explanation, or interpretation of the IEEE. 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. Recommendations to change the status of a stabilized standard should include a rationale as to why a revision or withdrawal is required.

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

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854
USA

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; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Introduction

This introduction is not part of IEEE Std 802.3ba-2010, IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements, Part 3: CSMA/CD Access Method and Physical Layer Specifications, Amendment 4: Media Access Control Parameters, Physical Layers and Management Parameters for 40 Gb/s and 100 Gb/s Operation.

IEEE Std 802.3™ was first published in 1985. Since the initial publication, many projects have added functionality or provided maintenance updates to the specifications and text included in the standard. Each IEEE 802.3 project/amendment is identified with a suffix (e.g., IEEE Std 802.3ba-2010).

The Media Access Control (MAC) protocol specified in IEEE Std 802.3 is Carrier Sense Multiple Access with Collision Detection (CSMA/CD). This MAC protocol was included in the experimental Ethernet developed at Xerox Palo Alto Research Center. While the experimental Ethernet had a 2.94 Mb/s data rate, IEEE Std 802.3-1985 specified operation at 10 Mb/s. Since 1985 new media options, new speeds of operation, and new capabilities have been added to IEEE Std 802.3.

Some of the major additions to IEEE Std 802.3 are identified in the marketplace with their project number. This is most common for projects adding higher speeds of operation or new protocols. For example, IEEE Std 802.3u™ added 100 Mb/s operation (also called Fast Ethernet), IEEE Std 802.3x™ specified full duplex operation and a flow control protocol, IEEE Std 802.3z™ added 1000 Mb/s operation (also called Gigabit Ethernet), IEEE Std 802.3ae™ added 10 Gb/s operation (also called 10 Gigabit Ethernet) and IEEE Std 802.3ah™ specified access network Ethernet (also called Ethernet in the First Mile). These major additions are all now included in, and are superseded by, IEEE Std 802.3-2008 and are not maintained as separate documents.

At the date of IEEE Std 802.3ba-2010 publication, IEEE Std 802.3 is comprised of the following documents:

IEEE Std 802.3-2008

Section One—Includes Clause 1 through Clause 20 and Annex A through Annex H and Annex 4A. Section One includes the specifications for 10 Mb/s operation and the MAC, frame formats and service interfaces used for all speeds of operation.

Section Two—Includes Clause 21 through Clause 33 and Annex 22A through Annex 33E. Section Two includes management attributes for multiple protocols and speed of operation as well as specifications for providing power over twisted pair cabling for multiple operational speeds. It also includes general information on 100 Mb/s operation as well as most of the 100 Mb/s Physical Layer specifications.

Section Three—Includes Clause 34 through Clause 43 and Annex 36A through Annex 43C. Section Three includes general information on 1000 Mb/s operation as well as most of the 1000 Mb/s Physical Layer specifications.

Section Four—Includes Clause 44 through Clause 55 and Annex 44A through Annex 55B. Section Four includes general information on 10 Gb/s operation as well as most of the 10 Gb/s Physical Layer specifications.

Section Five—Includes Clause 56 through Clause 74 and Annex 57A through Annex 74A. Clause 56 through Clause 67 and associated annexes specify subscriber access and other Physical Layers and sublayers for operation from 512 kb/s to 1000 Mb/s, and defines services and protocol elements that enable the exchange of IEEE Std 802.3 format frames between stations in a subscriber access network.

Clause 68 specifies a 10 Gb/s Physical Layer specification. Clause 69 through Clause 74 and associated annexes specify Ethernet operation over electrical backplanes at speeds of 1000 Mb/s and 10 Gb/s.

IEEE Std 802.3av™-2009

This amendment includes changes to IEEE Std 802.3-2008 and adds Clause 75 through Clause 77 and Annex 75A through Annex 76A. This amendment adds new Physical Layers for 10 Gb/s operation on point-to-multipoint passive optical networks.

IEEE Std 802.3bc™-2009

This amendment includes changes to IEEE Std 802.3-2008 and adds Clause 79. This amendment moves the Ethernet Organizationally Specific Type, Length, Value (TLV) information elements that were specified in IEEE Std 802.1AB™ to IEEE Std 802.3.

IEEE Std 802.3at™-2009

This amendment includes changes to IEEE Std 802.3-2008. This amendment augments the capabilities of IEEE Std 802.3-2008 with higher power levels and improved power management information.

IEEE Std 802.3-2008™/Cor 1-2009

This corrigendum corrects the PAUSE reaction timing delay value for the 10GBASE-T PHY type.

IEEE Std 802.3ba™-2010

This amendment includes changes to IEEE Std 802.3-2008 and adds Clause 80 through Clause 88 and Annex 83A through Annex 83C, Annex 85A, and Annex 86A. This amendment adds MAC parameters, Physical Layers, and management parameters for the transfer of IEEE 802.3 format frames at 40 Gb/s and 100 Gb/s.

IEEE Std 802.3 will continue to evolve. New Ethernet capabilities are anticipated to be added within the next few years as amendments to this standard.

Notice to users

Laws and regulations

Users of these documents should consult all applicable laws and regulations. Compliance with the provisions of this standard does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

This document is copyrighted by the IEEE. It is made available for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making this document

available for use and adoption by public authorities and private users, the IEEE does not waive any rights in copyright to this document.

Updating of IEEE documents

Users of IEEE standards should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect. In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE Standards Association website at <http://ieeexplore.ieee.org/xpl/standards.jsp>, or contact the IEEE at the address listed previously.

For more information about the IEEE Standards Association or the IEEE standards development process, visit the IEEE-SA website at <http://standards.ieee.org>.

Errata

Errata, if any, for this and all other standards can be accessed at the following URL: <http://standards.ieee.org/reading/ieee/updates/errata/index.html>. Users are encouraged to check this URL for errata periodically.

Downloads

Portions of this standard can be downloaded from the Internet. Materials include PICS tables, data tables, and code. URLs are listed in the text in the appropriate sections.

Interpretations

Current interpretations can be accessed at the following URL: <http://standards.ieee.org/reading/ieee/interp/index.html>.

Patents

Attention is called to the possibility that implementation of this amendment may require use of subject matter covered by patent rights. By publication of this amendment, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this amendment are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

The following individuals were officers and members of the IEEE 802.3 working group at the beginning of the working group ballot:

David J. Law, *IEEE 802.3 Working Group Chair*
Wael William Diab, *IEEE 802.3 Working Group Vice-Chair*
Steven B. Carlson, *IEEE 802.3 Working Group Executive Secretary*
Adam Healey, *IEEE 802.3 Working Group Secretary*
Bradley Booth, *IEEE 802.3 Working Group Treasurer*

John D’Ambrosia, *IEEE P802.3ba Task Force Chair*
Ilango S. Ganga, *IEEE P802.3ba Task Force Editor-in-Chief, Editor, Clauses 1,4,80, Annexes A, 4A, 31B*
Mark Gustlin, *IEEE P802.3ba “Logic” Sub-task Force Chair, Editor, Clauses 81, 82*
Chris Di Minico, *IEEE P802.3ba “Cu” Sub-task Force Chair, Editor, Clause 85, Annex 85A*
Pete Anslow, *IEEE P802.3ba “Optical” Sub-task Force Chair, Editor, Clause 88*
Hugh Barrass, *IEEE P802.3ba Task Force Editor, Clauses 30, 45*
Piers Dawe, *IEEE P802.3ba Task Force Editor, Clause 52, Clause 86, Annex 86A*
Jonathan King,^a *IEEE P802.3ba Task Force Editor, Clause 87*
Ryan Latchman, *IEEE P802.3ba Task Force Editor, Annex 83A, Annex 83B*
Arthur Marris, *IEEE P802.3ba Task Force Editor, Clauses 69, 73, 74, 84, Annex 69A, Annex 69B*
Steve Trowbridge, *IEEE P802.3ba Task Force Editor, Clause 83, Annex 83C*
George Oulundsen, *IEEE P802.3ba Task Force Secretary*
Frank Chang, *IEEE P802.3ba Task Force Web Master*

Ghani Abbas	Christopher R. Cole	Marek Hajduczenia
John Abbott	Doug Coleman	Hiroshi Hamano
Justin Abbott	Herbert V. Congdon	Bernie Hammond
Akira Agata	Charles Cook	Greg Hankins
Arne Alping	Fumio Daido	Robert Hays
Yehuda Alush	Yair Darshan	Kengo Hirano
Vittal Balasubramanian	Bill Delveaux	Ryan Hirth
Thananya Baldwin	Bryan Dietz	Keith Hopwood
Jaya Bandyopadhyay	Thomas Dineen	Rita Horner
Ozdal Barkan	Thuyen Dinh	Stanley Hronik
Jim Barnette	Dan Dove	Dean Huumala
Howard Baumer	Mike Dudek	Thong Huynh
Denis Beaudoin	Joseph Dupuis	Hiroki Ikeda
Christian Beia	Frank Effenberger	Kazuhiko Ishibe
Michael Bennett	George Eisler	Osamu Ishida
Ernest Bergmann	David Estes	Hideki Isono
Ralf-Peter Braun	John Ewen	John Jetzt
Dirk Breuer	Daniel Feldman	Jack L. Jewell
Alan Brown	Dongning Feng	Jessica Jiang
Robert Busse	Alan Flatman	Wenbin Jiang
Maurice Caldwell	Norbert Folkens	Thomas K. Joergensen
J. Martin Carroll	Howard Frazier	Chad Jones
Mandeep Chadha	Richard Frosch	Bheom-Soon Joo
David Chalupsky	Ali Ghiasi	Sanjay Kasturia
Sun-Hyok Chang	Dimitrios Giannakopoulos	Yasuaki Kawatsu
Jian Chen	Larry Green	Seung-Hwan Kim
Joseph Chou	Michael Grimwood	Yong Kim
Hwan-Seok Chung	Robert Grow	Mitsunobu Kimura
Terry Cobb	Paul Gyugyi	Scott Kipp

^aNot a member of the IEEE 802.3 working group at the beginning of the working group ballot.

Shoukei Kobayashi
David Koenen
Paul Kolesar
Seiji Kozaki
Glen Kramer
Joerg Kropp
Yasuyuki Kuroda
Toshihiko Kusano
Hans Lackner
Lowell Lamb
D. Matthew Landry
Jeff Lapak
Kyungh Lee
Andreas Lenkisch
Raymond W. K. Leung
Mike Peng Li
Ru Jian Lin
Robert Lingle
James Lott
Sharon Lutz
Eric Lynskey
Anthony Magee
Joseph Maggolino
Valerie Maguire
Jeffery J. Maki
Jeff Mandin
Carlo Mariotti
Arlon Martin
C. Phillip McClay
Michael S. McCormack
John McDonough
Jim McGrath
Greg McSorley
Richard Mellitz
Jeffrey Meyer
Tremont Miao
Andy Moorwood
Kazuyuki Mori
Shimon Muller
Angela Muscat
Gerard Nadeau
Takeshi Nagahori

Edward Nakamoto
Jay Neer
Gary Nicholl
George Noh
Takumi Nomura
Ronald Nordin
Ahmad Nouri
Mark Nowell
Satoshi Obara
David Ofelt
Gourgen Oganessyan
Padraig OMathuna
Akihiro Otaka
Tom Palkert
Sesha Panguluri
Gavin Parnaby
Shashi Patel
Martin Patoka
Petar Pepeljugoski
Gerald Pepper
John Petrilla
Velu Pillai
Rick Pimpinella
Scott Powell
Holger Quast
Rick Rabinovich
Randy Rannow
Duane Remein
Tamir Reshef
Michael Ressler
June-Koo (Kevin) Rhee
Poldi (Pavlick) Rimboim
Sam Sambasivan
Ramesh Sastry
Olindo Savi
Edward Sayre
Frederick Schindler
Thomas Schrans
Shawn Searles
Ted Seely
Khorvash (Kory) Sefidvash
Murat Serbay

Farhad Shafai
Masayuki Shigematsu
Jong-Yoon Shin
Larry Shorthill
Jesse Simsarian
Clayton Stanford
Henk Steenman
Christopher Stook
Olaf Storaasli
Alan Sugg
Ken-Ichi Suzuki
Naoki Suzuki
Steve Swanson
Andre Szczepanek
Dimitry Taich
Akio Tajima
Hidenori Takahashi
Motoyuki Takizawa
Keiji Tanaka
Sashi Thiagarajan
Geoffrey O. Thompson
Hidehiro Toyoda
Mario Traeber
Matthew Traverso
Shinji Tsuji
Eddie Tsumura
Brad Turner
Alexander Umnov
Sterling A. Vaden
Paul Vanderlaan
Albert Vareljian
Anoop Vetteth
Ionel Marius Vladan
Andrew Weitzner
Martin White
Bill Woodruff
Masaki Yasukawa
Tetsuya Yokomoto
George Young
George Zimmerman
Pavel Zivny

The following members of the individual balloting committee voted on this standard. Balloters may have voted for approval, disapproval, or abstention.

John Abbott	Marek Hajduczenia	Charles Moore
Thomas Alexander	Hiroshi Hamano	Charles Moorwood
Jon Anderson	John Hawkins	Jose Morales
Peter Anslow	Adam Healey	Shimon Muller
Butch Anton	Rita Horner	Michael S. Newman
Danilo Antonelli	Paul Isaacs	Paul Nikolich
Jacob Ben Ary	Akio Iso	Kevin Noll
John Barr	Atsushi Ito	Satoshi Obara
Hugh Barrass	Raj Jain	David Ofelt
Michael Bennett	Tal Kaitz	George Oulundsen
Bradley Booth	Shinkyō Kaku	Peter Pepeljugoski
Ralf-Peter Braun	Piotr Karocki	John Petrilla
Matthew Brown	Stuart J. Kerry	Rick Pimpinella
William Byrd	Yongbum Kim	Maximilian Riegel
Steven B. Carlson	Mitsunobu Kimura	Robert Robinson
Juan Carreon	Scott Kipp	Randall Safier
Jon Caswell	Paul Kolesar	Bartien Sayogo
David Chalupsky	Seiji Kozaki	Rich Seifert
Keith Chow	Glen Kramer	Farhad Shafai
Terry Cobb	Geoff Ladwig	Gil Shultz
Charles Cook	D. Matthew Landry	Kapil Sood
John D'Ambrosia	David J. Law	Matthew Squire
Piers Dawe	Li Li	Thomas Starai
Wael William Diab	Ru Lin	Walter Struppler
Chris Di Minico	Shen Loh	Steven Swanson
Thomas Dineen	William Lumpkins	Andre Szczepanek
Michael Dudek	G. Luri	Joseph Tardo
Sourav Dutta	Elvis Maculuba	William Taylor
Frank Effenberger	Valerie Maguire	Michael Johas Teener
Howard M. Frazier	Jeffery J. Maki	Patricia Thaler
Yukihiro Fujimoto	Mark Maloney	James Theodoras
Ilango S. Ganga	Arthur Marris	Geoffrey O. Thompson
Devon Gayle	David Martin	Michael Thompson
Ali Ghiasi	Peter Martini	Stephen Trowbridge
Joel Goergen	C. Phillip McClay	Edward J. Turner
Sergiu Goma	Richard Mellitz	Mark-Rene Uchida
Larry Green	Steven Methley	Prabodh Varshney
Randall Groves	Tremont Miao	Kunpeng Wu
Mark Gustlin	Brian Misek	George Young
C. Guy	Merrick Moeller	Oren Yuen
Stephen Haddock		Zhen Zhou

When the IEEE-SA Standards Board approved this standard on 17 June 2010, it had the following membership:

Robert M. Grow, *Chair*
Richard H. Hulett, *Vice Chair*
Steve M. Mills, *Past Chair*
Judith Gorman, *Secretary*

Karen Bartleson
Victor Berman
Ted Burse
Clint Chaplin
Andy Drozd
Alexander Gelman
Jim Hughes

Young Kyun Kim
Joseph L. Koepfinger*
John Kulick
David J. Law
Hung Ling
Oleg Logvinov
Ted Olsen

Ronald C. Petersen
Thomas Prevost
Jon Walter Rosdahl
Sam Sciacca
Mike Seavey
Curtis Siller
Don Wright

*Member Emeritus

Also included are the following nonvoting IEEE-SA Standards Board liaisons:

Satish K. Aggarwal, NRC Representative
Richard DeBlasio, DOE Representative
Michael Janezic, NIST Representative

Michelle Turner
IEEE Standards Program Manager, Document Development

Kathryn Bennett
IEEE Standards Program Manager, Technical Program Development

List of special symbols

For the benefit of those who have received this document by electronic means, what follows is a list of special symbols and operators. If any of these symbols or operators fail to print out correctly on your machine, the editors apologize, and hope that this table will at least help you to sort out the meaning of the resulting funny-shaped blobs and strokes

Special symbols and operators

Printed character	Meaning	Font
*	Boolean AND	Symbol
+	Boolean OR, arithmetic addition	Symbol
^	Boolean XOR	Times New Roman
!	Boolean NOT	Symbol
×	Multiplication	Symbol
<	Less than	Symbol
≤	Less than or equal to	Symbol
>	Greater than	Symbol
≥	Greater than or equal to	Symbol
=	Equal to	Symbol
≈	Approximately equal to	Symbol
≠	Not equal to	Symbol
←	Assignment operator	Symbol
∈	Indicates membership	Symbol
∉	Indicates nonmembership	Symbol
±	Plus or minus (a tolerance)	Symbol
°	Degrees	Symbol
∑	Summation	Symbol
√	Square root	Symbol
—	Big dash (em dash)	Times New Roman
-	Little dash (en dash), subtraction	Times New Roman
	Vertical bar	Times New Roman
†	Dagger	Times New Roman
‡	Double dagger	Times New Roman
α	Lower case alpha	Symbol
β	Lower case beta	Symbol
γ	Lower case gamma	Symbol
δ	Lower case delta	Symbol
ε	Lower case epsilon	Symbol
λ	Lambda	Symbol
μ	Micro	Times New Roman
Ω	Omega	Symbol

Contents

1.	Introduction.....	2
	1.2.3 Physical Layer and media notation.....	2
	1.3 Normative references.....	3
	1.4 Definitions.....	4
	1.5 Abbreviations.....	5
4.	Media Access Control.....	7
	4.4.2 MAC parameters.....	7
30.	Management.....	9
45.	Management Data Input/Output (MDIO) Interface.....	17
	45.2 MDIO Interface Registers.....	17
	45.2.1 PMA/PMD registers.....	17
	45.2.3 PCS registers.....	44
	45.2.7 Auto-Negotiation registers.....	62
	45.5 Protocol implementation conformance statement (PICS) proforma for Clause 45, MDIO interface.....	64
52.	Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-S (short wavelength serial), 10GBASE-L (long wavelength serial), and 10GBASE-E (extra long wavelength serial).....	71
	52.9.10 Transmitter and dispersion penalty measurement.....	71
69.	Introduction to Ethernet operation over electrical backplanes.....	73
	69.1 Overview.....	73
	69.1.1 Scope.....	73
	69.1.2 Objectives.....	73
	69.1.3 Relationship of Backplane Ethernet to the ISO OSI reference model.....	73
	69.2 Summary of Backplane Ethernet Sublayers.....	74
	69.2.1 Reconciliation Sublayer and media independent interfaces.....	74
	69.2.3 Physical Layer signaling systems.....	75
	69.2.5 Management.....	75
	69.3 Delay constraints.....	75
73.	Auto-Negotiation for backplane and copper cable assembly.....	77
	73.2 Relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model.....	77
	73.3 Functional specifications.....	77
	73.5 DME transmission.....	78
	73.5.1 DME electrical specifications.....	78
	73.6 Link codeword encoding.....	79
	73.6.4 Technology Ability Field.....	79
	73.6.5 FEC capability.....	79
	73.7 Receive function requirements.....	80
	73.7.1 DME page reception.....	80

73.7.2	Receive Switch function	80
73.7.6	Priority Resolution function.....	81
73.10	State diagrams and variable definitions	81
73.10.1	State diagram variables.....	81
73.10.2	State diagram timers	82
73.11	Protocol implementation conformance statement (PICS) proforma for Clause 73, Auto-Negotiation for backplane and copper cable assembly.....	84
73.11.1	Introduction.....	84
73.11.4	PICS proforma tables for Auto-Negotiation for backplane and copper cable assembly.....	84
74.	Forward Error Correction (FEC) sublayer for BASE-R PHYs	87
74.1	Overview.....	87
74.2	Objectives	87
74.3	Relationship to other sublayers.....	87
74.4	Inter-sublayer interfaces	88
74.4.1	Functional block diagram for 10GBASE-R PHYs.....	89
74.4.2	Functional block diagram for 40GBASE-R PHYs.....	89
74.4.3	Functional block diagram for 100GBASE-R PHYs	90
74.5	FEC service interface.....	91
74.5.1	10GBASE-R service primitives.....	92
74.5.2	40GBASE-R and 100GBASE-R service primitives.....	93
74.6	Delay constraints.....	94
74.7	FEC principle of operation	94
74.7.3	Composition of the FEC block	95
74.8	FEC MDIO function mapping	102
74.8.1	FEC capability	102
74.8.2	FEC enable.....	103
74.8.3	FEC Enable Error Indication	103
74.9	BASE-R PHY test-pattern mode	104
74.11	Protocol implementation conformance statement (PICS) proforma for Clause 74, Forward Error Correction (FEC) sublayer for BASE-R PHYs.....	105
74.11.1	Introduction.....	105
74.11.3	Major capabilities/options.....	106
74.11.5	FEC requirements	106
80.	Introduction to 40 Gb/s and 100 Gb/s networks	107
80.1	Overview.....	107
80.1.1	Scope.....	107
80.1.2	Objectives	107
80.1.3	Relationship of 40 Gigabit and 100 Gigabit Ethernet to the ISO OSI reference model.....	107
80.1.4	Nomenclature.....	109
80.1.5	Physical Layer signaling systems	109
80.2	Summary of 40 Gigabit and 100 Gigabit Ethernet sublayers	110
80.2.1	Reconciliation Sublayer (RS) and Media Independent Interface	110
80.2.2	Physical Coding Sublayer (PCS).....	110
80.2.3	Forward Error Correction (FEC) sublayer.....	111
80.2.4	Physical Medium Attachment (PMA) sublayer.....	111
80.2.5	Physical Medium Dependent (PMD) sublayer	111
80.2.6	Auto-Negotiation	111
80.2.7	Management interface (MDIO/MDC).....	111
80.2.8	Management.....	111

80.3	Service interface specification method and notation	112
80.3.1	Inter-sublayer service interface.....	112
80.3.2	Instances of the Inter-sublayer service interface.....	112
80.3.3	Semantics of inter-sublayer service interface primitives.....	115
80.4	Delay constraints.....	116
80.5	Skew constraints	117
80.6	State diagrams.....	121
80.7	Protocol implementation conformance statement (PICS) proforma.....	121
81.	Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation (XLGMII and CGMII).....	123
81.1	Overview.....	123
81.1.1	Summary of major concepts	124
81.1.2	Application.....	124
81.1.3	Rate of operation.....	124
81.1.4	Delay constraints.....	124
81.1.5	Allocation of functions	125
81.1.6	XLGMII/CGMII structure	125
81.1.7	Mapping of XLGMII/CGMII signals to PLS service primitives.....	126
81.2	XLGMII/CGMII data stream.....	128
81.2.1	Inter-frame <inter-frame>.....	129
81.2.2	Preamble <preamble> and start of frame delimiter <sfd>.....	129
81.2.3	Data <data>.....	130
81.2.4	End of frame delimiter <efd>	130
81.2.5	Definition of Start of Packet and End of Packet Delimiters	130
81.3	XLGMII/CGMII functional specifications	130
81.3.1	Transmit	130
81.3.2	Receive.....	134
81.3.3	Error and fault handling.....	137
81.3.4	Link fault signaling.....	138
81.4	Protocol implementation conformance statement (PICS) proforma for Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation	141
81.4.1	Introduction.....	141
81.4.2	Identification.....	141
81.4.3	PICS proforma tables for Reconciliation Sublayer and Media Independent Interface for 40 Gb/s and 100 Gb/s operation.....	142
82.	Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R	147
82.1	Overview.....	147
82.1.1	Scope.....	147
82.1.2	Relationship of 40GBASE-R and 100GBASE-R to other standards	147
82.1.3	Summary of 40GBASE-R and 100GBASE-R sublayers	148
82.1.4	Inter-sublayer interfaces	149
82.1.5	Functional block diagram	149
82.2	Physical Coding Sublayer (PCS).....	150
82.2.1	Functions within the PCS	150
82.2.2	Use of blocks	151
82.2.3	64B/66B transmission code	152
82.2.4	Transmit process	158
82.2.5	Scrambler.....	158
82.2.6	Block distribution	158
82.2.7	Alignment marker insertion	159

82.2.8	BIP calculations	161
82.2.9	PMA or FEC Interface	162
82.2.10	Test-pattern generators	163
82.2.11	Block synchronization	163
82.2.12	PCS lane deskew	163
82.2.13	PCS lane reorder	163
82.2.14	Alignment marker removal	164
82.2.15	Descrambler	164
82.2.16	Receive process.....	164
82.2.17	Test-pattern checker.....	164
82.2.18	Detailed functions and state diagrams	165
82.3	PCS Management	170
82.3.1	PMD MDIO function mapping.....	170
82.4	Loopback	171
82.5	Delay constraints.....	171
82.6	Auto-Negotiation	172
82.7	Protocol implementation conformance statement (PICS) proforma for Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R.....	178
82.7.1	Introduction.....	178
82.7.2	Identification.....	178
82.7.3	Major capabilities/options.....	179
82.7.4	PICS Proforma Tables for PCS, type 40GBASE-R and 100GBASE-R.....	180
82.7.5	Test-pattern modes.....	181
82.7.6	Management.....	182
83.	Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R.....	185
83.1	Overview.....	185
83.1.1	Scope.....	185
83.1.2	Position of the PMA in the 40GBASE-R or 100GBASE-R sublayers.....	185
83.1.3	Summary of functions.....	185
83.1.4	PMA sublayer positioning	185
83.2	PMA interfaces	188
83.3	PMA service interface	188
83.4	Service interface below PMA	190
83.5	Functions within the PMA	191
83.5.1	Per input-lane clock and data recovery	191
83.5.2	Bit-level multiplexing	191
83.5.3	Skew and Skew Variation.....	194
83.5.4	Delay constraints.....	195
83.5.5	Clocking architecture	195
83.5.6	Signal drivers	195
83.5.7	Link status	196
83.5.8	PMA local loopback mode	196
83.5.9	PMA remote loopback mode (optional)	196
83.5.10	PMA test patterns (optional).....	197
83.6	PMA MDIO function mapping	198
83.7	Protocol implementation conformance statement (PICS) proforma for Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R	203
83.7.1	Introduction.....	203
83.7.2	Identification.....	203
83.7.3	Major capabilities/options.....	204
83.7.4	Skew generation and tolerance	206
83.7.5	Test patterns	206

83.7.6	Loopback modes	207
84.	Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4.....	209
84.1	Overview.....	209
84.2	Physical Medium Dependent (PMD) service interface	209
84.3	PCS requirements for Auto-Negotiation (AN) service interface.....	210
84.4	Delay constraints.....	211
84.5	Skew constraints	211
84.6	PMD MDIO function mapping.....	211
84.7	PMD functional specifications.....	211
84.7.1	Link block diagram	211
84.7.2	PMD transmit function	211
84.7.3	PMD receive function.....	213
84.7.4	Global PMD signal detect function	213
84.7.5	PMD lane-by-lane signal detect function	213
84.7.6	Global PMD transmit disable function	213
84.7.7	PMD lane-by-lane transmit disable function	214
84.7.8	Loopback mode.....	214
84.7.9	PMD_fault function	214
84.7.10	PMD transmit fault function	214
84.7.11	PMD receive fault function.....	214
84.7.12	PMD control function	215
84.8	40GBASE-KR4 electrical characteristics	215
84.8.1	Transmitter characteristics.....	215
84.8.2	Receiver characteristics	215
84.9	Interconnect characteristics.....	215
84.10	Environmental specifications.....	215
84.10.1	General safety	215
84.10.2	Network safety	216
84.10.3	Installation and maintenance guidelines	216
84.10.4	Electromagnetic compatibility	216
84.10.5	Temperature and humidity.....	216
84.11	Protocol implementation conformance statement (PICS) proforma for Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4	217
84.11.1	Introduction.....	217
84.11.2	Identification.....	217
84.11.3	Major capabilities/options.....	218
84.11.4	PICS proforma tables for Clause 84, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-KR4.....	219
85.	Physical Medium Dependent sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10.....	223
85.1	Overview.....	223
85.2	Physical Medium Dependent (PMD) service interface	224
85.3	PCS requirements for Auto-Negotiation (AN) service interface.....	225
85.4	Delay constraints.....	225
85.5	Skew constraints	225
85.6	PMD MDIO function mapping.....	226
85.7	PMD functional specifications.....	228
85.7.1	Link block diagram	228
85.7.2	PMD Transmit function	229
85.7.3	PMD Receive function.....	229
85.7.4	Global PMD signal detect function	229

85.7.5	PMD lane-by-lane signal detect function	230
85.7.6	Global PMD transmit disable function	230
85.7.7	PMD lane-by-lane transmit disable function	230
85.7.8	Loopback mode.....	230
85.7.9	PMD_fault function	231
85.7.10	PMD transmit fault function	231
85.7.11	PMD receive fault function.....	231
85.7.12	PMD control function	231
85.8	MDI Electrical specifications for 40GBASE-CR4 and 100GBASE-CR10	231
85.8.1	Signal levels	231
85.8.2	Signal paths.....	232
85.8.3	Transmitter characteristics.....	232
85.8.4	Receiver characteristics at TP3 summary	241
85.9	Channel characteristics	245
85.10	Cable assembly characteristics	245
85.10.1	Characteristic impedance and reference impedance.....	245
85.10.2	Cable assembly insertion loss	245
85.10.3	Cable assembly insertion loss deviation (ILD).....	247
85.10.4	Cable assembly return loss	248
85.10.5	Cable assembly multiple disturber near-end crosstalk (MDNEXT) loss	249
85.10.6	Cable assembly multiple disturber far-end crosstalk (MDFEXT) loss.....	249
85.10.7	Cable assembly integrated crosstalk noise (ICN)	250
85.10.8	Cable assembly test fixture	252
85.10.9	Mated test fixtures	252
85.10.10	Shielding	256
85.10.11	Crossover function.....	256
85.11	MDI specification	257
85.11.1	40GBASE-CR4 MDI connectors.....	257
85.11.2	100GBASE-CR10 MDI connectors.....	260
85.11.3	Electronic keying	262
85.12	Environmental specifications.....	262
85.13	Protocol implementation conformance statement (PICS) proforma for Clause 85, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10.....	263
85.13.1	Introduction.....	263
85.13.2	Identification.....	263
85.13.3	PICS proforma tables for 40GBASE-CR4 and 100GBASE-CR10 PMDs and baseband medium	264
85.13.4	Major capabilities/options.....	264
86.	Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-SR4 and 100GBASE-SR10	271
86.1	Overview.....	271
86.2	Physical Medium Dependent (PMD) service interface	273
86.3	Delay and Skew	274
86.3.1	Delay constraints.....	274
86.3.2	Skew and Skew Variation constraints	274
86.4	PMD MDIO function mapping.....	274
86.5	PMD functional specifications.....	275
86.5.1	PMD block diagram.....	275
86.5.2	PMD transmit function	276
86.5.3	PMD receive function.....	276
86.5.4	PMD global signal detect function	276

86.5.5	PMD lane-by-lane signal detect function	277
86.5.6	PMD reset function	277
86.5.7	PMD global transmit disable function (optional)	277
86.5.8	PMD lane-by-lane transmit disable function (optional)	278
86.5.9	PMD fault function (optional)	278
86.5.10	PMD transmit fault function (optional)	278
86.5.11	PMD receive fault function (optional)	278
86.6	Lane assignments	278
86.7	PMD to MDI specifications for 40GBASE-SR4 or 100GBASE-SR10	278
86.7.1	Transmitter optical specifications	279
86.7.2	Characteristics of signal within, and at the receiving end of, a compliant optical channel	280
86.7.3	40GBASE-SR4 or 100GBASE-SR10 receiver optical specifications	281
86.7.4	40GBASE-SR4 or 100GBASE-SR10 illustrative link power budget	282
86.8	Definitions of optical and dual-use parameters and measurement methods	282
86.8.1	Test points and compliance boards	282
86.8.2	Test patterns and related subclauses	284
86.8.3	Parameters applicable to both electrical and optical signals	285
86.8.4	Optical parameter definitions	287
86.9	Safety, installation, environment, and labeling	290
86.9.1	General safety	290
86.9.2	Laser safety	290
86.9.3	Installation	290
86.9.4	Environment	290
86.9.5	PMD labeling	290
86.10	Optical channel	290
86.10.1	Fiber optic cabling model	290
86.10.2	Characteristics of the fiber optic cabling (channel)	291
86.10.3	Medium Dependent Interface (MDI)	292
86.11	Protocol implementation conformance statement (PICS) proforma for Clause 86, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-SR4 and 100GBASE-SR10	295
86.11.1	Introduction	295
86.11.2	Identification	295
86.11.3	Major capabilities/options	296
86.11.4	PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, types 40GBASE-SR4 and 100GBASE-SR10	297
87.	Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4	301
87.1	Overview	301
87.2	Physical Medium Dependent (PMD) service interface	301
87.3	Delay and Skew	303
87.3.1	Delay constraints	303
87.3.2	Skew constraints	303
87.4	PMD MDIO function mapping	303
87.5	PMD functional specifications	303
87.5.1	PMD block diagram	303
87.5.2	PMD transmit function	304
87.5.3	PMD receive function	305
87.5.4	PMD global signal detect function	305
87.5.5	PMD lane-by-lane signal detect function	306
87.5.6	PMD reset function	306
87.5.7	PMD global transmit disable function (optional)	306
87.5.8	PMD lane-by-lane transmit disable function	306

87.5.9	PMD fault function (optional)	307
87.5.10	PMD transmit fault function (optional)	307
87.5.11	PMD receive fault function (optional)	307
87.6	Wavelength-division-multiplexed lane assignments	307
87.7	PMD to MDI optical specifications for 40GBASE–LR4	307
87.7.1	40GBASE–LR4 transmitter optical specifications	308
87.7.2	40GBASE–LR4 receive optical specifications	309
87.7.3	40GBASE–LR4 illustrative link power budget	309
87.8	Definition of optical parameters and measurement methods	310
87.8.1	Test patterns for optical parameters	310
87.8.2	Skew and Skew Variation	310
87.8.3	Wavelength	311
87.8.4	Average optical power	311
87.8.5	Optical Modulation Amplitude (OMA)	311
87.8.6	Transmitter and dispersion penalty	312
87.8.7	Extinction ratio	313
87.8.8	Relative Intensity Noise (RIN20OMA)	313
87.8.9	Transmitter optical waveform (transmit eye)	313
87.8.10	Receiver sensitivity	314
87.8.11	Stressed receiver sensitivity	314
87.8.12	Receiver 3 dB electrical upper cutoff frequency	319
87.9	Safety, installation, environment, and labeling	319
87.9.1	General safety	319
87.9.2	Laser safety	319
87.9.3	Installation	320
87.9.4	Environment	320
87.9.5	PMD labeling requirements	320
87.10	Fiber optic cabling model	321
87.11	Characteristics of the fiber optic cabling (channel)	321
87.11.1	Optical fiber cable	322
87.11.2	Optical fiber connection	322
87.11.3	Medium Dependent Interface (MDI) requirements	322
87.12	Protocol implementation conformance statement (PICS) proforma for Clause 87, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4	323
87.12.1	Introduction	323
87.12.2	Identification	323
87.12.3	Major capabilities/options	324
87.12.4	PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4	325
88.	Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE–LR4 and 100GBASE–ER4	329
88.1	Overview	329
88.2	Physical Medium Dependent (PMD) service interface	329
88.3	Delay and Skew	331
88.3.1	Delay constraints	331
88.3.2	Skew constraints	331
88.4	PMD MDIO function mapping	331
88.5	PMD functional specifications	331
88.5.1	PMD block diagram	331
88.5.2	PMD transmit function	333
88.5.3	PMD receive function	333
88.5.4	PMD global signal detect function	333
88.5.5	PMD lane-by-lane signal detect function	334

88.5.6	PMD reset function	334
88.5.7	PMD global transmit disable function (optional)	334
88.5.8	PMD lane-by-lane transmit disable function	334
88.5.9	PMD fault function (optional)	335
88.5.10	PMD transmit fault function (optional)	335
88.5.11	PMD receive fault function (optional)	335
88.6	Wavelength-division-multiplexed lane assignments	335
88.7	PMD to MDI optical specifications for 100GBASE-LR4 and 100GBASE-ER4	335
88.7.1	100GBASE-LR4 and 100GBASE-ER4 transmitter optical specifications	336
88.7.2	100GBASE-LR4 and 100GBASE-ER4 receive optical specifications	338
88.7.3	100GBASE-LR4 and 100GBASE-ER4 illustrative link power budgets	339
88.8	Definition of optical parameters and measurement methods	339
88.8.1	Test patterns for optical parameters	339
88.8.2	Wavelength	339
88.8.3	Average optical power	340
88.8.4	Optical Modulation Amplitude (OMA)	340
88.8.5	Transmitter and dispersion penalty (TDP)	341
88.8.6	Extinction ratio	342
88.8.7	Relative Intensity Noise (RIN _{20OMA})	342
88.8.8	Transmitter optical waveform (transmit eye)	342
88.8.9	Receiver sensitivity	343
88.8.10	Stressed receiver sensitivity	343
88.8.11	Receiver 3 dB electrical upper cutoff frequency	343
88.9	Safety, installation, environment, and labeling	343
88.9.1	General safety	343
88.9.2	Laser safety	343
88.9.3	Installation	344
88.9.4	Environment	344
88.9.5	Electromagnetic emission	344
88.9.6	Temperature, humidity, and handling	344
88.9.7	PMD labeling requirements	344
88.10	Fiber optic cabling model	345
88.11	Characteristics of the fiber optic cabling (channel)	345
88.11.1	Optical fiber cable	346
88.11.2	Optical fiber connection	346
88.11.3	Medium Dependent Interface (MDI) requirements	346
88.12	Protocol implementation conformance statement (PICS) proforma for Clause 88, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4	347
88.12.1	Introduction	347
88.12.2	Identification	347
88.12.3	Major capabilities/options	348
88.12.4	PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, types 100GBASE-LR4 and 100GBASE-ER4	349
Annex A (informative) Bibliography		353
Annex 4A (normative) Simplified full duplex media access control		355
Annex 31B (normative) MAC Control PAUSE operation		357
31B.4	Protocol implementation conformance statement (PICS) proforma for MAC Control PAUSE operation	358
Annex 69A (normative) Interference tolerance testing		361

Annex 69B (informative) Interconnect characteristics	363
Annex 83A (normative) 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)	367
83A.1 Overview.....	367
83A.2 XLAUI/CAUI link block diagram	368
83A.3 XLAUI/CAUI electrical characteristics.....	370
83A.4 Interconnect characteristics.....	379
83A.5 Electrical parameter measurement methods	381
83A.6 Environmental specifications.....	382
83A.7 Protocol implementation conformance statement (PICS) proforma for Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)	383
Annex 83B (normative) Chip-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)	387
83B.1 Overview.....	387
83B.2 Compliance point specifications for chip-module XLAUI/CAUI.....	389
83B.3 Environmental specifications.....	395
83B.4 Protocol implementation conformance statement (PICS) proforma for Annex 83B, Chip- module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)	397
Annex 83C (informative) PMA sublayer partitioning examples.....	401
83C.1 Partitioning examples with FEC	401
83C.2 Partitioning examples without FEC.....	402
Annex 85A (informative) 40GBASE-CR4 and 100GBASE-CR10 TP0 and TP5 test point parameters	405
85A.1 Overview.....	405
85A.2 Transmitter characteristics at TP0	405
85A.3 Receiver characteristics at TP5	406
85A.4 Transmitter and receiver differential printed circuit board trace loss.....	406
85A.5 Channel insertion loss	407
85A.6 Channel return loss	408
85A.7 Channel insertion loss deviation (ILD).....	408
85A.8 Channel integrated crosstalk noise (ICN).....	409
Annex 86A (normative) Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPPI) and 100GBASE-SR10 (CPPI).....	411
86A.1 Overview.....	411
86A.2 Block diagram and test points.....	411
86A.3 Lane assignments	411
86A.4 Electrical specifications for nPPI.....	412
86A.5 Definitions of electrical parameters and measurement methods	415
86A.6 Recommended electrical channel	428
86A.7 Safety, installation, environment, and labeling.....	429
86A.8 Protocol implementation conformance statement (PICS) proforma for Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPPI) and 100GBASE-SR10 (CPPI)	430

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

Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications

Amendment 4: Media Access Control Parameters, Physical Layers, and Management Parameters for 40 Gb/s and 100 Gb/s Operation

IMPORTANT NOTICE: This standard is not intended to ensure safety, security, health, or environmental protection. Implementers of the standard are responsible for determining appropriate safety, security, environmental, and health practices or regulatory requirements.

This IEEE document is made available for use subject to important notices and legal disclaimers. These notices and disclaimers appear in all publications containing this document and may be found under the heading “Important Notice” or “Important Notices and Disclaimers Concerning IEEE Documents.” They can also be obtained on request from IEEE or viewed at <http://standards.ieee.org/IPR/disclaimers.html>.

[This amendment is based on IEEE Std 802.3-2008.]

NOTE—The editing instructions contained in this amendment define how to merge the material contained therein into the existing base standard and its amendments to form the comprehensive standard.

The editing instructions are shown in ***bold italic***. Four editing instructions are used: change, delete, insert, and replace. ***Change*** is used to make corrections in existing text or tables. The editing instruction specifies the location of the change and describes what is being changed by using ~~strikethrough~~ (to remove old material) and underline (to add new material). ***Delete*** removes existing material. ***Insert*** adds new material without disturbing the existing material. Insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. ***Replace*** is used to make changes in figures or equations by removing the existing figure or equation and replacing it with a new one. Editing instructions, change markings, and this NOTE will not be carried over into future editions because the changes will be incorporated into the base standard.¹

¹NOTES in text, tables, and figures are given for information only and do not contain requirements needed to implement the standard.

1. Introduction

1.1.3.2 Compatibility interfaces

Insert the following new compatibility interfaces to the end of the list as follows:

- i) *40 Gigabit Media Independent Interface (XLGMII)*. The XLGMII is designed to connect a 40 Gb/s capable MAC to a 40 Gb/s PHY. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 40 Gb/s speeds. The XLGMII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the XLGMII. The XLGMII is optional.
- j) *40 Gigabit Attachment Unit Interface (XLAUI)*. The XLAUI is a physical instantiation of the PMA service interface to extend the connection between 40 Gb/s capable PMAs. While conformance with implementation of this interface is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 40 Gb/s speeds. The XLAUI is intended for use as a chip-to-chip or a chip-to-module interface. No mechanical connector is specified for use with the XLAUI. The XLAUI is optional.
- k) *40 Gigabit Parallel Physical Interface (XLPPPI)*. The XLPPPI is provided as a physical instantiation of the PMD service interface for 40GBASE-SR4 and 40GBASE-LR4 PMDs. The XLPPPI has four lanes. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in connecting the 40GBASE-SR4 or 40GBASE-LR4 PMDs. The XLPPPI is intended for use as a chip-to-module interface. No mechanical connector is specified for use with the XLPPPI. The XLPPPI is optional.
- l) *100 Gigabit Media Independent Interface (CGMII)*. The CGMII is designed to connect a 100 Gb/s capable MAC to a 100 Gb/s PHY. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in intermixing PHYs and DTEs at 100 Gb/s speeds. The CGMII is a logical interconnection intended for use as an intra-chip interface. No mechanical connector is specified for use with the CGMII. The CGMII is optional.
- m) *100 Gigabit Attachment Unit Interface (CAUI)*. The CAUI is a physical instantiation of the PMA service interface to extend the connection between 100 Gb/s capable PMAs. While conformance with implementation of this interface is not necessary to ensure communication, it is recommended, since it allows maximum flexibility in intermixing PHYs and DTEs at 100 Gb/s speeds. The CAUI is intended for use as a chip-to-chip or a chip-to-module interface. No mechanical connector is specified for use with the CAUI. The CAUI is optional.
- n) *100 Gigabit Parallel Physical Interface (CPPI)*. The CPPI is provided as a physical instantiation of the PMD service interface for 100GBASE-SR10 PMDs. The CPPI has ten lanes. While conformance with implementation of this interface is not necessary to ensure communication, it allows flexibility in connecting the 100GBASE-SR10 PMDs. The CPPI is intended for use as a chip-to-module interface. No mechanical connector is specified for use with the CPPI. The CPPI is optional.

1.2.3 Physical Layer and media notation

Change the example in last paragraph of 1.2.3 as follows:

The data rate, if only a number, is in Mb/s, and if suffixed by a “G”, is in Gb/s. The modulation type (e.g., BASE) indicates how encoded data is transmitted on the medium. The additional distinction may identify characteristics of transmission or medium and, in some cases, the type of PCS encoding used (e.g., examples of additional distinctions are “T” for twisted pair, “S”-“B” for short wavelength bidirectional optics, and “X” for a block PCS coding used for that speed of operation). Expansions for defined Physical Layer types are included in 1.4.

1.3 Normative references

Insert the following references in alphanumerical order as follows:

IEC 60793-1-42:2007, Optical fibres—Part 1-42: Measurement methods and test procedures—Chromatic dispersion.²

IEC 60793-2-50:2008, Optical fibres—Part 2-50: Product specifications—Sectional specification for class B single-mode fibres.

IEC 61280-1-1:1998, Fibre optic communication subsystem basic test procedures—Part 1-1: Test procedures for general communication subsystems—Transmitter output optical power measurement for single-mode optical fibre cable.

IEC 61280-1-3:1998, Fibre optic communication subsystem basic test procedures—Part 1-3: Test procedures for general communication subsystems—Central wavelength and spectral width measurement.

IEC 61280-1-4:2009, Fibre optic communication subsystem test procedures—Part 1-4: General communication subsystems—Light source encircled flux measurement method.

IEC 61280-4-1:2009, Fibre-optic communication subsystem test procedures—Part 4-1: Installed cable plant—Multimode attenuation measurement.

IEC 61754-7: Fibre optic connector interfaces—Part 7: Type MPO connector family.

ITU-T Recommendation G.694.1—Spectral grids for WDM applications: DWDM frequency grid.³

ITU-T Recommendation G.694.2—Spectral grids for WDM applications: CWDM wavelength grid.

ITU-T Recommendation G.695, 2006—Optical interfaces for coarse wavelength division multiplexing applications.

ITU-T Recommendation G.959.1, 2008—Optical transport network physical layer interfaces.

SFF-8436, Rev 3.4, Nov. 12, 2009—Specification for QSFP+ Copper And Optical Modules.⁴

SFF-8642, Rev 2.4, Nov. 16, 2009—Specification for Mini Multilane Series: Shielded Integrated Connector.

TIA-455-127-A-2006, FOTP-127-A—Basic Spectral Characterization of Laser Diodes.⁵

TIA-492AAAD, Detail Specification for 850-nm Laser-Optimized, 50- μm core diameter/125- μm cladding diameter class Ia graded-index multimode optical fibers suitable for manufacturing OM4 cabled optical fiber.

²IEC publications are available from IEC Sales Department, Case Postale 131, 3 rue de Varembe, CH-1211, Genève 20, Switzerland/Suisse (<http://www.iec.ch/>). IEC publications are also available in the United States from the American National Standards Institute.

³ITU-T publications are available from the International Telecommunications Union, Place des Nations, CH-1211, Geneva 20, Switzerland (<http://www.itu.int/>).

⁴SFF specifications are available at <ftp://ftp.seagate.com/sff>.

⁵TIA publications are available from Global Engineering Documents, 15 Inverness Way East, Englewood, Colorado 80112, USA or online at <http://www.tiaonline.org>.

1.4 Definitions

Insert the following new definitions into the definitions list, in alphanumerical order:

1.4.x 40 Gigabit Attachment Unit Interface (XLAUI): A physical instantiation of the PMA service interface to extend the connection between 40 Gb/s capable PMAs, used for chip-to-chip or chip-to-module interconnections. (See IEEE Std 802.3, Annex 83A and Annex 83B.)

1.4.x 40 Gigabit Media Independent Interface (XLGMII): The interface between the Reconciliation Sublayer (RS) and the Physical Coding Sublayer (PCS) for 40 Gb/s operation. (See IEEE Std 802.3, Clause 81.)

1.4.x 40 Gigabit Parallel Physical Interface (XLPPPI): The interface between the Physical Medium Attachment (PMA) sublayer and the Physical Medium Dependent (PMD) sublayer for 40GBASE-SR4 and 40GBASE-LR4 PHYs. (See IEEE Std 802.3, Annex 86A.)

1.4.x 40GBASE-R: An IEEE 802.3 family of Physical Layer devices using the physical coding sublayer defined in Clause 82 for 40 Gb/s operation. (See IEEE Std 802.3, Clause 82.)

1.4.x 40GBASE-CR4: IEEE 802.3 Physical Layer specification for 40 Gb/s using 40GBASE-R encoding over four lanes of shielded balanced copper cabling, with reach up to at least 7 m. (See IEEE Std 802.3, Clause 85.)

1.4.x 40GBASE-KR4: IEEE 802.3 Physical Layer specification for 40 Gb/s using 40GBASE-R encoding over four lanes of an electrical backplane, with reach up to at least 1 m. (See IEEE Std 802.3, Clause 84.)

1.4.x 40GBASE-LR4: IEEE 802.3 Physical Layer specification for 40 Gb/s using 40GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km. (See IEEE Std 802.3, Clause 87.)

1.4.x 40GBASE-SR4: IEEE 802.3 Physical Layer specification for 40 Gb/s using 40GBASE-R encoding over four lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 86.)

1.4.x 100 Gigabit Attachment Unit Interface (CAUI): A physical instantiation of the PMA service interface to extend the connection between 100 Gb/s capable PMAs, used for chip-to-chip or chip-to-module interconnections. (See IEEE Std 802.3, Annex 83A and Annex 83B.)

1.4.x 100 Gigabit Media Independent Interface (CGMII): The interface between the Reconciliation Sublayer (RS) and the Physical Coding Sublayer (PCS) for 100 Gb/s operation. (See IEEE Std 802.3, Clause 81.)

1.4.x 100 Gigabit Parallel Physical Interface (CPPI): The interface between the Physical Medium Attachment (PMA) sublayer and the Physical Medium Dependent (PMD) sublayer for 100GBASE-SR10 PHYs. (See IEEE Std 802.3, Annex 86A.)

1.4.x 100GBASE-R: An IEEE 802.3 family of Physical Layer devices using the physical coding sublayer defined in Clause 82 for 100 Gb/s operation. (See IEEE Std 802.3, Clause 82.)

1.4.x 100GBASE-CR10: IEEE 802.3 Physical Layer specification for 100 Gb/s using 100GBASE-R encoding over ten lanes of shielded balanced copper cabling, with reach up to at least 7 m. (See IEEE Std 802.3, Clause 85.)

1.4.x 100GBASE-ER4: IEEE 802.3 Physical Layer specification for 100 Gb/s using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 40 km. (See IEEE Std 802.3, Clause 88.)

1.4.x 100GBASE-LR4: IEEE 802.3 Physical Layer specification for 100 Gb/s using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km. (See IEEE Std 802.3, Clause 88.)

1.4.x 100GBASE-SR10: IEEE 802.3 Physical Layer specification for 100 Gb/s using 100GBASE-R encoding over ten lanes of multimode fiber, with reach up to at least 100 m. (See IEEE Std 802.3, Clause 86.)

1.4.x BASE-R: An IEEE 802.3 family of Physical Layer devices using the 64B/66B encoding defined in Clause 49 or Clause 82. (See IEEE Std 802.3, Clause 49 and Clause 82.)

1.4.x PCS lane (PCSL): In 40GBASE-R and 100GBASE-R, the PCS distributes encoded data to multiple logical lanes, these logical lanes are called PCS lanes. One or more PCS lanes can be multiplexed and carried on a physical lane together at the PMA service interface. (See IEEE Std 802.3, Clause 83.)

1.4.x nPPI: The term “nPPI” denotes either XLPPi or CPPI or both. (See IEEE Std 802.3, Annex 86A.)

Change 1.4.311 as follows:

1.4.311 RMS spectral width: A measure of the optical wavelength range as defined by ~~ANSI/EIA/TIA 455-127-A-1991~~ (FOTP-127-A) ~~[B10]~~.

1.5 Abbreviations

Insert the following new abbreviations into the definitions list, in alphabetical order as follows:

CAUI	100 Gigabit Attachment Unit Interface
CGMII	100 Gigabit Media Independent Interface
CPPI	100 Gigabit Parallel Physical Interface
DIC	deficit idle count
HCB	Host Compliance Board
LSB	least significant bit
MCB	Module Compliance Board
MSB	most significant bit
OTN	Optical Transport Network
OPU3	Optical channel Payload Unit 3
PCSL	PCS lane
XLAUI	40 Gigabit Attachment Unit Interface
XLGMII	40 Gigabit Media Independent Interface
XLPPi	40 Gigabit Parallel Physical Interface

4. Media Access Control

4.4.2 MAC parameters

Insert a new column to Table 4-2 for 40 Gb/s and 100 Gb/s MAC data rates as follows:

Table 4–2—MAC parameters

Parameters	MAC data rate			
	Up to and including 100 Mb/s	1 Gb/s	10 Gb/s	<u>40 Gb/s and 100 Gb/s</u>
slotTime	512 bit times	4096 bit times	not applicable	<u>not applicable</u>
interPacketGap ^a	96 bits	96 bits	96 bits	<u>96 bits</u>
attemptLimit	16	16	not applicable	<u>not applicable</u>
backoffLimit	10	10	not applicable	<u>not applicable</u>
jamSize	32 bits	32 bits	not applicable	<u>not applicable</u>
maxBasicFrameSize	1518 octets	1518 octets	1518 octets	<u>1518 octets</u>
maxEnvelopeFrameSize	2000 octets	2000 octets	2000 octets	<u>2000 octets</u>
minFrameSize	512 bits (64 octets)	512 bits (64 octets)	512 bits (64 octets)	<u>512 bits (64 octets)</u>
burstLimit	not applicable	65 536 bits	not applicable	<u>not applicable</u>
ipgStretchRatio	not applicable	not applicable	104 bits	<u>not applicable</u>

^aReferences to interFrameGap or interFrameSpacing in other clauses (e.g., 13, 35, and 42) shall be interpreted as interPacketGap.

Insert the following note below Table 4-2 (above the warning box) for 40 Gb/s and 100 Gb/s MAC data rates, and renumber notes as appropriate:

NOTE 7—For 40 Gb/s and 100 Gb/s operation, the received interPacketGap (the spacing between two packets, from the last bit of the FCS field of the first packet to the first bit of the Preamble of the second packet) can have a minimum value of 8 BT (bit times), as measured at the XLGMII or CGMII receive signals at the DTE due to clock tolerance and lane alignment requirements.

30. Management

Insert two new objects into Table 31-1e, before aIdleErrorCount, as follows:

Table 30-1e—Capabilities

			DTE	Repeater	MAU			
			Basic Package (mandatory)					
			Mandatory Package (mandatory)					
			Recommended Package (optional)					
			Optional Package (optional)					
			Array Package (optional)					
			Excessive Deferral Package (optional)					
			Multiple PHY Package (optional)					
			PHY Error Monitor Capability (optional)					
			Basic Control Capability (mandatory)					
			Performance Monitor Capability (optional)					
			Address Tracking Capability (optional)					
			100/1000 Mb/s Monitor Capability (optional)					
			1000 Mb/s Burst Monitor Capability (optional)					
			Basic Package (mandatory)					
			MAU Control Package (optional)					
			Media Loss Tracking Package (conditional)					
			Broadband DTE MAU Package (conditional)					
			MII Capability (conditional)					
			PHY Error Monitor Capability (optional)					
			10GBASE-T Operating Margin package (conditional)					
			Forward Error Correction Package (conditional)					
			Auto-Negotiation Package (mandatory)					
oMAU managed object class (30.5.1)								
aBIPErrCount	ATTRIBUTE	GET						X
aLaneMapping	ATTRIBUTE	GET						X

Insert new PHY types in alphanumeric order into “APPROPRIATE SYNTAX” of 30.3.2.1.2:

30.3.2.1.2 aPhyType

APPROPRIATE SYNTAX:

- ...
- 40GBASE-R Clause 82 40 Gb/s multi-PCS lane 64B/66B
- 100GBASE-R Clause 82 100 Gb/s multi-PCS lane 64B/66B

Insert new PHY types in alphanumeric order into “APPROPRIATE SYNTAX” of 30.3.2.1.3:

30.3.2.1.3 aPhyTypeList

APPROPRIATE SYNTAX:

- ...
- 40GBASE-R Clause 82 40 Gb/s multi-PCS lane 64B/66B
- 100GBASE-R Clause 82 100 Gb/s multi-PCS lane 64B/66B

Change 30.3.2.1.5 as follows:

30.3.2.1.5 aSymbolErrorDuringCarrier

ATTRIBUTE

APPROPRIATE SYNTAX:

Generalized nonresetable counter. This counter has a maximum increment rate of 160 000 counts per second for 100 Mb/s implementations

BEHAVIOUR DEFINED AS:

For 100 Mb/s operation it is a count of the number of times when valid carrier was present and there was at least one occurrence of an invalid data symbol (see 23.2.1.4, 24.2.2.1.6, and 32.3.4.1).

For half duplex operation at 1000 Mb/s, it is a count of the number of times the receiving media is non-idle (a carrier event) for a period of time equal to or greater than slotTime (see 4.2.4), and during which there was at least one occurrence of an event that causes the PHY to indicate “Data reception error” or “Carrier Extend Error” on the GMII (see Table 35–2).

For full duplex operation at 1000 Mb/s, it is a count of the number of times the receiving media is non-idle (a carrier event) for a period of time equal to or greater than minFrameSize, and during which there was at least one occurrence of an event that causes the PHY to indicate “Data reception error” on the GMII (see Table 35–2).

For operation at 10 Gb/s, 40 Gb/s, and 100 Gb/s, it is a count of the number of times the receiving media is non-idle (the time between the Start of Packet Delimiter and the End of Packet Delimiter as defined by 46.2.5 and 81.2.5) for a period of time equal to or greater than minFrameSize, and during which there was at least one occurrence of an event that causes the PHY to indicate “Receive Error” on the XGMII (see Table 46-4), the XLGMII or the CGMII (see Table 81–3).

At all speeds this counter shall be incremented only once per valid CarrierEvent and if a collision is present this counter shall not increment.;

Insert new PHY types into “APPROPRIATE SYNTAX” before 802.9a (as modified by IEEE Std 802.3av) and change “BEHAVIOUR DEFINED AS” of 30.5.1.1.2 as follows:

30.5.1.1.2 aMAUType

APPROPRIATE SYNTAX:

...	
<u>40GBASE-R</u>	<u>Multi-lane PCS as specified in Clause 82 over undefined PMA/PMD</u>
<u>40GBASE-KR4</u>	<u>40GBASE-R PCS/PMA over an electrical backplane PMD as specified in Clause 84</u>
<u>40GBASE-CR4</u>	<u>40GBASE-R PCS/PMA over 4 lane shielded copper balanced cable PMD as specified in Clause 85</u>
<u>40GBASE-SR4</u>	<u>40GBASE-R PCS/PMA over 4 lane multimode fiber PMD as specified in Clause 86</u>
<u>40GBASE-LR4</u>	<u>40GBASE-R PCS/PMA over 4 WDM lane single mode fiber PMD, with long reach, as specified in Clause 87</u>
<u>100GBASE-R</u>	<u>Multi-lane PCS as specified in Clause 82 over undefined PMA/PMD</u>
<u>100GBASE-CR10</u>	<u>100GBASE-R PCS/PMA over 10 lane shielded copper balanced cable PMD as specified in Clause 85</u>
<u>100GBASE-SR10</u>	<u>100GBASE-R PCS/PMA over 10 lane multimode fiber PMD as specified in Clause 86</u>
<u>100GBASE-LR4</u>	<u>100GBASE-R PCS/PMA over 4 WDM lane single mode fiber PMD, with long reach, as specified in Clause 88</u>
<u>100GBASE-ER4</u>	<u>100GBASE-R PCS/PMA over 4 WDM lane single mode fiber PMD, with extended reach, as specified in Clause 88</u>

...

BEHAVIOUR DEFINED AS:

Returns a value that identifies the internal MAU type. If an AUI is to be identified to access an external MAU, the type “AUI” is returned. A SET operation to one of the possible enumerations indicated by aMAUTypeList will force the MAU into the new operating mode. If a Clause 22 MII or Clause 35 GMII is present, then this will map to the mode force bits specified in 22.2.4.1. If a Clause 45 MDIO Interface is present, then this will map to the PCS type selection bit(s) in the 10G WIS Control 2 register specified in 45.2.2.6.6, the ~~40G~~-PCS Control 2 register specified in 45.2.3.6.1, the PMA/PMD type selection bits in the ~~40G~~-PMA/PMD Control 2 register specified in 45.2.1.6.1, the PMA/PMD control 1 register specified in 45.2.1.1 and the PCS control 1 register 45.2.3.1. If Clause 28, Clause 37, or Clause 73 Auto-Negotiation is operational, then this will change the advertised ability to the single enumeration specified in the SET operation, and cause an immediate link renegotiation. A change in the MAU type will also be reflected in aPHYType.

The enumerations 1000BASE-X, 1000BASE-XHD, 1000BASE-XFD, 10GBASE-X, 10GBASE-R, ~~and~~ 10GBASE-W, 40GBASE-R and 100GBASE-R shall only be returned if the underlying PMD type is unknown.;

Change “BEHAVIOUR DEFINED AS” of 30.5.1.1.4 for high BER as follows:

30.5.1.1.4 aMediaAvailable**BEHAVIOUR DEFINED AS:**

If the MAU is a 10M b/s link or fiber type (FOIRL, 10BASE-T, 10BASE-F), then this is equivalent to the link test fail state/low light function. For an AUI, 10BASE2, 10BASE5, or 10BROAD36 MAU, this indicates whether or not loopback is detected on the DI circuit. The value of this attribute persists between packets for MAU types AUI, 10BASE5, 10BASE2, 10BROAD36, and 10BASE-FP.

At power-up or following a reset, the value of this attribute will be “unknown” for AUI, 10BASE5, 10BASE2, 10BROAD36, and 10BASE-FP MAUs. For these MAUs loopback will be tested on each transmission during which no collision is detected. If DI is receiving *input* when DO returns to IDL after a transmission and there has been no collision during the transmission, then loopback will be detected. The value of this attribute will only change during noncollided transmissions for AUI, 10BASE2, 10BASE5, 10BROAD36, and 10BASE-FP MAUs.

For 100BASE-T2 and 100BASE-T4 PHYs the enumerations match the states within the respective link integrity state diagrams, Figure 32–16 and Figure 23–12. For 100BASE-TX, 100BASE-FX, 100BASE-LX10 and 100BASE-BX10 PHYs the enumerations match the states within the link integrity state diagram Figure 24–15. Any MAU that implements management of Clause 28 or Clause 73 Auto-Negotiation will map remote fault indication to MediaAvailable “remote fault.” Any MAU that implements management of Clause 37 Auto-Negotiation will map the received RF1 and RF2 bits as specified in Table 37–2, as follows. Offline maps to the enumeration “offline,” Link_Failure maps to the enumeration “remote fault” and Auto-Negotiation Error maps to the enumeration “auto neg error.”

The enumeration “remote fault” applies to 10BASE-FB remote fault indication, the 100BASE-X far-end fault indication and nonspecified remote faults from a system running Clause 28 Auto-Negotiation. The enumerations “remote jabber”, “remote link loss”, or “remote test” should be used instead of “remote fault” where the reason for remote fault is identified in the remote signaling protocol.

Where a Clause 22 MII or Clause 35 GMII is present, a ~~logic~~ one in the remote fault bit (22.2.4.2.11) maps to the enumeration “remote fault,” a ~~logic~~ zero in the link status bit (22.2.4.2.13) maps to the enumeration “not available.” The enumeration “not available” takes precedence over “remote fault.”

For 40 Gb/s and 100 Gb/s the enumerations map to value of the link_fault variable (see 81.3.4) within

the Link Fault Signaling state diagram (see 81.3.4.1 and Figure 46-9) as follows: the value OK maps to the enumeration “available”, the value Local Fault maps to the enumeration “not available” and the value Remote Fault maps to the enumeration “remote fault.”

For 2BASE-TL and 10PASS-TS PHYs, the enumeration “unknown” maps to the condition where the PHY is initializing, the enumeration “ready” maps to the condition where at least one PME is available and is ready for handshake, the enumeration “available” maps to the condition where, at the PCS, at least one PME is operationally linked, the enumeration “not available” maps to the condition where the PCS is not operationally linked, the enumeration “available reduced” maps to the condition where a link fault is detected at the receive direction by one or more PMEs in the aggregation group and the enumeration “PMD link fault” maps to the condition where a link fault is detected at the receive direction by all of the PMA/PMDs in the aggregation group.

For 10 Gb/s the enumerations map to value of the link_fault variable within the Link Fault Signaling state diagram (Figure 46–9) as follows: the value OK maps to the enumeration “available”, the value Local Fault maps to the enumeration “not available” and the value Remote Fault maps to the enumeration “remote fault”. The enumeration “PMD link fault”, “WIS frame loss”, “WIS signal loss”, “PCS link fault”, “excessive BER” or “DXS link fault” should be used instead of the enumeration “not available” where the reason for the Local Fault state can be identified through the use of the Clause 45 MDIO Interface. Where multiple reasons for the Local Fault state can be identified only the highest precedence error should be reported. This precedence in descending order is as follows: “PXS link fault”, “PMD link fault”, “WIS frame loss”, “WIS signal loss”, “PCS link fault”, “excessive BER”, “DXS link fault”. Where a Clause 45 MDIO interface is present a logic zero in the PMA/PMD Receive link status bit (45.2.1.2.2) maps to the enumeration “PMD link fault”, a logic one in the LOF status bit (45.2.2.10.4) maps to the enumeration “WIS frame loss”, a logic one in the LOS status bit (45.2.2.10.5) maps to the enumeration “WIS signal loss”, a logic zero in the PCS Receive link status bit (45.2.3.2.2) maps to the enumeration “PCS link fault”, a logic one in the 10/40/100GBASE-R PCS Latched high BER status bit (45.2.3.12.2) maps to the enumeration “excessive BER”, a logic zero in the DTE XS receive link status bit (45.2.5.2.2) maps to the enumeration “DXS link fault” and a logic zero in the PHY XS transmit link status bit (45.2.4.2.2) maps to the enumeration “PXS link fault”;

Insert a new subclause after 30.5.1.1.10 as follows:

30.5.1.1.10a aBIPErrorCount

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized nonresetable counters. Each counter has a maximum increment rate of 10 000 counts per second for 40 Gb/s implementations and 5 000 counts per second for 100 Gb/s implementations.

BEHAVIOUR DEFINED AS:

For 40/100GBASE-R PHYs, an array of BIP error counters. The counters will not increment for other PHY types. The indices of this array (0 to n – 1) denote the PCS lane number where n is the number of PCS lanes in use. Each element of this array contains a count of BIP errors for that PCS lane.

Increment the counter by one for each BIP error detected during alignment marker removal in the PCS for the corresponding lane.

If a Clause 45 MDIO Interface to the PCS is present, then this attribute will map to the BIP error counters (see 45.2.3.36 and 45.2.3.37).;

30.5.1.1.10b aLaneMapping

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of INTEGERS.

BEHAVIOUR DEFINED AS:

For 40/100GBASE-R PHYs, an array of PCS lane identifiers. The indices of this array (0 to $n - 1$) denote the service interface lane number where n is the number of PCS lanes in use. Each element of this array contains the PCS lane number for the PCS lane that has been detected in the corresponding service interface lane.

If a Clause 45 MDIO Interface to the PCS is present, then this attribute will map to the Lane mapping registers (see 45.2.3.38 and 45.2.3.39).;

Change 30.5.1.1.14, 30.5.1.1.15, and 30.5.1.1.16 for FEC as follows:

30.5.1.1.14 aFECmode

ATTRIBUTE

APPROPRIATE SYNTAX:

A ENUMERATION that meets the requirement of the description below

unknown	initializing, true state not yet known
disabled	FEC disabled
enabled	FEC enabled

BEHAVIOUR DEFINED AS:

A read-write value that indicates the mode of operation of the optional FEC sublayer for forward error correction (see 65.2 and Clause 74).

A GET operation returns the current mode of operation of the PHY. A SET operation changes the mode of operation of the PHY to the indicated value. When Clause 73 Auto-Negotiation is enabled a SET operation is not allowed and a GET operation maps to the variable FEC enabled in Clause 74.

If a Clause 45 MDIO Interface to the PCS is present, then this attribute will map to the FEC control register (see 45.2.7.3) for 1000BASE-PX or FEC enable bit in 10GBASE-R FEC control register (see 45.2.1.86).;

30.5.1.1.15 aFECCorrectedBlocks

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized ~~Generalized~~ nonresetable counters. ~~This~~ Each counter has a maximum increment rate of 1 200 000 counts per second for 1000 Mb/s implementations, ~~and~~ 5 000 000 counts per second for 10 Gb/s and 40 Gb/s implementations, and 2 500 000 counts per second for 100 Gb/s implementations.

BEHAVIOUR DEFINED AS:

~~For 1000BASE-PX PHYs or 10/40/100GBASE-R PHYs, a count an array of corrected FEC blocks counters.~~ The counters will not increment for other PHY types. The indices of this array (0 to $N - 1$) denote the PCS lane number where N is the number of PCS lanes in use. The number of PCS lanes in use is set to one for PHYs that do not use PCS lanes. Each element of this array contains a count of corrected FEC blocks for that PCS lane.

Increment the counter by one for each received block that is corrected by the FEC function in the PHY for the corresponding lane.

If a Clause 45 MDIO Interface to the PCS is present, then this attribute will map to the FEC corrected blocks counter(s) (see 45.2.8.5, ~~and 45.2.1.87,~~ and 45.2.1.89).;

30.5.1.1.16 aFECUncorrectableBlocks

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized Generalized nonresetable counters. ThisEach counter has a maximum increment rate of 1 200 000 counts per second for 1000 Mb/s implementations, and 5 000 000 counts per second for 10 Gb/s and 40 Gb/s implementations, and 2 500 000 counts per second for 100 Gb/s implementations.

BEHAVIOUR DEFINED AS:

For 1000BASE-PX PHYs or 10/40/100GBASE-R PHYs, ~~a count an array~~ of uncorrectable FEC blocks counters. The counters will not increment for other PHY types. The indices of this array (0 to N – 1) denote the PCS lane number where N is the number of PCS lanes in use. The number of PCS lanes in use is set to one for PHYs that do not use PCS lanes. Each element of this array contains a count of uncorrectable FEC blocks for that PCS lane.

Increment the counter by one for each FEC block that is determined to be uncorrectable by the FEC function in the PHY for the corresponding lane.

If a Clause 45 MDIO Interface to the PCS is present, then this attribute will map to the FEC uncorrectable blocks counter(s) (see 45.2.8.6, ~~and 45.2.1.88,~~ and 45.2.1.90).;

Change 30.6.1.1.5 for autoneg ability and FEC request as follows:

30.6.1.1.5 aAutoNegLocalTechnologyAbility

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE that meets the requirements of the description below:

global	Reserved for future use
other	See 30.2.5
unknown	Initializing, true state or type not yet known
10BASE-T	10BASE-T half duplex as defined in Clause 14
10BASE-TFD	Full duplex 10BASE-T as defined in Clause 14 and Clause 31
100BASE-T4	100BASE-T4 half duplex as defined in Clause 23
100BASE-TX	100BASE-TX half duplex as defined in Clause 25
100BASE-TXFD	Full duplex 100BASE-TX as defined in Clause 25 and Clause 31
FDX PAUSE	PAUSE operation for full duplex links as defined in Annex 31B
FDX APAUSE	Asymmetric PAUSE operation for full duplex links as defined in Clause 37, Annex 28B, and Annex 31B
FDX SPAUSE	Symmetric PAUSE operation for full duplex links as defined in Clause 37, Annex 28B, and Annex 31B
FDX BPAUSE	Asymmetric and Symmetric PAUSE operation for full duplex links as defined in Clause 37, Annex 28B, and Annex 31B
100BASE-T2	100BASE-T2 half duplex as defined in Clause 32

100BASE-T2FD	Full duplex 100BASE-T2 as defined in Clause 31 and Clause 32
1000BASE-X	1000BASE-X half duplex as specified in Clause 36
1000BASE-XFD	Full duplex 1000BASE-X as specified in Clause 31 and Clause 36
1000BASE-T	1000BASE-T half duplex UTP PHY as specified in Clause 40
1000BASE-TFD	Full duplex 1000BASE-T UTP PHY as specified in Clause 31 and as specified in Clause 40
Rem Fault1	Remote fault bit 1 (RF1) as specified in Clause 37
Rem Fault2	Remote fault bit 2 (RF2) as specified in Clause 37
10GBASE-T	10GBASE-T PHY as specified in Clause 55
1000BASE-KXFD	Full duplex 1000BASE-KX as specified in Clause 70
10GBASE-KX4FD	Full duplex 10GBASE-KX4 as specified in Clause 71
10GBASE-KRFD	Full duplex 10GBASE-KR as specified in Clause 72
<u>40GBASE-KR4</u>	<u>40GBASE-KR4 as specified in Clause 84</u>
<u>40GBASE-CR4</u>	<u>40GBASE-CR4 as specified in Clause 85</u>
<u>100GBASE-CR10</u>	<u>100GBASE-CR10 as specified in Clause 85</u>
Rem Fault	Remote fault bit (RF) as specified in Clause 73
FEC Capable	FEC ability as specified in <u>Clause 73 (see 73.6.5) and Clause 74</u>
<u>FEC Requested</u>	<u>FEC requested as specified in Clause 73 (see 73.6.5) and Clause 74</u>
isoethernet	IEEE Std 802.9 ISLAN-16T

BEHAVIOUR DEFINED AS:

This indicates the technology ability of the local device, as defined in Clause 28, ~~and~~ Clause 37, and Clause 73.

45. Management Data Input/Output (MDIO) Interface

Change the indicated rows of Table 45–1 and Table 45–2 for new MMDs as follows:

45.2 MDIO Interface Registers

Table 45–1—MDIO Manageable Device addresses

Device address	MMD name
...	...
<u>8</u>	<u>Separated PMA (1)</u>
<u>9</u>	<u>Separated PMA (2)</u>
<u>10</u>	<u>Separated PMA (3)</u>
<u>11</u>	<u>Separated PMA (4)</u>
<u>12</u> through 28	Reserved
...	...

Table 45–2—Devices in package registers bit definitions

Bit(s) ^a	Name	Description	R/W ^b
...
m.5.15: <u>12</u>	Reserved	Ignore on read	RO
<u>m.5.11</u>	<u>Separated PMA (4) present</u>	<u>1 = Separated PMA^c (4) present in package</u> <u>0 = Separated PMA (4) not present in package</u>	<u>RO</u>
<u>m.5.10</u>	<u>Separated PMA (3) present</u>	<u>1 = Separated PMA (3) present in package</u> <u>0 = Separated PMA (3) not present in package</u>	<u>RO</u>
<u>m.5.9</u>	<u>Separated PMA (2) present</u>	<u>1 = Separated PMA (2) present in package</u> <u>0 = Separated PMA (2) not present in package</u>	<u>RO</u>
<u>m.5.8</u>	<u>Separated PMA (1) present</u>	<u>1 = Separated PMA (1) present in package</u> <u>0 = Separated PMA (1) not present in package</u>	<u>RO</u>
...

^am = Address of MMD accessed (see Table 45–1).

^bRO = Read only.

^cSeparated PMAs are defined in 45.2.1.

Change 45.2.1 for the general description of MMD 1, 8, 9, and 10 as follows:

45.2.1 PMA/PMD registers

For devices operating at 40 Gb/s or higher speeds, the PMA may be instantiated as multiple sublayers (see 83.1.4 for how MMD addresses are allocated to multiple PMA sublayers). A PMA sublayer that is packaged with the PMD is addressed as MMD 1. More addressable instances of PMA sublayers, each one separated

from lower addressable instances, may be implemented and addressed as MMD 8, 9, 10, and 11 where MMD 8 is the closest to the PMD and MMD 11 is the furthest from the PMD. The addresses and functions of all registers in MMD 8, 9, 10 and 11 are defined identically to MMD 1, except registers m.5 and m.6 as defined in Table 45–2.

The assignment of registers in the PMA/PMD is shown in Table 45–3.

Change the indicated rows of Table 45–3 (as modified by IEEE Std 802.3av) for 40 Gb/s and 100 Gb/s registers as follows:

Table 45–3—PMA/PMD registers

Register address	Register name	Clause
1.0	PMA/PMD control 1	<u>45.2.1.1</u>
1.1	PMA/PMD status 1	<u>45.2.1.2</u>
1.2, 1.3	PMA/PMD device identifier	<u>45.2.1.3</u>
1.4	PMA/PMD speed ability	<u>45.2.1.4</u>
1.5, 1.6	PMA/PMD devices in package	<u>45.2.1.5</u>
1.7	PMA/PMD control 2	<u>45.2.1.6</u>
1.8	10G -PMA/PMD status 2	<u>45.2.1.7</u>
1.9	10G -PMA/PMD transmit disable	<u>45.2.1.8</u>
1.10	10G -PMD receive signal detect	<u>45.2.1.9</u>
1.11	10G -PMA/PMD extended ability register	<u>45.2.1.10</u>
1.12	<u>10G-EPON PMA/PMD P2MP</u> ability register	<u>45.2.1.11</u>
<u>1.13</u>	<u>40G/100G PMA/PMD extended ability register</u>	<u>45.2.1.11a</u>
1.13	Reserved	
1.14, 1.15	PMA/PMD package identifier	<u>45.2.1.12</u>
...	...	
1.150	10GBASE-KR <u>BASE-R</u> PMD control	<u>45.2.1.77</u>
1.151	10GBASE-KR <u>BASE-R</u> PMD status	<u>45.2.1.78</u>
1.152	10GBASE-KR <u>BASE-R</u> LP coefficient update, <u>lane 0</u>	<u>45.2.1.79</u>
1.153	10GBASE-KR <u>BASE-R</u> LP status report, <u>lane 0</u>	<u>45.2.1.80</u>
1.154	10GBASE-KR <u>BASE-R</u> LD coefficient update, <u>lane 0</u>	<u>45.2.1.81</u>
1.155	10GBASE-KR <u>BASE-R</u> LD status report, <u>lane 0</u>	<u>45.2.1.82</u>
<u>1.156</u>	<u>BASE-R</u> PMD status 2	<u>45.2.1.82a</u>
<u>1.157</u>	<u>BASE-R</u> PMD status 3	<u>45.2.1.82b</u>
1.156 8 through 1.159	Reserved	
1.160	1000BASE-KX control	<u>45.2.1.83</u>
1.161	1000BASE-KX status	<u>45.2.1.84</u>
1.162 through 1.169	Reserved	
1.170	10GBASE-R FEC ability	<u>45.2.1.85</u>
1.171	10GBASE-R FEC control	<u>45.2.1.86</u>
1.172 through 1.173	10GBASE-R FEC corrected blocks counter	<u>45.2.1.87</u>

Table 45–3—PMA/PMD registers (continued)

Register address	Register name	Clause
1.174 through 1.175	10GBASE-R FEC uncorrected blocks counter	45.2.1.88
1.176 through 1.299	Reserved	
1.300 through 1.339	BASE-R FEC corrected blocks counter, lanes 0 through 19	45.2.1.89
1.340 through 1.699	Reserved	
1.700 through 1.739	BASE-R FEC uncorrected blocks counter, lanes 0 through 19	45.2.1.90
1.740 through 1.1099	Reserved	
1.1100	BASE-R LP coefficient update, lane 0 (copy)	45.2.1.79
1.1101 through 1.1109	BASE-R LP coefficient update, lanes 1 through 9	45.2.1.91
1.1110 through 1.1199	Reserved	
1.1200	BASE-R LP status report, lane 0 (copy)	45.2.1.80
1.1201 through 1.1209	BASE-R LP status report, lanes 1 through 9	45.2.1.92
1.1210 through 1.1299	Reserved	
1.1300	BASE-R LD coefficient update, lane 0 (copy)	45.2.1.81
1.1301 through 1.1309	BASE-R LD coefficient update, lanes 1 through 9	45.2.1.93
1.1310 through 1.1399	Reserved	
1.1400	BASE-R LD status report, lane 0 (copy)	45.2.1.82
1.1401 through 1.1409	BASE-R LD status report, lanes 1 through 9	45.2.1.94
1.1410 through 1.1499	Reserved	
1.1500	Test pattern ability	45.2.1.95
1.1501	PRBS pattern testing control	45.2.1.96
1.1502 through 1.1509	Reserved	
1.1510	Square wave testing control	45.2.1.97
1.1511 through 1.1599	Reserved	
1.1600 through 1.1609	PRBS Tx error counters, lane 0 through lane 9	45.2.1.98
1.1610 through 1.1699	Reserved	
1.1700 through 1.1709	PRBS Rx error counters, lane 0 through lane 9	45.2.1.99
1.176 1710 through 1.32 767	Reserved	

45.2.1.1 PMA/PMD control 1 register (Register 1.0)

Change Table 45–4 for 40 Gb/s and 100 Gb/s speed selection as follows:

Change 45.2.1.1.3 for 40 Gb/s and 100 Gb/s speed selection:

45.2.1.1.3 Speed selection (1.0.13, 1.0.6, 1.0.5:2)

For devices operating at 10 Mb/s, 100 Mb/s, or 1000 Mb/s the speed of the PMA/PMD may be selected using bits 13 and 6. The speed abilities of the PMA/PMD are advertised in the PMA/PMD speed ability register. These two bits use the same definition as the speed selection bits defined in Clause 22.

Table 45–4—PMA/PMD control 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.0.15	Reset	1 = PMA/PMD reset 0 = Normal operation	R/W SC
1.0.14	Reserved	Value always 0, writes ignored	R/W
1.0.13	Speed selection (LSB)	<u>1.0.6</u> <u>1.0.13</u> 1 1 = bits 5:2 select speed 1 0 = 1000 Mb/s 0 1 = 100 Mb/s 0 0 = 10 Mb/s	R/W
1.0.12	Reserved	Value always 0, writes ignored	R/W
1.0.11	Low power	1 = Low-power mode 0 = Normal operation	R/W
1.0.10:7	Reserved	Value always 0, writes ignored	R/W
1.0.6	Speed selection (MSB)	<u>1.0.6</u> <u>1.0.13</u> 1 1 = bits 5:2 select speed 1 0 = 1000 Mb/s 0 1 = 100 Mb/s 0 0 = 10 Mb/s	R/W
1.0.5:2	Speed selection	<u>5</u> <u>4</u> <u>3</u> <u>2</u> 1 x x x = Reserved x 1 x x = Reserved x x 1 x = Reserved <u>0 0 1 1 = 100 Gb/s</u> <u>0 0 1 0 = 40 Gb/s</u> 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s	R/W
1.0.4	Reserved	Value always 0, writes ignored	R/W
<u>1.0.1</u>	<u>PMA remote loopback</u>	<u>1 = Enable PMA remote loopback mode</u> <u>0 = Disable PMA remote loopback mode</u>	<u>R/W</u>
1.0.0	PMA <u>local</u> loopback	1 = Enable PMA <u>local</u> loopback mode 0 = Disable PMA <u>local</u> loopback mode	R/W

^aR/W = Read/Write, SC = Self-clearing

For devices not operating at 10 Mb/s, 100 Mb/s, or 1000 Mb/s, the speed of the PMA/PMD may be selected using bits 5 through 2. When bits 5 through 2 are set to 0000 the use of a 10G PMA/PMD is selected. More specific selection is performed using the PMA/PMD control 2 register (Register 1.7) (see 45.2.1.6). The speed abilities of the PMA/PMD are advertised in the PMA/PMD speed ability register. A PMA/PMD may ignore writes to the PMA/PMD speed selection bits that select speeds it has not advertised in the PMA/PMD speed ability register. It is the responsibility of the STA entity to ensure that mutually acceptable speeds are applied consistently across all the MMDs on a particular PHY.

The PMA/PMD speed selection defaults to a supported ability.

When set to 0001, bits 5:2 select the use of the 10PASS-TS or 2BASE-TL PMA/PMD. More specific mode selection is performed using the 10P/2B PMA control register (45.2.1.12).

When bits 5 through 2 are set to 0010 the use of a 40G PMA/PMD is selected; when set to 0011 the use of a 100G PMA/PMD is selected. More specific selection is performed using the PMA/PMD control 2 register (Register 1.7) (see 45.2.1.6.1).

Insert the following new subclause before 45.2.1.1.4:

45.2.1.1.3a PMA remote loopback (1.0.1)

The PMA shall be placed in a remote loopback mode of operation when bit 1.0.1 is set to a one. When bit 1.0.1 is set to a one, the PMA shall accept data on the receive path and return it on the transmit path.

The remote loopback function is optional for all port types, except 2BASE-TL and 10PASS-TS, which do not support loopback. A device's ability to perform the remote loopback function is advertised in the remote loopback ability bit of the related speed-dependent status register. A PMA that is unable to perform the remote loopback function shall ignore writes to this bit and shall return a value of zero when read. For 40/100 Gb/s operation, the remote loopback functionality is detailed in 83.5.9. For 40/100 Gb/s operation, the remote loopback ability bit is specified in the 40G/100G PMA/PMD extended ability register.

The default value of bit 1.0.1 is zero.

Change 45.2.1.1.4 (as modified by IEEE Std 802.3av) to distinguish from remote loopback as follows:

45.2.1.1.4 PMA local loopback (1.0.0)

The PMA shall be placed in a local loopback mode of operation when bit 1.0.0 is set to a one. When bit 1.0.0 is set to a one, the PMA shall accept data on the transmit path and return it on the receive path.

The local loopback function is mandatory for the 1000BASE-KX, 10GBASE-KR, ~~and~~ 10GBASE-X, 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10 port type and optional for all other port types, except 2BASE-TL, 10PASS-TS and 10/1GBASE-PRX, which do not support loopback. A device's ability to perform the local loopback function is advertised in the local loopback ability bit of the related speed-dependent status register. A PMA that is unable to perform the local loopback function shall ignore writes to this bit and shall return a value of zero when read. For 10 Gb/s operation, the local loopback functionality is detailed in 48.3.3 and 51.8, ~~and~~. For 40/100 Gb/s operation, the local loopback functionality is detailed in 83.5.8. For 10/40/100 Gb/s operation, the local loopback ability bit is specified in the ~~40G~~ PMA/PMD status 2 register.

The default value of bit 1.0.0 is zero.

NOTE—The signal path through the PMA that is exercised in the loopback mode of operation is implementation specific, but it is recommended that the signal path encompass as much of the PMA circuitry as is practical. The intention of providing this loopback mode of operation is to permit a diagnostic or self-test function to perform the transmission and reception of a PDU, thus testing the transmit and receive data paths. Other loopback signal paths may be enabled using loopback controls within other MMDs.

Change the indicated rows of Table 45–6 (as modified by IEEE Std 802.3av) for 40 Gb/s and 100 Gb/s speed ability:

Table 45–6—PMA/PMD speed ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.4.15:8 1.4.15:10	Reserved for future speeds	Value always 0, writes ignored	RO
<u>1.4.9</u>	<u>100G capable</u>	<u>1 = PMA/PMD is capable of operating at 100 Gb/s</u> <u>0 = PMA/PMD is not capable of operating as 100 Gb/s</u>	<u>RO</u>
<u>1.4.8</u>	<u>40G capable</u>	<u>1 = PMA/PMD is capable of operating at 40 Gb/s</u> <u>0 = PMA/PMD is not capable of operating as 40 Gb/s</u>	<u>RO</u>

^aRO = Read only

Change 45.2.1.2.1 for 10/40/100 Gb/s as follows:

45.2.1.2.1 Fault (1.1.7)

Fault is a global PMA/PMD variable. When read as a one, bit 1.1.7 indicates that either (or both) the PMA or the PMD has detected a fault condition on either the transmit or receive paths. When read as a zero, bit 1.1.7 indicates that neither the PMA nor the PMD has detected a fault condition. For 10/40/100 Gb/s operation, bit 1.1.7 is set to a one when either of the fault bits (1.8.11, 1.8.10) located in register 1.8 are set to a one. For 10PASS-TS or 2BASE-TL operations, when read as a one, a fault has been detected and more detailed information is conveyed in 45.2.1.16, 45.2.1.39, 45.2.1.40, and 45.2.1.55.

Insert 45.2.1.4.a and 45.2.1.4.b before 45.2.1.4.1 (inserted by IEEE Std 802.3av) as follows:

45.2.1.4.a 100G capable (1.4.9)

When read as a one, bit 1.4.9 indicates that the PMA/PMD is able to operate at a data rate of 100 Gb/s. When read as a zero, bit 1.4.9 indicates that the PMA/PMD is not able to operate at a data rate of 100 Gb/s.

45.2.1.4.b 40G capable (1.4.8)

When read as a one, bit 1.4.8 indicates that the PMA/PMD is able to operate at a data rate of 40 Gb/s. When read as a zero, bit 1.4.8 indicates that the PMA/PMD is not able to operate at a data rate of 40 Gb/s.

Change Table 45–7 (as modified by IEEE Std 802.3av) for 40 Gb/s and 100 Gb/s PMA/PMD type selections:

Table 45–7—PMA/PMD control 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.7.15:5 <u>1.7.15:6</u>	Reserved	Value always 0, writes ignored	R/W
1.7.4:0 <u>1.7.5:0</u>	PMA/PMD type selection	<u>5 4 3 2 1 0</u> <u>1 1 x x x x = reserved for future use</u> <u>1 0 1 1 x x = reserved for future use</u> <u>1 0 1 0 1 1 = 100GBASE-ER4 PMA/PMD type</u> <u>1 0 1 0 1 0 = 100GBASE-LR4 PMA/PMD type</u> <u>1 0 1 0 0 1 = 100GBASE-SR10 PMA/PMD type</u> <u>1 0 1 0 0 0 = 100GBASE-CR10 PMA/PMD type</u> <u>1 0 0 1 x x = reserved for future use</u> <u>1 0 0 0 1 1 = 40GBASE-LR4 PMA/PMD type</u> <u>1 0 0 0 1 0 = 40GBASE-SR4 PMA/PMD type</u> <u>1 0 0 0 0 1 = 40GBASE-CR4 PMA/PMD type</u> <u>1 0 0 0 0 0 = 40GBASE-KR4 PMA/PMD type</u> <u>0 1 1 1 x x = reserved</u> <u>0 1 1 0 1 1 = reserved</u> <u>0 1 1 0 1 0 = 10GBASE-PR-U3</u> <u>0 1 1 0 0 1 = 10GBASE-PR-U1</u> <u>0 1 1 0 0 0 = 10/1GBASE-PRX-U3</u> <u>0 1 0 1 1 1 = 10/1GBASE-PRX-U2</u> <u>0 1 0 1 1 0 = 10/1GBASE-PRX-U1</u> <u>0 1 0 1 0 1 = 10GBASE-PR-D3</u> <u>0 1 0 1 0 0 = 10GBASE-PR-D2</u> <u>0 1 0 0 1 1 = 10GBASE-PR-D1</u> <u>0 1 0 0 1 0 = 10/1GBASE-PRX-D3</u> <u>0 1 0 0 0 1 = 10/1GBASE-PRX-D2</u> <u>0 1 0 0 0 0 = 10/1GBASE-PRX-D1</u> <u>0 0 1 1 1 1 = 10BASE-T PMA/PMD type</u> <u>0 0 1 1 1 0 = 100BASE-TX PMA/PMD type</u> <u>0 0 1 1 0 1 = 1000BASE-KX PMA/PMD type</u> <u>0 0 1 1 0 0 = 1000BASE-T PMA/PMD type</u> <u>0 0 1 0 1 1 = 10GBASE-KR PMA/PMD type</u> <u>0 0 1 0 1 0 = 10GBASE-KX4 PMA/PMD type</u> <u>0 0 1 0 0 1 = 10GBASE-T PMA type</u> <u>0 0 1 0 0 0 = 10GBASE-LRM PMA/PMD type</u> <u>0 0 0 1 1 1 = 10GBASE-SR PMA/PMD type</u> <u>0 0 0 1 1 0 = 10GBASE-LR PMA/PMD type</u> <u>0 0 0 1 0 1 = 10GBASE-ER PMA/PMD type</u> <u>0 0 0 1 0 0 = 10GBASE-LX4 PMA/PMD type</u> <u>0 0 0 0 1 1 = 10GBASE-SW PMA/PMD type</u> <u>0 0 0 0 1 0 = 10GBASE-LW PMA/PMD type</u> <u>0 0 0 0 0 1 = 10GBASE-EW PMA/PMD type</u> <u>0 0 0 0 0 0 = 10GBASE-CX4 PMA/PMD type</u>	R/W

^aR/W = Read/Write

Change 45.2.1.6.1 for 40 Gb/s and 100 Gb/s PMA/PMD type selections as follows:

45.2.1.6.1 PMA/PMD type selection (1.7.35:0)

The PMA/PMD type of the PMA/PMD shall be selected using bits 35 through 0. The PMA/PMD type abilities of the PMA/PMD are advertised in bits 9 and 7 through 0 of the PMA/PMD status 2 register; ~~and bits 8 through 0 of the PMA/PMD extended ability register; and the 40G/100G PMA/PMD extended ability register.~~ A PMA/PMD shall ignore writes to the PMA/PMD type selection bits that select PMA/PMD types it has not advertised in the PMA/PMD status 2 register. It is the responsibility of the STA entity to ensure that mutually acceptable MMD types are applied consistently across all the MMDs on a particular PHY.

The PMA/PMD type selection defaults to a supported ability.

Change 45.2.1.7 and change the indicated row of Table 45–8 for naming as follows:

45.2.1.7 ~~40G~~ PMA/PMD status 2 register (Register 1.8)

The assignment of bits in the ~~40G~~ PMA/PMD status 2 register is shown in Table 45–8. All the bits in the ~~40G~~ PMA/PMD status 2 register are read only; a write to the ~~40G~~ PMA/PMD status 2 register shall have no effect.

Table 45–8—~~40G~~ PMA/PMD status 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.8.0	PMA <u>local</u> loopback ability	1 = PMA has the ability to perform a <u>local</u> loopback function 0 = PMA does not have the ability to perform a <u>local</u> loopback function	RO

^aRO = Read only, LH = Latching high

Change 45.2.1.7.4 and 45.2.1.7.5 as follows:

45.2.1.7.4 Transmit fault (1.8.11)

When read as a one, bit 1.8.11 indicates that the PMA/PMD has detected a fault condition on the transmit path. When read as a zero, bit 1.8.11 indicates that the PMA/PMD has not detected a fault condition on the transmit path. Detection of a fault condition on the transmit path is optional and the ability to detect such a condition is advertised by bit 1.8.13. A PMA/PMD that is unable to detect a fault condition on the transmit path shall return a value of zero for this bit. The description of the transmit fault function for the 10GBASE-KR PMD is given in 72.6.8, for 10GBASE-LRM serial PMDs in 68.4.8, and for other serial PMDs in 52.4.8. The description of the transmit fault function for WWDM PMDs is given in 53.4.10. The description of the transmit fault function for the 10GBASE-CX4 PMD is given in 54.5.10. The description of the transmit fault function for the 10GBASE-T PMA is given in 55.4.2.2. The description of the transmit fault function for the 10GBASE-KX4 PMD is given in 71.6.10. The description of the transmit fault function for the 40GBASE-KR4 PMD is given in 84.7.10. The description of the transmit fault function for the 40GBASE-CR4 and 100GBASE-CR10 PMDs is given in 85.7.10. The description of the transmit fault function for the 40GBASE-SR4 and 100GBASE-SR10 PMDs is given in 86.5.10. The description of the transmit fault function for the 40GBASE-LR4 PMD is given in 87.5.10. The description of the transmit fault function for the 100GBASE-LR4 and 100GBASE-ER4 PMDs is given in 88.5.10. The transmit fault bit shall be implemented with latching high behavior.

The default value of bit 1.8.11 is zero.

45.2.1.7.5 Receive fault (1.8.10)

When read as a one, bit 1.8.10 indicates that the PMA/PMD has detected a fault condition on the receive path. When read as a zero, bit 1.8.10 indicates that the PMA/PMD has not detected a fault condition on the receive path. Detection of a fault condition on the receive path is optional and the ability to detect such a condition is advertised by bit 1.8.12. A PMA/PMD that is unable to detect a fault condition on the receive path shall return a value of zero for this bit. The description of the receive fault function for the 10GBASE-KR PMD is given in 72.6.9, for 10GBASE-LRM serial PMDs in 68.4.9, and for other serial PMDs in 52.4.9. The description of the receive fault function for WWDM PMDs is given in 53.4.11. The description of the receive fault function for the 10GBASE-CX4 PMD is given in 54.5.11. The description of the receive fault function for the 10GBASE-T PMA is given in 55.4.2.4. The description of the receive fault function for the 10GBASE-KX4 PMD is given in 71.6.11. The description of the receive fault function for the 40GBASE-KR4 PMD is given in 84.7.11. The description of the receive fault function for the 40GBASE-CR4 and 100GBASE-CR10 PMDs is given in 85.7.11. The description of the receive fault function for the 40GBASE-SR4 and 100GBASE-SR10 PMDs is given in 86.5.11. The description of the receive fault function for the 40GBASE-LR4 PMD is given in 87.5.11. The description of the receive fault function for the 100GBASE-LR4 and 100GBASE-ER4 PMDs is given in 88.5.11. The receive fault bit shall be implemented with latching high behavior.

Change 45.2.1.7.15 for local loopback as follows:

45.2.1.7.15 PMA local loopback ability (1.8.0)

When read as a one, bit 1.8.0 indicates that the PMA is able to perform the local loopback function. When read as a zero, bit 1.8.0 indicates that the PMA is not able to perform the local loopback function. If a PMA is able to perform the local loopback function, then it is controlled using the PMA local loopback bit 1.0.0.

Change 45.2.1.8 for 40 Gb/s and 100 Gb/s PMA/PMD transmit disable as follows:

45.2.1.8 ~~40G~~ PMD transmit disable register (Register 1.9)

The assignment of bits in the ~~40G~~ PMD transmit disable register is shown in Table 45–9. The transmit disable functionality is optional and a PMD’s ability to perform the transmit disable functionality is advertised in the PMD transmit disable ability bit 1.8.8. A PMD that does not implement the transmit disable functionality shall ignore writes to the ~~40G~~ PMD transmit disable register and may return a value of zero for all bits. A PMD device that operates using a single ~~wavelength lane~~ and has implemented the transmit disable function shall use bit 1.9.0 to control the function. Such devices shall ignore writes to bits 1.9.410:1 and return a value of zero for those bits when they are read. The transmit disable function for the 10GBASE-KR PMD is described in 72.6.5, for 10GBASE-LRM serial PMDs in 68.4.7, and for other serial PMDs in 52.4.7. The transmit disable function for ~~wide wavelength division multiplexing (WWDM) PMDs~~ the 10GBASE-LX4 PMD is described in 53.4.7. The transmit disable function for the 10GBASE-CX4 PMD is described in 54.5.6. The transmit disable function for 10GBASE-KX4 is described in 71.6.6. The transmit disable function for the 10GBASE-T PMA is described in 55.4.2.3. The transmit disable function for 40GBASE-KR4 is described in 84.7.6. The transmit disable function for 40GBASE-CR4 and 100GBASE-CR10 is described in 85.7.6. The transmit disable function for 40GBASE-SR4 and 100GBASE-SR10 is described in 86.5.7. The transmit disable function for 40GBASE-LR4 is described in 87.5.7. The transmit disable function for 100GBASE-LR4 and 100GBASE-ER4 is described in 88.5.7.

NOTE—Disabling the transmitter on one or more lanes stops the entire link from carrying data.

Change the indicated rows of Table 45–9 for 10 lane transmit disables as follows:

Table 45–9—~~40G~~ PMD transmit disable register bit definitions

Bit(s)	Name	Description	R/W ^a
1.9.15: 5 11	Reserved	Value always 0, writes ignored	R/W
<u>1.9.10</u>	<u>PMD transmit disable 9</u>	<u>1 = Disable output on transmit lane 9</u> <u>0 = Enable output on transmit lane 9</u>	<u>R/W</u>
<u>1.9.9</u>	<u>PMD transmit disable 8</u>	<u>1 = Disable output on transmit lane 8</u> <u>0 = Enable output on transmit lane 8</u>	<u>R/W</u>
<u>1.9.8</u>	<u>PMD transmit disable 7</u>	<u>1 = Disable output on transmit lane 7</u> <u>0 = Enable output on transmit lane 7</u>	<u>R/W</u>
<u>1.9.7</u>	<u>PMD transmit disable 6</u>	<u>1 = Disable output on transmit lane 6</u> <u>0 = Enable output on transmit lane 6</u>	<u>R/W</u>
<u>1.9.6</u>	<u>PMD transmit disable 5</u>	<u>1 = Disable output on transmit lane 5</u> <u>0 = Enable output on transmit lane 5</u>	<u>R/W</u>
<u>1.9.5</u>	<u>PMD transmit disable 4</u>	<u>1 = Disable output on transmit lane 4</u> <u>0 = Enable output on transmit lane 4</u>	<u>R/W</u>

^aR/W = Read/Write

Insert two new subclauses before 45.2.1.8.1 for 10 lane transmit disable as follows:

45.2.1.8.a PMD transmit disable 9 (1.9.10)

When bit 1.9.10 is set to a one, the PMD shall disable output on lane 9 of the transmit path. When bit 1.9.10 is set to zero, the PMD shall enable output on lane 9 of the transmit path.

The default value for bit 1.9.10 is zero.

NOTE—Transmission will not be enabled when this bit is set to zero unless the global PMD transmit disable bit is also zero.

45.2.1.8.b PMD transmit disable 4, 5, 6, 7, 8 (1.9.5, 1.9.6, 1.9.7, 1.9.8, 1.9.9)

These bits are defined similarly to bit 1.9.10 for lanes 4, 5, 6, 7, and 8, respectively.

Change 45.2.1.9 for 40 Gb/s and 100 Gb/s PMA/PMD signal detect as follows:

45.2.1.9 ~~40G~~ PMD receive signal detect register (Register 1.10)

The assignment of bits in the ~~40G~~ PMD receive signal detect register is shown in Table 45–10. The ~~40G~~ PMD receive signal detect register is mandatory. PMD types that use only a single ~~wavelength lane~~ indicate the status of the receive signal detect using bit 1.10.0 and return a value of zero for bits 1.10.4:~~10:1~~ PMD types that use multiple wavelengths or lanes indicate the status of each lane in bits 1.10.4:~~10:1~~ and the logical AND of those bits in bit 1.10.0.

Change the indicated rows of Table 45–10 for 10 lane signal detect as follows:

Table 45–10—~~10G~~ PMD receive signal detect register bit definitions

Bit(s)	Name	Description	R/W ^a
1.10.15: 5 11	Reserved	Value always 0, writes ignored	RO
<u>1.10.10</u>	<u>PMD receive signal detect 9</u>	<u>1 = Signal detected on receive lane 9</u> <u>0 = Signal not detected on receive lane 9</u>	<u>RO</u>
<u>1.10.9</u>	<u>PMD receive signal detect 8</u>	<u>1 = Signal detected on receive lane 8</u> <u>0 = Signal not detected on receive lane 8</u>	<u>RO</u>
<u>1.10.8</u>	<u>PMD receive signal detect 7</u>	<u>1 = Signal detected on receive lane 7</u> <u>0 = Signal not detected on receive lane 7</u>	<u>RO</u>
<u>1.10.7</u>	<u>PMD receive signal detect 6</u>	<u>1 = Signal detected on receive lane 6</u> <u>0 = Signal not detected on receive lane 6</u>	<u>RO</u>
<u>1.10.6</u>	<u>PMD receive signal detect 5</u>	<u>1 = Signal detected on receive lane 5</u> <u>0 = Signal not detected on receive lane 5</u>	<u>RO</u>
<u>1.10.5</u>	<u>PMD receive signal detect 4</u>	<u>1 = Signal detected on receive lane 4</u> <u>0 = Signal not detected on receive lane 4</u>	<u>RO</u>

^aRO = Read only

Insert two new subclauses before 45.2.1.9.1 for 10 lane signal detect as follows:

45.2.1.9.a PMD receive signal detect 9 (1.10.10)

When bit 1.10.10 is read as a one, a signal has been detected on lane 9 of the PMD receive path. When bit 1.10.10 is read as a zero, a signal has not been detected on lane 9 of the PMD receive path.

45.2.1.9.b PMD receive signal detect 4, 5, 6, 7, 8 (1.10.5, 1.10.6, 1.10.7, 1.10.8, 1.10.9)

These bits are defined similarly to bit 1.10.10 for lanes 4, 5, 6, 7, and 8, respectively.

Change the indicated rows of Table 45–11 (as modified by IEEE Std 802.3av) in 45.2.1.10 for 40G/100G extended abilities as follows:

45.2.1.10 PMA/PMD extended ability register (Register 1.11)

Table 45–11—PMA/PMD extended ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.11.15:4011	Reserved	Ignore on read	RO
1.11.10	40G/100G extended abilities	1 = PMA/PMD has 40G/100G extended abilities listed in register 1.13 0 = PMA/PMD does not have 40G/100G extended abilities	RO

^aRO = Read only

Insert new subclause 45.2.1.11a and 45.2.1.11a.1 through 45.2.1.11a.9 after existing subclause 45.2.1.11 (this subclause was renumbered by IEEE Std 802.3av):

45.2.1.11a 40G/100G PMA/PMD extended ability register (Register 1.13)

The assignment of bits in the 40G/100G PMA/PMD extended ability register is shown in Table 45–12a. All of the bits in the PMA/PMD extended ability register are read only; a write to the PMA/PMD extended ability register shall have no effect.

Table 45–12a—40G/100G PMA/PMD extended ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.13.15	PMA remote loopback ability	1 = PMA has the ability to perform a remote loopback function 0 = PMA does not have the ability to perform a remote loopback function	RO
1.13.14:12	Reserved	Ignore on read	RO
1.13.11	100GBASE-ER4 ability	1 = PMA/PMD is able to perform 100GBASE-ER4 0 = PMA/PMD is not able to perform 100GBASE-ER4	RO
1.13.10	100GBASE-LR4 ability	1 = PMA/PMD is able to perform 100GBASE-LR4 0 = PMA/PMD is not able to perform 100GBASE-LR4	RO
1.13.9	100GBASE-SR10 ability	1 = PMA/PMD is able to perform 100GBASE-SR10 0 = PMA/PMD is not able to perform 100GBASE-SR10	RO
1.13.8	100GBASE-CR10 ability	1 = PMA/PMD is able to perform 100GBASE-CR10 0 = PMA/PMD is not able to perform 100GBASE-CR10	RO
1.13.7:4	Reserved	Ignore on read	RO
1.13.3	40GBASE-LR4 ability	1 = PMA/PMD is able to perform 40GBASE-LR4 0 = PMA/PMD is not able to perform 40GBASE-LR4	RO
1.13.2	40GBASE-SR4 ability	1 = PMA/PMD is able to perform 40GBASE-SR4 0 = PMA/PMD is not able to perform 40GBASE-SR4	RO
1.13.1	40GBASE-CR4 ability	1 = PMA/PMD is able to perform 40GBASE-CR4 0 = PMA/PMD is not able to perform 40GBASE-CR4	RO
1.13.0	40GBASE-KR4 ability	1 = PMA/PMD is able to perform 40GBASE-KR4 0 = PMA/PMD is not able to perform 40GBASE-KR4	RO

^aRO = Read only

45.2.1.11a.1 PMA remote loopback ability (1.13.15)

When read as a one, bit 1.13.15 indicates that the PMA is able to perform the remote loopback function. When read as a zero, bit 1.13.15 indicates that the PMA is not able to perform the remote loopback function. If a PMA is able to perform the remote loopback function, then it is controlled using the PMA remote loopback bit 1.0.1 (see 45.2.1.1.3a).

45.2.1.11a.2 100GBASE-ER4 ability (1.13.11)

When read as a one, bit 1.13.11 indicates that the PMA/PMD is able to operate as a 100GBASE-ER4 PMA/PMD type. When read as a zero, bit 1.13.11 indicates that the PMA/PMD is not able to operate as a 100GBASE-ER4 PMA/PMD type.

45.2.1.11a.3 100GBASE-LR4 ability (1.13.10)

When read as a one, bit 1.13.10 indicates that the PMA/PMD is able to operate as a 100GBASE-LR4 PMA/PMD type. When read as a zero, bit 1.13.10 indicates that the PMA/PMD is not able to operate as a 100GBASE-LR4 PMA/PMD type.

45.2.1.11a.4 100GBASE-SR10 ability (1.13.9)

When read as a one, bit 1.13.9 indicates that the PMA/PMD is able to operate as a 100GBASE-SR10 PMA/PMD type. When read as a zero, bit 1.13.9 indicates that the PMA/PMD is not able to operate as a 100GBASE-SR10 PMA/PMD type.

45.2.1.11a.5 100GBASE-CR10 ability (1.13.8)

When read as a one, bit 1.13.8 indicates that the PMA/PMD is able to operate as a 100GBASE-CR10 PMA/PMD type. When read as a zero, bit 1.13.8 indicates that the PMA/PMD is not able to operate as a 100GBASE-CR10 PMA/PMD type.

45.2.1.11a.6 40GBASE-LR4 ability (1.13.3)

When read as a one, bit 1.13.3 indicates that the PMA/PMD is able to operate as a 40GBASE-LR4 PMA/PMD type. When read as a zero, bit 1.13.3 indicates that the PMA/PMD is not able to operate as a 40GBASE-LR4 PMA/PMD type.

45.2.1.11a.7 40GBASE-SR4 ability (1.13.2)

When read as a one, bit 1.13.2 indicates that the PMA/PMD is able to operate as a 40GBASE-SR4 PMA/PMD type. When read as a zero, bit 1.13.2 indicates that the PMA/PMD is not able to operate as a 40GBASE-SR4 PMA/PMD type.

45.2.1.11a.8 40GBASE-CR4 ability (1.13.1)

When read as a one, bit 1.13.1 indicates that the PMA/PMD is able to operate as a 40GBASE-CR4 PMA/PMD type. When read as a zero, bit 1.13.1 indicates that the PMA/PMD is not able to operate as a 40GBASE-CR4 PMA/PMD type.

45.2.1.11a.9 40GBASE-KR4 ability (1.13.0)

When read as a one, bit 1.13.0 indicates that the PMA/PMD is able to operate as a 40GBASE-KR4 PMA/PMD type. When read as a zero, bit 1.13.0 indicates that the PMA/PMD is not able to operate as a 40GBASE-KR4 PMA/PMD type.

Change 45.2.1.77 (as renumbered by IEEE Std 802.3av) for register naming as follows:

45.2.1.77 ~~40GBASE-KR~~ BASE-R PMD control register (Register 1.150)

The BASE-R PMD control register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85. The assignment of bits in the BASE-R PMD control register is shown in Table 45-54.

Table 45-54—~~40GBASE-KR~~ BASE-R PMD control register

Bit(s)	Name	Description	R/W ^a
1.150.15:2	Reserved	Value always zero, writes ignored	RO
1.150.1	Training enable	1 = Enable the 40GBASE-KR <u>BASE-R</u> start-up protocol 0 = Disable the 40GBASE-KR <u>BASE-R</u> start-up protocol	R/W
1.150.0	Restart training	1 = Restart/Reset 40GBASE-KR <u>BASE-R</u> start-up protocol 0 = Normal operation	R/W SC

^aR/W = Read/Write, SC = Self-clearing, RO = Read only

Change 45.2.1.78 for register naming and multiple lanes as follows:

45.2.1.78 ~~40GBASE-KR~~ BASE-R PMD status register (Register 1.151)

The BASE-R PMD status register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85. The assignment of bits in the ~~40GBASE-KR~~ BASE-R PMD status register is shown in Table 45-55.

Table 45-55—~~40GBASE-KR~~ BASE-R PMD status register

Bit(s)	Name	Description	R/W ^a
1.151.15:4	Reserved	Value always zero, writes ignored	RO
<u>1.151.15</u>	<u>Training failure 3</u>	<u>1 = Training failure has been detected for lane 3</u> <u>0 = Training failure has not been detected for lane 3</u>	<u>RO</u>
<u>1.151.14</u>	<u>Start-up protocol status 3</u>	<u>1 = Start-up protocol in progress for lane 3</u> <u>0 = Start-up protocol complete for lane 3</u>	<u>RO</u>
<u>1.151.13</u>	<u>Frame lock 3</u>	<u>1 = Training frame delineation detected for lane 3</u> <u>0 = Training frame delineation not detected for lane 3</u>	<u>RO</u>
<u>1.151.12</u>	<u>Receiver status 3</u>	<u>1 = Receiver trained and ready to receive data for lane 3</u> <u>0 = Receiver training for lane 3</u>	<u>RO</u>
<u>1.151.11</u>	<u>Training failure 2</u>	<u>1 = Training failure has been detected for lane 2</u> <u>0 = Training failure has not been detected for lane 2</u>	<u>RO</u>
<u>1.151.10</u>	<u>Start-up protocol status 2</u>	<u>1 = Start-up protocol in progress for lane 2</u> <u>0 = Start-up protocol complete for lane 2</u>	<u>RO</u>
<u>1.151.9</u>	<u>Frame lock 2</u>	<u>1 = Training frame delineation detected for lane 2</u> <u>0 = Training frame delineation not detected for lane 2</u>	<u>RO</u>

Table 45–55—10GBASE-KR BASE-R PMD status register

1.151.8	<u>Receiver status 2</u>	<u>1 = Receiver trained and ready to receive data for lane 2</u> <u>0 = Receiver training for lane 2</u>	<u>RO</u>
1.151.7	<u>Training failure 1</u>	<u>1 = Training failure has been detected for lane 1</u> <u>0 = Training failure has not been detected for lane 1</u>	<u>RO</u>
1.151.6	<u>Start-up protocol status 1</u>	<u>1 = Start-up protocol in progress for lane 1</u> <u>0 = Start-up protocol complete for lane 1</u>	<u>RO</u>
1.151.5	<u>Frame lock 1</u>	<u>1 = Training frame delineation detected for lane 1</u> <u>0 = Training frame delineation not detected for lane 1</u>	<u>RO</u>
1.151.4	<u>Receiver status 1</u>	<u>1 = Receiver trained and ready to receive data for lane 1</u> <u>0 = Receiver training for lane 1</u>	<u>RO</u>
1.151.3	<u>Training failure 0</u>	<u>1 = Training failure has been detected for lane 0</u> <u>0 = Training failure has not been detected for lane 0</u>	<u>RO</u>
1.151.2	<u>Start-up protocol status 0</u>	<u>1 = Start-up protocol in progress for lane 0</u> <u>0 = Start-up protocol complete for lane 0</u>	<u>RO</u>
1.151.1	<u>Frame lock 0</u>	<u>1 = Training frame delineation detected for lane 0</u> <u>0 = Training frame delineation not detected for lane 0</u>	<u>RO</u>
1.151.0	<u>Receiver status 0</u>	<u>1 = Receiver trained and ready to receive data for lane 0</u> <u>0 = Receiver training for lane 0</u>	<u>RO</u>

^aRO = Read only

45.2.1.78.1 Receiver status 0 (1.151.0)

This bit maps to the state variable rx_trained as defined in 72.6.10.3.1.

45.2.1.78.2 Frame lock 0 (1.151.1)

This bit maps to the state variable frame_lock as defined in 72.6.10.3.1.

45.2.1.78.3 Start-up protocol status 0 (1.151.2)

This bit maps to the state variable training as defined in 72.6.10.3.1.

45.2.1.78.4 Training failure 0 (1.151.3)

This bit maps to the state variable training_failure as defined in 72.6.10.3.1.

45.2.1.78.5 Receiver status 1, 2, 3 (1.151.4, 1.151.8, 1.151.12)

These bits are defined identically to 1.151.0 for lanes 1, 2, and 3 respectively.

45.2.1.78.6 Frame lock 1, 2, 3 (1.151.5, 1.151.9, 1.151.13)

These bits are defined identically to 1.151.1 for lanes 1, 2, and 3 respectively.

45.2.1.78.7 Start-up protocol status 1, 2, 3 (1.151.6, 1.151.10, 1.151.14)

These bits are defined identically to 1.151.2 for lanes 1, 2, and 3 respectively.

45.2.1.78.8 Training failure 1, 2, 3 (1.151.7, 1.151.11, 1.151.15)

These bits are defined identically to 1.151.3 for lanes 1, 2, and 3 respectively.

Change 45.2.1.79 for register naming and for multi-lane as follows:

45.2.1.79 ~~40GBASE-KR~~ BASE-R LP coefficient update, lane 0 register (Register 1.152)

The BASE-R LP coefficient update, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85. The ~~40GBASE-KR~~ BASE-R LP coefficient update, lane 0 register reflects the contents of the first 16-bit word of the training frame most recently received from the control channel for lane 0 or for a single-lane PHY.

The assignment of bits in the ~~40GBASE-KR~~ BASE-R LP coefficient update, lane 0 register is shown in Table 45–56. Normally the bits in this register are read only; however, when training is disabled by setting low bit 1 in the ~~40GBASE-KR~~ BASE-R PMD control register, the ~~40GBASE-KR~~ BASE-R LP coefficient update, lane 0 register becomes writeable.

A copy of this register may be implemented at address 1.1100 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as accesses to the original register.

Change the title of Table 45-56 as follows:

Table 45-56 ~~40GBASE-KR~~ BASE-R LP coefficient update, lane 0 register bit definitions

Change 45.2.1.80 for register naming and for multi-lane as follows:

45.2.1.80 ~~40GBASE-KR~~ BASE-R LP status report, lane 0 register (Register 1.153)

The BASE-R LP status report, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, 84 or 85. The ~~40GBASE-KR~~ BASE-R LP status report, lane 0 register reflects the contents of the second 16-bit word of the training frame most recently received from the control channel for lane 0 or for a single-lane PHY.

The assignment of bits in the ~~40GBASE-KR~~ BASE-R LP status report, lane 0 register is shown in Table 45-57.

A copy of this register may be implemented at address 1.1200 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as the original register.

Change the title of Table 45-57 as follows:

Table 45-57 ~~40GBASE-KR~~ BASE-R LP status report, lane 0 register bit definitions

Change 45.2.1.81 for register naming and for multi-lane as follows:

45.2.1.81 ~~40GBASE-KR~~ BASE-R LD coefficient update, lane 0 register (Register 1.154)

The BASE-R LD coefficient update, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85. The ~~40GBASE-KR~~ BASE-R LD coefficient update, lane 0 register reflects the contents of the first 16-bit word of the outgoing training frame as defined by the LD receiver adaptation process in 72.6.10.2.5 for lane 0 or for a single-lane PHY.

The assignment of bits in the ~~40GBASE-KR~~ BASE-R LD coefficient update, lane 0 register is shown in Table 45-58.

A copy of this register may be implemented at address 1.1300 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as the original register.

Change the title of Table 45-58 as follows:

Table 45-58 ~~40GBASE-KR~~ BASE-R LD coefficient update, lane 0 register bit definitions

Change 45.2.1.82 for register naming and for multi-lane as follows:

45.2.1.82 ~~40GBASE-KR~~ BASE-R LD status report, lane 0 register (Register 1.155)

The BASE-R LD status report, lane 0 register is used for 10GBASE-KR and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85. The ~~40GBASE-KR~~ BASE-R LD status report, lane 0 register reflects the contents of the second 16-bit word of the current outgoing training frame, as defined in the training state diagram in Figure 72-5 for lane 0 or for a single-lane PHY.

The assignment of bits in the ~~40GBASE-KR~~ BASE-R LD status report, lane 0 register is shown in Table 45-59.

A copy of this register may be implemented at address 1.1400 to assist PHY access for devices using post-read-increment-address access for a multi-lane PCS. If implemented, all accesses to the copy have identical behavior as the original register.

Change the title of Table 45-59 as follows:

Table 45-59 ~~40GBASE-KR~~ BASE-R LD status report, lane 0 register bit definitions

Insert 45.2.1.82a and 45.2.1.82b for status 2 register and status 3 register as follows:

45.2.1.82a BASE-R PMD status 2 register (Register 1.156)

The BASE-R PMD status 2 register is used for 100GBASE-CR10 and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85 over more than 4 lanes. The assignment of bits in the BASE-R PMD status 2 register is shown in Table 45–59a.

Table 45–59a—BASE-R PMD status 2 register

Bit(s)	Name	Description	R/W ^a
1.156.15	Training failure 7	1 = Training failure has been detected for lane 7 0 = Training failure has not been detected for lane 7	RO
1.156.14	Start-up protocol status 7	1 = Start-up protocol in progress for lane 7 0 = Start-up protocol complete for lane 7	RO
1.156.13	Frame lock 7	1 = Training frame delineation detected for lane 7 0 = Training frame delineation not detected for lane 7	RO
1.156.12	Receiver status 7	1 = Receiver trained and ready to receive data for lane 7 0 = Receiver training for lane 7	RO
1.156.11	Training failure 6	1 = Training failure has been detected for lane 6 0 = Training failure has not been detected for lane 6	RO
1.156.10	Start-up protocol status 6	1 = Start-up protocol in progress for lane 6 0 = Start-up protocol complete for lane 6	RO
1.156.9	Frame lock 6	1 = Training frame delineation detected for lane 6 0 = Training frame delineation not detected for lane 6	RO
1.156.8	Receiver status 6	1 = Receiver trained and ready to receive data for lane 6 0 = Receiver training for lane 6	RO
1.156.7	Training failure 5	1 = Training failure has been detected for lane 5 0 = Training failure has not been detected for lane 5	RO
1.156.6	Start-up protocol status 5	1 = Start-up protocol in progress for lane 5 0 = Start-up protocol complete for lane 5	RO
1.156.5	Frame lock 5	1 = Training frame delineation detected for lane 5 0 = Training frame delineation not detected for lane 5	RO
1.156.4	Receiver status 5	1 = Receiver trained and ready to receive data for lane 5 0 = Receiver training for lane 5	RO
1.156.3	Training failure 4	1 = Training failure has been detected for lane 4 0 = Training failure has not been detected for lane 4	RO
1.156.2	Start-up protocol status 4	1 = Start-up protocol in progress for lane 4 0 = Start-up protocol complete for lane 4	RO
1.156.1	Frame lock 4	1 = Training frame delineation detected for lane 4 0 = Training frame delineation not detected for lane 4	RO
1.156.0	Receiver status 4	1 = Receiver trained and ready to receive data for lane 4 0 = Receiver training for lane 4	RO

^aRO = Read only

45.2.1.82a.1 Receiver status 4, 5, 6, 7 (1.156.0, 1.156.4, 1.156.8, 1.156.12)

These bits are defined identically to 1.151.0 for lanes 4, 5, 6, and 7 respectively.

45.2.1.82a.2 Frame lock 4, 5, 6, 7 (1.156.1, 1.156.5, 1.156.9, 1.156.13)

These bits are defined identically to 1.151.1 for lanes 4, 5, 6, and 7 respectively.

45.2.1.82a.3 Start-up protocol status 4, 5, 6, 7 (1.156.2, 1.156.6, 1.156.10, 1.156.14)

These bits are defined identically to 1.151.2 for lanes 4, 5, 6, and 7 respectively.

45.2.1.82a.4 Training failure 4, 5, 6, 7 (1.156.3, 1.156.7, 1.156.11, 1.156.15)

These bits are defined identically to 1.151.3 for lanes 4, 5, 6, and 7 respectively.

45.2.1.82b BASE-R PMD status 3 register (Register 1.157)

The BASE-R PMD status 3 register is used for 100GBASE-CR10 and other PHY types using the PMDs described in Clause 72, Clause 84, or Clause 85 over more than 8 lanes. The assignment of bits in the BASE-R PMD status 3 register is shown in Table 45–59b.

Table 45–59b—BASE-R PMD status 3 register

Bit(s)	Name	Description	R/W ^a
1.157.15:8	Reserved	Value always zero, writes ignored	RO
1.157.7	Training failure 9	1 = Training failure has been detected for lane 9 0 = Training failure has not been detected for lane 9	RO
1.157.6	Start-up protocol status 9	1 = Start-up protocol in progress for lane 9 0 = Start-up protocol complete for lane 9	RO
1.157.5	Frame lock 9	1 = Training frame delineation detected for lane 9 0 = Training frame delineation not detected for lane 9	RO
1.157.4	Receiver status 9	1 = Receiver trained and ready to receive data for lane 9 0 = Receiver training for lane 9	RO
1.157.3	Training failure 8	1 = Training failure has been detected for lane 8 0 = Training failure has not been detected for lane 8	RO
1.157.2	Start-up protocol status 8	1 = Start-up protocol in progress for lane 8 0 = Start-up protocol complete for lane 8	RO
1.157.1	Frame lock 8	1 = Training frame delineation detected for lane 8 0 = Training frame delineation not detected for lane 8	RO
1.157.0	Receiver status 8	1 = Receiver trained and ready to receive data for lane 8 0 = Receiver training for lane 8	RO

^aRO = Read only

45.2.1.82b.1 Receiver status 8, 9 (1.157.0, 1.157.4)

These bits are defined identically to 1.151.0 for lanes 8 and 9 respectively.

45.2.1.82b.2 Frame lock 8, 9 (1.157.1, 1.157.5)

These bits are defined identically to 1.151.1 for lanes 8 and 9 respectively.

45.2.1.82b.3 Start-up protocol status 8, 9 (1.157.2, 1.157.6)

These bits are defined identically to 1.151.2 for lanes 8 and 9 respectively.

45.2.1.82b.4 Training failure 8, 9 (1.157.3, 1.157.7)

These bits are defined identically to 1.151.3 for lanes 8 and 9 respectively.

Change 45.2.1.85 (as renumbered by IEEE Std 802.3av) for register naming:

45.2.1.85 40GBASE-R FEC ability register (Register 1.170)

The assignment of bits in the 40GBASE-R FEC ability register is shown in Table 45–62.

Table 45–62—40GBASE-R FEC ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.170.15:2	Reserved	Value always zero, writes ignored	RO
1.170.1	40GBASE-R FEC error indication ability	A read of 1 in this bit indicates that the 40GBASE-R PHY sublayer is able to report FEC decoding errors to the PCS layer	RO
1.170.0	40GBASE-R FEC ability	A read of 1 in this bit indicates that the 40GBASE-R PHY sublayer supports FEC	RO

^aRO Read only

45.2.1.85.1 40GBASE-R FEC ability (1.170.0)

When read as a one, this bit indicates that the ~~40GBASE-R PHY~~ sublayer supports forward error correction (FEC). When read as a zero, the ~~40GBASE-R PHY~~ sublayer does not support forward error correction.

45.2.1.85.2 40GBASE-R FEC error indication ability (1.170.1)

When read as a one, this bit indicates that the 40GBASE-R FEC sublayer is able to indicate decoding errors to the PCS layer (see 74.8.3). When read as a zero, the 40GBASE-R FEC sublayer is not able to indicate decoding errors to the PCS layer. 40GBASE-R FEC error indication is controlled by the FEC enable error indication bit in the BASE-R FEC control register (see 45.2.1.86.2).

Change 45.2.1.86 for register naming:

45.2.1.86 40GBASE-R FEC control register (Register 1.171)

The assignment of bits in the 40GBASE-R FEC control register is shown in Table 45–63.

Table 45–63—40GBASE-R FEC control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.171.15:2	Reserved	Value always zero, writes ignored	RO
1.171.1	FEC enable error indication	A write of 1 to this bit configures the FEC decoder to indicate errors to the PCS layer	R/W
1.171.0	FEC enable	A write of 1 to this bit enables 40GBASE-R FEC A write of 0 to this bit disables 40GBASE-R FEC	R/W

^aR/W = Read/Write, RO Read only

45.2.1.86.1 FEC enable (1.171.0)

When written as a one, this bit enables FEC for the 40GBASE-R PHY. When written as a zero, FEC is disabled in the 40GBASE-R PHY. This bit shall be set to zero upon execution of PHY reset.

45.2.1.86.2 FEC enable error indication (1.171.1)

This bit enables the 40GBASE-R FEC decoder to indicate decoding errors to the upper layers (PCS) through the sync bits for the 40GBASE-R PHY in the Local Device. When written as a one, this bit enables indication of decoding errors through the sync bits to the PCS layer. When written as zero the error indication function is disabled. Writes to this bit are ignored and reads return a zero if the 40GBASE-R FEC does not have the ability to indicate decoding errors to the PCS layer (see 45.2.1.85.2 and 74.8.3).

Change 45.2.1.87 and 88 for multi-lane:

45.2.1.87 10GBASE-R FEC corrected blocks counter (Register 1.172, 1.173)

The assignment of bits in the 10GBASE-R FEC corrected blocks counter register is shown in Table 45–64. See 74.8.4.1 for a definition of this register. These bits shall be reset to all ~~zeros~~zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.172, 1.173 are used to read the value of a 32-bit counter. When registers 1.172 and 1.173 are used to read the 32-bit counter value, the register 1.172 is read first, the value of the register 1.173 is latched when (and only when) register 1.172 is read and reads of register 1.173 returns the latched value rather than the current value of the counter.

Table 45–64—10GBASE-R FEC corrected blocks counter register bit definitions

Bit(s)	Name	Description	R/W ^a
1.172.15:0	FEC corrected blocks lower	FEC_corrected_blocks_counter[15:0]	RO, NR
1.173.15:0	FEC corrected blocks upper	FEC_corrected_blocks_counter[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over

For a multi-PCS lane PHY, this register may be a copy of register 1.300 BASE-R FEC corrected blocks counter, lane 0. If implemented, all accesses to the copy shall have identical behavior as the original register.

45.2.1.88 10GBASE-R FEC uncorrected blocks counter (Register 1.174, 1.175)

The assignment of bits in the 10GBASE-R FEC uncorrected blocks counter register is shown in Table 45–65. See 74.8.4.2 for a definition of this register. These bits shall be reset to all ~~zeros~~zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.174, 1.175 are used to read the value of a 32-bit counter. When registers 1.174 and 1.175 are used to read the 32-bit counter value, the register 1.174 is read first, the value of the register 1.175 is latched when (and only when) register 1.174 is read and reads of register 1.175 returns the latched value rather than the current value of the counter.

Table 45–65—10GBASE-R FEC uncorrected blocks counter register bit definitions

Bit(s)	Name	Description	R/W ^a
1.174.15:0	FEC uncorrected blocks lower	FEC_uncorrected_blocks_counter[15:0]	RO, NR
1.175.15:0	FEC uncorrected blocks upper	FEC_uncorrected_blocks_counter[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over

For a multi-PCS lane PHY, this register may be a copy of register 1.700 BASE-R FEC uncorrected blocks counter, lane 0. If implemented, all accesses to the copy shall have identical behavior as the original register.

Insert 45.2.1.89 and 45.2.1.90 for multi-lane FEC as follows:

45.2.1.89 BASE-R FEC corrected blocks counter, lanes 0 through 19

(Register 1.300, 1.301, 1.302, 1.303, 1.304, 1.305, 1.306, 1.307, 1.308, 1.309, 1.310, 1.311, 1.312, 1.313, 1.314, 1.315, 1.316, 1.317, 1.318, 1.319, 1.320, 1.321, 1.322, 1.323, 1.324, 1.325, 1.326, 1.327, 1.328, 1.329, 1.330, 1.331, 1.332, 1.333, 1.334, 1.335, 1.336, 1.337, 1.338, 1.339)

For multi-PCS lane BASE-R PHYs, the even-numbered registers in this set are defined similarly to register 1.172 (see 45.2.1.87) but for lanes 0 through 19 respectively of multi-PCS lane PHYs. The odd-numbered registers in this set are defined similarly to register 1.173 (see 45.2.1.87) but for lanes 0 through 19 respectively of multilane PHYs. Note that the lane numbers refer to the service interface lanes and do not necessarily correspond to the PCS lane numbers.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.90 BASE-R FEC uncorrected blocks counter, lanes 0 through 19

(Register 1.700, 1.701, 1.702, 1.703, 1.704, 1.705, 1.706, 1.707, 1.708, 1.709, 1.710, 1.711, 1.712, 1.713, 1.714, 1.715, 1.716, 1.717, 1.718, 1.719, 1.720, 1.721, 1.722, 1.723, 1.724, 1.725, 1.726, 1.727, 1.728, 1.729, 1.730, 1.731, 1.732, 1.733, 1.734, 1.735, 1.736, 1.737, 1.738, 1.739)

For multi-PCS lane BASE-R PHYs, the even-numbered registers in this set are defined similarly to register 1.174 (see 45.2.1.88) but for lanes 0 through 19 respectively of multi-PCS lane PHYs. The odd-numbered registers in this set are defined similarly to register 1.175 (see 45.2.1.88) but for lanes 0 through 19 respectively of multilane PHYs. Note that the lane numbers refer to the service interface lanes and do not necessarily correspond to the PCS lane numbers.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

Insert 45.2.1.91 through 45.2.1.94 for multi-lane coefficient exchange as follows:

45.2.1.91 BASE-R LP coefficient update register, lanes 1 through 9

(Register 1.1101, 1.1102, 1.1103, 1.1104, 1.1105, 1.1106, 1.1107, 1.1108, 1.1109)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.152 (which is used for lane 0, see 45.2.1.79) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.92 BASE-R LP status report register, lanes 1 through 9

(Register 1.1201, 1.1202, 1.1203, 1.1204, 1.1205, 1.1206, 1.1207, 1.1208, 1.1209)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.153 (which is used for lane 0, see 45.2.1.80) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.93 BASE-R LD coefficient update register, lanes 1 through 9

(Register 1.1301, 1.1302, 1.1303, 1.1304, 1.1305, 1.1306, 1.1307, 1.1308, 1.1309)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.154 (which is used for lane 0, see 45.2.1.81) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

45.2.1.94 BASE-R LD status report register, lanes 1 through 9

(Register 1.1401, 1.1402, 1.1403, 1.1404, 1.1405, 1.1406, 1.1407, 1.1408, 1.1409)

For multi-PCS lane BASE-R PHYs, these registers are defined similarly to register 1.155 (which is used for lane 0, see 45.2.1.82) but for lanes 1 through 9 respectively of multi-PCS lane PHYs.

Registers corresponding to lanes that are not used for the selected PHY shall return all zeros.

*Insert 45.2.1.95 for test-pattern ability register:***45.2.1.95 Test-pattern ability (Register 1.1500)**

The test-pattern ability register is used for PHY types that implement square wave testing and PRBS testing in the PMA. These functions are described in 83.5.10. The assignment of bits in the test-pattern ability register is shown in Table 45–65a.

Table 45–65a—Test-pattern ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.1500.15:13	Reserved	Value always zero, writes ignored	RO
1.1500.12	Square wave test ability	1 = Square wave testing supported 0 = Square wave testing not supported	RO
1.1500.11:6	Reserved	Value always zero, writes ignored	RO
1.1500.5	PRBS9 Tx generator ability	1 = PRBS9 transmit direction pattern generator supported 0 = PRBS9 transmit direction pattern generator not supported	RO
1.1500.4	PRBS9 Rx generator ability	1 = PRBS9 receive direction pattern generator supported 0 = PRBS9 receive direction pattern generator not supported	RO
1.1500.3	PRBS31 Tx generator ability	1 = PRBS31 Transmit direction pattern generator supported 0 = PRBS31 Tx direction pattern generator not supported	RO
1.1500.2	PRBS31 Tx checker ability	1 = PRBS31 Transmit direction pattern checker supported 0 = PRBS31 Tx direction pattern checker not supported	RO
1.1500.1	PRBS31 Rx generator ability	1 = PRBS31 Receive direction pattern generator supported 0 = PRBS31 Rx direction pattern generator not supported	RO
1.1500.0	PRBS31 Rx checker ability	1 = PRBS31 Receive direction pattern checker supported 0 = PRBS31 Rx direction pattern checker not supported	RO

^aR/W = Read/Write, RO = Read only

If square wave testing is supported and this register is implemented then bit 1.1500.12 shall be set to 1. The square wave test is controlled by register 1.1510 (see 45.2.1.97).

If PRBS9 pattern testing is supported and this register is implemented then bit 1.1500.5 shall indicate the generation ability in the transmit direction and bit 1.1500.4 shall indicate the generation ability in the receive direction. The PRBS pattern test is controlled by register 1.1501 (see 45.2.1.96).

If PRBS31 pattern testing is supported and this register is implemented then bit 1.1500.3 shall indicate the generation ability in the transmit direction and bit 1.1500.1 shall indicate the generation ability in the receive direction. Bit 1.1500.2 shall indicate the checker ability in the transmit direction and bit 1.1500.0 shall indicate the checker ability in the receive direction. The PRBS pattern test is controlled by register 1.1501 (see 45.2.1.96).

Insert 45.2.1.96 for PRBS pattern testing control as follows:

45.2.1.96 PRBS pattern testing control (Register 1.1501)

The PRBS pattern testing control register is used for PHY types that implement PRBS pattern testing in the PMA. This function is described in 83.5.10. The assignment of bits in the PRBS pattern testing control register is shown in Table 45–65b.

Table 45–65b—PRBS pattern testing control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.1501.15:8	Reserved	Value always zero, writes ignored	RO
1.1501.7	PRBS31 pattern enable	1 = Enable PRBS31 test-pattern 0 = Disable PRBS31 test-pattern	R/W
1.1501.6	PRBS9 pattern enable	1 = Enable PRBS9 test-pattern 0 = Disable PRBS9 test-pattern	R/W
1.1501.5:4	Reserved	Value always zero, writes ignored	RO
1.1501.3	Tx generator enable	1 = Enable transmit direction test-pattern generator 0 = Disable transmit direction test-pattern generator	R/W
1.1501.2	Tx checker enable	1 = Enable transmit direction test-pattern checker 0 = Disable transmit direction test-pattern checker	R/W
1.1501.1	Rx generator enable	1 = Enable receive direction test-pattern generator 0 = Disable receive direction test-pattern generator	R/W
1.1501.0	Rx checker enable	1 = Enable receive direction test-pattern checker 0 = Disable receive direction test-pattern checker	R/W

^aR/W = Read/Write, RO = Read only

Register 1.1501, bit 7 enables testing with the PRBS31 pattern defined in 83.5.10. Register 1.1501, bit 6 enables testing with the PRBS9 pattern defined in 83.5.10. The assertion of register 1.1501 bits 7 and 6 is mutually exclusive. If both bits are asserted the behavior is undefined. The assertion of register 1.1501, bits 7 and 6 works in conjunction with register 1.1501, bits 3:0. If none of the bits 3:0 are asserted then bits 7 and 6 have no effect.

Register 1.1501, bit 3 enables PRBS generation in the transmit direction. Register 1.1501, bit 2 enables PRBS checking in the transmit direction. Register 1.1501, bit 1 enables PRBS generation in the receive direction. Register 1.1501, bit 0 enables PRBS checking in the receive direction. If neither of the bits 7 and 6 are asserted then bits 3:0 have no effect.

Insert 45.2.1.97 for square wave pattern testing control:

45.2.1.97 Square wave testing control (Register 1.1510)

The square wave testing control register is used for PHY types that implement transmit square wave testing in the PMA. This function is described in 83.5.10. The assignment of bits in the square wave testing control register is shown in Table 45–65c.

Table 45–65c—Square wave testing control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.1510.15:10	Reserved	Value always zero, writes ignored	RO
1.1510.9	Lane 9 SW enable	1 = Enable square wave on lane 9 0 = Disable square wave on lane 9	R/W
1.1510.8	Lane 8 SW enable	1 = Enable square wave on lane 8 0 = Disable square wave on lane 8	R/W
1.1510.7	Lane 7 SW enable	1 = Enable square wave on lane 7 0 = Disable square wave on lane 7	R/W
1.1510.6	Lane 6 SW enable	1 = Enable square wave on lane 6 0 = Disable square wave on lane 6	R/W
1.1510.5	Lane 5 SW enable	1 = Enable square wave on lane 5 0 = Disable square wave on lane 5	R/W
1.1510.4	Lane 4 SW enable	1 = Enable square wave on lane 4 0 = Disable square wave on lane 4	R/W
1.1510.3	Lane 3 SW enable	1 = Enable square wave on lane 3 0 = Disable square wave on lane 3	R/W
1.1510.2	Lane 2 SW enable	1 = Enable square wave on lane 2 0 = Disable square wave on lane 2	R/W
1.1510.1	Lane 1 SW enable	1 = Enable square wave on lane 1 0 = Disable square wave on lane 1	R/W
1.1510.0	Lane 0 SW enable	1 = Enable square wave on lane 0 0 = Disable square wave on lane 0	R/W

^aR/W = Read/Write, RO = Read only

Register 1.1510, bits 0 through 9 enable square wave output on PMA lanes 0 through 9 respectively. Lanes for which a square wave pattern is not enabled act as determined by other registers.

Insert 45.2.1.98 and 45.2.1.99 for PRBS pattern testing counters as follows:

45.2.1.98 PRBS Tx pattern testing error counter (Register 1.1600, 1.1601, 1.1602, 1.1603, 1.1604, 1.1605, 1.1606, 1.1607, 1.1608, 1.1609)

The PRBS Tx pattern testing error counter registers are used for PHY types that implement PRBS Tx pattern testing in the PMA. This function is described in 83.5.10. The assignment of bits in the PRBS Tx pattern testing error counter registers are shown in Table 45–65d. Register 1.1600 contains the PRBS Tx pattern testing error counter for lane 0, register 1.1601 contains the PRBS Tx pattern testing error counter for lane 1,

and registers 1.1602 through 1.1609 contain the PRBS Tx pattern testing error counters for lanes 2 through 9 respectively. Counters corresponding to lanes that are not implemented in a PMA shall read all zeros.

Table 45–65d—PRBS Tx pattern testing error counter

Bit(s)	Name	Description	R/W ^a
1.1600 ^b .15:0	Error counter		RO, NR

^aRO = Read only, NR = Non Roll-over.

^bAll instances of address 1.1600 also apply to addresses 1.1601 through 1.1609.

The PRBS Tx pattern testing error counter is a 16-bit counter as defined in 83.5.10. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.1.99 PRBS Rx pattern testing error counter (Register 1.1700, 1.1701, 1.1702, 1.1703, 1.1704, 1.1705, 1.1706, 1.1707, 1.1708, 1.1709)

The PRBS Rx pattern testing error counter registers are used for PHY types that implement PRBS Rx pattern testing in the PMA. This function is described in 83.5.10. The assignment of bits in the PRBS Rx pattern testing error counter registers is identical to the PRBS Tx pattern testing error counter as shown in Table 45–65d. Register 1.1700 contains the PRBS Rx pattern testing error counter for lane 0, register 1.1701 contains the PRBS Rx pattern testing error counter for lane 1, and registers 1.1702 through 1.1709 contain the PRBS Rx pattern testing error counters for lanes 2 through 9 respectively. Counters corresponding to lanes that are not implemented in a PMA shall read all zeros.

The PRBS Rx pattern testing error counter is a 16-bit counter as defined in 83.5.10. These bits shall be reset to all zeros when the register is read by the management function or upon execution of the PMA reset. These bits shall be held at all ones in the case of overflow.

45.2.3 PCS registers

Change indicated rows of Table 45–83 (as modified and renumbered by IEEE Std 802.3av) for 40 Gb/s and 100 Gb/s PCS registers:

Table 45–83—PCS registers

Register address	Register name
3.0	PCS control 1
3.1	PCS status 1
3.2, 3.3	PCS device identifier
3.4	PCS speed ability
3.5, 3.6	PCS devices in package
3.7	40G PCS control 2
3.8	40G PCS status 2
3.9 through 3.13	Reserved
3.14, 3.15	PCS package identifier
3.16 through 3.23	Reserved
3.24	10GBASE-X PCS status
3.25	10GBASE-X PCS test control
3.26 through 3.31	Reserved
3.32	40G BASE-R and 10GBASE-T PCS status 1
3.33	40G BASE-R and 10GBASE-T PCS status 2
3.34 through 3.37	10GBASE-R PCS test pattern seed A
3.38 through 3.41	10GBASE-R PCS test pattern seed B
3.42	40G BASE-R PCS test pattern control
3.43	40G BASE-R PCS test pattern error counter
<u>3.44</u>	<u>BER high order counter</u>
<u>3.45</u>	<u>Errored blocks high order counter</u>
3.446 through 3.549	Reserved
<u>3.50 through 3.53</u>	<u>Multi-lane BASE-R PCS alignment status 1 through 4</u>
<u>3.54 through 3.59</u>	<u>Reserved</u>
...	
3.83 through 3.32-767 <u>199</u>	<u>Reserved</u>
<u>3.200 through 3.219</u>	<u>BIP error counters, lanes 0 through 19</u>
<u>3.220 through 3.399</u>	<u>Reserved</u>
<u>3.400 through 3.419</u>	<u>PCS lane mapping registers, lanes 0 through 19</u>
<u>3.420 through 3.32-767</u>	<u>Reserved</u>

Change the indicated row of Table 45–84 (as modified and renumbered by IEEE Std 802.3av) for 40 Gb/s and 100 Gb/s speed selection:.

Table 45–84—PCS control 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.0.5:2	Speed selection	5 4 3 2 1 x x x = Reserved 0 1 1 x = Reserved 0 1 0 1 = Reserved * 1 * * = Reserved * * 1 1 = Reserved 0 1 0 0 = 100 Gb/s 0 0 1 1 = 40 Gb/s 0 0 1 0 = 10/1 Gb/s 0 0 0 1 = 10PASS-TS/2BASE-TL 0 0 0 0 = 10 Gb/s	R/W

^aR/W = Read/Write, SC = Self-clearing

Change the indicated rows of Table 45–86 for 40 Gb/s and 100 Gb/s speed ability as follows:

Table 45–86—PCS speed ability register bit definitions

Bit(s)	Name	Description	R/W ^a
3.4.15:1 3.4.15:4	Reserved for future speeds	Value always 0, writes ignored	RO
<u>3.4.3</u>	<u>100G capable</u>	<u>1 = PCS is capable of operating at 100 Gb/s</u> <u>0 = PCS is not capable of operating at 100 Gb/s</u>	<u>RO</u>
<u>3.4.2</u>	<u>40G capable</u>	<u>1 = PCS is capable of operating at 40 Gb/s</u> <u>0 = PCS is not capable of operating at 40 Gb/s</u>	<u>RO</u>

^aRO = Read only

Change 45.2.3.2.2 for 40 Gb/s and 100 Gb/s link status as follows:

45.2.3.2.2 PCS receive link status (3.1.2)

When read as a one, bit 3.1.2 indicates that the PCS receive link is up. When read as a zero, bit 3.1.2 indicates that the PCS receive link is down. When a 10/40/100GBASE-R, 10GBASE-W, or 10GBASE-T mode of operation is selected for the PCS using the PCS type selection field (3.7.1:0), this bit is a latching low version of bit 3.32.12. When a 10GBASE-X mode of operation is selected for the PCS using the PCS type selection field (3.7.1:0), this bit is a latching low version of bit 3.24.12. The receive link status bit shall be implemented with latching low behavior.

Insert 45.2.3.4.3 and 45.2.3.4.4 as follows:

45.2.3.4.3 40G capable (3.4.2)

When read as a one, bit 3.4.2 indicates that the PCS is able to operate at a data rate of 40 Gb/s. When read as a zero, bit 3.4.2 indicates that the PCS is not able to operate at a data rate of 40 Gb/s.

45.2.3.4.4 100G capable (3.4.3)

When read as a one, bit 3.4.3 indicates that the PCS is able to operate at a data rate of 100 Gb/s. When read as a zero, bit 3.4.3 indicates that the PCS is not able to operate at a data rate of 100 Gb/s.

Change 45.2.3.6 for 40 Gb/s and 100 Gb/s PCS type select as follows:

45.2.3.6 40G PCS control 2 register (Register 3.7)

The assignment of bits in the ~~10G~~ PCS control 2 register is shown in Table 45–87. The default value for each bit of the ~~10G~~ PCS control 2 register should be chosen so that the initial state of the device upon power up or reset is a normal operational state without management intervention.

Table 45–87—~~10G~~ PCS control 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.7.15: 2 3	Reserved	Value always 0, writes ignored	R/W
3.7. 4 2:0	PCS type selection	$\begin{array}{l} \underline{2} \ \underline{1} \ \underline{0} \\ \underline{1} \ \underline{1} \ \underline{x} = \text{reserved} \\ \underline{1} \ \underline{0} \ \underline{1} = \text{Select 100GBASE-R PCS type} \\ \underline{1} \ \underline{0} \ \underline{0} = \text{Select 40GBASE-R PCS type} \\ \underline{0} \ \underline{1} \ \underline{1} = \text{Select 10GBASE-T PCS type} \\ \underline{0} \ \underline{1} \ \underline{0} = \text{Select 10GBASE-W PCS type} \\ \underline{0} \ \underline{0} \ \underline{1} = \text{Select 10GBASE-X PCS type} \\ \underline{0} \ \underline{0} \ \underline{0} = \text{Select 10GBASE-R PCS type} \end{array}$	R/W

^aR/W = Read/Write

45.2.3.6.1 PCS type selection (3.7.~~4~~2:0)

The PCS type shall be selected using bits ~~4~~2 through 0. The PCS type abilities of the ~~10G~~ PCS are advertised in bits 3.8.~~3~~2:0. A ~~10G~~ PCS shall ignore writes to the PCS type selection bits that select PCS types it has not advertised in the ~~10G~~ PCS status 2 register. It is the responsibility of the STA entity to ensure that mutually acceptable MMD types are applied consistently across all the MMDs on a particular PHY. The PCS type selection defaults to a supported ability.

Change 45.2.3.7 for 40 Gb/s and 100 Gb/s PCS type capability as follows:

45.2.3.7 40G PCS status 2 register (Register 3.8)

The assignment of bits in the ~~10G~~ PCS status 2 register is shown in Table 45–88. All the bits in the ~~10G~~ PCS status 2 register are read only; a write to the ~~10G~~ PCS status 2 register shall have no effect.

Change the indicated rows of Table 45–88 for 40 Gb/s and 100 Gb/s PCS ability as follows:

Insert the following two new subclauses before 45.2.3.7.4 for 40G/100G PCS abilities:

45.2.3.7.3a 100GBASE-R capable (3.8.5)

When read as a one, bit 3.8.5 indicates that the PCS is able to support the 100GBASE-R PCS type. When read as a zero, bit 3.8.5 indicates that the PCS is not able to support the 100GBASE-R PCS type.

Table 45–88—~~10G~~ PCS status 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.8.9:46	Reserved	Ignore when read	RO
<u>3.8.5</u>	<u>100GBASE-R capable</u>	<u>1 = PCS is able to support 100GBASE-R PCS type</u> <u>0 = PCS is not able to support 100GBASE-R PCS type</u>	<u>RO</u>
<u>3.8.4</u>	<u>40GBASE-R capable</u>	<u>1 = PCS is able to support 40GBASE-R PCS type</u> <u>0 = PCS is not able to support 40GBASE-R PCS type</u>	<u>RO</u>

^aRO = Read only, LH = Latching high

45.2.3.7.3b 40GBASE-R capable (3.8.4)

When read as a one, bit 3.8.4 indicates that the PCS is able to support the 40GBASE-R PCS type. When read as a zero, bit 3.8.4 indicates that the PCS is not able to support the 40GBASE-R PCS type.

Change 45.2.3.11 for naming as follows:

45.2.3.11 ~~40GBASE-R PCS~~ and 10GBASE-T PCS status 1 register (Register 3.32)

The assignment of bits in the ~~40GBASE-R~~ and 10GBASE-T PCS status 1 register is shown in Table 45–91. All the bits in the ~~40GBASE-R~~ and 10GBASE-T PCS status 1 register are read only; a write to the ~~40GBASE-R~~ and 10GBASE-T PCS status 1 register shall have no effect. A PCS device that does not implement ~~40GBASE-R~~ or the 10GBASE-T shall return a zero for all bits in the ~~40GBASE-R~~ and 10GBASE-T PCS status 1 register. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits. The contents of register 3.32 are undefined when the ~~40GBASE-R~~ PCS or the 10GBASE-T PCS is operating in seed test-pattern mode, PRBS31 test-pattern mode, or PRBS9 test-pattern mode.

Change 45.2.3.11.1 for naming as follows:

45.2.3.11.1 ~~40GBASE-R~~ and 10GBASE-T receive link status (3.32.12)

When read as a one, bit 3.32.12 indicates that the PCS is in a fully operational state. When read as a zero, bit 3.32.12 indicates that the PCS is not fully operational. This bit is a reflection of the PCS_status variable defined in 49.2.14.1 for 10GBASE-R, ~~and~~ in 55.3.6.1 for 10GBASE-T and in 82.3.1 for 40/100GBASE-R.

Change 45.2.3.11.4 and 45.2.3.11.5 for naming as follows:

45.2.3.11.4 ~~40GBASE-R~~ and 10GBASE-T PCS high BER (3.32.1)

For ~~40GBASE-R~~, when read as a one, bit 3.32.1 indicates that the 64B/66B receiver is detecting a BER of $\geq 10^{-4}$. When read as a zero, bit 3.32.1 indicates that the 64B/66B receiver is detecting a BER of $< 10^{-4}$. This bit is a direct reflection of the state of the hi_ber variable in the 64B/66B state diagram and is defined in 49.2.13.2.2 for 10GBASE-R and in 82.2.18.2.2 for 40/100GBASE-R.

For 10GBASE-T, when read as a one, bit 3.32.1 indicates that the 64B/65B receiver is detecting a BER of $\geq 10^{-4}$. When read as a zero, bit 3.32.1 indicates that the 64B/65B receiver is detecting a BER of $< 10^{-4}$. This bit is a direct reflection of the state of the hi_lfer variable in the 64B/65B state diagram and is defined in 55.3.6.1.

Table 45–91—~~40GBASE-R and 10GBASE-T PCS status 1 register bit definitions~~

Bit(s)	Name	Description	R/W ^a
3.32.15:13	Reserved	Value always 0, writes ignored	RO
3.32.12	40GBASE-R and 10GBASE-T receive link status	1 = 40GBASE-R or 10GBASE-T PCS receive link up 0 = 40GBASE-R or 10GBASE-T PCS receive link down	RO
3.32.11:4	Reserved	Ignore when read	RO
3.32.3	<u>10GBASE-R PRBS9 pattern testing ability</u>	1 = PCS is able to support PRBS9 pattern testing 0 = PCS is not able to support PRBS9 pattern testing	RO
3.32.2	<u>10GBASE-R PRBS31 pattern testing ability</u>	1 = PCS is able to support PRBS31 pattern testing 0 = PCS is not able to support PRBS31 pattern testing	RO
3.32.1	40GBASE-R and 10GBASE-T PCS high BER	1 = 40GBASE-R or 10GBASE-T PCS reporting a high BER 0 = 40GBASE-R or 10GBASE-T PCS not reporting a high BER	RO
3.32.0	40GBASE-R and 10GBASE-T PCS block lock	1 = 40GBASE-R or 10GBASE-T PCS locked to received blocks 0 = 40GBASE-R or 10GBASE-T PCS not locked to received blocks	RO

^aRO = Read only

45.2.3.11.5 ~~40GBASE-R and 10GBASE-T PCS block lock (3.32.0)~~

When read as a one, bit 3.32.0 indicates that the 64B/66B receiver for ~~40GBASE-R~~ or the 64B/65B receiver for the 10GBASE-T has block lock. When read as a zero, bit 3.32.0 indicates that the 64B/66B receiver for ~~40GBASE-R~~ or the 64B/65B receiver for the 10GBASE-T has not ~~got achieved~~ block lock. This bit is a direct reflection of the state of the `block_lock` variable in the 64B/66B state diagram and is defined in 49.2.13.2.2 for 10GBASE-R PCS and in 82.2.18.2.2 for 40/100GBASE-R. For the 10GBASE-T PCS, the `block_lock` variable in the 64B/65B state diagram is defined in 55.3.2.3. For a multi-lane PCS, this bit indicates that the receiver has both block lock and alignment for all lanes and is identical to 3.50.12 (see 45.2.3.16c.1).

Change 45.2.3.12 for naming as follows:

45.2.3.12 ~~40GBASE-R and 10GBASE-T PCS status 2 register (Register 3.33)~~

The assignment of bits in the ~~40GBASE-R and 10GBASE-T PCS status 2 register~~ is shown in Table 45–92. All the bits in the ~~40GBASE-R and 10GBASE-T PCS status 2 register~~ are read only; a write to the ~~40GBASE-R and 10GBASE-T PCS status 2 register~~ shall have no effect. A PCS device that does not implement ~~40GBASE-R or 10GBASE-T~~ shall return a zero for all bits in the ~~40GBASE-R and 10GBASE-T PCS status 2 register~~. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits. The contents of register 3.33 are undefined when the

10GBASE-R or the 10GBASE-T PCS is operating seed test-pattern mode, PRBS31 test-pattern mode, or PRBS9 test-pattern mode.

Table 45–92—~~10GBASE-R~~ and 10GBASE-T PCS status 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.33.15	Latched block lock	1 = 10GBASE-R or 10GBASE-T PCS has block lock 0 = 10GBASE-R or 10GBASE-T PCS does not have block lock	RO/LL
3.33.14	Latched high BER	1 = 10GBASE-R or 10GBASE-T PCS has reported a high BER 0 = 10GBASE-R or 10GBASE-T PCS has not reported a high BER	RO/LH
3.33.13:8	BER	BER counter	RO/NR
3.33.7:0	Errored blocks	Errored blocks counter	RO/NR

^aRO = Read only, LL = Latching low, LH = Latching high, NR = Non Roll-over

45.2.3.12.1 Latched block lock (3.33.15)

When read as a one, bit 3.33.15 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has achieved block lock. When read as a zero, bit 3.33.15 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has lost block lock.

The latched block lock bit shall be implemented with latching low behavior.

This bit is a latching low version of the 10/40/100GBASE-R and 10GBASE-T PCS block lock status bit (3.32.0).

45.2.3.12.2 Latched high BER (3.33.14)

When read as a one, bit 3.33.14 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has detected a high BER. When read as a zero, bit 3.33.14 indicates that the 10/40/100GBASE-R or the 10GBASE-T PCS has not detected a high BER.

The latched high BER bit shall be implemented with latching high behavior.

This bit is a latching high version of the 10/40/100GBASE-R and 10GBASE-T PCS high BER status bit (3.32.1).

45.2.3.12.3 BER (3.33.13:8)

The BER counter is a six bit count as defined by the ber_count variable in 49.2.14.2 and 82.2.18.2.4 for 10/40/100GBASE-R and defined by the lfer_count variable in 55.3.6.2 for 10GBASE-T. These bits shall be reset to all zeros when the ~~10GBASE-R~~ and 10GBASE-T PCS status 2 register is read by the management function or upon execution of the PCS reset. If the BER high order counter, 3.44 (see 45.2.3.16a) is not implemented then tThese bits shall be held at all ones in the case of overflow.

45.2.3.12.4 Errored blocks (3.33.7:0)

The errored blocks counter is an eight bit count defined by the errored_block_count counter specified in 49.2.14.2 for 10GBASE-R, in 82.3.1 for 40/100GBASE-R and defined by the errored_block_count variable

in 55.3.6.2 for 10GBASE-T. These bits shall be reset to all zeros when the errored blocks count is read by the management function or upon execution of the PCS reset. If the Errored blocks high order counter, 3.45 (see 45.2.3.16b) is not implemented then tThese bits shall be held at all ones in the case of overflow.

Change 45.2.3.15 for naming as follows:

45.2.3.15 ~~40~~GBASE-R PCS test-pattern control register (Register 3.42)

The assignment of bits in the ~~40~~GBASE-R PCS test-pattern control register is shown in Table 45–95. This register is only required when the ~~40~~GBASE-R capability is supported. If both ~~40~~GBASE-R and 10GBASE-W capability is supported, then this register may either ignore writes and return zeros for reads when in 10GBASE-W mode or may function as defined for ~~40~~GBASE-R. PRBS9, PRBS31, pseudo random, and square wave test patterns are defined for 10GBASE-R PCS only. Scrambled idle test patterns are defined for 40/100GBASE-R PCS only. The test-pattern methodology is described in 49.2.8 and 82.2.10.

Table 45–95—~~40~~GBASE-R PCS test-pattern control register bit definitions

Bit(s)	Name	Description	R/W ^a
3.42.15:7 8	Reserved	Value always 0, writes ignored	R/W
<u>3.42.7</u>	<u>Scrambled idle test-pattern enable</u>	<u>1 = Enable scrambled idle test-pattern mode</u> <u>0 = Disable scrambled idle test-pattern mode</u>	<u>R/W</u>
3.42.6	<u>10GBASE-R</u> PRBS9 transmit test-pattern enable	1 = Enable PRBS9 test-pattern mode on the transmit path 0 = Disable PRBS9 test-pattern mode on the transmit path	R/W
3.42.5	<u>10GBASE-R</u> PRBS31 receive test-pattern enable	1 = Enable PRBS31 test-pattern mode on the receive path 0 = Disable PRBS31 test-pattern mode on the receive path	R/W
3.42.4	<u>10GBASE-R</u> PRBS31 transmit test-pattern enable	1 = Enable PRBS31 test-pattern mode on the transmit path 0 = Disable PRBS31 test-pattern mode on the transmit path	R/W
3.42.3	Transmit test-pattern enable	1 = Enable transmit test pattern 0 = Disable transmit test pattern	R/W
3.42.2	Receive test-pattern enable	1 = Enable receive test-pattern testing 0 = Disable receive test-pattern testing	R/W
3.42.1	Test-pattern select	1 = Square wave test pattern 0 = Pseudo random test pattern	R/W
3.42.0	Data pattern select	1 = Zeros data pattern 0 = LF data pattern	R/W

^aR/W = Read/Write

Insert new subclause 45.2.3.15.a as follows before 45.2.3.15.1:

45.2.3.15.a Scrambled idle test-pattern enable (3.42.7)

When bit 3.42.7 is set to a one, scrambled idle pattern testing is enabled. When bit 3.42.7 is set to zero, scrambled idle pattern testing is disabled.

The default value for bit 3.42.3 is zero.

Change subclause headings for 45.2.3.15.1, 45.2.3.15.2, and 45.2.3.15.3 for naming as follows:

45.2.3.15.1 10GBASE-R PRBS9 transmit test-pattern enable (3.42.6)

45.2.3.15.2 10GBASE-R PRBS31 receive test-pattern enable (3.42.5)

45.2.3.15.3 10GBASE-R PRBS31 transmit test-pattern enable (3.42.4)

Change 45.2.3.16 as follows:

45.2.3.16 ~~40~~GBASE-R PCS test-pattern error counter register (Register 3.43)

The assignment of bits in the ~~40~~GBASE-R PCS test-pattern error counter register is shown in Table 45–96. This register is only required when the ~~40~~GBASE-R capability is supported. If both ~~40~~GBASE-R and 10GBASE-W capability is supported, then this register may either ignore writes and return zeros for reads when in 10GBASE-W mode, or may function as defined for ~~40~~GBASE-R.

The test-pattern error counter is a sixteen bit counter that contains the number of errors received during a pattern test. These bits shall be reset to all zeros when the test-pattern error counter is read by the management function or upon execution of the PCS reset. These bits shall be held at all ones in the case of overflow. The test-pattern methodology is described in 49.2.12 and 82.2.10. This counter will count either block errors or bit errors dependent on the test mode (see 49.2.12).

Change the title of Table 45-96 as follows:

Table 45-96 ~~40~~GBASE-R PCS test-pattern error counter register bit definitions

Insert subclauses 45.2.3.16a, 45.2.3.16b, 45.2.3.16c, 45.2.3.16d, 45.2.3.16e, and 45.2.3.16f after 45.2.3.16:

45.2.3.16a BER high order counter (Register 3.44)

The assignment of bits in the BER high order counter register is shown in Table 45–96a. This register is mandatory when the 40/100GBASE-R capability is supported and optional for other PHY types that implement register 3.33.

Table 45–96a—BER high order counter register bit definitions

Bit(s)	Name	Description	R/W ^a
3.44.15:0	BER high order	Bits 21:6 of BER counter	RO

^aRO = Read only

Bits 15:0 of this register are concatenated with bits 13:8 of register 3.33 (see 45.2.3.12) to read the value of a twenty-two bit BER counter. When registers 3.44 and 3.33 are used to read the 22-bit counter value, the register 3.33 is read first (indicating the lower 6 bits). The remaining 16 bits shall be latched when register 3.33 is read and reads to register 3.44 return the latched value. The counter continues to function as defined in 45.2.3.12.3 regardless of the state of the latched register. The 22-bit counter shall be reset to all zeros when register 3.33 is read or upon PCS reset. The 22-bit counter shall be held at all ones in the case of overflow.

45.2.3.16b Errored blocks high order counter (Register 3.45)

The assignment of bits in the Errored blocks high order counter register is shown in Table 45–96b. This register is mandatory when the 40/100GBASE-R capability is supported and optional for other PHY types that implement register 3.33. If this register is implemented then bit 3.45.15 shall be set to 1.

Table 45–96b—Errored blocks high order counter register bit definitions

Bit(s)	Name	Description	R/W ^a
3.45.15	High order counter present	Always reads as 1 if this register is implemented	RO
3.45.14	Reserved	Value always 0	RO
3.45.13:0	Errored blocks high order	Bits 21:8 of errored blocks counter	RO

^aRO = Read only

Bits 13:0 of this register are concatenated with bits 7:0 of register 3.33 (see 45.2.3.12) to read the value of a twenty-two bit errored blocks counter. When registers 3.45 and 3.33 are used to read the 22-bit counter value, the register 3.33 is read first (indicating the lower 8 bits). The remaining 14 bits shall be latched when register 3.33 is read and reads to register 3.45 return the latched value. The counter continues to function as defined in 45.2.3.12.4 regardless of the state of the latched register. The 22-bit counter shall be reset to all zeros when register 3.33 is read or upon PCS reset. The 22-bit counter shall be held at all ones in the case of overflow.

45.2.3.16c Multi-lane BASE-R PCS alignment status 1 register (Register 3.50)

The assignment of bits in the multi-lane BASE-R PCS alignment status 1 register is shown in Table 45–96c. All the bits in the multi-lane BASE-R PCS alignment status 1 register are read only; a write to the multi-lane BASE-R PCS alignment status 1 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 1 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 1 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

45.2.3.16c.1 Multi-lane BASE-R PCS alignment status (3.50.12)

When read as a one, bit 3.50.12 indicates that the PCS has locked and aligned all receive lanes. When read as a zero, bit 3.50.12 indicates that the PCS has not locked and aligned all receive lanes.

45.2.3.16c.2 Block 7 lock (3.50.7)

When read as a one, bit 3.50.7 indicates that the PCS receiver has achieved block lock for service interface lane 7. When read as a zero, bit 3.50.7 indicates that the PCS receiver lane 7 has not achieved block lock. This bit reflects the state of `block_lock[7]` (see 82.2.18.2.2).

45.2.3.16c.3 Block 6 lock (3.50.6)

When read as a one, bit 3.50.6 indicates that the PCS receiver has achieved block lock for service interface lane 6. When read as a zero, bit 3.50.6 indicates that the PCS receiver lane 6 has not achieved block lock. This bit reflects the state of `block_lock[6]` (see 82.2.18.2.2).

Table 45–96c—Multi-lane BASE-R PCS alignment status 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.50.15:13	Reserved	Ignore when read	RO
3.50.12	PCS lane alignment status	1 = PCS receive lanes locked and aligned 0 = PCS receive lanes not locked and aligned	RO
3.50.11:8	Reserved	Ignore when read	RO
3.50.7	Block 7 lock	1 = Lane 7 is locked 0 = Lane 7 is not locked	RO
3.50.6	Block 6 lock	1 = Lane 6 is locked 0 = Lane 6 is not locked	RO
3.50.5	Block 5 lock	1 = Lane 5 is locked 0 = Lane 5 is not locked	RO
3.50.4	Block 4 lock	1 = Lane 4 is locked 0 = Lane 4 is not locked	RO
3.50.3	Block 3 lock	1 = Lane 3 is locked 0 = Lane 3 is not locked	RO
3.50.2	Block 2 lock	1 = Lane 2 is locked 0 = Lane 2 is not locked	RO
3.50.1	Block 1 lock	1 = Lane 1 is locked 0 = Lane 1 is not locked	RO
3.50.0	Block 0 lock	1 = Lane 0 is locked 0 = Lane 0 is not locked	RO

^aRO = Read only

45.2.3.16c.4 Block 5 lock (3.50.5)

When read as a one, bit 3.50.5 indicates that the PCS receiver has achieved block lock for service interface lane 5. When read as a zero, bit 3.50.5 indicates that the PCS receiver lane 5 has not achieved block lock. This bit reflects the state of `block_lock[5]` (see 82.2.18.2.2).

45.2.3.16c.5 Block 4 lock (3.50.4)

When read as a one, bit 3.50.4 indicates that the PCS receiver has achieved block lock for service interface lane 4. When read as a zero, bit 3.50.4 indicates that the PCS receiver lane 4 has not achieved block lock. This bit reflects the state of `block_lock[4]` (see 82.2.18.2.2).

45.2.3.16c.6 Block 3 lock (3.50.3)

When read as a one, bit 3.50.3 indicates that the PCS receiver has achieved block lock for service interface lane 3. When read as a zero, bit 3.50.3 indicates that the PCS receiver lane 3 has not achieved block lock. This bit reflects the state of `block_lock[3]` (see 82.2.18.2.2).

45.2.3.16c.7 Block 2 lock (3.50.2)

When read as a one, bit 3.50.2 indicates that the PCS receiver has achieved block lock for service interface lane 2. When read as a zero, bit 3.50.2 indicates that the PCS receiver lane 2 has not achieved block lock. This bit reflects the state of `block_lock[2]` (see 82.2.18.2.2).

45.2.3.16c.8 Block 1 lock (3.50.1)

When read as a one, bit 3.50.1 indicates that the PCS receiver has achieved block lock for service interface lane 1. When read as a zero, bit 3.50.1 indicates that the PCS receiver lane 1 has not achieved block lock. This bit reflects the state of `block_lock[1]` (see 82.2.18.2.2).

45.2.3.16c.9 Block 0 lock (3.50.0)

When read as a one, bit 3.50.0 indicates that the PCS receiver has achieved block lock for service interface lane 0. When read as a zero, bit 3.50.0 indicates that the PCS receiver lane 0 has not achieved block lock. This bit reflects the state of `block_lock[0]` (see 82.2.18.2.2).

45.2.3.16d Multi-lane BASE-R PCS alignment status 2 register (Register 3.51)

The assignment of bits in the multi-lane BASE-R PCS alignment status 2 register is shown in Table 45–96d. All the bits in the multi-lane BASE-R PCS alignment status 2 register are read only; a write to the multi-lane BASE-R PCS alignment status 2 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 2 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 2 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

45.2.3.16d.1 Block 19 lock (3.51.11)

When read as a one, bit 3.51.11 indicates that the PCS receiver has achieved block lock for service interface lane 19. When read as a zero, bit 3.51.11 indicates that the PCS receiver lane 19 has not achieved block lock. This bit reflects the state of `block_lock[19]` (see 82.2.18.2.2).

45.2.3.16d.2 Block 18 lock (3.51.10)

When read as a one, bit 3.51.10 indicates that the PCS receiver has achieved block lock for service interface lane 18. When read as a zero, bit 3.51.10 indicates that the PCS receiver lane 18 has not achieved block lock. This bit reflects the state of `block_lock[18]` (see 82.2.18.2.2).

45.2.3.16d.3 Block 17 lock (3.51.9)

When read as a one, bit 3.51.9 indicates that the PCS receiver has achieved block lock for service interface lane 17. When read as a zero, bit 3.51.9 indicates that the PCS receiver lane 17 has not achieved block lock. This bit reflects the state of `block_lock[17]` (see 82.2.18.2.2).

45.2.3.16d.4 Block 16 lock (3.51.8)

When read as a one, bit 3.51.8 indicates that the PCS receiver has achieved block lock for service interface lane 16. When read as a zero, bit 3.51.8 indicates that the PCS receiver lane 16 has not achieved block lock. This bit reflects the state of `block_lock[16]` (see 82.2.18.2.2).

Table 45–96d—Multi-lane BASE-R PCS alignment status 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.51.15:12	Reserved	Ignore when read	RO
3.51.11	Block 19 lock	1 = Lane 19 is locked 0 = Lane 19 is not locked	RO
3.51.10	Block 18 lock	1 = Lane 18 is locked 0 = Lane 18 is not locked	RO
3.51.9	Block 17 lock	1 = Lane 17 is locked 0 = Lane 17 is not locked	RO
3.51.8	Block 16 lock	1 = Lane 16 is locked 0 = Lane 16 is not locked	RO
3.51.7	Block 15 lock	1 = Lane 15 is locked 0 = Lane 15 is not locked	RO
3.51.6	Block 14 lock	1 = Lane 14 is locked 0 = Lane 14 is not locked	RO
3.51.5	Block 13 lock	1 = Lane 13 is locked 0 = Lane 13 is not locked	RO
3.51.4	Block 12 lock	1 = Lane 12 is locked 0 = Lane 12 is not locked	RO
3.51.3	Block 11 lock	1 = Lane 11 is locked 0 = Lane 11 is not locked	RO
3.51.2	Block 10 lock	1 = Lane 10 is locked 0 = Lane 10 is not locked	RO
3.51.1	Block 9 lock	1 = Lane 9 is locked 0 = Lane 9 is not locked	RO
3.51.0	Block 8 lock	1 = Lane 8 is locked 0 = Lane 8 is not locked	RO

^aRO = Read only

45.2.3.16d.5 Block 15 lock (3.51.7)

When read as a one, bit 3.51.7 indicates that the PCS receiver has achieved block lock for service interface lane 15. When read as a zero, bit 3.51.7 indicates that the PCS receiver lane 15 has not achieved block lock. This bit reflects the state of `block_lock[15]` (see 82.2.18.2.2).

45.2.3.16d.6 Block 14 lock (3.51.6)

When read as a one, bit 3.51.6 indicates that the PCS receiver has achieved block lock for service interface lane 14. When read as a zero, bit 3.51.6 indicates that the PCS receiver lane 14 has not achieved block lock. This bit reflects the state of `block_lock[14]` (see 82.2.18.2.2).

45.2.3.16d.7 Block 13 lock (3.51.5)

When read as a one, bit 3.51.5 indicates that the PCS receiver has achieved block lock for service interface lane 13. When read as a zero, bit 3.51.5 indicates that the PCS receiver lane 13 has not achieved block lock. This bit reflects the state of block_lock[13] (see 82.2.18.2.2).

45.2.3.16d.8 Block 12 lock (3.51.4)

When read as a one, bit 3.51.4 indicates that the PCS receiver has achieved block lock for service interface lane 12. When read as a zero, bit 3.51.4 indicates that the PCS receiver lane 12 has not achieved block lock. This bit reflects the state of block_lock[12] (see 82.2.18.2.2).

45.2.3.16d.9 Block 11 lock (3.51.3)

When read as a one, bit 3.51.3 indicates that the PCS receiver has achieved block lock for service interface lane 11. When read as a zero, bit 3.51.3 indicates that the PCS receiver lane 11 has not achieved block lock. This bit reflects the state of block_lock[11] (see 82.2.18.2.2).

45.2.3.16d.10 Block 10 lock (3.51.2)

When read as a one, bit 3.51.2 indicates that the PCS receiver has achieved block lock for service interface lane 10. When read as a zero, bit 3.51.2 indicates that the PCS receiver lane 10 has not achieved block lock. This bit reflects the state of block_lock[10] (see 82.2.18.2.2).

45.2.3.16d.11 Block 9 lock (3.51.1)

When read as a one, bit 3.51.1 indicates that the PCS receiver has achieved block lock for service interface lane 9. When read as a zero, bit 3.51.1 indicates that the PCS receiver lane 9 has not achieved block lock. This bit reflects the state of block_lock[9] (see 82.2.18.2.2).

45.2.3.16d.12 Block 8 lock (3.51.0)

When read as a one, bit 3.51.0 indicates that the PCS receiver has achieved block lock for service interface lane 8. When read as a zero, bit 3.51.0 indicates that the PCS receiver lane 8 has not achieved block lock. This bit reflects the state of block_lock[8] (see 82.2.18.2.2).

45.2.3.16e Multi-lane BASE-R PCS alignment status 3 register (Register 3.52)

The assignment of bits in the multi-lane BASE-R PCS alignment status 3 register is shown in Table 45–96e. All the bits in the multi-lane BASE-R PCS alignment status 3 register are read only; a write to the multi-lane BASE-R PCS alignment status 3 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 3 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 3 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

45.2.3.16e.1 Lane 7 aligned (3.52.7)

When read as a one, bit 3.52.7 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 7. When read as a zero, bit 3.52.7 indicates that the PCS receiver lane 7 has not achieved alignment marker lock. This bit reflects the state of am_lock[7] (see 82.2.18.2.2).

Table 45–96e—Multi-lane BASE-R PCS alignment status 3 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.52.15:8	Reserved	Ignore when read	RO
3.52.7	Lane 7 aligned	1 = Lane 7 alignment marker is locked 0 = Lane 7 alignment marker is not locked	RO
3.52.6	Lane 6 aligned	1 = Lane 6 alignment marker is locked 0 = Lane 6 alignment marker is not locked	RO
3.52.5	Lane 5 aligned	1 = Lane 5 alignment marker is locked 0 = Lane 5 alignment marker is not locked	RO
3.52.4	Lane 4 aligned	1 = Lane 4 alignment marker is locked 0 = Lane 4 alignment marker is not locked	RO
3.52.3	Lane 3 aligned	1 = Lane 3 alignment marker is locked 0 = Lane 3 alignment marker is not locked	RO
3.52.2	Lane 2 aligned	1 = Lane 2 alignment marker is locked 0 = Lane 2 alignment marker is not locked	RO
3.52.1	Lane 1 aligned	1 = Lane 1 alignment marker is locked 0 = Lane 1 alignment marker is not locked	RO
3.52.0	Lane 0 aligned	1 = Lane 0 alignment marker is locked 0 = Lane 0 alignment marker is not locked	RO

^aRO = Read only

45.2.3.16e.2 Lane 6 aligned (3.52.6)

When read as a one, bit 3.52.6 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 6. When read as a zero, bit 3.52.6 indicates that the PCS receiver lane 6 has not achieved alignment marker lock. This bit reflects the state of `am_lock[6]` (see 82.2.18.2.2).

45.2.3.16e.3 Lane 5 aligned (3.52.5)

When read as a one, bit 3.52.5 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 5. When read as a zero, bit 3.52.5 indicates that the PCS receiver lane 5 has not achieved alignment marker lock. This bit reflects the state of `am_lock[5]` (see 82.2.18.2.2).

45.2.3.16e.4 Lane 4 aligned (3.52.4)

When read as a one, bit 3.52.4 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 4. When read as a zero, bit 3.52.4 indicates that the PCS receiver lane 4 has not achieved alignment marker lock. This bit reflects the state of `am_lock[4]` (see 82.2.18.2.2).

45.2.3.16e.5 Lane 3 aligned (3.52.3)

When read as a one, bit 3.52.3 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 3. When read as a zero, bit 3.52.3 indicates that the PCS receiver lane 3 has not achieved alignment marker lock. This bit reflects the state of `am_lock[3]` (see 82.2.18.2.2).

45.2.3.16e.6 Lane 2 aligned (3.52.2)

When read as a one, bit 3.52.2 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 2. When read as a zero, bit 3.52.2 indicates that the PCS receiver lane 2 has not achieved alignment marker lock. This bit reflects the state of `am_lock[2]` (see 82.2.18.2.2).

45.2.3.16e.7 Lane 1 aligned (3.52.1)

When read as a one, bit 3.52.1 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 1. When read as a zero, bit 3.52.1 indicates that the PCS receiver lane 1 has not achieved alignment marker lock. This bit reflects the state of `am_lock[1]` (see 82.2.18.2.2).

45.2.3.16e.8 Lane 0 aligned (3.52.0)

When read as a one, bit 3.52.0 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 3.52.0 indicates that the PCS receiver lane 0 has not achieved alignment marker lock. This bit reflects the state of `am_lock[0]` (see 82.2.18.2.2).

45.2.3.16f Multi-lane BASE-R PCS alignment status 4 register (Register 3.53)

The assignment of bits in the multi-lane BASE-R PCS alignment status 4 register is shown in Table 45–96f. All the bits in the multi-lane BASE-R PCS alignment status 4 register are read only; a write to the multi-lane BASE-R PCS alignment status 4 register shall have no effect. A PCS device that does not implement multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 4 register. A device that implements multi-lane BASE-R PCS shall return a zero for all bits in the multi-lane BASE-R PCS alignment status 4 register that are not required for the PCS configuration. It is the responsibility of the STA management entity to ensure that a port type is supported by all MMDs before interrogating any of its status bits.

45.2.3.16f.1 Lane 19 aligned (3.53.11)

When read as a one, bit 3.53.11 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 19. When read as a zero, bit 3.53.11 indicates that the PCS receiver lane 19 has not achieved alignment marker lock. This bit reflects the state of `am_lock[19]` (see 82.2.18.2.2).

45.2.3.16f.2 Lane 18 aligned (3.53.10)

When read as a one, bit 3.53.10 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 18. When read as a zero, bit 3.53.10 indicates that the PCS receiver lane 18 has not achieved alignment marker lock. This bit reflects the state of `am_lock[18]` (see 82.2.18.2.2).

45.2.3.16f.3 Lane 17 aligned (3.53.9)

When read as a one, bit 3.53.9 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 17. When read as a zero, bit 3.53.9 indicates that the PCS receiver lane 17 has not achieved alignment marker lock. This bit reflects the state of `am_lock[17]` (see 82.2.18.2.2).

45.2.3.16f.4 Lane 16 aligned (3.53.8)

When read as a one, bit 3.53.8 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 16. When read as a zero, bit 3.53.8 indicates that the PCS receiver lane 16 has not achieved alignment marker lock. This bit reflects the state of `am_lock[16]` (see 82.2.18.2.2).

Table 45–96f—Multi-lane BASE-R PCS alignment status 4 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.53.15:12	Reserved	Ignore when read	RO
3.53.11	Lane 19 aligned	1 = Lane 19 alignment marker is locked 0 = Lane 19 alignment marker is not locked	RO
3.53.10	Lane 18 aligned	1 = Lane 18 alignment marker is locked 0 = Lane 18 alignment marker is not locked	RO
3.53.9	Lane 17 aligned	1 = Lane 17 alignment marker is locked 0 = Lane 17 alignment marker is not locked	RO
3.53.8	Lane 16 aligned	1 = Lane 16 alignment marker is locked 0 = Lane 16 alignment marker is not locked	RO
3.53.7	Lane 15 aligned	1 = Lane 15 alignment marker is locked 0 = Lane 15 alignment marker is not locked	RO
3.53.6	Lane 14 aligned	1 = Lane 14 alignment marker is locked 0 = Lane 14 alignment marker is not locked	RO
3.53.5	Lane 13 aligned	1 = Lane 13 alignment marker is locked 0 = Lane 13 alignment marker is not locked	RO
3.53.4	Lane 12 aligned	1 = Lane 12 alignment marker is locked 0 = Lane 12 alignment marker is not locked	RO
3.53.3	Lane 11 aligned	1 = Lane 11 alignment marker is locked 0 = Lane 11 alignment marker is not locked	RO
3.53.2	Lane 10 aligned	1 = Lane 10 alignment marker is locked 0 = Lane 10 alignment marker is not locked	RO
3.53.1	Lane 9 aligned	1 = Lane 9 alignment marker is locked 0 = Lane 9 alignment marker is not locked	RO
3.53.0	Lane 8 aligned	1 = Lane 8 alignment marker is locked 0 = Lane 8 alignment marker is not locked	RO

^aRO = Read only

45.2.3.16f.5 Lane 15 aligned (3.53.7)

When read as a one, bit 3.53.7 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 15. When read as a zero, bit 3.53.7 indicates that the PCS receiver lane 15 has not achieved alignment marker lock. This bit reflects the state of `am_lock[15]` (see 82.2.18.2.2).

45.2.3.16f.6 Lane 14 aligned (3.53.6)

When read as a one, bit 3.53.6 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 14. When read as a zero, bit 3.53.6 indicates that the PCS receiver lane 14 has not achieved alignment marker lock. This bit reflects the state of `am_lock[14]` (see 82.2.18.2.2).

45.2.3.16f.7 Lane 13 aligned (3.53.5)

When read as a one, bit 3.53.5 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 13. When read as a zero, bit 3.53.5 indicates that the PCS receiver lane 13 has not achieved alignment marker lock. This bit reflects the state of `am_lock[13]` (see 82.2.18.2.2).

45.2.3.16f.8 Lane 12 aligned (3.53.4)

When read as a one, bit 3.53.4 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 12. When read as a zero, bit 3.53.4 indicates that the PCS receiver lane 12 has not achieved alignment marker lock. This bit reflects the state of `am_lock[12]` (see 82.2.18.2.2).

45.2.3.16f.9 Lane 11 aligned (3.53.3)

When read as a one, bit 3.53.3 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 11. When read as a zero, bit 3.53.3 indicates that the PCS receiver lane 11 has not achieved alignment marker lock. This bit reflects the state of `am_lock[11]` (see 82.2.18.2.2).

45.2.3.16f.10 Lane 10 aligned (3.53.2)

When read as a one, bit 3.53.2 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 10. When read as a zero, bit 3.53.2 indicates that the PCS receiver lane 10 has not achieved alignment marker lock. This bit reflects the state of `am_lock[10]` (see 82.2.18.2.2).

45.2.3.16f.11 Lane 9 aligned (3.53.1)

When read as a one, bit 3.53.1 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 9. When read as a zero, bit 3.53.1 indicates that the PCS receiver lane 9 has not achieved alignment marker lock. This bit reflects the state of `am_lock[9]` (see 82.2.18.2.2).

45.2.3.16f.12 Lane 8 aligned (3.53.0)

When read as a one, bit 3.53.8 indicates that the PCS receiver has achieved alignment marker lock for service interface lane 0. When read as a zero, bit 3.53.8 indicates that the PCS receiver lane 0 has not achieved alignment marker lock. This bit reflects the state of `am_lock[8]` (see 82.2.18.2.2).

Insert the following new subclauses after 45.2.3.35 (inserted by IEEE Std 802.3av) for BIP error counters and lane mapping:

45.2.3.36 BIP error counter lane 0 (Register 3.200)

The assignment of bits in the BIP error counter lane 0 is shown in Table 45–115a. The multi-lane PCS described in Clause 82 calculates a BIP value for each PCS lane (see 82.2.8, 82.2.14). Errors detected in PCS lane 0 are counted and shown in register 3.200.15:0. The 16-bit counter shall be reset to all zeros when register 3.200 is read or upon PCS reset. The 16-bit counter shall be held at all ones in the case of overflow.

Table 45–115a—BIP error counter, lane 0 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.200.15:0	BIP error counter, lane 0	Errors detected by BIP in PCS lane 0	RO

^aRO = Read only

45.2.3.37 BIP error counter, lanes 1 through 19 (Registers 3.201 through 3.219)

The behavior of the BIP error counters, lanes 1 through 19 is identical to that described for PCS lane 0 in 45.2.3.36. Errors detected in each PCS lane are counted and shown in register bits 15:0 in the corresponding register. PCS lane 1 is shown in register 3.201; PCS lane 2 is shown in register 3.202; through register 3.219 for PCS lane 19.

45.2.3.38 Lane 0 mapping register (Register 3.400)

The assignment of bits in the Lane 0 mapping register is shown in Table 45–115b. When the multi-lane PCS described in Clause 82 detects and locks the alignment marker for service interface lane 0, the detected PCS lane number is recorded in this register. The contents of the Lane 0 mapping register is valid when Lane 0 aligned bit (3.52.0) is set to one and is invalid otherwise.

Table 45–115b—Lane 0 mapping register bit definitions

Bit(s)	Name	Description	R/W ^a
3.400.15:6	Reserved	Ignore when read	RO
3.400.5:0	Lane 0 mapping	PCS lane received in service interface lane 0	RO

^aRO = Read only

45.2.3.39 Lanes 1 through 19 mapping registers (Registers 3.401 through 3.419)

The definition of lanes 1 through 19 mapping registers is identical to that described for lane 0 in 45.2.3.38. The lane mapping for lane 1 is in register 3.401; lane 2 is in register 3.402; etc.

45.2.7 Auto-Negotiation registers

Change the following row of Table 45–141 (as renumbered by IEEE Std 802.3av) for Backplane, BASE-R copper naming:

Table 45–141— Auto-Negotiation MMD registers

Register address	Register name
7.48	BP <u>Backplane Ethernet, BASE-R copper</u> status

Change 45.2.7.12 for Backplane Ethernet, BASE-R copper, and naming as follows:

45.2.7.12 Backplane Ethernet, BASE-R copper status (Register 7.48)

The assignment of bits in the Backplane Ethernet, BASE-R copper status register is shown in Table 45–150.

Table 45–150— Backplane Ethernet, BASE-R copper status register (Register 7.48) bit definitions

Bit(s)	Name	Description	RO ^a
7.48.15:4 9	Reserved	Ignore on read	RO
<u>7.48.8</u>	<u>100GBASE-CR10</u>	<u>1 = PMA/PMD is negotiated to perform 100GBASE-CR10</u> <u>0 = PMA/PMD is not negotiated to perform 100GBASE-CR10</u>	<u>RO</u>
<u>7.48.7</u>	<u>Reserved</u>	<u>Ignore on read</u>	<u>RO</u>
<u>7.48.6</u>	<u>40GBASE-CR4</u>	<u>1 = PMA/PMD is negotiated to perform 40GBASE-CR4</u> <u>0 = PMA/PMD is not negotiated to perform 40GBASE-CR4</u>	<u>RO</u>
<u>7.48.5</u>	<u>40GBASE-KR4</u>	<u>1 = PMA/PMD is negotiated to perform 40GBASE-KR4</u> <u>0 = PMA/PMD is not negotiated to perform 40GBASE-KR4</u>	<u>RO</u>
7.48.4	10GBASE-KR <u>BASE-R FEC</u> negotiated	1 = PMA/PMD is negotiated to perform 10GBASE-KR <u>BASE-R</u> FEC 0 = PMA/PMD is not negotiated to perform 10GBASE-KR <u>BASE-R</u> FEC	RO
7.48.3	10GBASE-KR	1 = PMA/PMD is negotiated to perform 10GBASE-KR 0 = PMA/PMD is not negotiated to perform 10GBASE-KR	RO
7.48.2	10GBASE-KX4	1 = PMA/PMD is negotiated to perform 10GBASE-KX4 <u>or CX4</u> 0 = PMA/PMD is not negotiated to perform 10GBASE-KX4 <u>CX4</u>	RO
7.48.1	1000BASE-KX	1 = PMA/PMD is negotiated to perform 1000BASE-KX 0 = PMA/PMD is not negotiated to perform 1000BASE-KX	RO
7.48.0	BP AN ability	If a 1000BASE-KX, 10GBASE-KX4 or 10GBASE-KR <u>Backplane, BASE-R copper</u> PHY type is implemented, this bit is set to 1	RO

^aRO = Read only

45.2.7.12.1 ~~10GBASE-KR~~ BASE-R FEC negotiated (7.48.4)

When the Auto-Negotiation process has completed as indicated by the AN complete bit (7.1.5), bit 7.48.4 indicates that ~~10GBASE-KR~~ BASE-R FEC operation has been negotiated. This bit is set only if ~~10GBASE-KR~~ a BASE-R PHY supporting FEC operation has also been negotiated.

45.2.7.12.2 Negotiated Port Type (7.48.1, 7.48.2, 7.48.3, ~~7.48.5, 7.48.6, 7.48.8~~)

When the AN process has been completed as indicated by the AN complete bit, ~~one of the three~~ these bits (1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, ~~40GBASE-KR4, 40GBASE-CR4, 100GBASE-CR10~~) indicates the negotiated port type. Only one of ~~the three~~ these bits is set depending on the priority resolution function. System developers need to distinguish between parallel detection of 10GBASE-KX4 and 10GBASE-CX4 based on the MDI and media type present.

45.2.7.12.3 Backplane Ethernet, BASE-R copper AN ability (7.48.0)

If a ~~1000BASE-KX, 10GBASE-KX4 or 10GBASE-KR~~ Backplane, BASE-R copper PHY type is implemented, this bit shall be set to 1.

When read as a one, bit 7.48.0 indicates that the PMA/PMD has the ability to perform Backplane Ethernet, BASE-R copper AN. When read as a zero, bit 7.48.0 indicates that the PMA/PMD lacks the ability to perform Backplane Ethernet, BASE-R copper AN.

45.5 Protocol implementation conformance statement (PICS) proforma for Clause 45, MDIO interface⁶

45.5.3.2 PMA/PMD MMD options

Change the table in 45.5.3.2 as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
*ALB	Implementation of PMA <u>local</u> loopback function	45.2.1.1.4		PMA:O	Yes [] No [] N/A []
*RLB	<u>Implementation of PMA remote loopback function</u>	45.2.1.1.3a		PMA:O	Yes [] No [] N/A []
*PLF	Implementation of fault detection	45.2.1.7		PMA:O	Yes [] No [] N/A []
*ALP	<u>Implementation of PMA/PMD low power ability</u>	45.2.1.2.3		PMA:O	Yes [] No [] N/A []
*PTD	Implementation of transmit disable function	45.2.1.8		PMA:O	Yes [] No [] N/A []
*10T	Implementation of the 10GBASE-T PMA	45.2.1.6		PMA:O	Yes [] No []
*40G	<u>Implementation of 40 Gb/s PMA/PMD</u>	45.2.1.4		PMA:O	Yes [] No []
*100G	<u>Implementation of 100 Gb/s PMA/PMD</u>	45.2.1.4		PMA:O	Yes [] No []
*10M	<u>Implementation of 10 Mb/s PMA/PMD</u>	45.2.1.4		PMA:O	Yes [] No []
*100M	<u>Implementation of 100 Mb/s PMA/PMD</u>	45.2.1.4		PMA:O	Yes [] No []
*1G	<u>Implementation of 1000 Mb/s PMA/PMD</u>	45.2.1.4		PMA:O	Yes [] No []
*10P	Implementation of the 10PASS-TS PMA/PMD	45.2.1.4		PMA:O	Yes [] No []
*2B	Implementation of the 2BASE-TL PMA/PMD	45.2.1.4		PMA:O	Yes [] No []
*10G	<u>Implementation of 10 Gb/s PMA/PMD</u>	45.2.1.4		PMA:O	Yes [] No []

⁶Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

Item	Feature	Subclause	Value/Comment	Status	Support
*KX	Implementation of 1000BASE-KX PMA/PMD	<u>45.2.1.6</u>		PMA:O	Yes [] No []
*KX4	Implementation of 10GBASE-KX4 PMA/PMD	<u>45.2.1.6</u>		PMA:O	Yes [] No []
*KR	Implementation of 10GBASE-KR PMA/PMD	<u>45.2.1.6</u>		PMA:O	Yes [] No []
*FEC-R	Implementation of 10GBASE-R FEC	45.2.1.85		PMA:O	Yes [] No []
*PXAR	<u>Implementation of PMA/PMD Extended Ability Register</u>	<u>45.2.1.10</u>	<u>Required for certain abilities</u>	<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
*40XAR	<u>Implementation of 40G/100G PMA/PMD Extended Ability Register</u>	<u>45.2.1.11a</u>	<u>Required for certain abilities</u>	<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
*MMD8	<u>Implementation of separately addressable instance, address 8</u>	<u>45.2.1</u>	<u>All register definitions as for address 1</u>	<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
*MMD9	<u>Implementation of separately addressable instance, address 9</u>	<u>45.2.1</u>	<u>All register definitions as for address 1</u>	<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
*MMD10	<u>Implementation of separately addressable instance, address 10</u>	<u>45.2.1</u>	<u>All register definitions as for address 1</u>	<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
*MMD11	<u>Implementation of separately addressable instance, address 11</u>	<u>45.2.1</u>	<u>All register definitions as for address 1</u>	<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>

45.5.3.3 PMA/PMD management functions

Change the table in 45.5.3.3 as follows, unchanged rows are not shown:

Item	Feature	Subclause	Value/Comment	Status	Support
MM14	Speed selection bits 13 and 6 are written as one	45.2.1.1.3		PMA:M	Yes [] N/A []
MM16	PMA is set into <u>local</u> loopback mode when bit 0 is set to a one	45.2.1.1.4		PMA*ALB:M	Yes [] N/A []
MM17	PMA transmit data is returned on receive path when in <u>local</u> loopback	45.2.1.1.4		PMA*ALB:M	Yes [] N/A []
MM18	PMA ignores writes to this bit if it does not support <u>local</u> loopback.	45.2.1.1.4		PMA*!ALB:M	Yes [] N/A []
MM19	PMA returns a value of zero when read if it does not support <u>local</u> loopback.	45.2.1.1.4		PMA*!ALB:M	Yes [] N/A []
<u>MM19a</u>	<u>PMA is set into remote loopback mode when bit 1 is set to a one</u>	<u>45.2.1.1.3a</u>		<u>RLB:M</u>	<u>Yes []</u> <u>N/A []</u>
<u>MM19b</u>	<u>PMA receive data is returned on transmit path when in remote loopback</u>	<u>45.2.1.1.3a</u>		<u>RLB:M</u>	<u>Yes []</u> <u>N/A []</u>
<u>MM19c</u>	<u>PMA ignores writes to this bit if it does not support remote loopback.</u>	<u>45.2.1.1.3a</u>		<u>PMA*!RLB:M</u>	<u>Yes []</u> <u>N/A []</u>
<u>MM19d</u>	<u>PMA returns a value of zero when read if it does not support remote loopback.</u>	<u>45.2.1.1.3a</u>		<u>PMA*!RLB:M</u>	<u>Yes []</u> <u>N/A []</u>
MM23	40G PMA/PMD type is selected using bits 53:0	45.2.1.6.1		PMA:M	Yes [] N/A []
MM24	40G PMA/PMD ignores writes to type selection bits that select types that it has not advertised	45.2.1.6.1		PMA:M	Yes [] N/A []
MM32	Single wavelength device ignores writes to bits 4—4 10:1 and returns a value of zero for them	45.2.1.8		PMA*PTD:M	Yes [] N/A []
MM33	Setting bit 4 to a one disables transmission on lane 3	45.2.1.8.1		PMA*PTD:M	Yes [] No [] N/A []
MM34	Setting bit 4 to a zero enables transmission on lane 3	45.2.1.8.1		PMA*PTD:M	Yes [] No [] N/A []
MM35	Setting bit 3 to a one disables transmission on lane 2	45.2.1.8.2		PMA*PTD:M	Yes [] No [] N/A []

Item	Feature	Subclause	Value/Comment	Status	Support
MM36	Setting bit 3 to a zero enables transmission on lane 2	45.2.1.8.2		PMA*PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM37	Setting bit 2 to a one disables transmission on lane 1	45.2.1.8.3		PMA*PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM38	Setting bit 2 to a zero enables transmission on lane 1	45.2.1.8.3		PMA*PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM39	Bits 1 to 10, set to 1 disables transmission on corresponding lane	45.2.1.8		PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM39	Setting bit 1 to a one disables transmission on lane 0	45.2.1.8.4		PMA*PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM40	Bits 1 to 10, set to 0 enables transmission on corresponding lane	45.2.1.8		PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM40	Setting bit 1 to a zero enables transmission on lane 0	45.2.1.8.4		PMA*PTD:M	Yes <input type="checkbox"/> No <input type="checkbox"/> N/A <input type="checkbox"/>
MM42a	Receive signal detect register behaves as described	45.2.1.9		PMA:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM45a	Writes to the 40G/100G PMA/PMD extended ability register have no effect	45.2.1.11a		PMA:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
*MM47a	PRBS pattern testing implemented	45.2.1.95	Ability indicated in register 1.1500	PMA:O	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM47b	Square wave testing implemented	45.2.1.95	Bit 1.1500.12 is set to one	PMA:O	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM47c	Counters corresponding to lanes that are not implemented return all zeros	45.2.1.98		MM47a:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM47d	Counters are reset to zero by read or PMA reset	45.2.1.98		MM47a:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM47d	Counters are held at all ones in the case of overflow	45.2.1.98		MM47a:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM114	FEC corrected blocks counters is are reset when read or upon PHY reset.	45.2.1.87, 45.2.1.89		FEC-R:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM115	FEC corrected blocks counters is are held at all ones in the case of overflow	45.2.1.87, 45.2.1.89		FEC-R:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM116	FEC uncorrected blocks counters is are reset when read or upon PHY reset.	45.2.1.88, 45.2.1.90		FEC-R:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>
MM117	FEC uncorrected blocks counters is are held at all ones in the case of overflow	45.2.1.88, 45.2.1.90		FEC-R:M	Yes <input type="checkbox"/> N/A <input type="checkbox"/>

45.5.3.6 PCS options

Change the table in 45.5.3.6 as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
*100CR	<u>Implementation of 100GBASE-R PCS</u>	<u>45.2.3.7</u>		<u>PCS:O</u>	<u>Yes []</u> <u>No []</u> <u>N/A []</u>
*40CR	<u>Implementation of 40GBASE-R PCS</u>	<u>45.2.3.7</u>		<u>PCS:O</u>	<u>Yes []</u> <u>No []</u> <u>N/A []</u>
*CR	Implementation of 10GBASE-R PCS	45.2.3.7		PCS:O	Yes [] No [] N/A []
*XCR	<u>Implementation of 10/40/100GBASE-R PCS</u>	<u>45.2.3.7</u>		<u>100CR:M</u> <u>40CR:M</u> <u>10CR:M</u>	<u>Yes []</u> <u>No []</u> <u>N/A []</u>
*CT	Implementation of the 10GBASE-T PCS	45.2.3.7		PCS:O	Yes [] No [] N/A []
*CX	Implementation of 10GBASE-X PCS	45.2.3.7		PCS:O	Yes [] No [] N/A []
*XP	Implementation of 10GBASE-X pattern testing	45.2.3		PCS *CX:O	Yes [] No [] N/A []
*PPT	Implementation of PRBS31 pattern testing	45.2.3		PCS:O	Yes [] No [] N/A []
*PTT	Implementation of PRBS9 pattern testing	45.2.3		PCS:O	Yes [] No [] N/A []
*EPC	Implementation of the 10BASE-TS/2BASE-TL PCS	45.2.3.17		PCS:O	Yes [] No [] N/A []
*PAF	Implementation of the PME aggregation function	45.2.3.17		PCS *EPC:O	Yes [] No [] N/A []

45.5.3.7 PCS management functions

Change the table in 45.5.3.7 as follows, unchanged rows are not shown:

Item	Feature	Subclause	Value/Comment	Status	Support
RM35	Writes to 40GBASE-R and <u>10GBASE-T</u> PCS status 1 register have no effect	45.2.3.11		PCS*CR:M PCS*CT:M	Yes [] N/A []
RM36	Reads from 40GBASE-R and <u>10GBASE-T</u> PCS status 1 register return zero for PCS that does not support 10GBASE-R	45.2.3.11		PCS*CR:M PCS*CT:M	Yes [] N/A []
RM37	Writes to 40GBASE-R and <u>10GBASE-T</u> PCS status 2 register have no effect	45.2.3.12		PCS*XCR:M PCS*CT:M	Yes [] N/A []
RM38	Reads from 40GBASE-R and <u>10GBASE-T</u> PCS status 2 register return zero for PCS that does not support 40GBASE-R <u>10/40/100GBASE-R</u> or <u>10GBASE-T</u>	45.2.3.12		PCS*XCR:M PCS*CT:M	Yes [] N/A []
RM39	Latched block lock implemented with latching low behavior	45.2.3.12.1		PCS*XCR:M PCS*CT:M	Yes [] N/A []
RM40	Latched high BER implemented with latching high behavior	45.2.3.12.2		PCS*XCR:M PCS*CT:M	Yes [] N/A []
RM41	BER counter clears to <u>resets to</u> <u>all zeros</u> on read or reset	45.2.3.12.3		PCS*XCR:M PCS*CT:M	Yes [] N/A []
RM42	BER counter holds at all ones at overflow	45.2.3.12.3		PCS*CR:M PCS*CT:M <u>!RM50a:M</u>	Yes [] N/A []
RM43	Errored blocks counter implemented as a non roll over counter	45.2.3.12.4		PCS*CR:M PCS*CT:M <u>!RM50f:M</u>	Yes [] N/A []
RM44	Errored blocks counter clears to <u>resets to</u> <u>all zeros</u> on read	45.2.3.12.4		PCS*XCR:M PCS*CT:M	Yes [] N/A []
RM49	Test-pattern error counter clears to <u>resets to</u> <u>all zeros</u> on read or reset	45.2.3.16		PCS*XCR:M	Yes [] N/A []
RM50	Test-pattern error counter holds at all ones at overflow	45.2.3.16		PCS*XCR:M	Yes [] N/A []
<u>*RM50a</u>	<u>BER high order counter implemented</u>	<u>45.2.3.16a</u>		<u>CR:O</u> <u>40CR:M</u> <u>100CR:M</u>	<u>Yes []</u> <u>N/A []</u>
<u>RM50c</u>	<u>High order bits latched on read to 3.33</u>	<u>45.2.3.16a</u>		<u>RM50a:M</u>	<u>Yes []</u> <u>N/A []</u>
<u>RM50d</u>	<u>Counter reset on read to 3.33 or PCS reset</u>	<u>45.2.3.16a</u>		<u>RM50a:M</u>	<u>Yes []</u> <u>N/A []</u>
<u>RM50e</u>	<u>Counter held at all ones in the case of an overflow</u>	<u>45.2.3.16a</u>		<u>RM50a:M</u>	<u>Yes []</u> <u>N/A []</u>

Item	Feature	Subclause	Value/Comment	Status	Support
*RM50f	<u>Errored blocks high order counter implemented</u>	<u>45.2.3.16b</u>		<u>CR:O</u> <u>40CR:M</u> <u>100CR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM50g	<u>Register bit 3.45.15 set to 1</u>	<u>45.2.3.16b</u>		<u>RM50f:M</u>	<u>Yes []</u> <u>N/A []</u>
RM50h	<u>High order bits latched on read to 3.33</u>	<u>45.2.3.16b</u>		<u>RM50f:M</u>	<u>Yes []</u> <u>N/A []</u>
RM50i	<u>Counter reset on read to 3.33 or PCS reset</u>	<u>45.2.3.16b</u>		<u>RM50f:M</u>	<u>Yes []</u> <u>N/A []</u>
RM50j	<u>Counter held at all ones in the case of an overflow</u>	<u>45.2.3.16b</u>		<u>RM50f:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52a	<u>Writes to multi-lane BASE-R PCS alignment status 1 register have no effect</u>	<u>45.2.3.16c</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52b	<u>Non multi-lane BASE-R device shall return all zeros</u>	<u>45.2.3.16c</u>		<u>!XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52c	<u>Device shall return a zero for lanes not required</u>	<u>45.2.3.16c</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52c	<u>Writes to multi-lane BASE-R PCS alignment status 2 register have no effect</u>	<u>45.2.3.16d</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52d	<u>Non multi-PCS lane BASE-R device shall return all zeros</u>	<u>45.2.3.16d</u>		<u>!XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52e	<u>Device shall return a zero for lanes not required</u>	<u>45.2.3.16d</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52f	<u>Writes to multi-lane BASE-R PCS alignment status 3 register have no effect</u>	<u>45.2.3.16e</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52g	<u>Non multi-PCS lane BASE-R device shall return all zeros</u>	<u>45.2.3.16e</u>		<u>!XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52h	<u>Device shall return a zero for lanes not required</u>	<u>45.2.3.16e</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52i	<u>Writes to multi-lane BASE-R PCS alignment status 4 register have no effect</u>	<u>45.2.3.16f</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52j	<u>Non multi-PCS lane BASE-R device shall return all zeros</u>	<u>45.2.3.16f</u>		<u>!XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52k	<u>Device shall return a zero for lanes not required</u>	<u>45.2.3.16f</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52m	<u>Counters reset on read to 3.200 through 3.219 or PCS reset</u>	<u>45.2.3.36</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>
RM52n	<u>Counters held at all ones in the case of an overflow</u>	<u>45.2.3.36</u>		<u>XCR:M</u>	<u>Yes []</u> <u>N/A []</u>

52. Physical Medium Dependent (PMD) sublayer and baseband medium, type 10GBASE-S (short wavelength serial), 10GBASE-L (long wavelength serial), and 10GBASE-E (extra long wavelength serial)

52.9.10 Transmitter and dispersion penalty measurement

Change last paragraph of 52.9.10 as follows (note that there is no change to Figure 52-12):

The transmitter and dispersion penalty (TDP) measurement tests for transmitter impairments with chromatic effects for a transmitter to be used with single-mode fiber, and for transmitter impairments with modal (not chromatic) dispersion effects for a transmitter to be used with multimode fiber.

The setup for measurement of transmitter and dispersion penalty is shown in Figure 52–12 and consists of a reference transmitter, the transmitter under test, a controlled optical reflection, an optical attenuator, a test fiber, a reference receiver, a transversal filter for 10GBASE-S, and a bit-error rate tester. For 10GBASE-S, the polarization rotator shown in Figure 52–12 is removed from the setup and the single-mode fiber replaced with a multimode fiber. All BER and sensitivity measurements are made with the test patterns in 52.9.1.

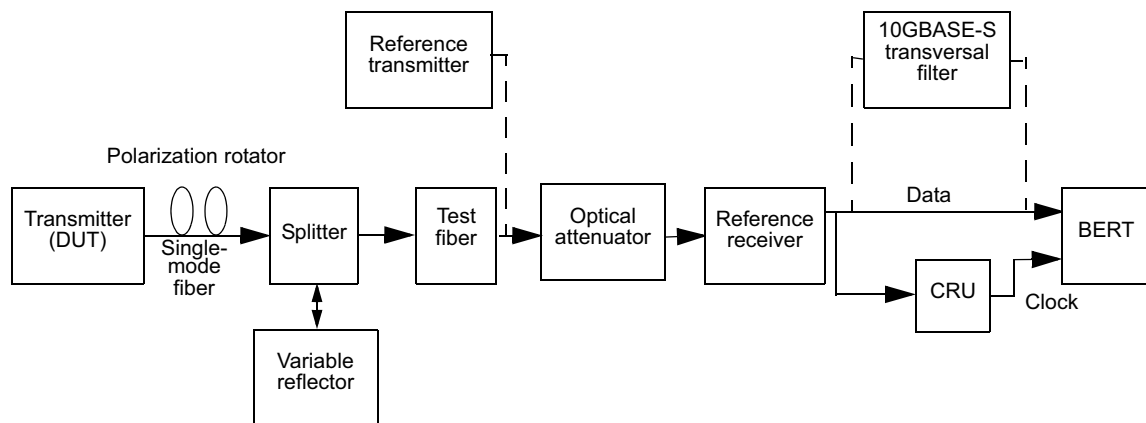


Figure 52–12—Test setup for measurement of transmitter and dispersion penalty

69. Introduction to Ethernet operation over electrical backplanes

69.1 Overview

69.1.1 Scope

Change the second paragraph as follows:

Backplane Ethernet supports the IEEE 802.3 MAC operating at 1000 Mb/s, ~~or~~ 10 Gb/s, or 40 Gb/s. For 1000 Mb/s operation, the family of 1000BASE-X Physical Layer signaling systems is extended to include 1000BASE-KX. For 10 Gb/s operation, two Physical Layer signaling systems are defined. For operation over four logical lanes, the 10GBASE-X family is extended to include 10GBASE-KX4. For serial operation, the 10GBASE-R family is extended to include 10GBASE-KR. For 40 Gb/s operation, there is 40GBASE-KR4 that operates over four lanes.

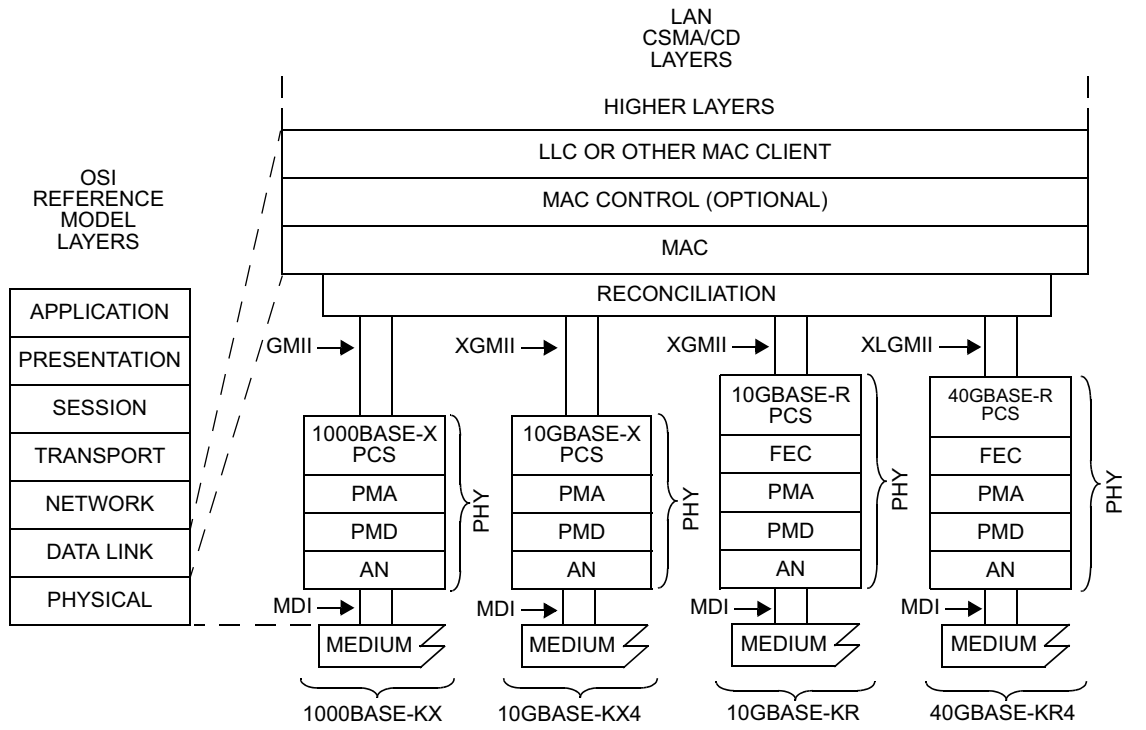
69.1.2 Objectives

Change item d) as follows:

- d) Support operation of the following PHY over differential, controlled impedance traces on a printed circuit board with 2 connectors and total length up to at least 1 m consistent with the guidelines of Annex 69B.
 - i) a 1 Gb/s PHY
 - ii) a 4-lane 10 Gb/s PHY
 - iii) a single-lane 10 Gb/s PHY
 - iv) a 4-lane 40 Gb/s PHY

69.1.3 Relationship of Backplane Ethernet to the ISO OSI reference model

Replace Figure 69-1 with the following:



AN = AUTO-NEGOTIATION
 FEC = FORWARD ERROR CORRECTION
 GMII = GIGABIT MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE
 XLGMII = 40 GIGABIT MEDIA INDEPENDENT INTERFACE

Figure 69–1—Architectural positioning of Backplane Ethernet

Change item f) as follows:

- f) The MDI as specified in Clause 70 for 1000BASE-KX, Clause 71 for 10GBASE-KX4, ~~or~~ Clause 72 for 10GBASE-KR, or Clause 84 for 40GBASE-KR4.

69.2 Summary of Backplane Ethernet Sublayers

Change 69.2.1 as follows:

69.2.1 Reconciliation sSublayer and media independent interfaces

The Clause 35 RS and GMII, ~~and~~ the Clause 46 RS and XGMII, ~~and the Clause 81 RS and XLGMII~~ are ~~both~~ employed for the same purpose in Backplane Ethernet, that being the interconnection between the MAC sublayer and the PHY.

69.2.3 Physical Layer signaling systems

Change first sentence of third paragraph as follows:

Finally, Backplane Ethernet also extends the family of 10GBASE-R Physical Layer signaling systems to include the 10GBASE-KR.

Insert a new paragraph after the third paragraph of the subclause as follows:

Backplane Ethernet also specifies 40GBASE-KR4. This embodiment employs the PCS defined in Clause 82, the PMA defined in Clause 83, and the PMD defined in Clause 84 and specifies 40 Gb/s operation over four differential paths in each direction for a total of eight pairs.

Replace Table 69-1 with the following:

Table 69–1—Nomenclature and clause correlation

Nomenclature	Clause																	
	35	36	46		48	49	51	70	71	72	73	74	81		82	83	83A	84
	RS/GMI	1000BASE-X PCS/PMA	RS	XGMII	10GBASE-X PCS/PMA	10GBASE-R PCS	Serial PMA	1000BASE-KX PMD	10GBASE-KX4 PMD	10GBASE-KR PMD	Auto-Negotiation	BASE-R FEC	RS	XLGMII	40GBASE-R PCS	40GBASE-R PMA	XLAUI	40GBASE-KR4 PMD
1000BASE-KX	O ^a	M ^a	M					M			M							
10GBASE-KX4				O	M				M		M							
10GBASE-KR						M	M			M	M	O						
40GBASE-KR4											M	O	M	O	M	M	O	M

^aO = Optional, M = Mandatory

69.2.5 Management

Change the paragraph as follows:

Managed objects, attributes, and actions are defined for all Backplane Ethernet components. Clause 30 consolidates all IEEE 802.3 management specifications so that ~~10 Mb/s, 100 Mb/s, 1000 Mb/s and 10 Gb/s~~ agents of different speed can be managed by existing network management stations with little or no modification to the agent code.

69.3 Delay constraints

Insert a new paragraph at the end of the subclause as follows:

For 40GBASE-KR4, normative delay specifications may be found in 81.1.4, 82.5, 83.5.4, and 84.4 and also referenced in 80.4.

Change the title of Clause 73 as follows:

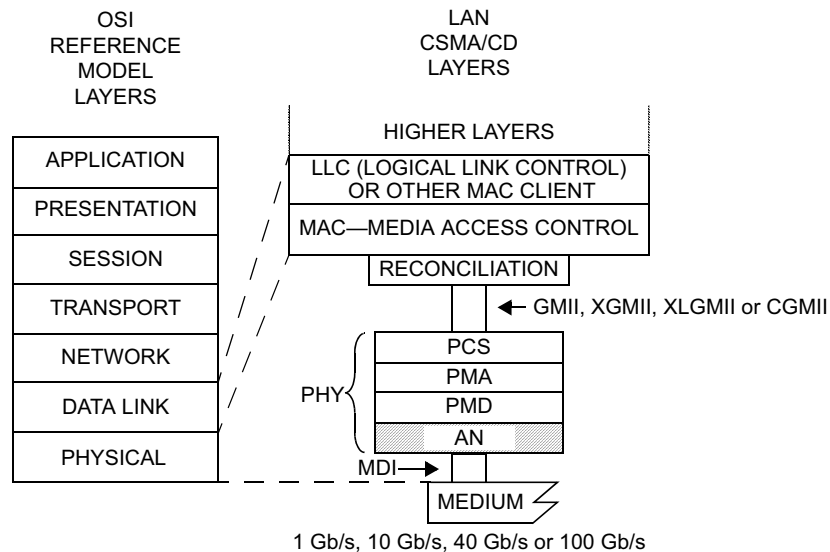
73. Auto-Negotiation for ~~B~~ackplane Ethernet and copper cable assembly

Insert the following text under the clause title:

Note that although the Auto-Negotiation defined in this clause was originally intended for use with Backplane Ethernet PHYs, it is also specified for use with 40GBASE-CR4 and 100GBASE-CR10 PHYs.

73.2 Relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model

Replace Figure 73-1 with the following:



- AN = AUTO-NEGOTIATION
- CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
- GMII = GIGABIT MEDIA INDEPENDENT INTERFACE
- MDI = MEDIUM DEPENDENT INTERFACE
- PCS = PHYSICAL CODING SUBLAYER
- PHY = PHYSICAL LAYER DEVICE
- PMA = PHYSICAL MEDIUM ATTACHMENT
- PMD = PHYSICAL MEDIUM DEPENDENT
- XGMII = 10 Gb/s MEDIA INDEPENDENT INTERFACE
- XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 73–1—Location of Auto-Negotiation function within the ISO/IEC OSI reference model

73.3 Functional specifications

Change third and fourth paragraphs as follows:

These functions shall comply with the state diagrams from Figure 73-9 through Figure 73-11. The Auto-Negotiation functions shall interact with the technology-dependent PHYs through the Technology-Dependent interface (see 73.9). Technology-Dependent PHYs include 1000BASE-KX, 10GBASE-KX4, and 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10.

When the MDI supports multiple lanes (e.g., for operation of 10GBASE-KX4), then lane 0 of the MDI shall be used for Auto-Negotiation and for connection of any single-lane PHYs (e.g., 1000BASE-KX or 10GBASE-KR).

Change 73.5.1 so it becomes electrical specifications as follows and delete existing 73.5.1.1 while preserving Table 73-1:

73.5.1 DME ~~page encoding~~ electrical specifications

DME pages can be transmitted by local devices capable of operating in 1 Gb/s (1000BASE-KX) mode, 10 Gb/s over 4 lane (10GBASE-KX4) mode or 10 Gb/s over 1 lane (10GBASE-KR) mode.

Transmitter characteristics shall meet the specifications in Table 73-1 at TP1 while transmitting DME pages. Receiver characteristics shall meet the specifications in Table 73-1 at TP4 while receiving DME pages.

Table 73-1—DME electrical characteristics

Parameter	Value	Units
Transmit differential peak-to-peak output voltage	600 to 1200	mV
Receive differential peak-to-peak input voltage	200 to 1200	mV

For any multi-lane PHY, DME pages shall be transmitted only on lane 0. The transmitters on other lanes should be disabled as specified in 71.6.7, 84.7.7, or 85.7.7.

~~73.5.1.1 DME electrical specifications~~

~~Transmitter characteristics shall meet the specifications in Table 73-1 at TP1 while transmitting DME pages. Receiver characteristics shall meet the specifications in Table 73-1 at TP4 while receiving DME pages~~

~~When the PHY has 10GBASE-KX4 capability, DME pages shall be transmitted only on lane 0. The lane 1 to lane 3 transmitters should be disabled as specified in 71.6.7.~~

73.6.4 Technology Ability Field

Change Table 73-4 as follows:

Table 73—4—Technology Ability Field encoding

Bit	Technology
A0	1000BASE-KX
A1	10GBASE-KX4
A2	10GBASE-KR
<u>A3</u>	<u>40GBASE-KR4</u>
<u>A4</u>	<u>40GBASE-CR4</u>
<u>A5</u>	<u>100GBASE-CR10</u>
A3 <u>A6</u> through A24	Reserved for future technology

Insert a new paragraph as second to last paragraph:

40GBASE-CR4 and 40GBASE-KR4 shall not be advertised simultaneously as their physical interfaces are different.

Change last paragraph as follows:

The fields A[24:3 6] are reserved for future use. Reserved fields shall be sent as zero and ignored on receive.

73.6.5 FEC capability

Change text as follows:

FEC (F0:F1) is encoded in bits D46:D47 of the base link codeword. The two FEC bits are used as follows:

- a) F0 is FEC ability.
- b) F1 is FEC requested.

When the FEC ability bit is set to logical one, it indicates that the ~~10GBASE-KR~~ PHY has FEC ability (see Clause 74). When the FEC requested bit is set to logical one, it indicates a request to enable FEC on the link.

Since the local device and the link partner may have set the FEC capability bits differently ~~and this FEC capability is only used with 10GBASE-KR~~, the priority resolution function is used to enable FEC in the respective PHYs. The FEC function shall be enabled on the link if 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4 or 100GBASE-CR10 is the HCD technology (see 73.7.6), both devices advertise FEC ability on the F0 bits, and at least one device requests FEC on the F1 bits. ~~If 10GBASE-KR is not the HCD technology, FEC shall not be enabled. If either device does not have FEC ability, FEC shall not be enabled. If neither device requests FEC, FEC shall not be enabled even if both devices have FEC ability; otherwise FEC shall not be enabled.~~

73.7 Receive function requirements

Change text as follows:

The Receive function detects the DME page sequence, decodes the information contained within, and stores the data in rx_link_code_word[48:1]. The receive function incorporates a receive switch to control connection to the 1000BASE-KX, 10GBASE-KX4, ~~or 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, or 100GBASE-CR10~~ PHYs.

73.7.1 DME page reception

Change text as follows:

To be able to detect the DME bits, the receiver should have the capability to receive DME signals sent with the electrical specifications of ~~any IEEE 802.3 Backplane Ethernet~~ the PHY (1000BASE-KX, 10GBASE-KX4, ~~or 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, or 100GBASE-CR10~~). The DME transmit signal level and receive sensitivity are specified in 73.5.1.1.

73.7.2 Receive Switch function

Change text as follows:

The Receive Switch function shall enable the receive path from the MDI to a single technology-dependent PHY once a highest common denominator choice has been made and Auto-Negotiation has completed.

During Auto-Negotiation, the Receive Switch function shall connect the DME page receiver controlled by the Receive state diagram to the MDI and the Receive Switch function shall also connect the 1000BASE-KX, 10GBASE-KX4, ~~and 10GBASE-KR, 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10~~ PMA receivers to the MDI if the PMAs are present.

73.7.4.1 Parallel Detection function

Change second paragraph as follows:

A local device shall provide Parallel Detection for 1000BASE-KX and 10GBASE-KX4 if it supports those PHYs. ~~Parallel Detection is not performed for 10GBASE-KR. Additionally, parallel detection may be used for 10GBASE-CX4. Parallel detection of 10GBASE-CX4 will be indicated by the setting of the Negotiated Port Type to “10GBASE-KX4 or 10GBASE-CX4” in the management register bit 7.48.2. The means to distinguish between 10GBASE-KX4 and 10GBASE-CX4 is implementation dependent.~~ Parallel Detection shall be performed by directing the MDI receive activity to the PHY. This detection may be done in sequence between detection of DME pages and detection of each supported PHY. If at least one of the 1000BASE-KX, or 10GBASE-KX4 establishes link_status=OK, the LINK STATUS CHECK state is entered and the autoneg_wait_timer is started. If exactly one link_status=OK indication is present when the autoneg_wait_timer expires, then Auto-Negotiation shall set link_control=ENABLE for the PHY indicating link_status=OK. If a PHY is enabled, the Arbitration function shall set link_control=DISABLE to all other PHYs and indicate that Auto-Negotiation has completed. On transition to the AN GOOD CHECK state from the LINK STATUS CHECK state, the Parallel Detection function shall set the bit in the AN LP base page ability registers (see 45.2.7.7) corresponding to the technology detected by the Parallel Detection function.

73.7.6 Priority Resolution function

Change Table 73-5 as follows:

Table 73–5—Priority Resolution

Priority	Technology	Capability
<u>1</u>	<u>100GBASE-CR10</u>	<u>100 Gb/s 10 lane, highest priority</u>
<u>2</u>	<u>40GBASE-CR4</u>	<u>40 Gb/s 4 lane</u>
<u>3</u>	<u>40GBASE-KR4</u>	<u>40 Gb/s 4 lane</u>
4	10GBASE-KR	10 Gb/s 1 lane, highest priority
5	10GBASE-KX4	10 Gb/s 4 lane, second highest priority
6	1000BASE-KX	1 Gb/s 1 lane, third highest lowest priority

73.10.1 State diagram variables

Change second paragraph as follows:

A variable with “_[x]” appended to the end of the variable name indicates a variable or set of variables as defined by “x”. “x” may be as follows:

- all; represents all specific technology-dependent PMAs supported in the local device.
- 1GKX; represents that the 1000BASE-KX PMA is the signal source.
- 10GKR; represents that the 10GBASE-KR PMA is the signal source.
- 10GKX4; represents that the 10GBASE-KX4 or 10GBASE-CX4 PMA is the signal source.
- 40GKR4; represents that the 40GBASE-KR4 PMD is the signal source.
- 40GCR4; represents that the 40GBASE-CR4 PMD is the signal source.
- 100GCR10; represents that the 100GBASE-CR10 PMD is the signal source.
- HCD; represents the single technology-dependent PMA chosen by Auto-Negotiation as the highest common denominator technology through the Priority Resolution or parallel detection function.
- notHCD; represents all technology-dependent PMAs not chosen by Auto-Negotiation as the highest common denominator technology through the Priority Resolution or parallel detection function.
- PD; represents all of the following that are present: 1000BASE-KX PMA, 10GBASE-KX4 PMA or 10GBASE-CX4 PMA, and 10GBASE-KR PMA, 40GBASE-KR4 PMD, 40GBASE-CR4 PMD, and 100GBASE-CR10 PMD.

Change first sentence of `mr_parallel_detection_fault` description as follows:

`mr_parallel_detection_fault`

Error condition indicating that while performing parallel detection, either `DMEan_receive_idle` = false, or zero or more than one of the following indications were present when the `autoneg_wait_timer` expired.

Change `single_link_ready` as follows:

`single_link_ready`

Status indicating that `DMEan_receive_idle` = true and only one the of the following indications is being received:

- 1) `link_status_[1GKX]` = OK
- 2) `link_status_[10GKX4]` = OK
- 3) `link_status_[10GKR]` = OK
- 4) `link_status_[40GKR4]` = OK
- 5) `link_status_[40GCR4]` = OK
- 6) `link_status_[100GCR10]` = OK

Values: false; either zero or more than one of the above ~~three~~ indications are true or `an_receive_idle` = false.
true; Exactly one of the above ~~three~~ indications is true and `an_receive_idle` = true.

NOTE—This variable is set by this variable definition; it is not set explicitly in the state diagrams.

73.10.2 State diagram timers***Change `autoneg_wait_timer` as follows:***

`autoneg_wait_timer`

Timer for the amount of time to wait before evaluating the number of link integrity test functions with `link_status=OK` asserted. The `autoneg_wait_timer` shall expire 25 ms to 50 ms from the assertion of `link_status=OK` from the ~~1000BASE-KX PCS, 10GBASE-KX4 PCS, or 10GBASE-KR PCS.~~

Change `link_fail_inhibit_timer` as follows:

`link_fail_inhibit_timer`

Timer for qualifying a `link_status=FAIL` indication or a `link_status=OK` indication when a specific technology link is first being established. A link will only be considered “failed” if the `link_fail_inhibit_timer` has expired and the link has still not gone into the `link_status=OK` state. The `link_fail_inhibit_timer` shall expire 40 ms to 50 ms after entering the AN LINK GOOD CHECK state when the link is ~~not 10GBASE-KR~~ 1000BASE-KX or 10GBASE-KX4. ~~Otherwise the~~ The `link_fail_inhibit_timer` shall expire 500 ms to 510 ms after entering the AN LINK GOOD CHECK state ~~when the link is 10GBASE-KR.~~

Change Table 73-7 as follows:

Table 73–7—Timer min/max value summary

Parameter	Min	Value and tolerance	Max	Units
autoneg_wait_timer	25		50	ms
break_link_timer	60		75	ms
clock_detect_min_timer	4.8		6.2	ns
clock_detect_max_timer	6.6		8.0	ns
data_detect_min_timer	1.6		3.0	ns
data_detect_max_timer	3.4		4.8	ns
interval_timer		3.2 ± 0.01%		ns
link_fail_inhibit_timer (when the link is 10GBASE-KR neither <u>1000BASE-KX</u> nor <u>10GBASE-KX4</u>)	500		510	ms
link_fail_inhibit_timer (when the link is not 10GBASE-KR- <u>1000BASE-KX</u> or <u>10GBASE-KX4</u>)	40		50	ms
page_test_min_timer	305		330	ns
page_test_max_timer	350		375	ns

Change title of 73.11 as follows:

73.11 Protocol implementation conformance statement (PICS) proforma for Clause 73, Auto-Negotiation for ~~B~~ackplane ~~E~~thernet and copper cable assembly⁷

73.11.1 Introduction

Change first sentence as follows:

The supplier of a protocol implementation that is claimed to conform to IEEE Std 802.3-2008 as modified by IEEE Std 802.3ba-2010, Clause 73, Auto-Negotiation for ~~B~~ackplane ~~E~~thernet and copper cable assembly, shall complete the following protocol implementation conformance statement (PICS) proforma.

73.11.2.2 Protocol summary

Change Identification of protocol standard as follows:

IEEE Std 802.3-2008 as modified by IEEE Std 802.3ba-2010, Clause 73, Auto-Negotiation for ~~B~~ackplane ~~E~~thernet and copper cable assembly

Change the title of 73.11.4 as follows:

73.11.4 PICS proforma tables for Auto-Negotiation for ~~B~~ackplane ~~E~~thernet and copper cable assembly

Change DT2 and DT3 as follows:

73.11.4.2 DME transmission

Item	Feature	Subclause	Value/Comment	Status	Support
DT2	10GBASE-KX4 Multi-lane DME pages	73.5.1.1 73.5.1	10GBASE-KX4 Multi-lane DME pages shall be transmitted on Lane 0	M	Yes []
DT3	DME electrical characteristics	73.5.1.1 73.5.1	Meet requirements of Table 73-1	M	Yes []

⁷Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

Insert LE17 as follows:

73.11.4.3 Link codeword encoding

Item	Feature	Subclause	Value/Comment	Status	Support
<u>LE17</u>	<u>Incompatible abilities</u>	<u>73.6.4</u>	<u>40GBASE-CR4 and 40GBASE-KR4 shall not be advertised simultaneously</u>	<u>M</u>	<u>Yes []</u>

Change SD15 as follows:

73.11.4.7 State diagrams and variable definitions

Item	Feature	Subclause	Value/Comment	Status	Support
SD15	link_fail_inhibit_timer	73.10.2	500 to 510 ms when the link is not 1000BASE-KX or 10GBASE-KX4 40GBASE-KR and 40 ms to 50 ms otherwise when the link is not 10GBASE-KR	M	Yes []

74. Forward Error Correction (FEC) sublayer for ~~40~~GBASE-R PHYs

Change 74.1 as follows:

74.1 Overview

This clause specifies an optional Forward Error Correction (FEC) sublayer for 10GBASE-R and other BASE-R PHYs. The FEC sublayer can be placed in between the PCS and PMA sublayers of the 10GBASE-R and other BASE-R Physical Layer implementations as shown in Figure 74-2, Figure 74-2a, and Figure 74-2b. For a PHY with a multi-lane BASE-R PCS, the FEC sublayer is instantiated for each PCS lane and operates autonomously on a per PCS lane basis. The FEC provides coding gain to increase the link budget and BER performance. The 10GBASE-KR and 40GBASE-KR4 PHYs described in Clause 72 and Clause 84 optionally uses the FEC sublayer to increase the performance on a broader set of backplane channels as than are defined in Clause 69. The FEC sublayer provides additional margin to account for variations in manufacturing and environmental conditions. The 40GBASE-CR4 and 100GBASE-CR10 PHYs described in Clause 85 optionally use the FEC sublayer to improve the BER performance beyond 10^{-12} .

Change 74.2 as follows:

74.2 Objectives

The following are the objectives for the FEC:

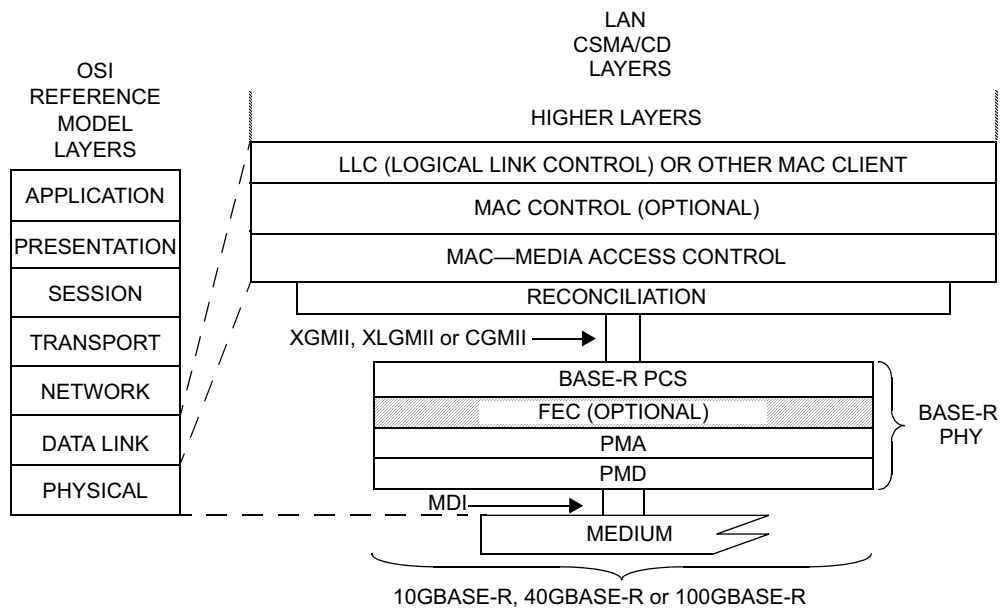
- a) To support forward error correction mechanism for ~~40~~GBASE-R PHYs.
- b) To support the full duplex mode of operation of the Ethernet MAC.
- c) To support the PCS, PMA, and PMD sublayers defined for 10GBASE-R, 40GBASE-R, and 100GBASE-R.
- d) To provide a ~~10.3125 Gb/s effective data rate~~ GBd on each BASE-R PCS lane at the service interface ~~presented by the PMA sublayer, for 10G and 40G; and 5.15625 GBd for 100G~~.
- e) To support operations over links consistent with differential, controlled impedance traces on a printed circuit board with two connectors and total length up to at least 1 m meeting the guidelines of Annex 69B.
- f) To support a BER objective of 10^{-12} or better.

Change 74.3 as follows:

74.3 Relationship to other sublayers

Figure 74-1 depicts the relationships among the ~~40~~GBASE-R FEC (shown shaded), the ~~10~~ Gb/s MAC and Reconciliation Sublayers, the ~~40~~GBASE-R PCS, PMA and PMD, the ISO/IEC 8802-2 LLC, and the ISO/IEC Open System Interconnection (OSI) reference model.

Replace Figure 74-1 with the following:



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE PMA = PHYSICAL MEDIUM ATTACHMENT
 FEC = FORWARD ERROR CORRECTION PMD = PHYSICAL MEDIUM DEPENDENT
 MDI = MEDIUM DEPENDENT INTERFACE XGMII = 10 Gb/s MEDIA INDEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
 PHY = PHYSICAL LAYER DEVICE

Figure 74-1—BASE-R FEC relationship to ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 CSMA/CD LAN model

Change 74.4 as follows:

74.4 Inter-sublayer interfaces

An FEC service interface is provided to allow the FEC sublayer to transfer information to and from the 10GBASE-R PCS, which is the sole FEC client. An abstract service model is used to define the operation of this interface. For 10GBASE-R, the FEC service interface directly maps to the PMA service interface of the 10GBASE-R PCS defined in Clause 49 and, in addition, the lower FEC sublayer interface maps to utilizes the service interface provided by the serial PMA sublayer defined in Clause 51 to transfer information to and from the PMA. For 40GBASE-R and 100GBASE-R, the FEC service interface is an instance of the inter-sublayer service interface defined in 80.3 as is the PMA service interface defined in 83.2.

For 40GBASE-R and 100GBASE-R the FEC service interface can either connect to the PCS as illustrated in Figure 74-1 or the PMA as illustrated in Figure 83-2 where the FEC and PCS are in separate devices connected by XLAUI/CAUI.

This standard defines these interfaces in terms of bits, octets, data-group, data units, and signals; however, implementors may choose other data-path widths and other control mechanisms for implementation convenience, provided that the implementation adheres to the logical model of the service interface.

Change the title of 74.4.1 to be explicitly for 10GBASE-R as follows:

74.4.1 Functional block diagram for 10GBASE-R PHYs

Replace the title of Figure 74–2 with the following:

Figure 74–2—Functional block diagram for 10GBASE-R PHYs

Insert a new subclause 74.4.2 to show the block diagram for 40GBASE-R as follows:

74.4.2 Functional block diagram for 40GBASE-R PHYs

Figure 74–2a shows the functional block diagram of FEC for 40GBASE-R PHYs and the relationship between the PCS and PMA sublayers.

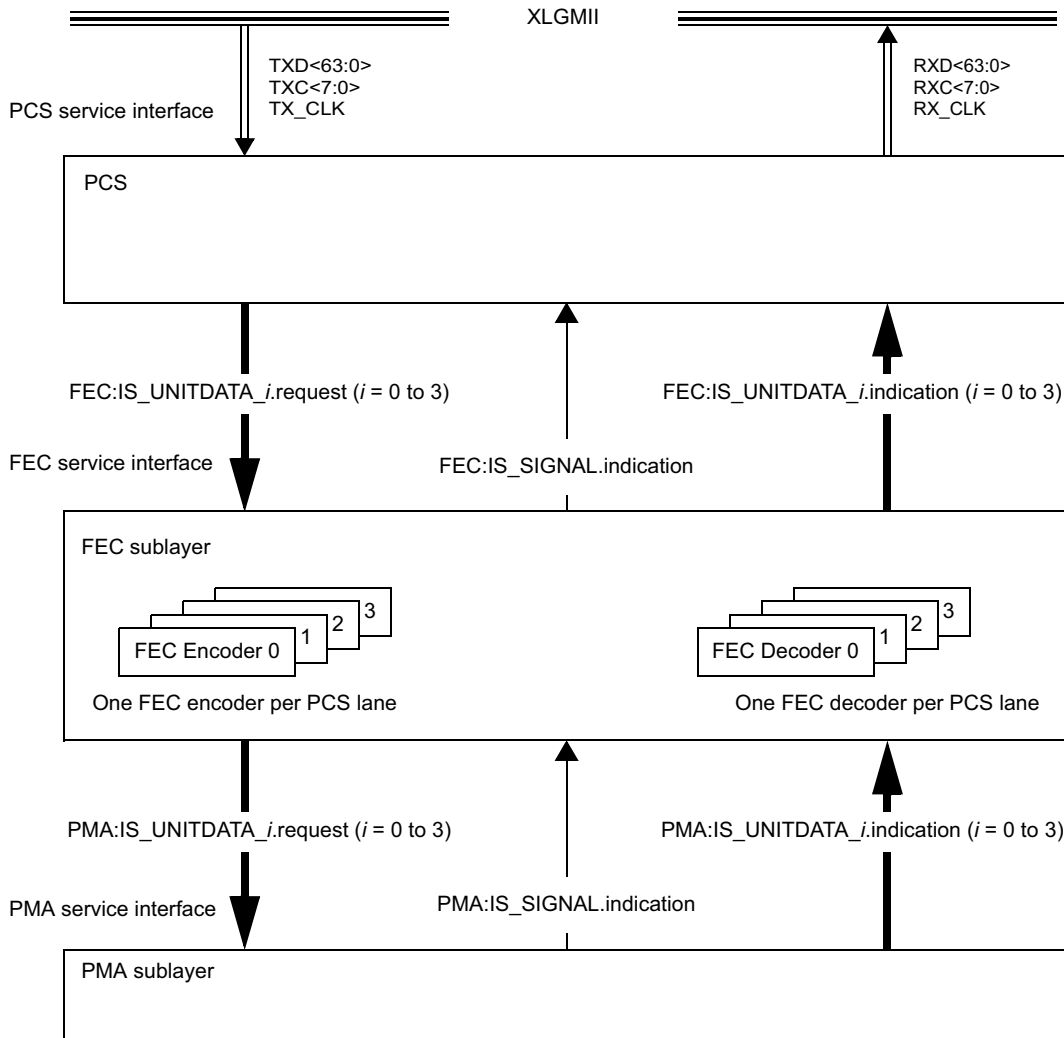


Figure 74–2a—Functional block diagram for 40GBASE-R PHY

Insert a new subclause 74.4.3 to show the block diagram for 100GBASE-R as follows:

74.4.3 Functional block diagram for 100GBASE-R PHYs

Figure 74–2b shows the functional block diagram of FEC for 100GBASE-R PHYs and the relationship between the PCS and PMA sublayers.

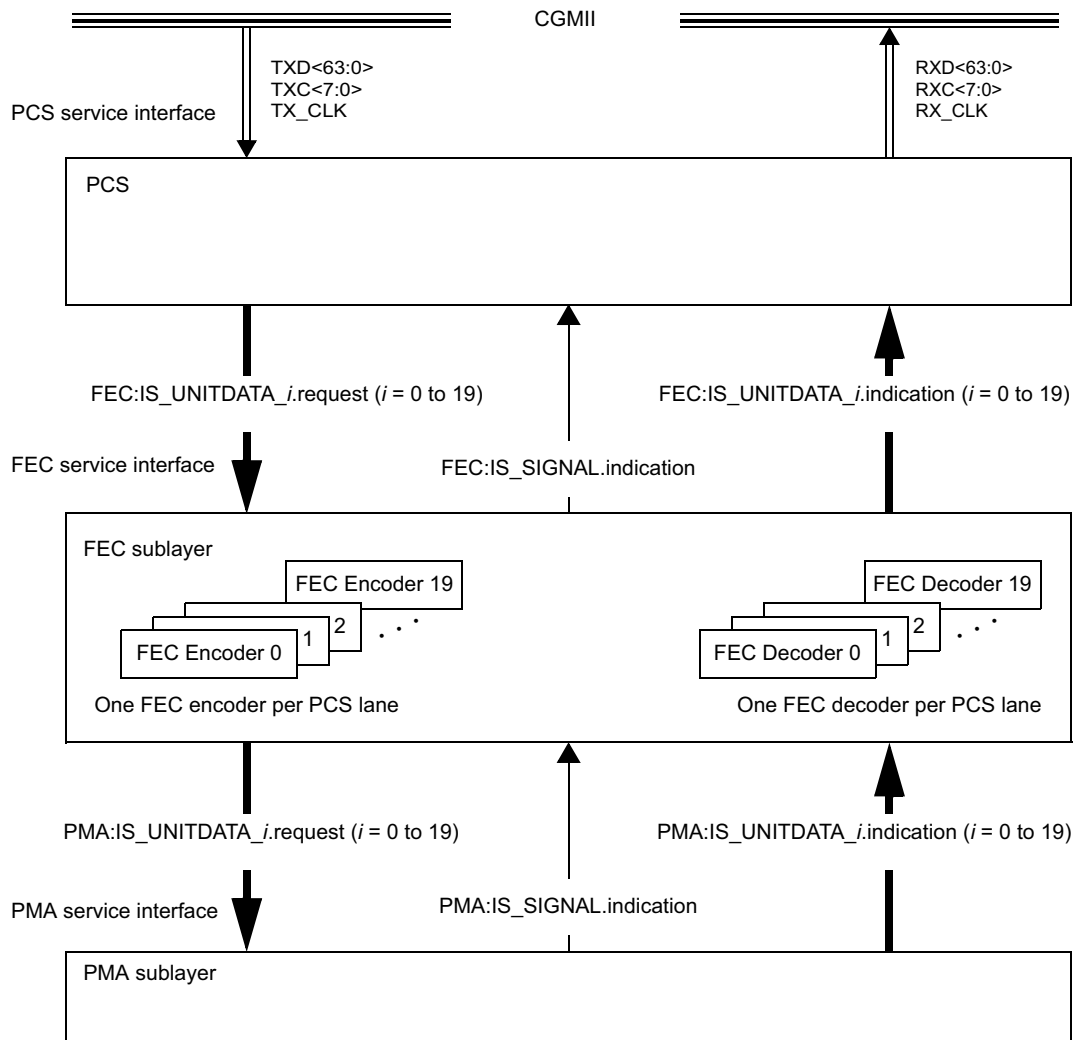


Figure 74–2b—Functional block diagram for 100GBASE-R PHY

Replace 74.5 with the following:

74.5 FEC service interface

The FEC service interface is provided to allow the PCS to transfer information to and from the FEC. The FEC service interface is equivalent to the PMA service interface for 10GBASE-R and an instance of the inter-sublayer service interface defined in 80.3 for 40GBASE-R and 100GBASE-R. These services are defined in an abstract manner and do not imply any particular implementation. The FEC service interface supports exchange of data units between PCS entities on either side of a link using request and indication primitives. Data units are mapped into FEC blocks by the FEC and passed to the PMA, and vice versa.

The service primitives are defined differently for 10GBASE-R and for 40GBASE-R and 100GBASE-R.

Optional physical instantiations of the PMA service interface have been defined (see Clause 51, Annex 83A, and Annex 83B). There is XSBI (10 Gigabit Sixteen-Bit Interface) for 10GBASE-R, XLAUI for 40GBASE-R and CAUI for 100GBASE-R. These physical instantiations, with a PMA if required, may also be used for the FEC service interface.

74.5.1 10GBASE-R service primitives

The following primitives are defined within the FEC service interface:

- a) FEC_UNITDATA.request(tx_data-group<15:0>)
- b) FEC_UNITDATA.indication(rx_data-group<15:0>)
- c) FEC_SIGNAL.indication(SIGNAL_OK)

The FEC service interface directly maps to the PMA service interface of the 10GBASE-R PCS defined in Clause 49. The FEC_UNITDATA.request maps to the PMA_UNITDATA.request primitive, the FEC_UNITDATA.indication maps to the PMA_UNITDATA.indication primitive, and the FEC_SIGNAL.indication maps to the PMA_SIGNAL.indication primitive of the 10GBASE-R PCS.

74.5.1.1 FEC_UNITDATA.request

This primitive defines the transfer of data in the form of constant-width data units from the PCS to the FEC. The data supplied via FEC_UNITDATA.request is mapped by the FEC Transmit process into the payload capacity of the outgoing FEC block stream.

74.5.1.1.1 Semantics of the service primitive

FEC_UNITDATA.request(tx_data-group<15:0>)

The data conveyed by FEC_UNITDATA.request is a 16-bit vector representing a single data unit that has been prepared for transmission by the 10GBASE-R PCS Transmit process.

74.5.1.1.2 When generated

The 10GBASE-R PCS sends tx_data-group<15:0> to the FEC at a nominal rate of 644.53125 MHz, corresponding to the 10GBASE-R signaling rate of 10.3125 Gb/s.

74.5.1.1.3 Effect of receipt

Upon receipt of this primitive, the FEC Transmit process maps the data conveyed by the tx_data unit<15:0> parameter into the payload of the transmitted FEC block stream, adds FEC overhead as required, scrambles the data, and transfers the result to the PMA via the PMA_UNITDATA.request primitives.

74.5.1.2 FEC_UNITDATA.indication

This primitive defines the transfer of received data in the form of constant-width data units from the FEC to the PCS. FEC_UNITDATA.indication is generated by the FEC Receive process in response to FEC block data received from the PMA.

74.5.1.2.1 Semantics of the service primitive

FEC_UNITDATA.indication(rx_data-group<15:0>)

The rx_data-group<15:0> parameter is a 16-bit vector that represents the data unit transferred by the FEC to the 10GBASE-R PCS.

74.5.1.2.2 When generated

The FEC sends one rx_data-group<15:0> to the 10GBASE-R PCS for each 16 bits received from the PMA sublayer. The nominal rate of generation of the FEC_UNITDATA.indication primitive is 644.53125 Mtransfers/s.

74.5.1.2.3 Effect of receipt

The effect of receipt of this primitive by the FEC client is unspecified by the FEC sublayer.

74.5.1.3 FEC_SIGNAL.indication

This primitive is sent by the FEC to the PCS to indicate the status of the Receive process. FEC_SIGNAL.indication is generated by the FEC Receive process in order to propagate the detection of severe error conditions (e.g., no valid signal being received from the PMA sublayer) to the PCS.

74.5.1.3.1 Semantics of the service primitive

FEC_SIGNAL.indication(SIGNAL_OK)

The SIGNAL_OK parameter can take one of two values: OK or FAIL. A value of OK denotes that the FEC Receive process is successfully delineating valid payload information from the incoming data stream received from the PMA sublayer indicated by the fec_signal_ok variable equal to true, and this payload information is being presented to the PCS via the FEC_UNITDATA.indication primitive. A value of FAIL denotes that errors have been detected by the Receive process indicated by the fec_signal_ok variable equal to false, that prevent valid data from being presented to the PCS, in this case the FEC_UNITDATA.indication primitive and its associated rx_data-group<15:0> parameter are meaningless.

74.5.1.3.2 When generated

The FEC generates the FEC_SIGNAL.indication primitive to the 10GBASE-R PCS whenever there is a change in the value of the SIGNAL_OK parameter and FEC block synchronization is achieved.

74.5.1.3.3 Effect of receipt

The effect of receipt of this primitive by the FEC client is unspecified by the FEC sublayer.

74.5.2 40GBASE-R and 100GBASE-R service primitives

The FEC service interface for 40GBASE-R and 100GBASE-R is an instance of the inter-sublayer service interface defined in 80.3. The FEC service interface primitives are summarized as follows:

```
FEC:IS_UNITDATA_i.request
FEC:IS_UNITDATA_i.indication
FEC:IS_SIGNAL.indication
```

The 40GBASE-R FEC has four parallel bit streams, hence $i = 0$ to 3 for 40GBASE-R and the 100GBASE-R FEC has twenty parallel bit streams, hence $i = 0$ to 19 for 100GBASE-R.

The PCS (or PMA) continuously sends four or twenty parallel bit streams to the FEC, one per PCS lane, each at a nominal signaling rate of 10.3125 GBd for 40GBASE-R and 5.15625 GBd for 100GBASE-R.

This FEC:IS_SIGNAL.indication primitive is sent by the FEC to the PCS (or PMA) to indicate the status of the Receive process. FEC:IS_SIGNAL.indication is generated by the FEC Receive process in order to propagate the detection of severe error conditions (e.g., no valid signal being received from the PMA sublayer) to the PCS (or PMA).

The SIGNAL_OK parameter in FEC:IS_SIGNAL.indication can take one of two values: OK or FAIL. A value of OK denotes that the FEC Receive process is successfully delineating valid payload information from all of the incoming data streams received from the PMA sublayer indicated by the fec_signal_ok variable equal to true, for all data streams, and this payload information is being presented to the PCS (or PMA) via the FEC:IS_UNITDATA_i.indication primitive. A value of FAIL denotes that errors have been detected by the Receive process indicated by the fec_signal_ok variable equal to false, in any of the data streams, that prevent valid data from being presented to the PCS (or PMA), in this case the FEC:IS_UNITDATA_i.indication primitive is a direct pass through of the PMA:IS_UNITDATA_i.indication from the PMA.

Change 74.6 as follows:

74.6 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementors must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The sum of transmit and receive maximum delay contributed by the 10GBASE-R FEC (sum of transmit and receive delays at one end of the link) shall be no more than 6144 BT (or 12 pause quanta or 614.4 ns).

The maximum delay contributed by the 40GBASE-R FEC (sum of transmit and receive delays at one end of the link) shall be no more than 24576 BT (or 48 pause quanta or 614.4 ns).

The maximum delay contributed by the 100GBASE-R FEC (sum of transmit and receive delays at one end of the link) shall be no more than 122880 BT (or 240 pause quanta or 1228.8 ns).

A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 80.4 and its references.

Change 74.7 as follows:

74.7 FEC principle of operation

On transmission, the FEC sublayer receives data from the ~~10GBASE-R~~ PCS, transcodes 64B/66B words, performs the FEC coding/framing, scrambles and sends the data to the PMA. On reception, the FEC sublayer receives data from the PMA, performs descrambling, achieves FEC framing synchronization, decodes the FEC code, correcting data where necessary and possible, re-codes 64B/66B words, and sends the data to the ~~10GBASE-R~~ PCS.

Change the second and third paragraphs of 74.7.3 as follows:

74.7.3 Composition of the FEC block

The 10GBASE-R 64B/66B PCS maps 64 bits of scrambled payload and 2 bits of unscrambled sync header into 66-bit encoded blocks. The 2-bit sync header allows establishment of 64B/66B block boundaries by the PCS sync process. The sync header is 01 for data blocks and 10 for control blocks; the sync header is the only position in the PCS block that always contain a transition and this feature of the code is used to establish 64B/66B block boundaries.

The FEC sublayer compresses the 2 bits of the sync header to 1 transcode bit. The transcode bit carries the state of 10GBASE-R sync bits for the associated payload. This is achieved by eliminating the first bit in 64B/66B block, which is also the first sync bit, and preserving the second bit. The value of the second bit defines the value of the removed first bit uniquely, since it is always an inversion of the first bit. The transcode bits are further scrambled (as explained in 74.7.4.2) to ensure DC balance.

Delete the last paragraph of 74.7.3 as it is redundant.

~~The 16-bit data transmitted from the PCS function is encoded by the FEC encoder and sent to the PMA sublayer; similarly, the 16-bit data received from the PMA sublayer is decoded by the FEC decoder. The resulting 64B/66B blocks are sent to the PCS sublayer.~~

74.7.4.1 Reverse gearbox function

Insert new subheading keeping original text from 74.7.4.1 as follows:

74.7.4.1.1 Reverse gearbox function for 10GBASE-R

The reverse gearbox function adapts between the 66-bit width of the 64B/66B blocks and the 16-bit width of the PCS interface. It receives the 16-bit stream from the PCS interface and converts them back to 66-bit encoded blocks for the FEC Encoder to process. The reverse gearbox function operates in the same manner as the block sync function defined in 49.2.9.

The reverse gearbox function receives data via 16-bit FEC_UNITDATA.request primitive. It will form a bit stream from the primitives by concatenating requests with the bits of each primitive in order to form tx_data-group<0> to tx_data-group<15> (see Figure 49–6). It obtains lock to the 66-bit blocks in the bit stream using the sync headers and outputs 66-bit blocks. Lock is obtained as specified in the block lock state diagram shown in Figure 49–12.

The reverse gearbox functionality is necessary only when the optional PMA compatibility interface named XSBI is implemented between the PCS and FEC functions, since that interface passes data via a 16-bit wide path. When the XSBI is not implemented, the internal data-path width between the PCS and FEC is an implementation choice. Depending on the path width, the reverse gearbox function may not be necessary.

Insert 74.7.4.1.2 for reverse gearbox for 40GBASE-R and 100GBASE-R as follows:

74.7.4.1.2 Reverse gearbox function for 40GBASE-R and 100GBASE-R

The reverse gearbox function adapts between the 66-bit width of the 64B/66B blocks and the 1-bit wide lane of the 40GBASE-R or 100GBASE-R PCS to FEC interface (or PMA to FEC interface). It receives the 1-bit

stream from the FEC service interface (or PMA service interface) and converts it back to 66-bit encoded blocks for the FEC Encoder to process. The reverse gearbox function, if implemented, shall operate in the same manner as the lane block sync function defined in 82.2.11.

The reverse gearbox function receives data via the 40GBASE-R and 100GBASE-R FEC:IS_UNITDATA_*i*.request primitive see 74.5.2 (or via the PMA:IS_UNITDATA_*i*.request primitive). It obtains lock to the 66-bit blocks in each bit stream using the sync headers and outputs 66-bit blocks to the FEC encoder function (see 74.7.4.4). PCS lane lock is obtained as specified in the PCS lane lock state diagram shown in Figure 82–10.

The internal data-path width from the PCS or PMA is an implementation choice. Depending on the path width, the reverse gearbox function may not be necessary.

74.7.4.3 FEC transmission bit ordering

Replace Figure 74-3 with the following:

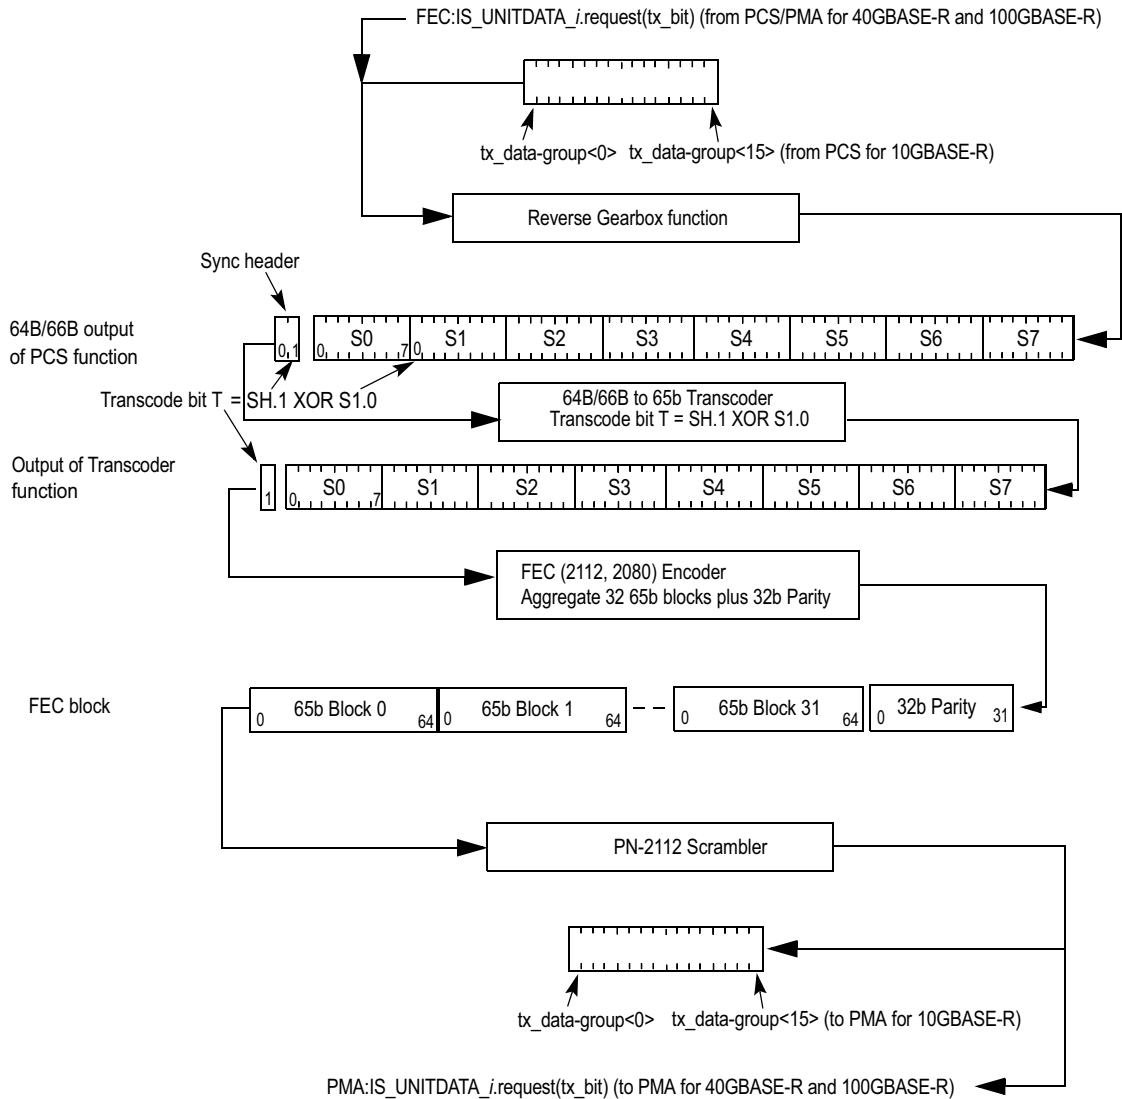


Figure 74-3—FEC Transmit bit ordering

74.7.4.4 FEC (2112, 2080) encoder

Replace Figure 74-4 with the following:

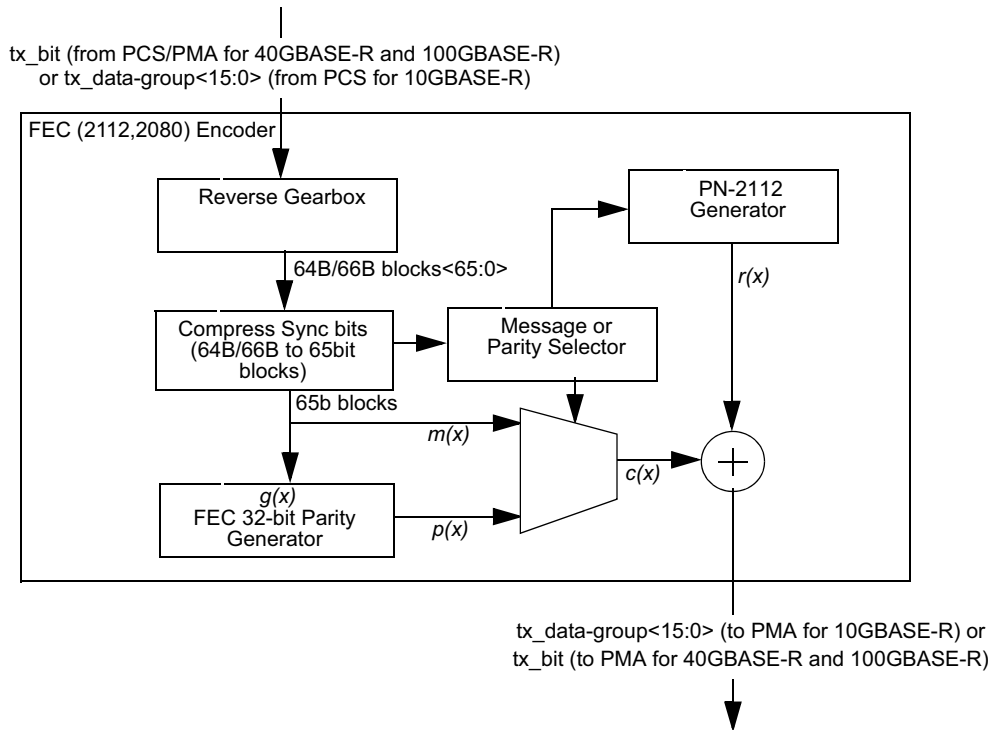


Figure 74-4—FEC (2112,2080) encoding

Change 74.7.4.5 as follows:

74.7.4.5 FEC decoder

The FEC decoder establishes FEC block synchronization based on repeated decoding of the received sequence. Decoding and error correction is performed after FEC synchronization is achieved. There is an option for the FEC decoder to indicate any decoding errors to the upper layer.

The FEC decoder recovers and extracts the information bits using the parity-check data. In case of successful decoding the decoder restores the sync bits in each of the 64B/66B blocks sent to the PCS function, by first performing an XOR operation of the received transcode bit with the associated data bit 8 and then generating the two sync bits.

When the decoder for 10GBASE-R is configured to indicate decoding error, the decoder indicates error to the PCS by means of setting both sync bits to the value 11 in the 1st, 9th, 17th, 25th, and 32nd of the 32 decoded 64B/66B blocks from the corresponding errored FEC block, thus forcing the PCS sublayer to consider this block as invalid.

When the decoder for 40GBASE-R or 100GBASE-R is configured to indicate decoding error, the decoder needs to mark errors in more of the 64B/66B blocks to ensure that detected errors are signaled to the MAC

for every frame containing an error. The FEC sublayers for 40GBASE-R and 100GBASE-R set both sync bits to the value 11 in all thirty-two 64B/66B blocks to indicate error to the PCS.

The FEC Synchronization process continuously monitors `PMA_SIGNAL.indication(SIGNAL_OK)` or `PMA:IS_SIGNAL.indication(SIGNAL_OK)`. When `SIGNAL_OK` indicates OK, the FEC Synchronization process accepts data units via the `PMA_UNITDATA.indication` or the `PMA:IS_UNITDATA.i.indication` primitives. It attains block synchronization based on the decoding of FEC blocks and conveys received 64B/66B blocks to the PCS Receive process. The FEC Synchronization process sets the `sync_status` flag to the PCS function to indicate whether the FEC has obtained synchronization.

Change 74.7.4.5.1 as follows and replace Figure 74-6:

74.7.4.5.1 FEC (2112,2080) decoding

The FEC decoding function block diagram is shown in Figure 74-6. The decoder processes the 16-bit `rx_data-group` stream received from the PMA sublayer and descrambles the data using the PN-2112 pseudo-noise sequence as described in 74.7.4.4.1.

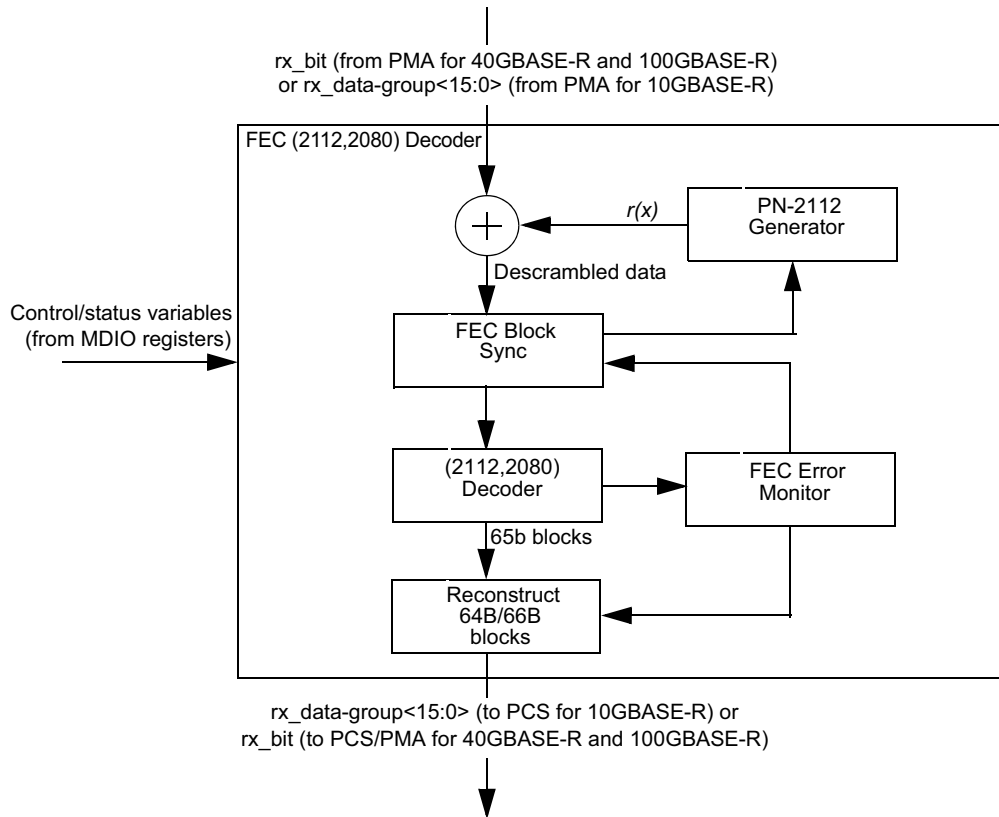


Figure 74-6—FEC (2112,2080) decoding

The synchronization of the 2112 bit FEC block is established using FEC decoding as described in 74.7.4.7. Each of the 32 65-bit data words is extracted from the recovered FEC block and the 2-bit sync is reconstructed for the 64B/66B codes from the transcode bit as shown in Figure 74-7. The FEC decoder provides an option to indicate decoding errors in the reconstructed sync bits. The sync bits {SH.0, SH.1} take the value as described in the following:

- a) If decoding is successful (by either the parity match or the FEC block is correctable) and the descrambled received transcode bit (T) is 1 then the sync bits take a value of {SH.0,SH.1} = 01 or if the descrambled received transcode bit (T) is 0 then the sync bits take a value of {SH.0,SH.1} = 10.
- b) If the variable FEC_Enable_Error_to_PCS is set to 1 to indicate error to PCS layer and the received FEC block has uncorrectable errors then the sync bits for the 1st, 9th, 17th, 25th, and 32nd of the 32 decoded 64B/66B blocks take a value of {SH.0,SH.1} = 11 for the 10GBASE-R PHY. For the 40GBASE-R and 100GBASE-R PHYs, sync bits in all thirty-two 64B/66B decoded 64B/66B blocks take a value of {SH.0,SH.1} = 11. The sync bits for all other 64B/66B blocks take a value as described in item a) above.
- c) If the variable FEC_Enable_Error_to_PCS is set to 0 and the received FEC block has uncorrectable errors then the sync bits take a value as described in item a) above.

This information corresponds to one complete (2112,2080) FEC block that is equal to 32 64B/66B code blocks.

The FEC code (2112, 2080) and its performance is specified in 74.7.1. The FEC (2112, 2080) decoder implementations shall be able to correct up to a minimum of 11-bit burst errors per FEC block.

74.7.4.6 FEC receive bit ordering

Replace Figure 74-7 with the following:

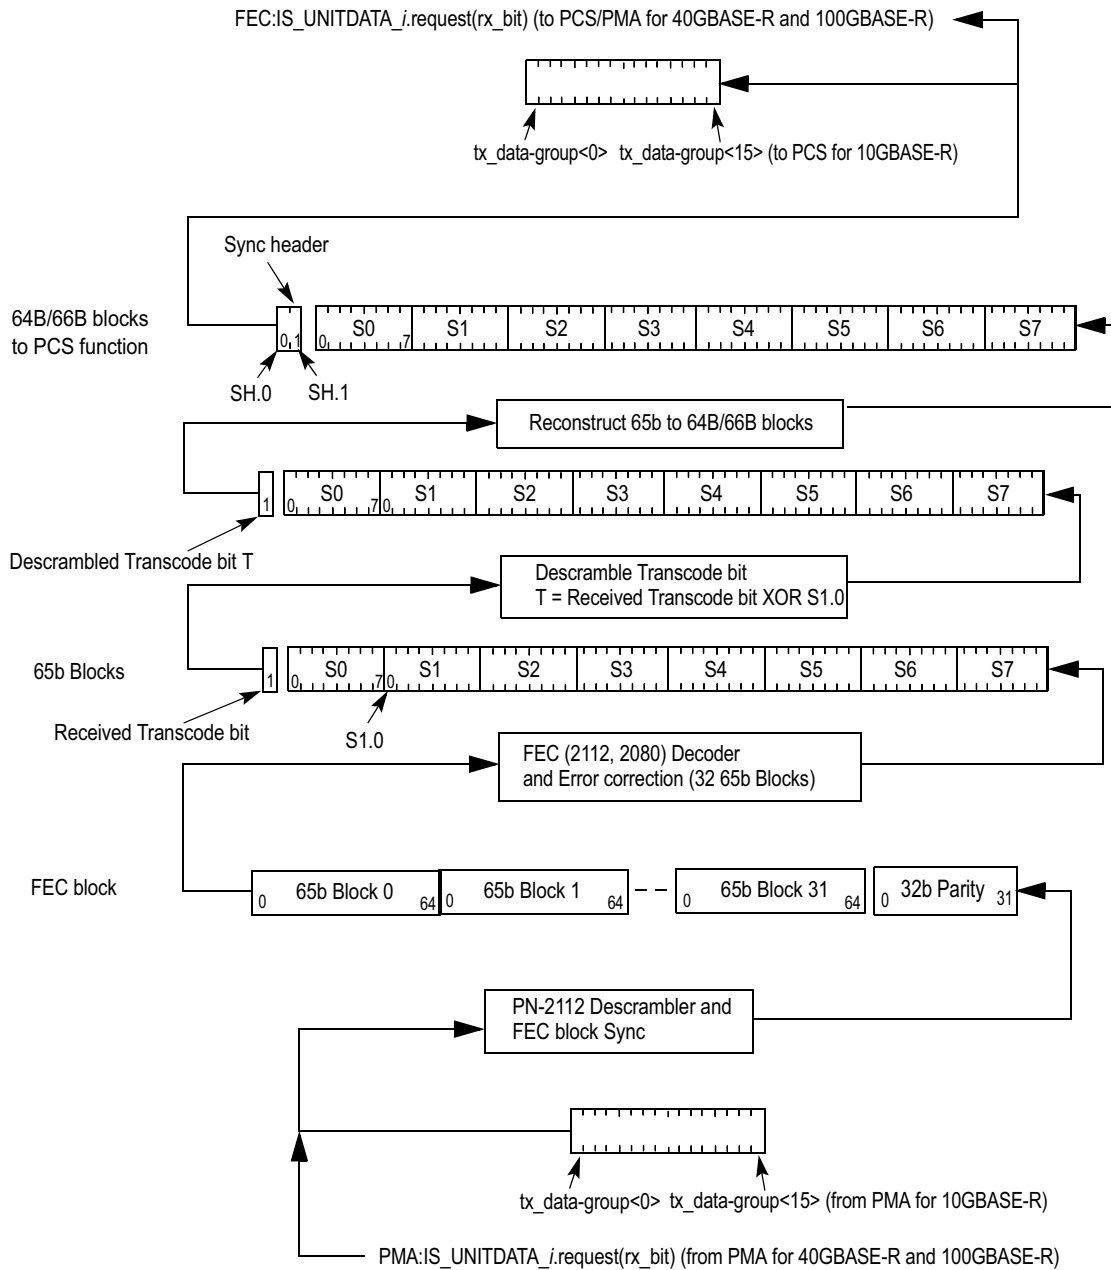


Figure 74-7—FEC Receive bit ordering

74.8 FEC MDIO function mapping

Change Table 74-2 as follows:

Table 74–2—MDIO/FEC variable mapping

MDIO variable	PMA/PMD register name	Register/bit number	FEC variable
10GBASE-R FEC ability	10GBASE-R FEC ability register	1.170.0	FEC_ability
10GBASE-R FEC Error Indication ability	10GBASE-R FEC ability register	1.170.1	FEC_Error_Indication_ability
FEC Enable	10GBASE-R FEC control register	1.171.0	FEC_Enable
FEC Enable Error Indication	10GBASE-R FEC control register	1.171.1	FEC_Enable_Error_to_PCS
FEC corrected blocks	10GBASE-R FEC corrected blocks counter register	1.172, 1.173	FEC_corrected_blocks_counter
FEC uncorrected blocks	10GBASE-R FEC uncorrected blocks counter register	1.174, 1.175	FEC_uncorrected_blocks_counter
<u>FEC corrected blocks, lanes 0 through 19</u>	<u>BASE-R FEC corrected blocks counter register, lanes 0 through 19</u>	<u>1.300 through 1.339</u>	<u>FEC corrected blocks counter <i>i</i></u>
<u>FEC uncorrected blocks, lanes 0 through 19</u>	<u>BASE-R FEC uncorrected blocks counter register, lanes 0 through 19</u>	<u>1.700 through 1.739</u>	<u>FEC uncorrected blocks counter <i>i</i></u>

Change 74.8.1 as follows:

74.8.1 FEC capability

Since the FEC is an optional sublayer, the FEC ability is indicated by the variable FEC_ability for each of the ~~10GBASE-R~~ PHY types. An MDIO interface or an equivalent management interface shall be provided to access the variable FEC_ability for the ~~10GBASE-R~~ PHY type (refer to ~~45.2.1.84~~ ~~10GBASE-R~~ ~~FEC ability register 1.170~~). The FEC_ability variable bit is set to a one to indicate that the ~~10GBASE-R~~ PHY supports FEC sublayer, it defaults to zero otherwise.

The FEC_ability variable for the ~~10GBASE-R~~ PHY is mapped to register bit 1.170.0 (refer to 45.2.1.85.1).

~~For the 10GBASE-KR PHY type, the~~ The FEC capability between the link partners can be negotiated using the Clause 73 Auto-Negotiation as defined in 73.6.5. The FEC function is enabled on the link only if both the link partners advertise they have FEC ability and either one of them requests to enable FEC through the Auto-Negotiation function

Change 74.8.2 as follows:

74.8.2 FEC enable

The FEC sublayer shall have capability to enable or disable the FEC function. An MDIO interface or an equivalent management interface shall be provided to access the variable FEC_Enable for the ~~10~~GBASE-R PHY (refer to 45.2.1.86 register bit 1.171.0). When FEC_Enable variable bit is set to a one, this enables the FEC for the ~~10~~GBASE-R PHY. When the variable is set to zero, the FEC is disabled in the ~~10~~GBASE-R PHY. This variable shall be set to zero upon execution of PHY reset. When the FEC function is disabled, the PHY shall have a mechanism to bypass the FEC Encode and Decode functions so as not to cause additional latency associated with encoding or decoding functions.

Change 74.8.3 as follows:

74.8.3 FEC Enable Error Indication

The FEC sublayer may have the option to enable the ~~10~~GBASE-R FEC decoder to indicate decoding errors to the upper layers (PCS) through the sync bits for the ~~10~~GBASE-R PHY as defined in 74.7.4.5, if this ability is supported. An MDIO interface or an equivalent management interface shall be provided to access the variable FEC_Enable_Error_to_PCS. When the variable is set to one, this enables indication of decoding errors through the sync bits to the PCS layer. When set to zero, the error indication function is disabled.

Change 74.8.3.1 as follows:

74.8.3.1 FEC Error Indication ability

The FEC error indication ability shall be indicated by the variable FEC_Error_Indication_ability. The variable is set to one to indicate that the ~~10~~GBASE-R FEC has the ability to indicate decoding errors to the PCS layer. The variable is set to zero if this ability is not supported by the ~~10~~GBASE-R FEC. An MDIO interface or an equivalent management interface shall be provided to access the variable

Change second paragraph of 74.8.4.1 as follows:

74.8.4.1 FEC_corrected_blocks_counter

FEC_corrected_blocks_counter (for single-lane PHYs) or FEC_corrected_blocks_counter_{*i*} (for multi-PCS-lane PHYs, where $i = 0$ to 3 for 40 Gb/s and $i = 0$ to 19 for 100 Gb/s,) counts once for each corrected FEC blocks processed when FEC_SIGNAL.indication or FEC:IS_SIGNAL.indication is OK. These are 32-bit counters. These variables are accessed through a management interface that may be mapped to the registers defined in 45.2.1.87 register (1.172, 1.173) for single-lane PHYs and 45.2.1.89 (1.300 to 1.339) for multi-lane PHYs.

Change second paragraph of 74.8.4.2 as follows:

74.8.4.2 FEC_uncorrected_blocks_counter

FEC_uncorrected_blocks_counter (for single-lane PHYs) or FEC_uncorrected_blocks_counter_{*i*} (for multi-PCS-lane PHYs, where *i* = 0 to 3 for 40 Gb/s and *i* = 0 to 19 for 100 Gb/s.) counts once for each uncorrected FEC blocks processed when FEC_SIGNAL.indication or FEC:IS_SIGNAL.indication is OK. This is a
These are 32-bit counters. This variable is provided by These variables are accessed through a management interface that may be mapped to the registers defined in 45.2.1.88 ~~register (1.174, 1.175)~~ for single-lane PHYs and 45.2.1.90 (1.700 to 1.739) for multi-lane PHYs.

Change 74.9 as follows:

74.9 10GBASE-R PHY test-pattern mode

The 10GBASE-R PCS provides test-pattern functionality and the PCS transmit channel and receive channel can each operate in normal mode or test-pattern mode (see 49.2.2). When the 10GBASE-R PHY is configured for test-pattern mode, the FEC function may be disabled by setting the FEC Enable variable to zero, so the test-pattern from the 10GBASE-R PCS can be sent to the PMA service interface, bypassing the FEC Encode and Decode functions.

The Clause 82 PCS can also operate in test pattern mode (see 82.2.10); however, the scrambled idle test pattern does not require bypassing FEC encode and decode.

Change 74.10.2.2 fec_signal_ok variable definition as follows:

74.10.2.2 Variables

fec_signal_ok

Boolean variable that is set based on the most recently received value of PMA_UNITDATA.indication(SIGNAL_OK) or PMA:IS_SIGNAL.indication(SIGNAL_OK) and fec_block_lock. It is set to true if the ~~fec_block_lock~~ fec_block_lock value is true and PMA_UNITDATA.indication(SIGNAL_OK) or PMA:IS_SIGNAL.indication(SIGNAL_OK) value was OK and set to false otherwise. The value is sent to the PCS layer through the primitive FEC_SIGNAL.indication or FEC:IS_SIGNAL.indication as specified in ~~74.5.3~~ 74.5.1.3 or 74.5.2.

Change clause title as follows:

74.11 Protocol implementation conformance statement (PICS) proforma for Clause 74, Forward Error Correction (FEC) sublayer for 40GBASE-R PHYs⁸

74.11.1 Introduction

Change first sentence as follows:

The supplier of a protocol implementation that is claimed to conform to IEEE Std 802.3-2008 as modified by IEEE Std 802.3ba-2010, Clause 74 Forward Error Correction (FEC) sublayer for 40GBASE-R PHYs, shall complete the following protocol implementation conformance statement (PICS) proforma.

74.11.2.2 Protocol summary

Change Identification of protocol standard as follows:

IEEE Std 802.3-2008 as modified by IEEE Std 802.3ba-2010, Clause 74, Forward Error Correction (FEC) sublayer for 40GBASE-R PHYs.

⁸*Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

74.11.3 Major capabilities/options*Change delay constraint item as follows:*

Item	Feature	Subclause	Value/Comment	Status	Support
DC	FEC Delay Constraints	74.6	Device implements FEC delay constraints, no more than 6144 BT for the sum of transmit and receive path delays as specified in 74.6 Sum of transmit and receive. No more than 12 <u>Device implements FEC delay constraints, no more than 6144 BT for the sum of transmit and receive path delays as specified in 74.6 Sum of transmit and receive. No more than 12</u> <u>pause quanta for 10GBASE-R, 48 pause quanta for 40GBASE-R, and 240 pause quanta for 100GBASE-R</u>	M	Yes []
*MD	MDIO Interface	45, 74.8.2, 74.8.4	Device implements MDIO registers and interface	O	Yes [] No []
EF	FEC_Enable	74.8.2	The device has the capability to enable/disable the FEC function	M	Yes []
*EIA	FEC Error Indication ability	74.8.3, 74.8.3.1	The device has ability to indicate FEC decoding errors to the PCS layer as specified in 74.8.3	O	Yes [] No []
BF	Bypass FEC function	74.8.2	The device has mechanism to bypass FEC encode/decode functions to reduce latency	M	Yes []
*XSBI	PMA compatibility interface XSBI	51, 74.7.4.1	Optional PMA compatibility interface named XSBI is implemented between the PCS and FEC functions	O	Yes [] No []

74.11.5 FEC requirements*Change FE3 item as follows:*

Item	Feature	Subclause	Value/Comment	Status	Support
FE3	Reverse Gear-Box gearbox function for <u>10GBASE-R</u>	<u>74.7.4.1.1</u>	Reverse Gear-Box gearbox function implemented	XSBI:M	N/A [] Yes []

Add PICS item FE3a after FE3 as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
FE3a	Reverse gearbox function for 40GBASE-R and 100GBASE-R	74.7.4.1.2	Reverse gearbox function meets the requirements of 82.2.11	O	Yes [] No []

Insert the following new subclauses and corresponding annexes as follows:

80. Introduction to 40 Gb/s and 100 Gb/s networks

80.1 Overview

80.1.1 Scope

40 Gigabit and 100 Gigabit Ethernet uses the IEEE 802.3 MAC sublayer, connected through a Media Independent Interface to 40 Gb/s and 100 Gb/s Physical Layer entities such as those specified in Table 80–1.

40 Gigabit and 100 Gigabit Ethernet is defined for full duplex operation only.

80.1.2 Objectives

The following are the objectives of 40 Gigabit and 100 Gigabit Ethernet:

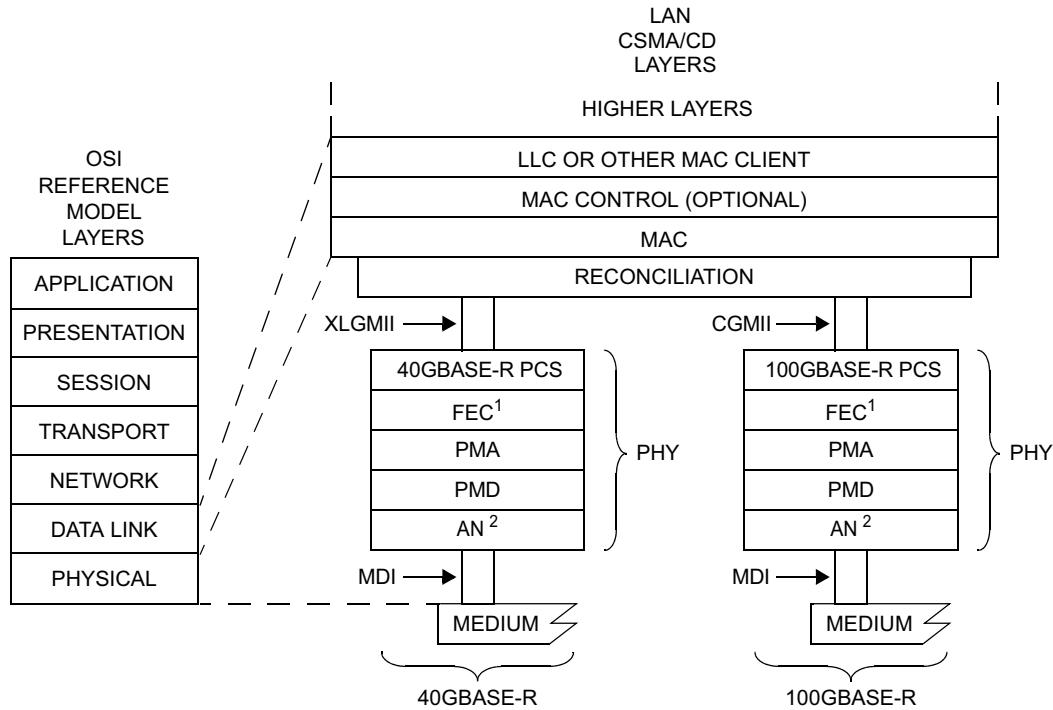
- a) Support the full duplex Ethernet MAC.
- b) Preserve the IEEE 802.3 Ethernet frame format utilizing the IEEE 802.3 MAC.
- c) Preserve minimum and maximum frame size of IEEE Std 802.3.
- d) Support a BER better than or equal to 10^{-12} at the MAC/PLS service interface.
- e) Provide appropriate support for Optical Transport Network (OTN).
- f) Support a MAC data rate of 40 Gb/s.
- g) Provide Physical Layer specifications that support 40 Gb/s operation over up to the following:
 - 1) At least 10 km on single-mode fiber (SMF)
 - 2) At least 100 m on OM3 multimode fiber (MMF)
 - 3) At least 7 m over a copper cable assembly
 - 4) At least 1 m over a backplane
- h) Support a MAC data rate of 100 Gb/s.
- i) Provide Physical Layer specifications that support 100 Gb/s operation over up to the following:
 - 1) At least 40 km on single-mode fiber (SMF)
 - 2) At least 10 km on single-mode fiber (SMF)
 - 3) At least 100 m on OM3 multimode fiber (MMF)
 - 4) At least 7 m over a copper cable assembly

80.1.3 Relationship of 40 Gigabit and 100 Gigabit Ethernet to the ISO OSI reference model

40 Gigabit and 100 Gigabit Ethernet couples the IEEE 802.3 MAC to a family of 40 Gb/s and 100 Gb/s Physical Layers. The relationships among 40 Gigabit and 100 Gigabit Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in Figure 80–1.

While this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience. The only exceptions are as follows:

- a) The XLGMII and CGMII, which, when implemented as a logical interconnection port between the MAC sublayer and the Physical Layer (PHY), uses a 64-bit wide data path as specified in Clause 81.
- b) The management interface, which, when physically implemented as the MDIO/MDC (Management Data Input/Output and Management Data Clock) at an observable interconnection port, uses a bit-wide data path as specified in Clause 45.
- c) The PMA service interface, which, when physically implemented as XLAUI (40 Gigabit Attachment Unit Interface) at an observable interconnection port, uses a 4 lane data path as specified in Annex 83A or Annex 83B.



AN = AUTO-NEGOTIATION

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE

FEC = FORWARD ERROR CORRECTION

LLC = LOGICAL LINK CONTROL

MAC = MEDIA ACCESS CONTROL

MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE

PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT

XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE

NOTE 2—CONDITIONAL BASED ON PHY TYPE

Figure 80–1—Architectural positioning of 40 Gigabit and 100 Gigabit Ethernet

- d) The PMA service interface, which, when physically implemented as CAUI (100 Gigabit Attachment Unit Interface) at an observable interconnection port, uses a 10 lane data path as specified in Annex 83A or Annex 83B.
- e) The PMD service interface, which, when physically implemented as XLPPi (40 Gigabit Parallel Physical Interface) at an observable interconnection port, uses a 4 lane data path as specified in Annex 86A.
- f) The PMD service interface, which, when physically implemented as CPPI (100 Gigabit Parallel Physical Interface) at an observable interconnection port, uses a 10 lane data path as specified in Annex 86A.
- g) The MDIs as specified in Clause 84 for 40GBASE-KR4, in Clause 85 for 40GBASE-CR4, in Clause 86 for 40GBASE-SR4, Clause 87 for 40GBASE-LR4 and in Clause 88 for 100GBASE-LR4 and 100GBASE-ER4 all use a 4 lane data path.
- h) The MDIs as specified in Clause 85 for 100GBASE-CR10, and in Clause 86 for 100GBASE-SR10 use a 10 lane data path.

80.1.4 Nomenclature

The nomenclature employed by the 40 Gigabit and 100 Gigabit Physical Layers is explained as follows.

The alpha-numeric prefix 40GBASE in the port type (e.g., 40GBASE-R) represents a family of Physical Layer devices operating at a speed of 40 Gb/s. The alpha-numeric prefix 100GBASE in the port type (e.g., 100GBASE-R) represents a family of Physical Layer devices operating at a speed of 100 Gb/s.

40GBASE-R or 100GBASE-R represents a family of Physical Layer devices using a physical coding sublayer for 40 Gb/s or 100 Gb/s operation over multiple PCS lanes based on 64B/66B block encoding (see Clause 82).

Physical Layer devices listed in Table 80–1 are defined for operation at 40 Gb/s and 100 Gb/s.

Table 80–1—40 Gb/s and 100 Gb/s PHYs

Name	Description
40GBASE-KR4	40 Gb/s PHY using 40GBASE-R encoding over four lanes of an electrical backplane, with reach up to at least 1 m (see Clause 84)
40GBASE-CR4	40 Gb/s PHY using 40GBASE-R encoding over four lanes of shielded balanced copper cabling, with reach up to at least 7 m (see Clause 85)
40GBASE-SR4	40 Gb/s PHY using 40GBASE-R encoding over four lanes of multimode fiber, with reach up to at least 100 m (see Clause 86)
40GBASE-LR4	40 Gb/s PHY using 40GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 87)
100GBASE-CR10	100 Gb/s PHY using 100GBASE-R encoding over ten lanes of shielded balanced copper cabling, with reach up to at least 7 m (see Clause 85)
100GBASE-SR10	100 Gb/s PHY using 100GBASE-R encoding over ten lanes of multimode fiber, with reach up to at least 100 m (see Clause 86)
100GBASE-LR4	100 Gb/s PHY using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 10 km (see Clause 88)
100GBASE-ER4	100 Gb/s PHY using 100GBASE-R encoding over four WDM lanes on single-mode fiber, with reach up to at least 40 km (see Clause 88)

80.1.5 Physical Layer signaling systems

This standard specifies a family of Physical Layer implementations. The generic term 40 Gigabit and 100 Gigabit Ethernet refers to any use of the 40 Gb/s and 100 Gb/s IEEE 802.3 MAC (the 40 Gigabit and 100 Gigabit Ethernet MAC) coupled with any IEEE 802.3 40GBASE or 100GBASE Physical Layer implementations. Table 80–2 specifies the correlation between nomenclature and clauses. Implementations conforming to one or more nomenclatures must meet the requirements of the corresponding clauses.

Table 80–2—Nomenclature and clause correlation

Nomenclature	Clause ^a																							
	73	74	81		82		83		83A		83B		84	85		86		86A		87	88			
	Auto-Negotiation	BASE-R FEC	RS	XLGMII	CGMII	40GBASE-R PCS	100GBASE-R PCS	40GBASE-R PMA	100GBASE-R PMA	XL/AUI	CAUI	XL/AUI	CAUI	40GBASE-KR4 PMD	40GBASE-CR4 PMD	100GBASE-CR10 PMD	40GBASE-SR4 PMD	100GBASE-SR10 PMD	XL/PPI	CPPI	40GBASE-LR4 PMD	100GBASE-LR4 PMD	100GBASE-ER4 PMD	
40GBASE-KR4	M	O	M	O		M	M		O				M											
40GBASE-CR4	M	O	M	O		M	M		O					M										
40GBASE-SR4			M	O		M	M		O		O					M		O						
40GBASE-LR4			M	O		M	M		O		O							O		M				
100GBASE-CR10	M	O	M		O	M	M		O						M									
100GBASE-SR10			M		O	M	M		O		O						M		O					
100GBASE-LR4			M		O	M	M		O		O												M	
100GBASE-ER4			M		O	M	M		O		O													M

^aO = Optional, M = Mandatory.

80.2 Summary of 40 Gigabit and 100 Gigabit Ethernet sublayers

80.2.1 Reconciliation Sublayer (RS) and Media Independent Interface

The Media Independent Interface (Clause 81) provides a logical interconnection between the MAC sublayer and Physical Layer entities (PHY). The Media Independent Interface is not intended to be physically instantiated, rather it can logically connect layers within a device.

The XLGMII supports 40 Gb/s and CGMII supports 100 Gb/s operation through its 64-bit-wide transmit and receive data paths. The Reconciliation Sublayer (RS) provides a mapping between the signals provided at the Media Independent Interface (XLGMII and CGMII) and the MAC/PLS service definition.

While XLGMII and CGMII are optional interfaces, they are used extensively in this standard as a basis for functional specification and provides a common service interface for the physical coding sublayers defined in Clause 82.

80.2.2 Physical Coding Sublayer (PCS)

The terms 40GBASE-R and 100GBASE-R refer to a specific family of Physical Layer implementations based upon the 64B/66B data coding method specified in Clause 82 and the PMA specification defined in Clause 83. The 40GBASE-R and 100GBASE-R PCSs perform encoding (decoding) of data from (to) the XLGMII/CGMII to 64B/66B code blocks, distribute the data to multiple lanes, and transfer the encoded data to the PMA.

80.2.3 Forward Error Correction (FEC) sublayer

The Forward Error Correction sublayer is an optional sublayer for 40GBASE-R and 100GBASE-R copper and backplane PHYs. The FEC sublayer can be placed in between the PCS and PMA sublayers or between two PMA sublayers, is instantiated for each PCS lane, and operates autonomously on a per PCS lane basis.

The FEC sublayer is specified in Clause 74.

80.2.4 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of a range of physical media. The 40GBASE-R and 100GBASE-R PMAs perform the mapping of transmit and receive data streams between the PCS and PMA via the PMA service interface, and the mapping and multiplexing of transmit and receive data streams between the PMA and PMD via the PMD service interface. In addition, the PMAs perform retiming of the received data stream when appropriate, optionally provide data loopback at the PMA or PMD service interface, and optionally provide test pattern generation and checking.

The 40GBASE-R and 100GBASE-R PMAs are specified in Clause 83.

80.2.5 Physical Medium Dependent (PMD) sublayer

The Physical Medium Dependent sublayer is responsible for interfacing to the transmission medium. The PMD is located just above the Medium Dependent Interface (MDI). The MDI, logically subsumed within each PMD subclause, is the actual medium attachment for the various supported media.

The 40GBASE-R and 100GBASE-R PMDs and their corresponding media are specified in Clause 84 through Clause 88.

80.2.6 Auto-Negotiation

Auto-Negotiation provides a linked device with the capability to detect the abilities (modes of operation) supported by the device at the other end of the link, determine common abilities, and configure for joint operation.

Clause 73 Auto-Negotiation is used by the 40 Gb/s backplane PHY (40GBASE-KR4, see Clause 84) and the 40 Gb/s and 100 Gb/s copper PHYs (40GBASE-CR4 and 100GBASE-CR10, see Clause 85).

80.2.7 Management interface (MDIO/MDC)

The optional MDIO/MDC management interface (Clause 45) provides an interconnection between MDIO Manageable Devices (MMDs) and Station Management (STA) entities.

80.2.8 Management

Managed objects, attributes, and actions are defined for all 40 Gigabit and 100 Gigabit Ethernet components. Clause 30 consolidates all IEEE 802.3 management specifications so that 10/100/1000 Mb/s, 10 Gb/s, 40 Gb/s, and 100 Gb/s agents can be managed by existing network management stations with little or no modification to the agent code.

80.3 Service interface specification method and notation

The service interface specification for 40GBASE-R and 100GBASE-R Physical Layers is as per the definition in 1.2.2. Note that the 40GBASE-R and 100GBASE-R inter-sublayer service interfaces use multiple scalar REQUEST and INDICATION primitives, to indicate the transfer of multiple independent streams of data units, as explained in 80.3.1 through 80.3.3.

80.3.1 Inter-sublayer service interface

The inter-sublayer service interface is described in an abstract manner and does not imply any particular implementation. The inter-sublayer service interface primitives are defined as follows:

```
IS_UNITDATA_i.request
IS_UNITDATA_i.indication
IS_SIGNAL.indication
```

The IS_UNITDATA_i.request (where $i = 0$ to $n - 1$, and n is the number of streams of data units) primitive is used to define the transfer of multiple streams of data units from a sublayer N to the next lower sublayer $N - 1$. The IS_UNITDATA_i.indication (where $i = 0$ to $n - 1$, and n is the number of streams of data units) primitive is used to define the transfer of multiple streams of data units from a sublayer $N - 1$ to the next higher sublayer N . The IS_SIGNAL.indication primitive is used to define the transfer of signal status from a sublayer $N - 1$ to the next higher sublayer N .

80.3.2 Instances of the Inter-sublayer service interface

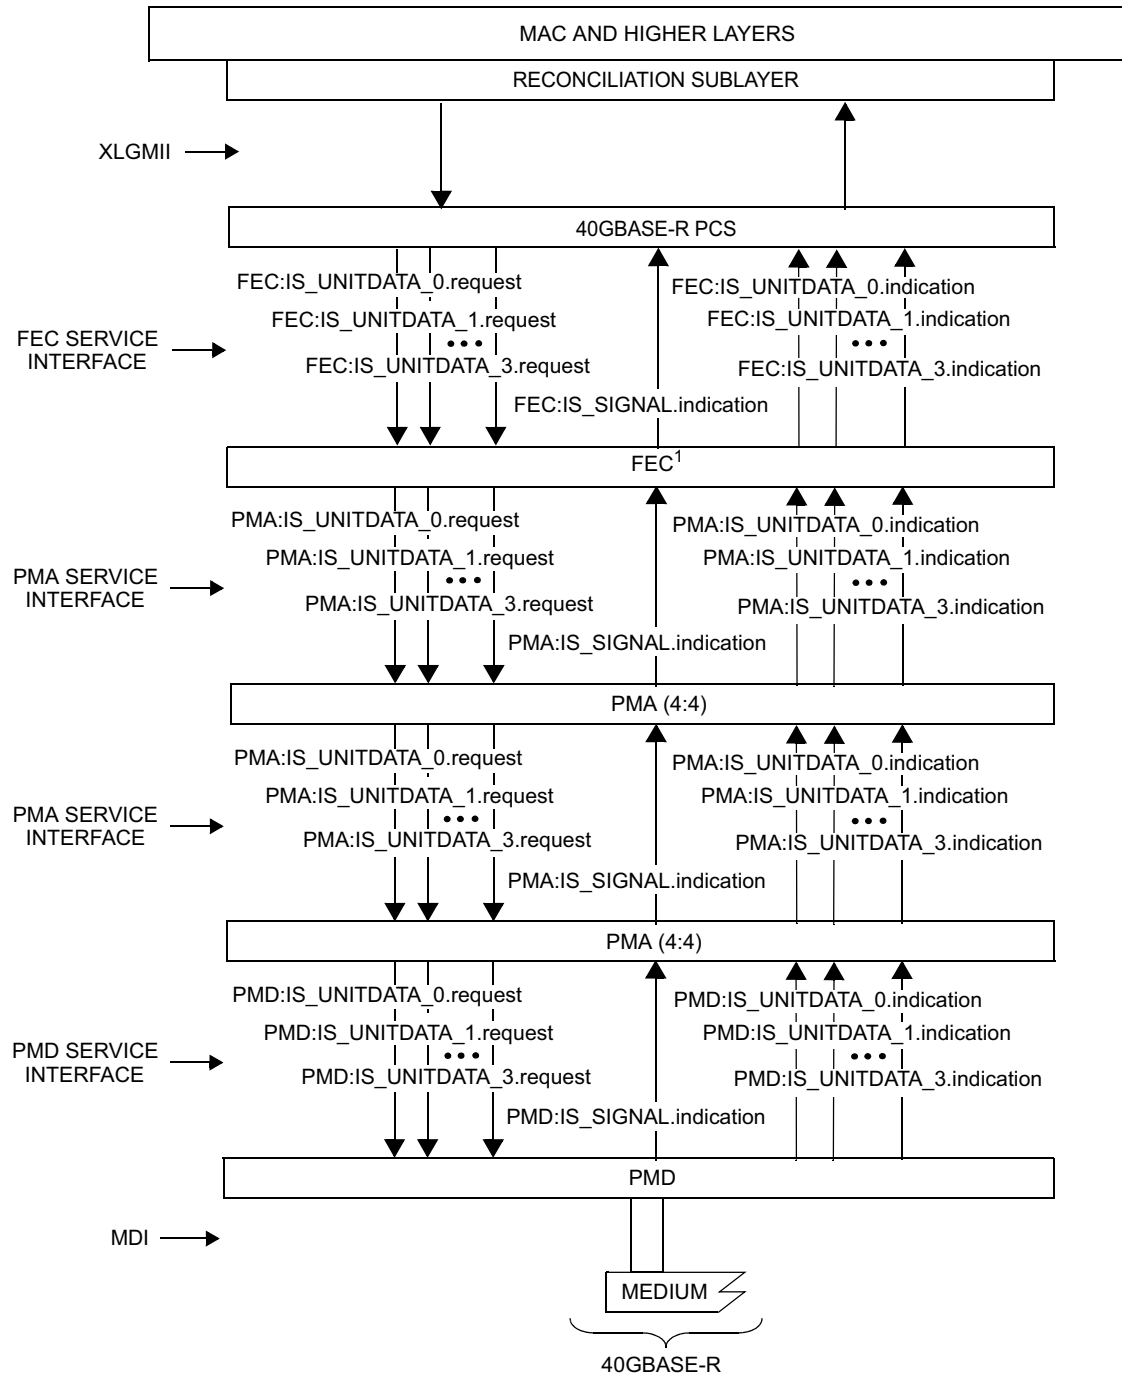
The inter-sublayer interface can be instantiated between different sublayers, hence a prefix notation is defined to identify a specific instance of an inter-sublayer service interface. The following prefixes are defined:

- a) PMD:—for primitives issued on the interface between the PMD sublayer and the PMA sublayer called the PMD service interface.
- b) PMA:—for primitives issued on the interface between the PMA sublayer and the PCS (or the FEC) sublayer called the PMA service interface.
- c) FEC:—for primitives issued on the interface between the FEC sublayer and the PCS (or the PMA) sublayer called the FEC service interface.

Examples of inter-sublayer service interfaces for 40GBASE-R and 100GBASE-R with their corresponding instance names are illustrated in Figure 80–2 and Figure 80–3. For example, the primitives for one instance of the inter-sublayer service interface, named the PMD service interface, are identified as follows:

```
PMD:IS_UNITDATA_i.request
PMD:IS_UNITDATA_i.indication
PMD:IS_SIGNAL.indication.
```

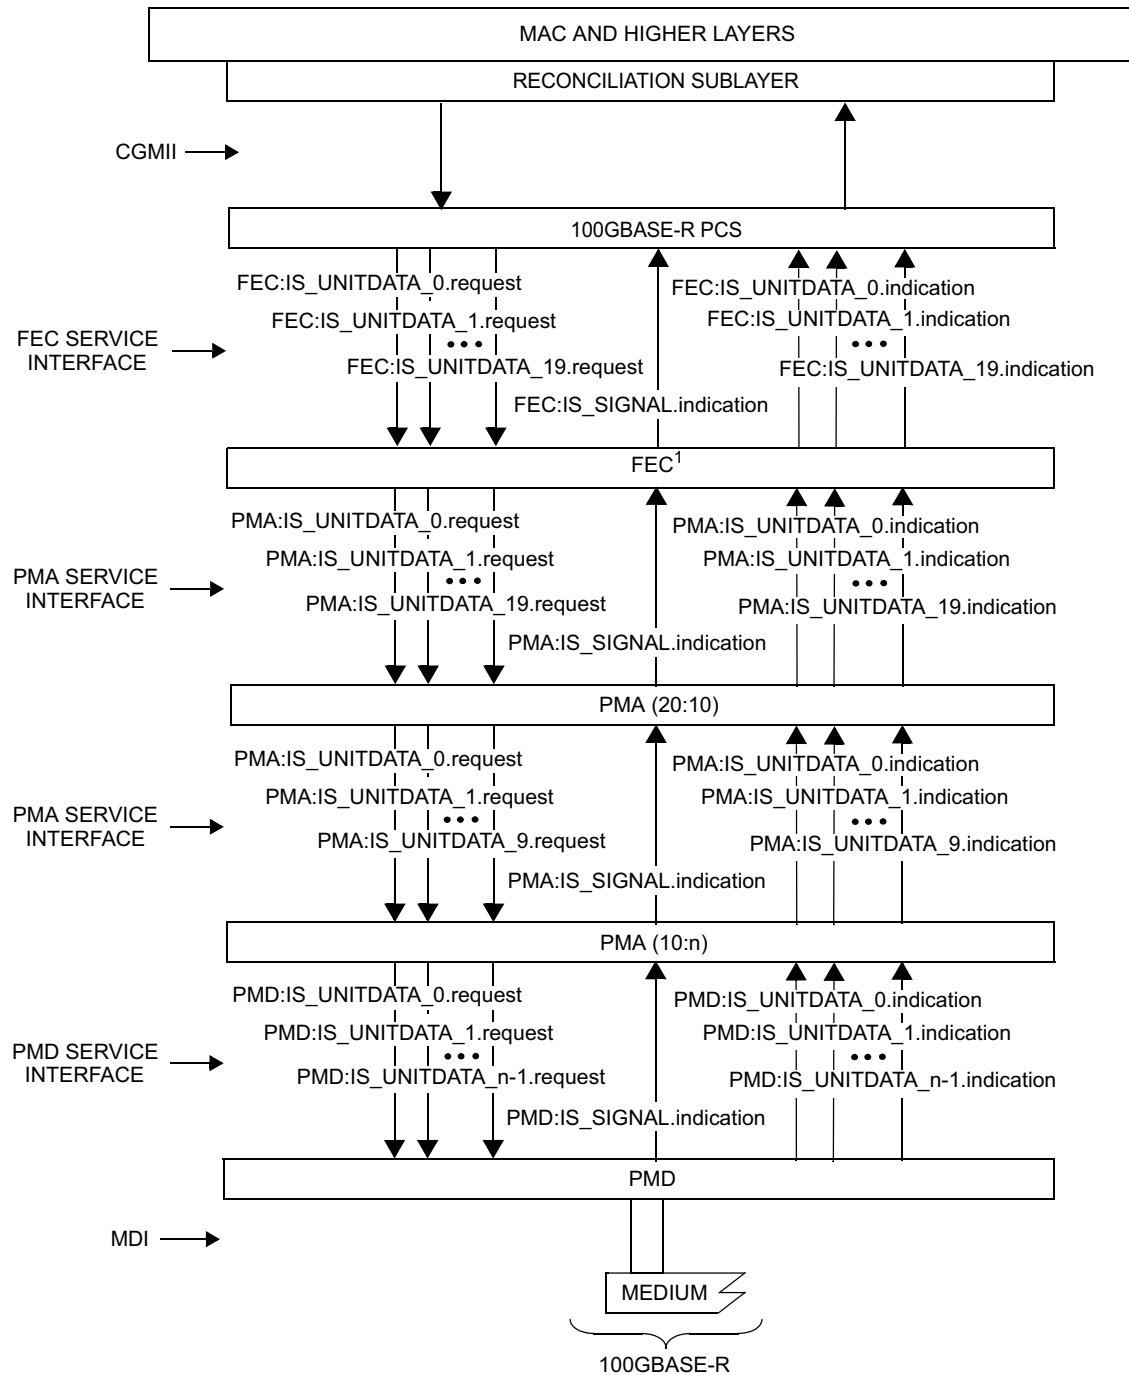
Primitives for other instances, of inter-sublayer interfaces, are represented in a similar manner as described above.



FEC = FORWARD ERROR CORRECTION
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
 NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE

Figure 80-2—40GBASE-R inter-sublayer service interfaces



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE PMD = PHYSICAL MEDIUM DEPENDENT
 FEC = FORWARD ERROR CORRECTION n = NUMBER OF PARALLEL STREAMS OF DATA UNITS
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
 PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT

Figure 80-3—100GBASE-R inter-sublayer service interfaces

80.3.3 Semantics of inter-sublayer service interface primitives

The semantics of the inter-sublayer service interface primitives for the 40GBASE-R and 100GBASE-R sublayers are described in 80.3.3.1 through 80.3.3.3.

80.3.3.1 IS_UNITDATA_*i*.request

The IS_UNITDATA_*i*.request (where $i = 0$ to $n - 1$) primitive is used to define the transfer of multiple streams of data units from a sublayer N to the next lower sublayer $N - 1$, where n is the number of parallel streams of data units.

80.3.3.1.1 Semantics of the service primitive

```
IS_UNITDATA_0.request(tx_bit)
IS_UNITDATA_1.request(tx_bit)
...
IS_UNITDATA_n-1.request(tx_bit)
```

The data conveyed by IS_UNITDATA_0.request to IS_UNITDATA_n-1.request consists of n parallel continuous streams of encoded bits, one stream for each lane. Each of the tx_bit parameters can take one of two values: one or zero.

80.3.3.1.2 When generated

The sublayer N continuously sends n parallel bit streams IS_UNITDATA_*i*.request(tx_bit) to the next lower sublayer $N - 1$, each at a nominal signaling rate defined by a specific instance of the inter-sublayer service interface.

80.3.3.1.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

80.3.3.2 IS_UNITDATA_*i*.indication

The IS_UNITDATA_*i*.indication (where $i = 0$ to $n - 1$) primitive is used to define the transfer of multiple streams of data units from the sublayer $N - 1$ to the next higher sublayer N , where n is the number of parallel streams of data units.

80.3.3.2.1 Semantics of the service primitive

```
IS_UNITDATA_0.indication(rx_bit)
IS_UNITDATA_1.indication(rx_bit)
...
IS_UNITDATA_n-1.indication(rx_bit)
```

The data conveyed by IS_UNITDATA_0.indication to IS_UNITDATA_n-1.indication consists of n parallel continuous streams of encoded bits, one stream for each lane. Each of the rx_bit parameters can take one of two values: one or zero.

80.3.3.2.2 When generated

The sublayer N-1 continuously sends n parallel bit streams IS_UNITDATA_i.indication(rx_bit) to the next higher sublayer N, each at a nominal signaling rate defined by a specific instance of the inter-sublayer service interface.

80.3.3.2.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

80.3.3.3 IS_SIGNAL.indication

The IS_SIGNAL.indication primitive is generated by the sublayer N - 1 to the next higher sublayer N to indicate the status of the receive process. This primitive is generated by the receive process to propagate the detection of severe error conditions (e.g., no valid signal being received by the sublayer that generates this primitive) to the next higher sublayer N.

80.3.3.3.1 Semantics of the service primitive

IS_SIGNAL.indication(SIGNAL_OK)

The SIGNAL_OK parameter can take on one of two values: OK or FAIL. A value of FAIL denotes that invalid data is being presented (rx_bit parameters undefined) by the sublayer N - 1 to the next higher sublayer N. A value of OK does not guarantee valid data is being presented by the sublayer N - 1 to the next higher sublayer N.

80.3.3.3.2 When generated

The sublayer N - 1 generates the IS_SIGNAL.indication primitive to the next higher sublayer N whenever there is change in the value of the SIGNAL_OK parameter.

80.3.3.3.3 Effect of receipt

The effect of receipt of this primitive is defined by the sublayer that receives this primitive.

80.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementors must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices. Table 80-3 contains the values of maximum sublayer delay (sum of transmit and receive delays at one end of the link) in bit times as specified in 1.4 and pause_quanta as specified in 31B.2. If a PHY contains an Auto-Negotiation sublayer, the delay of the Auto-Negotiation sublayer is included within the delay of the PMD and medium.

See 44.3 for the calculation of bit time per meter of fiber or electrical cable.

See 31B.3.7 for PAUSE reaction timing constraints for stations at operating speeds of 40 Gb/s and 100 Gb/s.

Table 80–3—Sublayer delay constraints

Sublayer	Maximum (bit time) ^a	Maximum (pause_quanta) ^b	Maximum (ns)	Notes ^c
40G MAC, RS, and MAC Control	16384	32	409.6	See 81.1.4.
40GBASE-R PCS	11264	22	281.6	See 82.5.
40GBASE-R FEC	24576	48	614.4	See 74.6.
40GBASE-R PMA	4096	8	102.4	See 83.5.4.
40GBASE-KR4 PMD	2048	4	51.2	Includes delay of one direction through backplane medium. See 84.4.
40GBASE-CR4 PMD	6144	12	153.6	Includes delay of one direction through cable medium. See 85.4.
40GBASE-SR4 PMD	1024	2	25.6	Includes 2 m of fiber. See 86.3.1.
40GBASE-LR4 PMD	1024	2	25.6	Includes 2 m of fiber. See 87.3.1.
100G MAC, RS, and MAC Control	24576	48	245.76	See 81.1.4.
100GBASE-R PCS	35328	69	353.28	See 82.5.
100GBASE-R FEC	122880	240	1228.8	See 74.6.
100GBASE-R PMA	9216	18	92.16	See 83.5.4.
100GBASE-CR10 PMD	14848	29	148.48	Includes delay of one direction through cable medium. See 85.4.
100GBASE-SR10 PMD	2048	4	20.48	Includes 2 m of fiber. See 86.3.1.
100GBASE-LR4 PMD	2048	4	20.48	Includes 2 m of fiber. See 88.3.1.
100GBASE-ER4 PMD	2048	4	20.48	Includes 2 m of fiber. See 88.3.1.

^a Note that for 40GBASE-R, 1 bit time (BT) is equal to 25 ps and for 100GBASE-R, 1 bit time (BT) is equal to 10 ps. (See 1.4.81 for the definition of bit time.)

^b Note that for 40GBASE-R, 1 pause_quantum is equal to 12.8 ns and for 100GBASE-R, 1 pause_quantum is equal to 5.12 ns. (See 31B.2 for the definition of pause_quanta.)

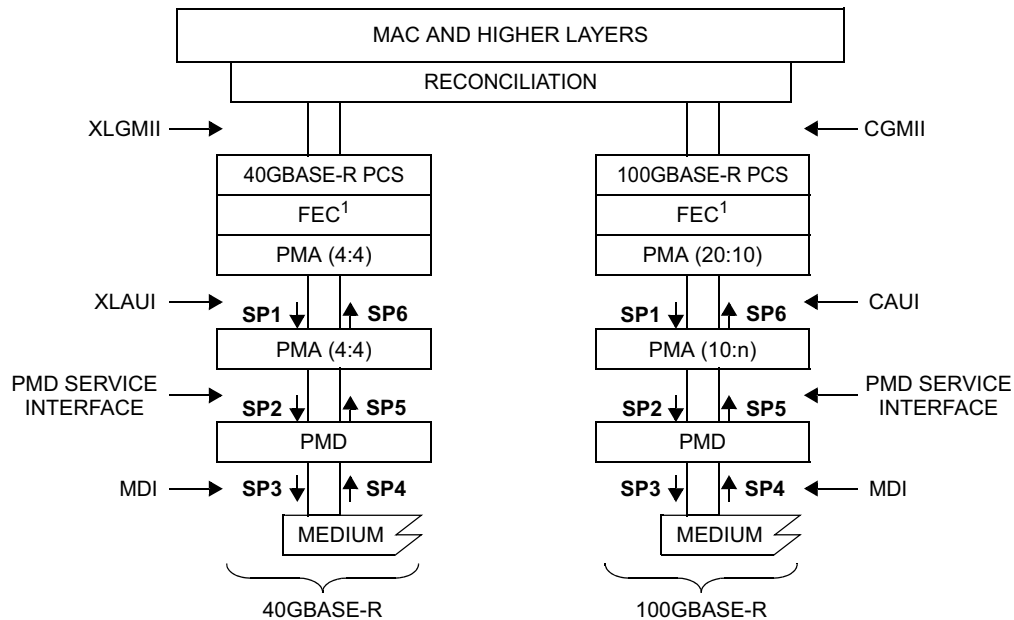
^c Should there be a discrepancy between this table and the delay requirements of the relevant sublayer clause, the sublayer clause prevails.

80.5 Skew constraints

Skew (or relative delay) can be introduced between lanes by both active and passive elements of a 40GBASE-R or 100GBASE-R link. Skew is defined as the difference between the times of the earliest PCS lane and latest PCS lane for the one to zero transition of the alignment marker sync bits. The PCS deskew

function (see 82.2.12) compensates for all lane-to-lane Skew observed at the receiver. The Skew between the lanes must be kept within limits as shown in Table 80–4 so that the transmitted information on the lanes can be reassembled by the receive PCS.

Skew Variation may be introduced due to variations in electrical, thermal or environmental characteristics. Skew Variation is defined as the change in Skew between any PCS lane and any other PCS lane over the entire time that the link is in operation. From the time the link is brought up, Skew Variation must be limited to ensure that each PCS lane always traverses the same lane between any pair of adjacent sublayers while the link remains in operation.



CAUI = 100 Gb/s ATTACHMENT UNIT INTERFACE
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT
 XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
 n = 4 or 10

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE

Figure 80–4—40GBASE-R and 100GBASE-R Skew points for single XLAUI or CAUI

The maximum Skew and Skew Variation at physically instantiated interfaces is specified at Skew points SP1, SP2, and SP3 for the transmit direction and SP4, SP5, and SP6 for the receive direction as illustrated in Figure 80–4 (single XLAUI or CAUI interface) and Figure 80–5 (multiple XLAUI or CAUI interfaces).

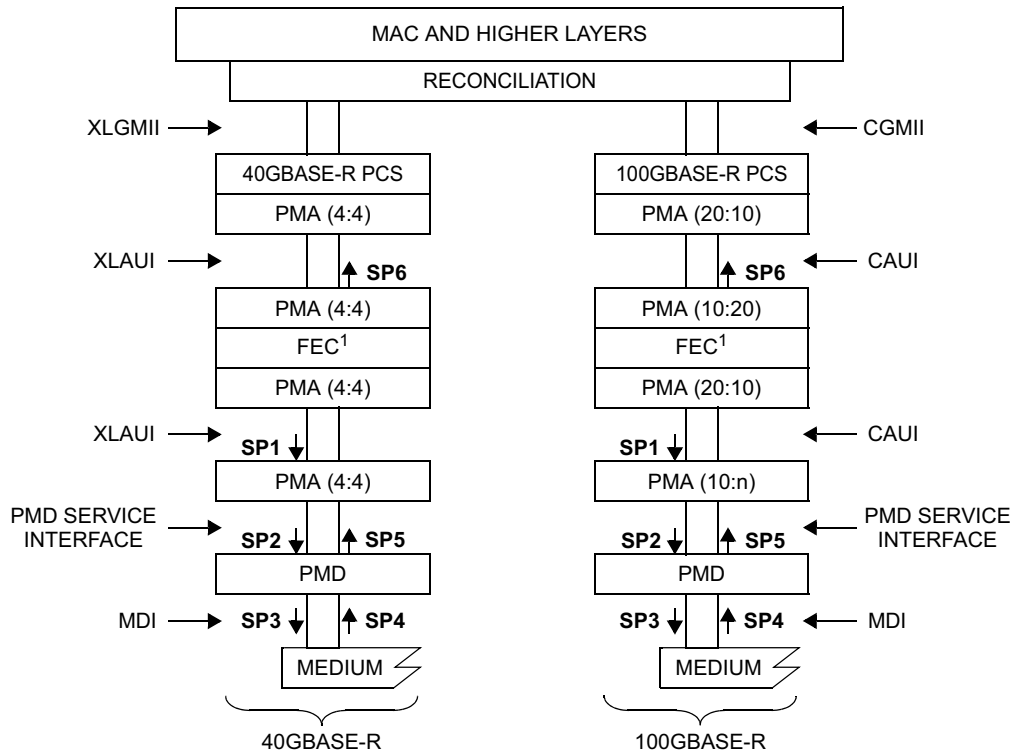
In the transmit direction, the Skew points are defined in the following locations:

- SP1 on the XLAUI/CAUI interface, at the input of the PMA closest to the PMD;
- SP2 on the PMD service interface, at the input of the PMD;
- SP3 at the output of the PMD, at the MDI.

In the receive direction, the Skew points are defined in the following locations:

- SP4 at the MDI, at the input of the PMD;
- SP5 on the PMD service interface, at the output of the PMD;
- SP6 on the XLAUI/CAUI interface, at the output of the PMA closest to the PCS.

The allowable limits for Skew are shown in Table 80–4 and the allowable limits for Skew Variation are shown in Table 80–5.



CAUI = 100 Gb/s ATTACHMENT UNIT INTERFACE
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIA DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT
 XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
 n = 4 or 10

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE

Figure 80–5—40GBASE-R and 100GBASE-R Skew points for multiple XLAUI or CAUI

The Skew requirements for the PCS, PMA and PMD sublayers are specified in the respective clauses as noted in Table 80–4 and Table 80–5.

Table 80–4—Summary of Skew constraints

Skew points	Maximum Skew (ns) ^a	Maximum Skew for 40GBASE-R PCS lane (UI) ^b	Maximum Skew for 100GBASE-R PCS lane (UI) ^c	Notes ^d
SP1	29	≈ 299	≈ 150	See 83.5.3.1
SP2	43	≈ 443	≈ 222	See 83.5.3.3 or 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP3	54	≈ 557	≈ 278	See 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP4	134	≈ 1382	≈ 691	See 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP5	145	≈ 1495	≈ 748	See 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP6	160	≈ 1649	≈ 824	See 83.5.3.5
At PCS receive	180	≈ 1856	≈ 928	See 82.2.12

^aThe Skew limit includes 1 ns allowance for PCB traces that are associated with the Skew points.

^bNote that ≈ indicates approximate equivalent of maximum Skew in UI for 40GBASE-R, based on 1 UI equals 96.969697 ps at PCS lane signaling rate of 10.3125 GBd.

^cNote that ≈ indicates approximate equivalent of maximum Skew in UI for 100GBASE-R, based on 1 UI equals 193.939394 ps at PCS lane signaling rate of 5.15625 GBd.

^dShould there be a discrepancy between this table and the Skew requirements of the relevant sublayer clause, the sub-layer clause prevails.

Table 80–5—Summary of Skew Variation constraints

Skew points	Maximum Skew Variation (ns)	Maximum Skew Variation for 10.3125 GBd PMD lane (UI) ^a	Maximum Skew Variation for 25.78125 GBd PMD lane (UI) ^b	Notes ^c
SP1	0.2	≈ 2	N/A	See 83.5.3.1
SP2	0.4	≈ 4	≈ 10	See 83.5.3.3 or 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2.
SP3	0.6	≈ 6	≈ 15	See 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP4	3.4	≈ 35	≈ 88	See 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP5	3.6	≈ 37	≈ 93	See 83.5.3.4 or 84.5 or 85.5 or 86.3.2 or 87.3.2 or 88.3.2
SP6	3.8	≈ 39	N/A	See 83.5.3.5
At PCS receive	4	≈ 41	N/A	See 82.2.12

^aNote that ≈ indicates approximate equivalent of maximum Skew Variation in UI for 40GBASE-R, based on 1 UI equals 96.969697 ps at PMD lane signaling rate of 10.3125 GBd.

^bNote that ≈ indicates approximate equivalent of maximum Skew Variation in UI for 100GBASE-R, based on 1 UI equals 38.787879 ps at PMD lane signaling rate of 25.78125 GBd.

^cShould there be a discrepancy between this table and the Skew requirements of the relevant sublayer clause, the sub-layer clause prevails.

80.6 State diagrams

State diagrams take precedence over text.

The conventions of 1.2 are adopted, along with the extensions listed in 21.5.

Multiple states of a function that have a transition to a common state utilizing different qualifiers (for example, multiple exit conditions to an IDLE or WAIT state) may be indicated by a shared arrow. An exit transition arrow must connect to the shared arrow, and the qualifier must be met prior to termination of the transition arrow on the shared arrow. The shared arrow has no qualifier.

80.7 Protocol implementation conformance statement (PICS) proforma

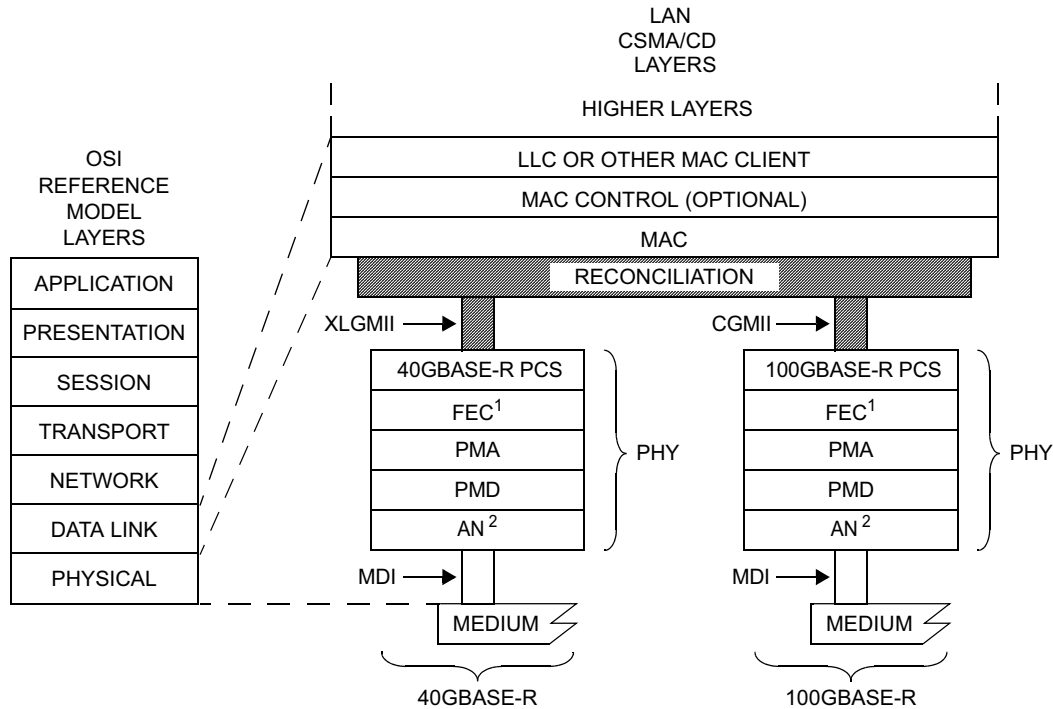
The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, Clause 45, Clause 73, Clause 74, Clause 81 through Clause 88, and related annexes demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma.

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. A PICS is included at the end of each clause as appropriate. Each of the 40 Gigabit and 100 Gigabit Ethernet PICS conforms to the same notation and conventions used in 21.6.

81. Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation (XLGMII and CGMII)

81.1 Overview

This clause defines the characteristics of the Reconciliation Sublayer (RS) and the Media Independent Interface between CSMA/CD media access controllers and various PHYs. Figure 81–1 shows the relationship of the RS and Media Independent Interface to the ISO/IEC OSI reference model. Note that there are two variants of the Media Independent Interface in this clause, the 40 Gb/s Media Independent Interface (XLGMII) and the 100 Gb/s Media Independent Interface (CGMII).



AN = AUTO-NEGOTIATION

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE

FEC = FORWARD ERROR CORRECTION

LLC = LOGICAL LINK CONTROL

MAC = MEDIA ACCESS CONTROL

MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE

PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT

XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE

NOTE 2—CONDITIONAL BASED ON PHY TYPE

Figure 81–1—RS and MII relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 CSMA/CD LAN model

The XLGMII and the CGMII are optional logical interfaces between the MAC sublayer and the Physical Layer (PHY).

The RS adapts the bit serial protocols of the MAC to the parallel format of the PCS service interface. Though the XLGMII/CGMII is an optional interface, it is used in this standard as a basis for specification. The Physical Coding Sublayer (PCS) is specified to the XLGMII/CGMII, so if not implemented, a conforming implementation shall behave functionally as if the RS and XLGMII/CGMII were implemented.

The XLGMII/CGMII has the following characteristics:

- a) The XLGMII interface supports a speed of 40 Gb/s.
- b) The CGMII interface supports a speed of 100 Gb/s.
- c) Data and delimiters are synchronous to a clock reference.
- d) It provides independent 64-bit wide transmit and receive data paths.
- e) It supports full duplex operation only.

81.1.1 Summary of major concepts

The following are the major concepts of the XLGMII/CGMII:

- a) The XLGMII/CGMII is functionally similar to other media independent interfaces that have been defined for lower speeds, as they all define an interface allowing independent development of MAC and PHY logic.
- b) The RS converts between the MAC serial data stream and the parallel data paths of the XLGMII/CGMII.
- c) The RS maps the signal set provided at the XLGMII/CGMII to the PLS service primitives provided at the MAC.
- d) Each direction of data transfer is independent and serviced by data, control, and clock signals.
- e) The RS generates continuous data or control characters on the transmit path and expects continuous data or control characters on the receive path.
- f) The RS participates in link fault detection and reporting by monitoring the receive path for status reports that indicate an unreliable link, and generating status reports on the transmit path to report detected link faults to the DTE on the remote end of the connecting link.

81.1.2 Application

This clause applies to the interface between the MAC and PHY. This logical interface is used to provide media independence so that an identical media access controller may be used with supported PHY types.

81.1.3 Rate of operation

The XLGMII is specified to support 40 Gb/s operation and the CGMII is specified to support 100 Gb/s operation.

81.1.4 Delay constraints

The maximum cumulative MAC Control, MAC, and RS delay (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 81–1. A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

Table 81–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
40 Gb/s MAC, RS, and MAC Control	16384	32	409.6
100 Gb/s MAC, RS, and MAC Control	24576	48	245.76

81.1.5 Allocation of functions

The allocation of functions at the XLGMII/CGMII balances the need for media independence with interface simplicity. The XLGMII and CGMII maximize media independence by separating the Data Link and Physical Layers of the OSI seven-layer reference model.

81.1.6 XLGMII/CGMII structure

The XLGMII/CGMII is composed of independent transmit and receive paths. Each direction uses 64 data signals (TXD<63:0> and RXD<63:0>), 8 control signals (TXC<7:0> and RXC<7:0>), and a clock (TX_CLK and RX_CLK). Figure 81–2 depicts a schematic view of the RS inputs and outputs.




Figure 81–2—Reconciliation Sublayer (RS) inputs and outputs

The 64 TXD and 8 TXC signals shall be organized into eight data lanes, as shall the 64 RXD and 8 RXC signals (see Table 81–2). The eight lanes in each direction share a common clock, TX_CLK for transmit and RX_CLK for receive. The eight lanes are used in round-robin sequence to carry an octet stream. On transmit, each eight PLS_DATA.request transactions represent an octet transmitted by the MAC. The first octet is aligned to lane 0, the second to lane 1, the third to lane 2, the fourth to lane 3, the fifth to lane 4, the sixth to lane 5, the seventh to lane 6 and the eighth to lane 7, then repeating with the ninth to lane 0, etc. Delimiters and interframe idle characters are encoded on the TXD and RXD signals with the control code indicated by assertion of TXC and RXC, respectively.

Table 81–2—Transmit and receive lane associations

TXD, RXD	TXC, RXC	Lane
<7:0>	<0>	0
<15:8>	<1>	1
<23:16>	<2>	2
<31:24>	<3>	3
<39:32>	<4>	4
<47:40>	<5>	5
<55:48>	<6>	6
<63:56>	<7>	7

81.1.7 Mapping of XLGMII/CGMII signals to PLS service primitives

The Reconciliation Sublayer (RS) shall map the signals provided at the XLGMII/CGMII to the PLS service primitives defined in Clause 6. The PLS service primitives provided by the RS and described here behave in exactly the same manner as defined in Clause 6. Full duplex operation only is implemented at 40 Gb/s and 100 Gb/s; therefore, PLS service primitives supporting CSMA/CD operation are not mapped through the RS to the XLGMII/CGMII.

Mappings for the following primitives are defined for 40 Gb/s and 100 Gb/s operation:

- PLS_DATA.request
- PLS_DATA.indication
- PLS_CARRIER.indication
- PLS_SIGNAL.indication
- PLS_DATA_VALID.indication

81.1.7.1 Mapping of PLS_DATA.request

81.1.7.1.1 Function

The RS maps the primitive PLS_DATA.request to the XLGMII/CGMII signals TXD<63:0>, TXC<7:0>, and TX_CLK.

81.1.7.1.2 Semantics of the service primitive

PLS_DATA.request(OUTPUT_UNIT)

The OUTPUT_UNIT parameter can take one of three values: one, zero, or DATA_COMPLETE. One or zero represents a single data bit. The DATA_COMPLETE value signifies that the MAC sublayer has no more data to output.

81.1.7.1.3 When generated

This primitive is generated by the MAC sublayer to request the transmission of a single data bit on the physical medium or to stop transmission.

81.1.7.1.4 Effect of receipt

The OUTPUT_UNIT values are conveyed to the PHY by the signals TXD<63:0> and TXC<7:0> on each TX_CLK rising edge. Each PLS_DATA.request transaction shall be mapped to a TXD signal in sequence (TXD<0>, TXD<1>, ... TXD<63>, TXD<0>) as described in 81.2, and for every eight transactions, a TXC signal in sequence (TXC<0>, TXC<1>, ... TXC<7>, TXC<0>) is generated, as described in 81.3.1.2. After 64 PLS_DATA.request transactions from the MAC sublayer (eight octets of eight PLS_DATA.request transactions each), the RS requests transmission of 64 data bits by the PHY. The first octet of preamble shall be converted to a Start control character and aligned to lane 0. The TXD<63:0> and TXC<7:0> shall be generated by the RS every 64 bit-times of the MAC sublayer.

The DATA_COMPLETE value shall be mapped to a Terminate control character encoded on the next eight TXD signals in sequence after the last data octet; and is transferred to the PHY at the next TX_CLK rising edge. This may be on the same TX_CLK rising edge as the last data octet or the subsequent TX_CLK rising edge. When the Terminate control character is in lane 0, 1, 2, 3, 4, 5, or 6, the lanes following in sequence are encoded with an Idle control character.

81.1.7.2 Mapping of PLS_DATA.indication

81.1.7.2.1 Function

The RS maps the XLGMII/CGMII signals RXD<63:0>, RXC<7:0>, and RX_CLK to the primitive PLS_DATA.indication.

81.1.7.2.2 Semantics of the service primitive

PLS_DATA.indication (INPUT_UNIT)

The INPUT_UNIT parameter can take one of two values: one or zero. It represents a single data bit.

81.1.7.2.3 When generated

The INPUT_UNIT values are derived from the signals RXC<7:0> and RXD<63:0> received from the PHY on each rising edge of RX_CLK. Each primitive generated to the MAC sublayer entity corresponds to a PLS_DATA.request issued by the MAC at the remote end of the link connecting two DTEs. For each RXD<63:0> during frame reception, the RS shall generate 64 PLS_DATA.indication transactions until the end of frame (Terminate control character), where 0, 8, 16, 24, 32, 40, 48, or 56 PLS_DATA.indication transactions will be generated from the RXD<63:0> containing the Terminate. During frame reception, each RXD signal shall be mapped in sequence into a PLS_DATA.indication transaction (RXD<0>, RXD<1>, ... RXD<63>, RXD<0>) as described in 81.2.

The RS shall convert a valid Start control character to a preamble octet prior to generation of the associated PLS_DATA.indication transactions. The RS shall not generate any PLS_DATA.indication primitives for a Terminate control character. To assure robust operation, the value of the data transferred to the MAC may be changed by the RS as required by XLGMII/CGMII error indications (see 81.3.3). Sequence ordered_sets are not indicated to the MAC (see 81.3.4).

81.1.7.2.4 Effect of receipt

The effect of receipt of this primitive by the MAC sublayer is unspecified.

81.1.7.3 Mapping of PLS_CARRIER.indication

40 Gb/s and 100 Gb/s operation supports full duplex operation only. The RS never generates this primitive.

81.1.7.4 Mapping of PLS_SIGNAL.indication

40 Gb/s and 100 Gb/s operation supports full duplex operation only. The RS never generates this primitive.

81.1.7.5 Mapping of PLS_DATA_VALID.indication

81.1.7.5.1 Function

The RS maps the XLGMII/CGMII signals RXC<7:0> and RXD<63:0> to the primitive PLS_DATA_VALID.indication.

81.1.7.5.2 Semantics of the service primitive

PLS_DATA_VALID.indication (DATA_VALID_STATUS)

The `DATA_VALID_STATUS` parameter can take one of two values: `DATA_VALID` or `DATA_NOT_VALID`. The `DATA_VALID` value indicates that the `INPUT_UNIT` parameter of the `PLS_DATA.indication` primitive contains valid data of an incoming frame. The `DATA_NOT_VALID` value indicates that the `INPUT_UNIT` parameter of the `PLS_DATA.indication` primitive does not contain valid data of an incoming frame.

81.1.7.5.3 When generated

The `PLS_DATA_VALID.indication` service primitive shall be generated by the RS whenever the `DATA_VALID_STATUS` parameter changes from `DATA_VALID` to `DATA_NOT_VALID` or vice versa.

`DATA_VALID_STATUS` shall assume the value `DATA_VALID` when a `PLS_DATA.indication` transaction is generated in response to reception of a Start control character on lane 0 if the prior `RXC<7:0>` and `RXD<63:0>` contained eight Idle characters, a Sequence ordered set, or a Terminate character. `DATA_VALID_STATUS` shall assume the value `DATA_NOT_VALID` when `RXC` of the current lane in sequence is asserted for anything except an Error control character. In the absence of errors, `DATA_NOT_VALID` is caused by a Terminate control character. When `DATA_VALID_STATUS` changes from `DATA_VALID` to `DATA_NOT_VALID` because of a control character other than Terminate, the RS shall ensure that the MAC will detect a `FrameCheckError` prior to indicating `DATA_NOT_VALID` to the MAC (see 81.3.3.1).

81.1.7.5.4 Effect of receipt

The effect of receipt of this primitive by the MAC sublayer is unspecified.

81.2 XLGMII/CGMII data stream

Packets transmitted through the XLGMII/CGMII shall be transferred within the XLGMII/CGMII data stream. A data stream is a sequence of bytes, where each byte conveys either a data octet or control character. The parts of the data stream are shown in Figure 81–3.




Figure 81–3—XLGMII/CGMII data stream

For the XLGMII/CGMII, transmission and reception of each bit and mapping of data octets to lanes shall be as shown in Figure 81-4.

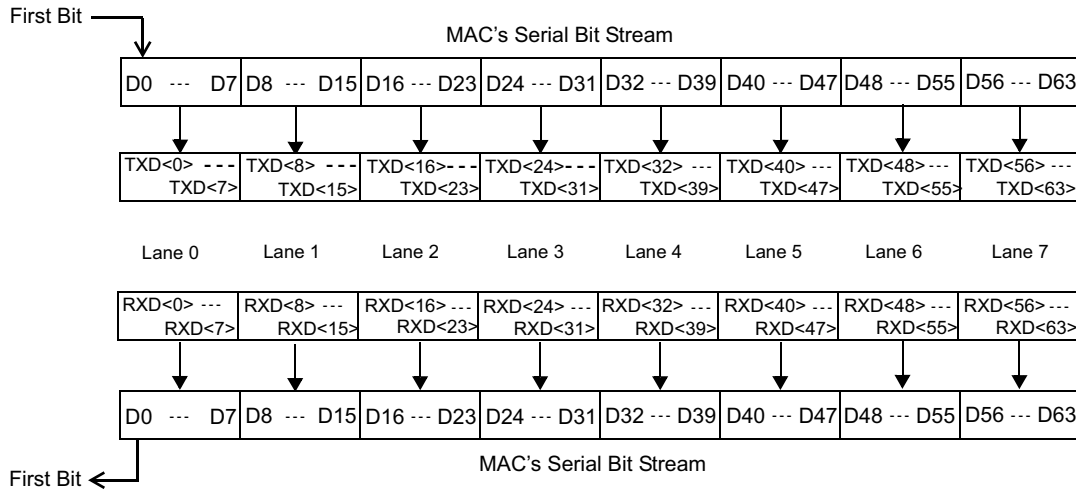


Figure 81-4—Relationship of data lanes to MAC serial bit stream

81.2.1 Inter-frame <inter-frame>

The inter-frame <inter-frame> period on an XLGMII/CGMII transmit or receive path is an interval during which no frame data activity occurs. The <inter-frame> corresponding to the MAC interpacket gap begins with the Terminate control character, continues with Idle control characters and ends with the Idle control character prior to a Start control character. The length of the interpacket gap may be changed between the transmitting MAC and receiving MAC by one or more functions (e.g., RS lane alignment or PHY clock rate compensation). The minimum IPG at the XLGMII/CGMII of the receiving RS is one octet.

The signaling of link status information logically occurs in the <inter-frame> period (see 81.3.4). Frame processing when signaling of link status information is initiated or terminated is described in 81.3.3.

81.2.2 Preamble <preamble> and start of frame delimiter <sfd>

The preamble <preamble> begins a frame transmission by a MAC as specified in 4.2.5 and when generated by a MAC consists of 7 octets with the following bit values:

10101010 10101010 10101010 10101010 10101010 10101010 10101010

The Start control character indicates the beginning of MAC data on the XLGMII/CGMII. On transmit, the RS converts the first data octet of preamble transferred from the MAC into a Start control character. On receive, the RS converts the Start control character into a preamble data octet. The Start control character is aligned to lane 0 of the XLGMII/CGMII by the RS on transmit and by the PHY on receive.

The start of frame delimiter <sfd> indicates the start of a frame and immediately follows the preamble. The bit value of <sfd> at the XLGMII/CGMII is the same as the Start Frame Delimiter (SFD) specified in 4.2.6 and is equal to the following:

10101011

The preamble and SFD are shown previously with their bits ordered for serial transmission from left to right. As shown, the left-most bit of each octet is the LSB of the octet and the right-most bit of each octet is the MSB of the octet.

The preamble and SFD are transmitted through the XLGMII/CGMII as octets sequentially ordered on the lanes of the XLGMII/CGMII. The first preamble octet is replaced with a Start control character and it is aligned to lane 0, the second octet on lane 1, the third on lane 2, the fourth on lane 3, the fifth on lane 4, the sixth on lane 5, the seventh on lane 6, and the SFD on lane 7, and the eight octets are transferred on the next rising edge of TX_CLK. The ninth octet is assigned to lane 0 with subsequent octets sequentially assigned to the lanes. The XLGMII/CGMII <preamble> and <sfd> are as follows:

Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7
Start	10101010	10101010	10101010	10101010	10101010	10101010	10101011

81.2.3 Data <data>

The data <data> in a frame shall consist of a set of data octets.

81.2.4 End of frame delimiter <efd>

Assertion of TXC with the appropriate Terminate control character encoding of TXD on a lane constitutes an end of frame delimiter <efd> for the transmit data stream. Similarly, assertion of RXC with the appropriate Terminate control character encoding of RXD constitutes an end of frame delimiter for the receive data stream. The XLGMII/CGMII shall recognize the end of frame delimiter on any of the eight lanes of the XLGMII/CGMII.

81.2.5 Definition of Start of Packet and End of Packet Delimiters

For the purposes of Clause 30, the Start of Packet Delimiter is defined as the Start control character, and the End of Packet delimiter is defined as the end of the last sequential data octet preceding the Terminate control character or other control character causing a change from DATA_VALID to DATA_NOT_VALID (see 81.1.7.5.2 and 30.3.2.1.5).

81.3 XLGMII/CGMII functional specifications

The XLGMII/CGMII is designed to make the differences among the various media and transceiver combinations transparent to the MAC sublayer. The selection of logical control signals and the functional procedures are all designed to this end.

NOTE—No XLGMII/CGMII loopback is defined, but XLGMII/CGMII signals are specified such that transmit signals may be connected to receive signals to create a loopback path. To do this, TXD<0> is connected to RXD<0> ... TXD<63> to RXD<63>, TXC<0> to RXC<0> ... TXC<7> to RXC<7>, and TXCLK to RXCLK. Such a loopback does not test the Link Fault Signaling state diagram, nor any of the error handling functions of the receive RS.

81.3.1 Transmit

81.3.1.1 TX_CLK

TX_CLK is a continuous clock used for operation at the appropriate frequency. TX_CLK provides the timing reference for the transfer of the TXC<7:0> and TXD<63:0> signals from the RS to the PHY. The values of TXC<7:0> and TXD<63:0> shall be sampled by the PHY on the rising edge of TX_CLK. TX_CLK is sourced by the RS.

The TX_CLK frequency shall be one-sixty-fourth of the MAC transmit data rate.

81.3.1.2 TXC<7:0> (transmit control)

TXC<7:0> indicate that the RS is presenting either data or control characters on the XLGMII/CGMII for transmission. The TXC signal for a lane shall be de-asserted when a data octet is being sent on the corresponding lane and asserted when a control character is being sent. In the absence of errors, the TXC signals are de-asserted by the RS for each octet of the preamble (except the first octet that is replaced with a Start control character) and remain de-asserted while all octets to be transmitted are presented on the lanes of the XLGMII/CGMII. TXC<7:0> are driven by the RS and shall transition synchronously with respect to the rising edge of TX_CLK. Table 81–3 specifies the permissible encodings of TXD and TXC for an XLGMII/CGMII transmit lane. Additional requirements apply for proper code sequences and in which lanes particular codes are valid (e.g., Start control character is to be aligned to lane 0).

Table 81–3—Permissible encodings of TXC and TXD

TXC	TXD	Description	PLS_DATA.request parameter
0	0x00 through 0xFF	Normal data transmission	Zero, one (eight bits)
1	0x00 through 0x06	Reserved	—
1	0x07	Idle	No applicable parameter (normal inter-frame)
1	0x08 through 0x9B	Reserved	—
1	0x9C	Sequence (only valid in lane 0)	No applicable parameter (inter-frame status signal)
1	0x9D through 0xFA	Reserved	—
1	0xFB	Start (only valid in lane 0)	No applicable parameter, replaces first eight zero, one of a frame (preamble octet)
1	0xFC	Reserved	—
1	0xFD	Terminate	DATA_COMPLETE
1	0xFE	Error	No applicable parameter
1	0xFF	Reserved	—

81.3.1.3 TXD<63:0> (transmit data)

TXD is a bundle of 64 data signals organized into eight lanes of eight signals each (TXD<7:0>, TXD<15:8>, TXD<23:16>, TXD<31:24>, TXD<39:32>, TXD<47:40>, TXD<55:48>, and TXD<63:56>) that are driven by the RS. Each lane is associated with a TXC signal as shown in Table 81–2 and shall be encoded as shown in Table 81–3. TXD<63:0> shall transition synchronously with respect to the rising edge of TX_CLK. For each high TX_CLK transition, data and/or control are presented on TXD<63:0> to the PHY for transmission. TXD<0> is the least significant bit of lane 0, TXD<8> the least significant bit of lane 1, TXD<16> the least significant bit of lane 2, TXD<24> the least significant bit of lane 3, TXD<32> the least significant bit of lane 4, TXD<40> the least significant bit of lane 5, TXD<48> the least significant bit of lane 6, and TXD<56> the least significant bit of lane 7.

Assertion on a lane of appropriate TXD values when TXC is asserted will cause the PHY to generate code-groups associated with either Idle, Start, Terminate, Sequence, or Error control characters. While the TXC of a lane is de-asserted, TXD of the lane is used to request the PHY to generate code-groups corresponding to the data octet value of TXD. An example of normal frame transmission is illustrated in Figure 81–5.

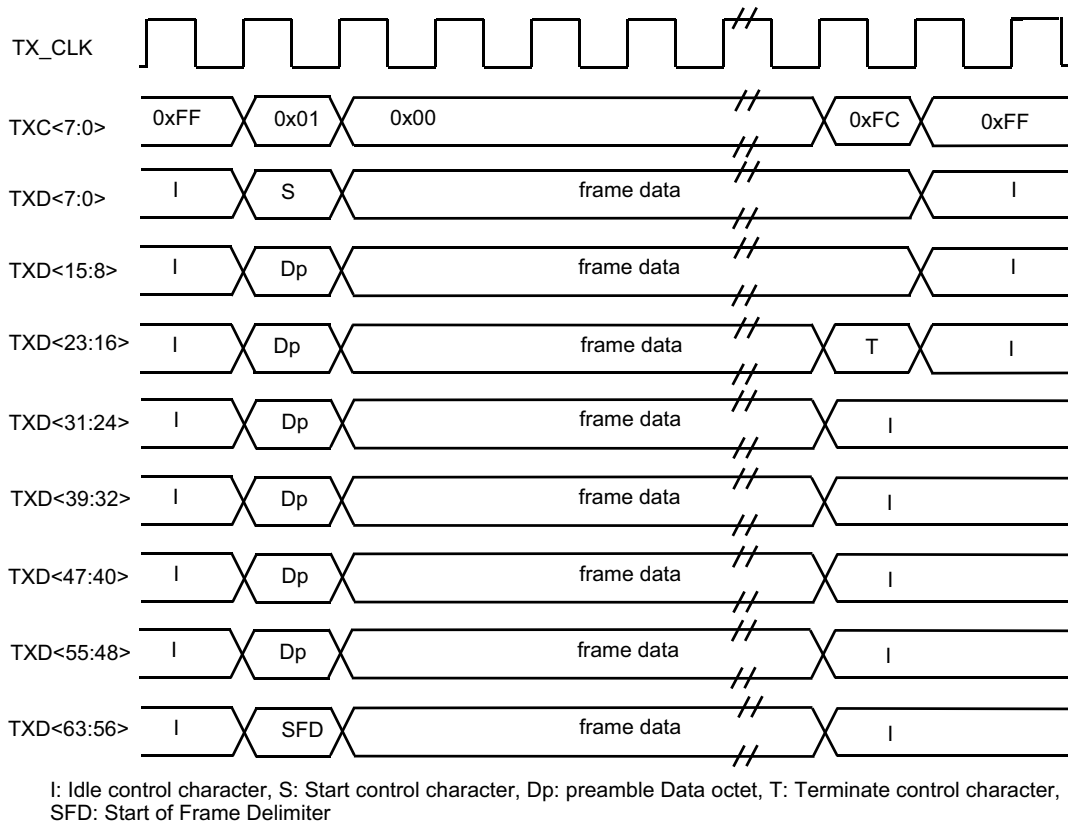


Figure 81-5—Normal frame transmission

Figure 81–6 shows the behavior of TXD and TXC during an example transmission of a frame propagating an error.

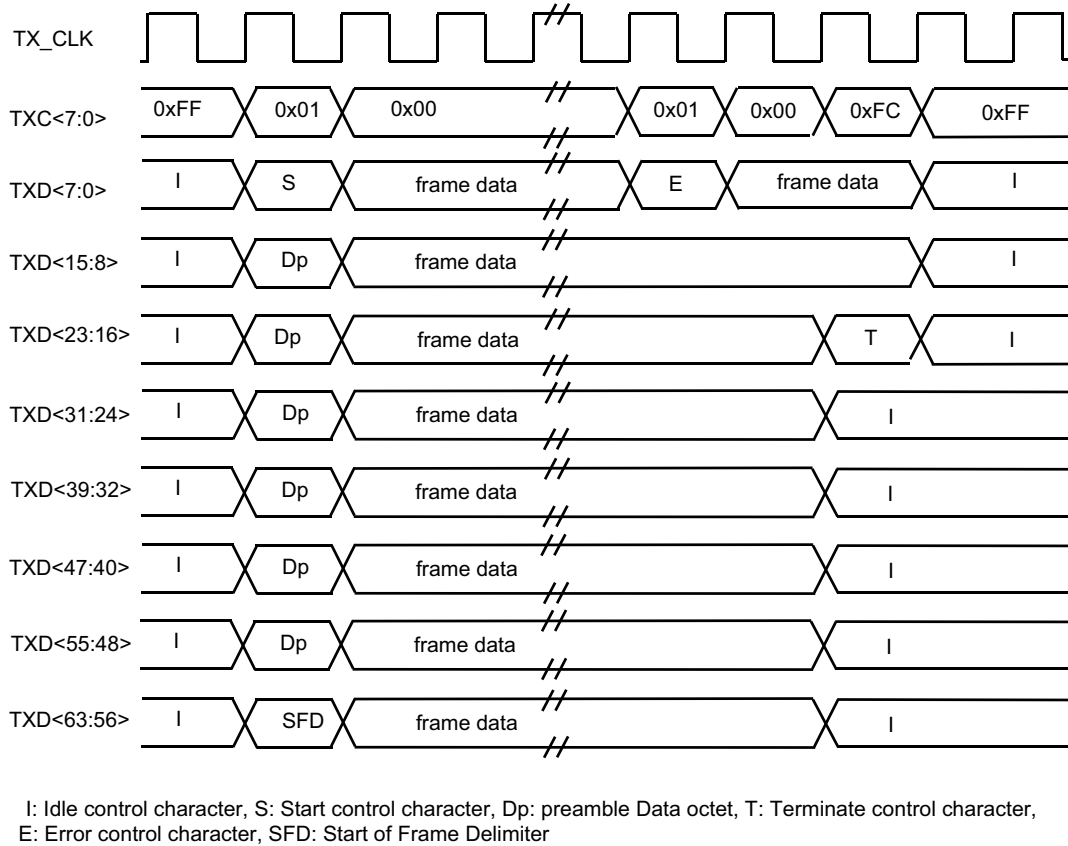


Figure 81–6—Transmit Error Propagation

81.3.1.4 Start control character alignment

On transmit, it may be necessary for the RS to modify the length of the <inter-frame> in order to align the Start control character (first octet of preamble) on lane 0. This shall be accomplished in one of the following two ways:

- a) A MAC implementation may incorporate this RS function into its design and always insert additional idle characters to align the start of preamble on an eight byte boundary. Note that this will reduce the effective data rate for certain packet sizes separated with minimum inter-frame spacing.
- b) Alternatively, the RS may maintain the effective data rate by sometimes inserting and sometimes deleting idle characters to align the Start control character. When using this method the RS must maintain a Deficit Idle Count (DIC) that represents the cumulative count of idle characters deleted or inserted. The DIC is incremented for each idle character deleted, decremented for each idle character inserted, and the decision of whether to insert or delete idle characters is constrained by bounding the DIC to a minimum value of zero and maximum value of seven. Note that this may result in inter-frame spacing observed on the transmit XLGMII/CGMII that is up to seven octets shorter than the minimum transmitted inter-frame spacing specified in Clause 4; however, the frequency of shortened inter-frame spacing is constrained by the DIC rules. The DIC is only reset at

initialization and is applied regardless of the size of the IPG transmitted by the MAC sublayer. An equivalent technique may be employed to control RS alignment of the Start control character provided that the result is the same as if the RS implemented the DIC as described.

81.3.2 Receive

81.3.2.1 RX_CLK (receive clock)

RX_CLK is a continuous clock that provides the timing reference for the transfer of the RXC<7:0> and RXD<63:0> signals from the PHY to the RS. RXC<7:0> and RXD<63:0> shall be sampled by the RS on the rising edge of RX_CLK. RX_CLK is sourced by the PHY.

The frequency of RX_CLK may be derived from the received data or it may be that of a nominal clock (e.g., TX_CLK). When the received data rate at the PHY is within tolerance, the RX_CLK frequency shall be one-sixty-fourth of the MAC receive data rate.

There is no need to transition between the recovered clock reference and a nominal clock reference on a frame-by-frame basis. If loss of received signal from the medium causes a PHY to lose the recovered RX_CLK reference, the PHY shall source the RX_CLK from a nominal clock reference.

NOTE—This standard neither requires nor assumes a guaranteed phase relationship between the RX_CLK and TX_CLK signals.

81.3.2.2 RXC<7:0> (receive control)

RXC<7:0> indicate that the PHY is presenting either recovered and decoded data or control characters on the XLGMII/CGMII. The RXC signal for a lane shall be de-asserted when a data octet is being received on the corresponding lane and asserted when a control character is being received. In the absence of errors, the RXC signals are de-asserted by the PHY for each octet of the preamble (except the first octet that is replaced with a Start control character) and remain de-asserted while all octets to be received are presented on the lanes of the XLGMII/CGMII. RXC<7:0> are driven by the PHY and shall transition synchronously with respect to the rising edge of RX_CLK. Table 81–4 specifies the permissible encodings of RXD and RXC for an XLGMII/CGMII receive lane. Additional requirements apply for proper code sequences and in which lanes particular codes are valid (e.g., Start control character is to be aligned to lane 0).

Table 81–4—Permissible lane encodings of RXD and RXC

RXC	RXD	Description	PLS_DATA.indication parameter
0	0x00 through 0xFF	Normal data reception	Zero, one (eight bits)
1	0x00 through 0x06	Reserved	—
1	0x07	Idle	No applicable parameter (Normal inter-frame)
1	0x08 through 0x9B	Reserved	—
1	0x9C	Sequence (only valid in lane 0)	No applicable parameter (Inter-frame status signal)
1	0x9D through 0xFA	Reserved	—
1	0xFB	Start (only valid in lane 0)	No applicable parameter, first eight zero, one of a frame (a preamble octet)
1	0xFC	Reserved	—
1	0xFD	Terminate	No applicable parameter (start of inter-frame)
1	0xFE	Error	No applicable parameter
1	0xFF	Reserved	—

Figure 81–7 shows the behavior of RXC<7:0> during an example frame reception with no errors.

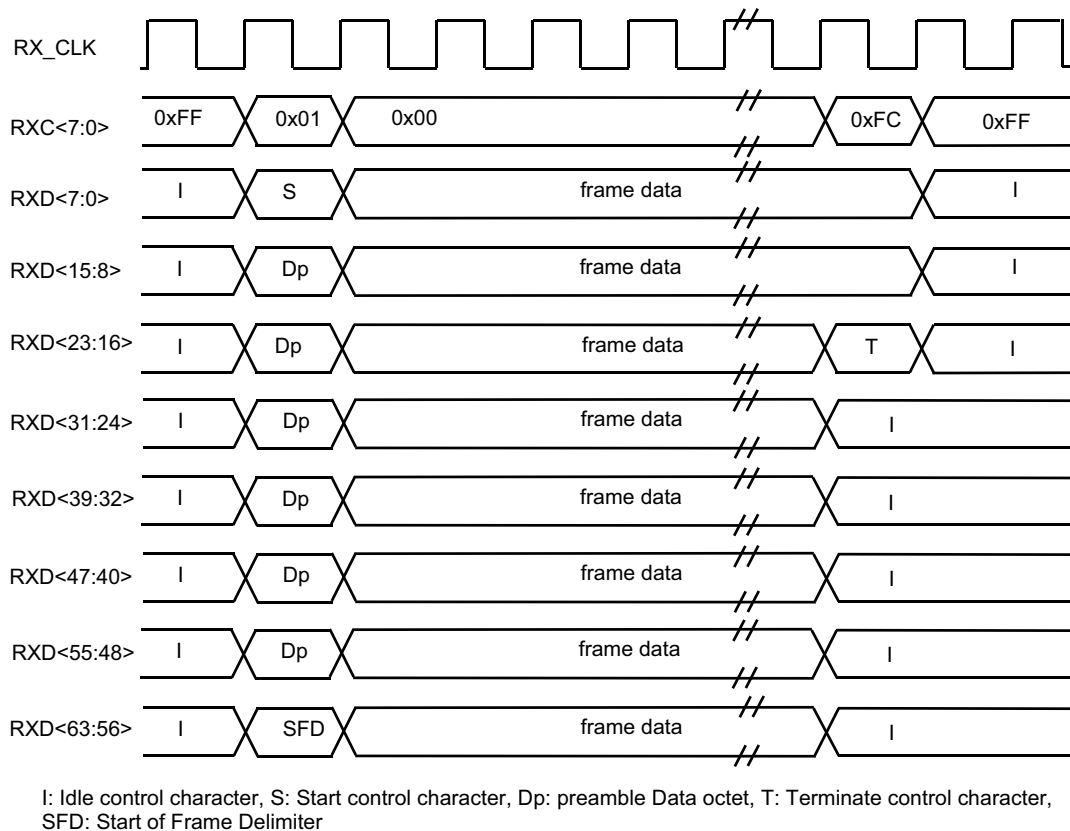


Figure 81–7—Frame reception without error

81.3.2.3 RXD (receive data)


RXD is a bundle of 64 data signals (RXD<63:0>) organized into eight lanes of eight signals each (RXD<7:0>, RXD<15:8>, RXD<23:16>, RXD<31:24>, RXD<39:32>, RXD<47:40>, RXD<55:48>, and RXD<63:56>) that are driven by the PHY. Each lane is associated with an RXC signal as shown in Table 81–2 and shall be decoded by the RS as shown in Table 81–4. RXD<63:0> shall transition synchronously with respect to the rising edge of RX_CLK. For each rising RX_CLK transition, received data and/or control are presented on RXD<63:0> for mapping by the RS. RXD<0> is the least significant bit of lane 0, RXD<8> the least significant bit of lane 1, RXD<16> the least significant bit of lane 2, RXD<24> the least significant bit of lane 3, RXD<32> the least significant bit of lane 4, RXD<40> the least significant bit of lane 5, RXD<48> the least significant bit of lane 6, and RXD<56> the least significant bit of lane 7. Figure 81–7 shows the behavior of RXD<63:0> during frame reception.

While the RXC of a lane is de-asserted, RXD of the lane is used by the RS to generate PLS_DATA indications. Assertion on a lane of appropriate RXD values when RXC is asserted indicates to the RS the Start control character, Terminate control character, Sequence control character, or Error control character that drive its mapping functions.

RXC of a lane is asserted with the appropriate Error control character encoding on RXD of the lane to indicate that an error was detected somewhere in the frame presently being transferred from the PHY to the

RS (e.g., a coding error, or any error that the PHY is capable of detecting, and that may otherwise be undetectable at the MAC sublayer).

The effect of an Error control character on the RS is defined in 81.3.3.1. Figure 81–8 shows the behavior of RXC and RXD during the reception of an example frame with an error.



I: Idle control character, S: Start control character, Dp: preamble Data octet, T: Terminate control character, E: Error control character, SFD: Start of Frame Delimiter

Figure 81–8—Reception with error

81.3.3 Error and fault handling

81.3.3.1 Response to error indications by the XLGMII/CGMII

If, during frame reception (i.e., when DATA_VALID_STATUS = DATA_VALID), a control character other than a Terminate control character is signaled on a received lane, the RS shall ensure that the MAC will detect a FrameCheckError in that frame. This requirement may be met by incorporating a function in the RS that produces a received frame data sequence delivered to the MAC sublayer that is guaranteed to not yield a valid CRC result, as specified by the frame check sequence algorithm (see 3.2.8). This data sequence may be produced by substituting data delivered to the MAC. The RS generates eight PLS_DATA.indication primitives for each Error control character received within a frame, and may generate eight PLS_DATA.indication primitives to ensure FrameCheckError when a control character other than Terminate causes the end of the frame.

Other techniques may be employed to respond to a received Error control character provided that the result is that the MAC sublayer behaves as though a FrameCheckError occurred in the received frame.

81.3.3.2 Conditions for generation of transmit Error control characters

If, during the process of transmitting a frame, it is necessary to request that the PHY deliberately corrupt the contents of the frame in such a manner that a receiver will detect the corruption with the highest degree of probability, then an Error control character may be asserted on a transmit lane by the appropriate encoding of the lane's TXD and TXC signals.

81.3.3.3 Response to received invalid frame sequences

The RS shall not indicate DATA_VALID to the MAC for a Start control character received on any lane other than lane 0. Error free operation will not change the SFD alignment in lane 7. A MAC/RS implementation is not required to process a packet that has an SFD in a position other than lane 7 in the XLGMII/CGMII transfer containing the Start control character.

81.3.4 Link fault signaling

Link fault signaling operates between the remote RS and the local RS. Faults detected between the remote RS and the local RS are received by the local RS as Local Fault. Only an RS originates Remote Fault signals. The behavior of the fault signaling is the same as it is for Clause 46 with the exception that the ordered sets are aligned to eight byte boundaries, padding lanes 4 to 7 with 0x00.

Clause 46 uses the term column when describing data transfers on the XGMII. The eight lanes of data and control transferred per clock cycle on XLGMII/CGMII are equivalent to a column in the following description of link fault signaling.

Sublayers within the PHY are capable of detecting faults that render a link unreliable for communication. Upon recognition of a fault condition, a PHY sublayer indicates Local Fault status on the data path. When this Local Fault status reaches an RS, the RS stops sending MAC data, and continuously generates a Remote Fault status on the transmit data path (possibly truncating a MAC frame being transmitted). When Remote Fault status is received by an RS, the RS stops sending MAC data, and continuously generates Idle control characters. When the RS no longer receives fault status messages, it returns to normal operation, sending MAC data. Note that this behavior only supports bidirectional operation.

Status is signaled in an eight byte Sequence ordered_set as shown in Table 81–5. The PHY indicates Local Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1, 2, 4, 5, 6, and 7 plus a data character of 0x01 in lane 3. The RS indicates a Remote Fault with a Sequence control character in lane 0 and data characters of 0x00 in lanes 1, 2, 4, 5, 6, and 7 plus a data character of 0x02 in lane 3. Though most fault detection is on the receive data path of a PHY, in some specific sublayers, faults can be detected on the transmit side of the PHY. This is also indicated by the PHY with a Local Fault status. All other values in lanes 1 to 3 not shown in Table 81–5 are reserved. The link fault signaling state diagram allows future standardization of reserved Sequence ordered sets for functions other than link fault indications.

Table 81–5—Sequence ordered_sets

Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Description
Sequence	0x00	0x00	0x00	0x00	0x00	0x00	0x00	Reserved
Sequence	0x00	0x00	0x01	0x00	0x00	0x00	0x00	Local Fault
Sequence	0x00	0x00	0x02	0x00	0x00	0x00	0x00	Remote Fault

The RS reports the fault status of the link. Local Fault indicates a fault detected on the receive data path between the remote RS and the local RS. Remote Fault indicates a fault on the transmit path between the local RS and the remote RS. The RS shall implement the link fault signaling state diagram (see Figure 81–9).

81.3.4.1 Variables and counters

The Link Fault Signaling state diagram uses the following variables and counters:

col_cnt

A count of the number of columns received not containing a fault_sequence. This counter increments at RX_CLK rate (on the rising clock transitions) unless reset.

fault_sequence

A new column received on RXC<7:0> and RXD<63:0> comprising a Sequence ordered_set of eight bytes and consisting of a Sequence control character in lane 0 and a seq_type in lanes 1, 2, 3, 4, 5, 6, and 7 indicating either Local Fault or Remote Fault.

last_seq_type

The seq_type of the previous Sequence ordered_set received

Values: Local Fault; 0x00 in lane 1, 0x00 in lane 2, 0x01 in lane 3, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.

Remote Fault; 0x00 in lane 1, 0x00 in lane 2, 0x02 in lane 3, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.

link_fault

An indicator of the fault status.

Values: OK; No fault.

Local Fault; fault detected by the PHY.

Remote Fault; fault detection signaled by the remote RS.

reset

Condition that is true until such time as the power supply for the device that contains the RS has reached the operating region.

Values: FALSE: The device is completely powered and has not been reset (default).

TRUE: The device has not been completely powered or has been reset.

seq_cnt

A count of the number of received Sequence ordered_sets of the same type.

seq_type

The value received in the current Sequence ordered_set

Values: Local Fault; 0x00 in lane 1, 0x00 in lane 2, 0x01 in lane 3, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.

Remote Fault; 0x00 in lane 1, 0x00 in lane 2, 0x02 in lane 3, 0x00 in lane 4, 0x00 in lane 5, 0x00 in lane 6, 0x00 in lane 7.

81.3.4.2 State Diagram

The Link Fault Signaling state diagram specifies the RS monitoring of RXC<7:0> and RXD<63:0> for Sequence ordered_sets. The variable link_fault is set to indicate the value of a received Sequence ordered_set when four fault_sequences containing the same fault value have been received with fault sequence separated by less than 128 columns and no intervening fault_sequences of a different fault value. The state diagram is shown in Figure 81–9.

The variable link_fault is set to OK following any interval of 128 columns not containing a Remote Fault or Local Fault Sequence ordered_set.

The RS output onto TXC<7:0> and TXD<63:0> is controlled by the variable link_fault.

- a) link_fault = OK
The RS shall send MAC frames as requested through the PLS service interface. In the absence of MAC frames, the RS shall generate Idle control characters.
- b) link_fault = Local Fault
The RS shall continuously generate Remote Fault Sequence ordered_sets.
- c) link_fault = Remote Fault
The RS shall continuously generate Idle control characters.

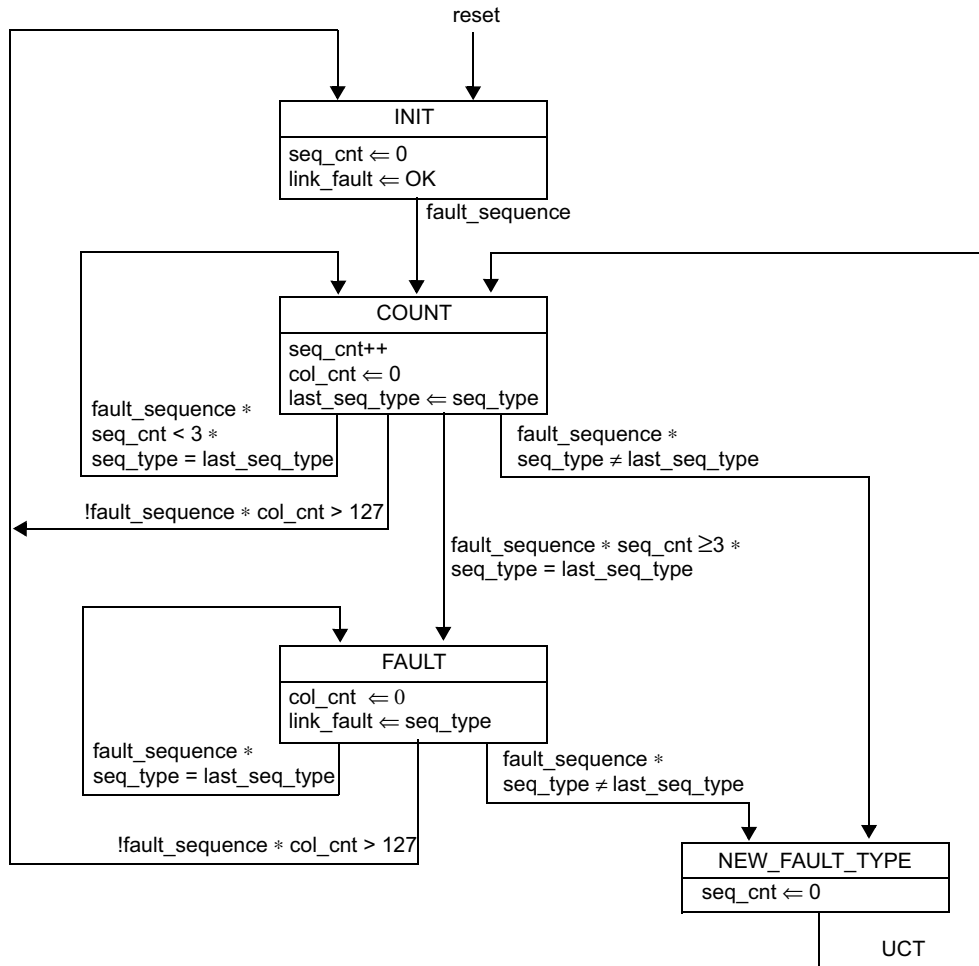


Figure 81-9—Link Fault Signaling state diagram

81.4 Protocol implementation conformance statement (PICS) proforma for Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation⁹

81.4.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

81.4.2 Identification

81.4.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

81.4.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 81, Reconciliation Sublayer (RS) and Media Independent Interface for 40 Gb/s and 100 Gb/s operation
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

⁹Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

81.4.2.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PHY40	PHY support of XLGMII	81.2, 81.3		O	Yes [] No []
*PHY100	PHY support of CGMII	81.2, 81.3		O	Yes [] No []
*RS40	Reconciliation Sublayer support of XLGMII	81.2, 81.3		O	Yes [] No []
*RS100	Reconciliation Sublayer support of CGMII	81.2, 81.3		O	Yes [] No []

81.4.3 PICS proforma tables for Reconciliation Sublayer and Media Independent Interface for 40 Gb/s and 100 Gb/s operation

81.4.3.1 General

Item	Feature	Subclause	Value/Comment	Status	Support
G1	PHY support of MAC data rate	81.1.3	Support MAC data rate of 40 Gb/s	PHY40:M	Yes [] N/A []
G2	PHY support of MAC data rate	81.1.3	Support MAC data rate of 100 Gb/s	PHY100:M	Yes [] N/A []
G3	Cumulative MAC Control, MAC and RS delay	81.1.4	Per Table 81–1 for 40 Gb/s	RS40:M	Yes [] N/A []
G4	Cumulative MAC Control, MAC and RS delay	81.1.4	Per Table 81–1 for 100 Gb/s	RS100:M	Yes [] N/A []
G5	Lane structure	81.1.6	Per Table 81–2	M	Yes []

81.4.3.2 Mapping of PLS service primitives

Item	Feature	Subclause	Value/Comment	Status	Support
PL1	Mapping to Clause 6	81.1.7	RS implements mapping to Clause 6 PLS service primitives	RS:M	Yes [] N/A []
PL2	Mapping of PLS_DATA.requests	81.1.7.1.4	In sequence TXD<0> to TXD<63>	RS:M	Yes [] N/A []
PL3	Start control character creation	81.1.7.1.4	First octet of preamble converted to Start control character	RS:M	Yes [] N/A []
PL4	TXD and TXC generation	81.1.7.1.4	For each 64 PLS_DATA.requests	RS:M	Yes [] N/A []

Item	Feature	Subclause	Value/Comment	Status	Support
PL5	Terminate control character creation	81.1.7.1.4	DATA_COMPLETE causes creation of Terminate control character in next lane in sequence	RS:M	Yes [] N/A []
PL6	Mapping RXD to PLS_DATA.indications	81.1.7.2.3	Create PLS_DATA.indications in sequence from RXD<0> to RXD<63>	RS:M	Yes [] N/A []
PL7	PLS_DATA.indication generation	81.1.7.2.3	Generate 64 PLS_DATA.indications for each RXD<0:63> until Terminate then generating 0, 8, 16, 24, 32, 40, 48, or 56	RS:M	Yes [] N/A []
PL8	Start control character conversion	81.1.7.2.3	Convert valid Start control character to preamble before generating PLS_DATA.indications	RS:M	Yes [] N/A []
PL9	Terminate control character	81.1.7.2.3	No PLS_DATA.indications generated	RS:M	Yes [] N/A []
PL10	PLS_DATA_VALID.indication generation	81.1.7.5.3	On change of value of DATA_VALID_STATUS	RS:M	Yes [] N/A []
PL11	DATA_VALID_STATUS	81.1.7.5.3	Value of DATA_VALID on a lane 0 Start control character preceded by eight idles, a Sequence ordered set, or a Terminate character	RS:M	Yes [] N/A []
PL12	DATA_VALID_STATUS	81.1.7.5.3	Value of DATA_NOT_VALID on any control character but Error	RS:M	Yes [] N/A []
PL13	Frame not ending with Terminate control character	81.1.7.5.3	Ensure MAC detects CRC error	RS:M	Yes [] N/A []

81.4.3.3 Data stream structure

Item	Feature	Subclause	Value/Comment	Status	Support
DS1	Frame transfer	81.2	Within XLGMII/CGMII data stream	RS:M	Yes [] N/A []
DS2	Bit mapping	81.2	Per Figure 81-4	RS:M	Yes [] N/A []
DS3	Content of <data>	81.2.3	Consist of data octets	RS:M	Yes [] N/A []
DS4	Recognition of <efd>	81.2.4	Terminate recognized in any lane	RS:M	Yes [] N/A []

81.4.3.4 XLGMII/CGMII signal functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	TX_CLK active edge	81.3.1.1	TXD and TXC sampled on the rising edge of TX_CLK	XGE:M	Yes [] N/A []
FS2	TX_CLK frequency	81.3.1.1	One-sixty-fourth of the MAC transmit data rate	XGE:M	Yes [] N/A []
FS3	TXC assertion and de-assertion	81.3.1.2	De-asserted for data, asserted for control character	RS:M	Yes [] N/A []
FS4	TXC clock	81.3.1.2	Synchronous to TX_CLK	XGE:M	Yes [] N/A []
FS5	TXD encoding	81.3.1.3	Per Table 81–3	RS:M	Yes [] N/A []
FS6	TXD clock	81.3.1.3	Synchronous to TX_CLK	XGE:M	Yes [] N/A []
FS7	Start alignment	81.3.1.4	Start control character aligned to lane 0	RS:M	Yes [] N/A []
FS8	RX_CLK active edge	81.3.2.1	RXD and RXC sampled on the rising edge of RX_CLK	XGE:M	Yes [] N/A []
FS9	RX_CLK frequency	81.3.2.1	One-sixty-fourth of the MAC receive data rate	XGE:M	Yes [] N/A []
FS10	Loss of receive signal	81.3.2.1	Source RX_CLK from nominal clock	PHY:M	Yes [] N/A []
FS11	RXC assertion and de-assertion	81.3.2.2	De-asserted for data, asserted for control character	PHY:M	Yes [] N/A []
FS12	RXC clock	81.3.2.2	Synchronous to RX_CLK	XGE:M	Yes [] N/A []
FS13	RXD decoding	81.3.2.3	Per Table 81–4	RS:M	Yes [] N/A []
FS14	RXD clock	81.3.2.3	Synchronous to RX_CLK	XGE:M	Yes [] N/A []
FS15	Received Error control character	81.3.3.1	RS cause MAC FrameCheck-Error	RS:M	Yes [] N/A []
FS16	DATA_VALID assertion	81.3.3.3	RS not assert DATA_VALID unless Start control character in lane 0	RS:M	Yes [] N/A []

81.4.3.5 Link fault signaling state diagram

Item	Feature	Subclause	Value/Comment	Status	Support
LF1	Link fault signaling state diagram	81.3.4	Implement per Figure 81-9	RS:M	Yes [] N/A []
LF2	link_fault = OK and MAC frames	81.3.4.2	RS services MAC frame transmission requests	RS:M	Yes [] N/A []
LF3	link_fault = OK and no MAC frames	81.3.4.2	In absence of MAC frames, RS transmits Idle control characters	RS:M	Yes [] N/A []
LF4	link_fault = Local Fault	81.3.4.2	RS transmits continuous Remote Fault Sequence ordered_sets	RS:M	Yes [] N/A []
LF5	link_fault = Remote Fault	81.3.4.2	RS transmits continuous Idle control characters	RS:M	Yes [] N/A []

82. Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R

82.1 Overview

82.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) that is common to two families of (40 Gb/s and 100 Gb/s) Physical Layer implementations, known as 40GBASE-R and 100GBASE-R. The 40GBASE-R PCS is a sublayer of the 40 Gb/s PHYs listed in Table 80–1. The 100GBASE-R PCS is a sublayer of the 100 Gb/s PHYs listed in Table 80–1. The terms 40GBASE-R and 100GBASE-R are used when referring generally to Physical Layers using the PCS defined in this clause.

Both 40GBASE-R and 100GBASE-R are based on a 64B/66B code. The 64B/66B code supports transmission of data and control characters, while maintaining robust error detection. Data distribution is introduced to support multiple lanes in the Physical Layer. Part of the distribution includes the periodic insertion of an alignment marker, which allows the receive PCS to align data from multiple lanes.

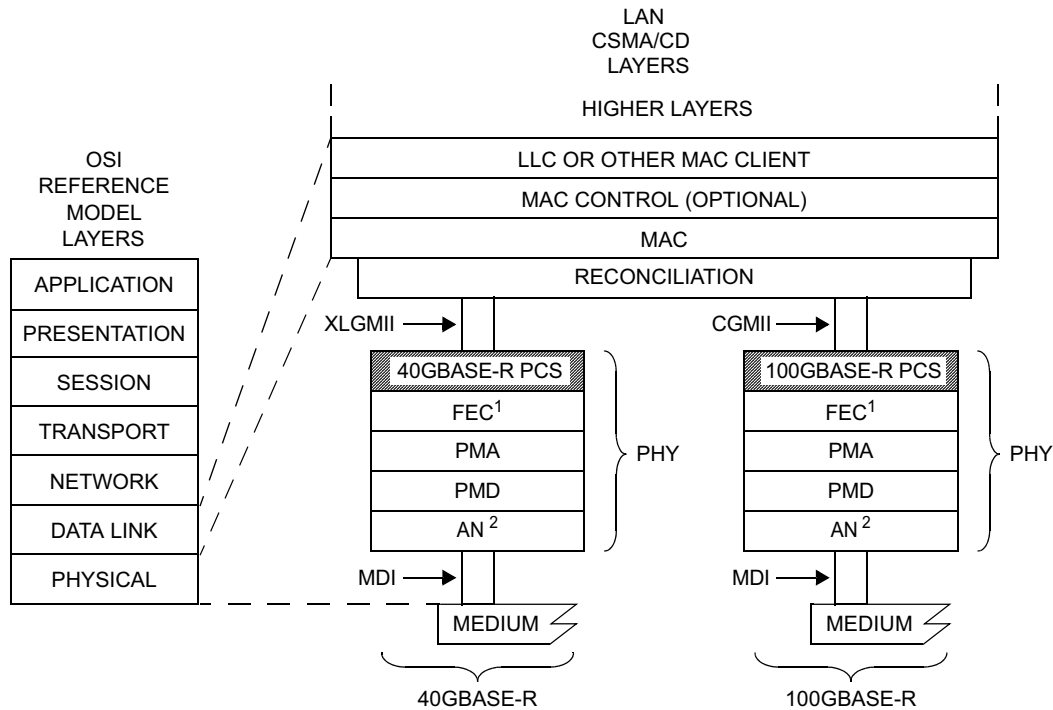
82.1.2 Relationship of 40GBASE-R and 100GBASE-R to other standards

Figure 82–1 depicts the relationships among the 40GBASE-R and 100GBASE-R sublayers (shown shaded), the Ethernet MAC and Reconciliation Sublayers, and the higher layers.

This clause borrows heavily from Clause 49. 64B/66B encoding is reused with appropriate changes made to support 8 byte alignment versus 4 byte alignment in Clause 49. In addition to 64B/66B encoding, there is a methodology to add alignment markers and distribute data to multiple lanes.

82.1.3 Summary of 40GBASE-R and 100GBASE-R sublayers

Figure 82–1 shows the relationship of the 40GBASE-R PCS and 100GBASE-R PCS sublayers (shown shaded) with other sublayers to the ISO Open System Interconnection (OSI) reference model.



AN = AUTO-NEGOTIATION
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
 NOTE 2—CONDITIONAL BASED ON PHY TYPE

Figure 82–1—40GBASE-R and 100GBASE-R PCS relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 CSMA/CD LAN model

82.1.3.1 Physical Coding Sublayer (PCS)

The PCS service interface is the Media Independent Interface (XLGMII/CGMII), which is defined in Clause 81. The 40 Gb/s variant of this is called the 40 Gb/s Media Independent Interface (XLGMII) and the 100 Gb/s variant of this interface is called the 100 Gb/s Media Independent Interface (CGMII). The XLGMII/CGMII provides a uniform interface to the Reconciliation Sublayer for all 40 Gb/s and 100 Gb/s PHY implementations.

The 40GBASE-R and 100GBASE-R PCSs provide all services required by the XLGMII/CGMII, including the following:

- a) Encoding (decoding) of eight XLGMII/CGMII data octets to (from) 66-bit blocks (64B/66B).
- b) Transferring encoded data to (from) the PMA.
- c) Compensation for any rate differences caused by the insertion or deletion of alignment markers or due to any rate difference between the XLGMII/CGMII and PMA through the insertion or deletion of idle control characters.
- d) Determining when a functional link has been established and informing the management entity via the MDIO when the PHY is ready for use.

82.1.4 Inter-sublayer interfaces

The upper interface of the PCS may connect to the Reconciliation Sublayer through the XLGMII/CGMII. The lower interface of the PCS connects to the PMA sublayer to support a PMD. If the optional FEC sublayer is implemented (see Clause 74) and an optional physical instantiation, i.e., XLAUI or CAUI, is not implemented directly below the PCS sublayer, then the lower interface connects to the FEC sublayer. The 40GBASE-R PCS has a nominal rate at the PMA/FEC service interface of 10.3125 Gtransfers/s per PCS lane, which provides capacity for the MAC data rate of 40 Gb/s. The 100GBASE-R PCS has a nominal rate at the PMA/FEC service interface of 5.15625 Gtransfers/s per PCS lane, which provides capacity for the MAC data rate of 100 Gb/s.

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience.

82.1.4.1 PCS service interface (XLGMII/CGMII)

The PCS service interface allows the 40GBASE-R or 100GBASE-R PCS to transfer information to and from a PCS client. The PCS client is the Reconciliation Sublayer. The PCS Service Interface is precisely defined as the Media Independent Interface (XLGMII/CGMII) in Clause 81.

82.1.4.2 Physical Medium Attachment (PMA) service interface

The PMA or FEC service interface for the PCS is described in an abstract manner and does not imply any particular implementation. The PMA/FEC Service Interface supports the exchange of encoded data between the PCS and PMA or FEC sublayer. The PMA or FEC service interface is defined in 83.3 and is an instance of the inter-sublayer service interface definition in 80.3.

82.1.5 Functional block diagram

Figure 82–2 provides a functional block diagram of the 40GBASE-R PCS and 100GBASE-R PCS.

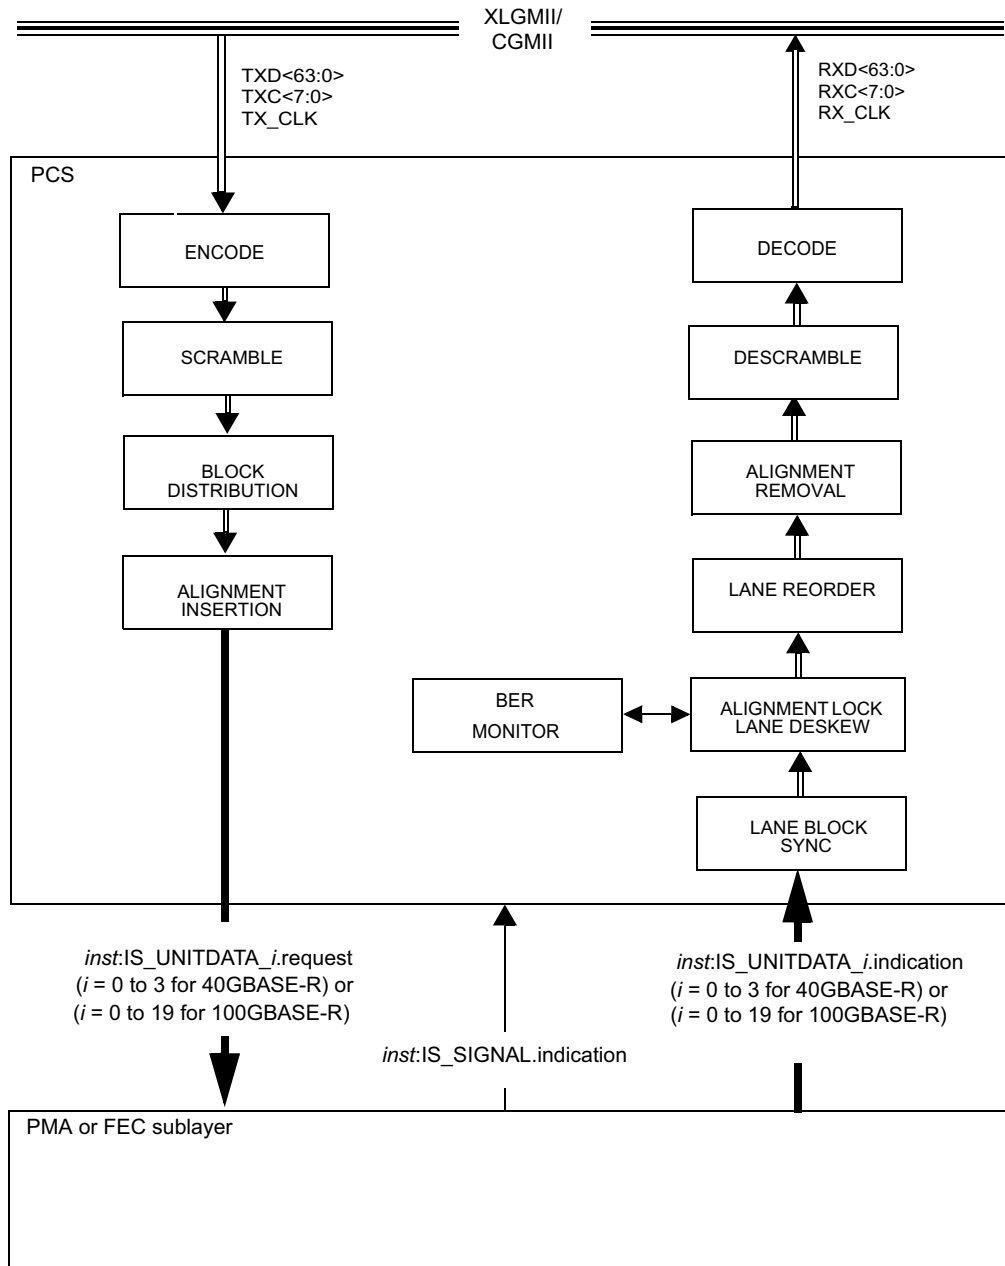


Figure 82-2—Functional block diagram

82.2 Physical Coding Sublayer (PCS)

82.2.1 Functions within the PCS

The 40GBASE-R and 100GBASE-R PCSs comprise the PCS Transmit and PCS Receive processes for each rate of operation. The PCS shields the Reconciliation Sublayer (and MAC) from the specific nature of the underlying channel. The PCS transmit channel and receive channel can each operate in normal mode or test-pattern mode.

When communicating with the XLGMII/CGMII, the PCS uses an eight octet-wide, synchronous data path, with packet delimiting being provided by transmit control signals ($TXC_{<n>} = 1$) and receive control signals ($RXC_{<n>} = 1$). When communicating with the PMA, the PCS uses multiple serial streams, 4 encoded bit streams for a 40GBASE-R PCS, or 20 encoded bit streams for a 100GBASE-R PCS. Alignment to 64B/66B blocks is performed in the PCS. The PMA sublayer operates independently of block and packet boundaries. The PCS provides the functions necessary to map packets between the XLGMII/CGMII format and the PMA/FEC service interface format.

Note that these serial streams originate from a common clock in each direction, but may vary in phase and Skew dynamically.

When the transmit channel is in normal mode, the PCS Transmit process continuously generates blocks based upon the $TXD_{<63:0>}$ and $TXC_{<7:0>}$ signals on the XLGMII/CGMII. The blocks are scrambled and then distributed to individual PCS lanes. After distribution, alignment marker blocks are periodically added to each PCS lane. Transmit data-units are sent to the service interface via the *inst:IS_UNITDATA_i.request* primitive. In these primitives, the *inst* can be either FEC or PMA depending on which sublayer is adjacent to the PCS sublayer.

When the transmit channel is in test-pattern mode, a test pattern is packed into the transmit data-units that are sent to the PMA/FEC service interface via the *inst:IS_UNITDATA_i.request* primitive.

When the receive channel is in normal or test-pattern mode, the PCS Synchronization process continuously monitors *inst:IS_SIGNAL.indication(SIGNAL_OK)*. When SIGNAL_OK indicates OK, then the PCS Synchronization process accepts data-units via the *inst:IS_UNITDATA_i.indication* primitive. It attains block synchronization based on the 2-bit synchronization headers on each one of the PCS lanes. Once block synchronization is found on a PCS lane, then alignment marker lock can be attained by searching for valid alignment markers. After alignment markers are found on all PCS lanes, the PCS lanes can be reordered and deskewed. Note that a particular transmit PCS lane can be received on any receive lane of the service interface due to the Skew and multiplexing that occurs in the path.

The PCS deskew process conveys received blocks to the PCS Receive process. The PCS deskew process deskews and aligns the individual PCS lanes, removes the alignment markers, forms a single stream, and sets the *align_status* flag to indicate whether the PCS has obtained alignment.

When the PCS deskew process has obtained alignment, the BER monitor process monitors the signal quality asserting *hi_ber* if excessive errors are detected. When *align_status* is asserted and *hi_ber* is de-asserted, the PCS Receive process continuously accepts blocks and generates $RXD_{<63:0>}$ and $RXC_{<7:0>}$ on the XLGMII/CGMII.

When the receive channel is in test-pattern mode, the BER monitor process is disabled. The Receive process will be held in the *RX_INIT* state. The received bits will be compared to the test pattern and errors counted.

The PCS shall provide transmit test-pattern mode for the scrambled idle pattern (see 82.2.10), and shall provide receive test-pattern mode for the scrambled idle pattern (see 82.2.10). Test-pattern mode is activated separately for transmit and receive. The PCS shall support transmit test-pattern mode and receive test-pattern mode operating simultaneously so as to support loopback testing.

82.2.2 Use of blocks

The PCS maps XLGMII/CGMII signals into 66-bit blocks, and vice versa, using a 64B/66B coding scheme. The synchronization headers of the blocks allow establishment of block boundaries by the PCS Synchronization process. The PCS functions ENCODE and DECODE generate, manipulate, and interpret blocks as defined in 82.2.3.

82.2.3 64B/66B transmission code

The PCS uses a transmission code to improve the transmission characteristics of information to be transferred across the link and to support transmission of control and data characters. The encodings defined by the transmission code ensure that sufficient transitions are present in the PHY bit stream to make clock recovery possible at the receiver. The encoding also preserves the likelihood of detecting any single or multiple bit errors that may occur during transmission and reception of information. In addition, the synchronization headers of the code enable the receiver to achieve block alignment on the incoming PHY bit stream. The 64B/66B transmission code specified for use in this standard is a run-length-limited code.¹⁰

The relationship of block bit positions relative to XLGMII/CGMII, PMA, and other PCS functions is illustrated in Figure 82–3 for transmit and Figure 82–4 for receive. These figures illustrate the processing of a block containing 8 data octets. See 82.2.3.3 for information on how blocks containing control characters are mapped into 66-bit blocks. Note that the sync header is generated by the encoder and bypasses the scrambler.

82.2.3.1 Notation conventions

For values shown as binary, the leftmost bit is the first transmitted bit.

64B/66B encodes 8 data octets or control characters into a block. Blocks containing control characters also contain a block type field. Data octets are labeled D_0 to D_7 . The control characters /I/ and /E/ are labeled C_0 to C_7 . The control characters, /Q/ and /Fsig/, for ordered_sets are labeled as O_0 since they are only valid on the first octet of the XLGMII/CGMII. The control character for start is labeled as S_0 for the same reason. The control character for terminate is labeled as T_0 to T_7 . The four trailing zero data octets in ordered_sets are labeled as Z_4 to Z_7 .

One XLGMII/CGMII transfer provides eight characters that are encoded into one 66-bit transmission block. The subscript in the above labels indicates the position of the character in the eight characters from the XLGMII/CGMII transfer.

Contents of block type fields, data octets, and control characters are shown as hexadecimal values. The LSB of the hexadecimal value represents the first transmitted bit. For instance, the block type field $0x1E$ is sent from left to right as 01111000 . The bits of a transmitted or received block are labeled $TxB\langle 65:0 \rangle$ and $RxB\langle 65:0 \rangle$, respectively, where $TxB\langle 0 \rangle$ and $RxB\langle 0 \rangle$ represent the first transmitted bit. The value of the sync header is shown as a binary value. Binary values are shown with the first transmitted bit (the LSB) on the left.

82.2.3.2 Transmission order

Block bit transmission order shall be as illustrated in Figure 82–3 and Figure 82–4. Note that these figures show the mapping from XLGMII/CGMII to 64B/66B block for a block containing eight data characters.

¹⁰In 10GBASE-R the run length is limited to 66 bits, but in 40GBASE-R and 100GBASE-R, when multiplexing occurs in the PMA, this guaranteed run length limit increases. For example, if two PCS lanes are multiplexed in the PMA, then the possible run length would double.




Figure 82-3—PCS Transmit bit ordering




Figure 82-4—PCS Receive bit ordering

82.2.3.3 Block structure

Blocks consist of 66 bits. The first two bits of a block are the synchronization header (sync header). Blocks are either data blocks or control blocks. The sync header is 01 for data blocks and 10 for control blocks. Thus, there is always a transition between the first two bits of a block. The remainder of the block contains the payload. The payload is scrambled and the sync header bypasses the scrambler. Therefore, the sync header is the only position in the block that is always guaranteed to contain a transition. This feature of the code is used to obtain block synchronization.

Data blocks contain eight data characters. Control blocks begin with an 8-bit block type field that indicates the format of the remainder of the block. For control blocks containing a Start, Terminate character, or ordered set, that character is implied by the block type field. Other control characters are encoded in a 7-bit control code. Each control block encodes eight characters.

The format of the blocks is as shown in Figure 82–5. In the figure, the column labeled Input Data shows, in abbreviated form, the eight characters used to create the 66-bit block. These characters are either data characters or control characters and, when transferred across the XLGMII/CGMII interface, the corresponding TXC or RXC bit is set accordingly. Within the Input Data column, D_0 through D_7 are data octets and are transferred with the corresponding TXC or RXC bit set to zero. All other characters are control characters and are transferred with the corresponding TXC or RXC bit set to one. The single bit fields (thin rectangles with no label in the figure) are sent as zero and ignored upon receipt.

All unused values of block type field¹¹ are invalid; they shall not be transmitted and shall be considered an error if received.

¹¹The block type field values have been chosen to have a 4-bit Hamming distance between them. There are four unused values that maintain this Hamming distance: 0x00, 0x2D, 0x33, and 0x66.

Input Data	S y n c	Block Payload									
Bit Position:	0 1 2	65									
Data Block Format:											
D ₀ D ₁ D ₂ D ₃ D ₄ D ₅ D ₆ D ₇	01	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇		
Control Block Formats:		Block Type Field									
C ₀ C ₁ C ₂ C ₃ C ₄ C ₅ C ₆ C ₇	10	0x1E	C ₀	C ₁	C ₂	C ₃	C ₄	C ₅	C ₆	C ₇	
S ₀ D ₁ D ₂ D ₃ D ₄ D ₅ D ₆ D ₇	10	0x78	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆	D ₇		
O ₀ D ₁ D ₂ D ₃ Z ₄ Z ₅ Z ₆ Z ₇	10	0x4B	D ₁	D ₂	D ₃	O ₀	0x000_0000				
T ₀ C ₁ C ₂ C ₃ C ₄ C ₅ C ₆ C ₇	10	0x87			C ₁	C ₂	C ₃	C ₄	C ₅	C ₆	C ₇
D ₀ T ₁ C ₂ C ₃ C ₄ C ₅ C ₆ C ₇	10	0x99	D ₀			C ₂	C ₃	C ₄	C ₅	C ₆	C ₇
D ₀ D ₁ T ₂ C ₃ C ₄ C ₅ C ₆ C ₇	10	0xAA	D ₀	D ₁			C ₃	C ₄	C ₅	C ₆	C ₇
D ₀ D ₁ D ₂ T ₃ C ₄ C ₅ C ₆ C ₇	10	0xB4	D ₀	D ₁	D ₂			C ₄	C ₅	C ₆	C ₇
D ₀ D ₁ D ₂ D ₃ T ₄ C ₅ C ₆ C ₇	10	0xCC	D ₀	D ₁	D ₂	D ₃			C ₅	C ₆	C ₇
D ₀ D ₁ D ₂ D ₃ D ₄ T ₅ C ₆ C ₇	10	0xD2	D ₀	D ₁	D ₂	D ₃	D ₄			C ₆	C ₇
D ₀ D ₁ D ₂ D ₃ D ₄ D ₅ T ₆ C ₇	10	0xE1	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅			C ₇
D ₀ D ₁ D ₂ D ₃ D ₄ D ₅ D ₆ T ₇	10	0xFF	D ₀	D ₁	D ₂	D ₃	D ₄	D ₅	D ₆		

Figure 82–5—64B/66B block formats

WARNING

The mapping of 40GBASE-R PCS blocks into OPU3 specified in ITU-T G.709[Bx1] depends on the set of control block types shown in Figure 82–5. Any deviation from the coding specified in Figure 82–5 will break the mapping and may prevent 40GBASE-R PCS blocks from being mapped into OPU3 (see ITU-T G.709 [Bx1] for more details).

82.2.3.4 Control codes

The same set of control characters are supported by the XLGMII/CGMII and the PCS. The representations of the control characters are the control codes. XLGMII/CGMII encodes a control character into an octet (an eight bit value). The 40GBASE-R and 100GBASE-R PCS encode the start and terminate control characters implicitly by the block type field. The 40GBASE-R and 100GBASE-R PCS encode the ordered_set control codes using the block type field. The 40GBASE-R and 100GBASE-R PCS encode each of the other control characters into a 7-bit control code.

The control characters and their mappings to 40GBASE-R and 100GBASE-R control codes and XLGMII/CGMII control codes are specified in Table 82–1. All XLGMII/CGMII, 40GBASE-R, and 100GBASE-R control code values that do not appear in the table shall not be transmitted and shall be considered an error if received.

82.2.3.5 Valid and invalid blocks

A block is invalid if any of the following conditions exists:

- a) The sync header has a value of 00 or 11.
- b) The block type field contains an invalid value (one not included in Figure 82–5).
- c) Any control character contains a value not included in Table 82–1.
- d) The set of eight XLGMII/CGMII characters does not have a corresponding block format in Figure 82–5.

82.2.3.6 Idle (/I/)

Idle control characters (/I/) are transmitted when idle control characters are received from the XLGMII/CGMII. Idle control characters may be added or deleted by the PCS to adapt between clock rates. /I/ insertion and deletion shall occur in groups of 8. /I/s may be added following idle control characters or ordered sets. They shall not be added while data is being received.

Table 82–1—Control codes

Control character	Notation	XLGMII/ CGMII control code	40/100GBASE-R O code	40GBASE-R and 100GBASE-R control code
idle	/I/	0x07		0x00
start	/S/	0xFB		Encoded by block type field
terminate	/T/	0xFD		Encoded by block type field
error	/E/	0xFE		0x1E
Sequence ordered_set	/Q/	0x9C	0x0	Encoded by block type 0x4B plus O code, control codes are set to 0x00
Signal ordered_set ^a	/Fsig/	0x5C	0xF	Encoded by block type 0x4B plus O code, control codes are set to 0x00

^aReserved for INCITS T11 Fibre Channel use.

82.2.3.7 Start (/S/)

The start control character (/S/) indicates the start of a packet. This delimiter is only valid on the first octet of the XLGMII/CGMII (TXD<0:7> and RXD<0:7>). Receipt of an /S/ on any other octet of TXD indicates an error. Block type field values implicitly encode an /S/ as the first character of the block.

82.2.3.8 Terminate (/T/)

The terminate control character (/T/) indicates the end of a packet. Since packets vary in length, the /T/ can occur on any octet of the XLGMII/CGMII interface and within any character of the block. The location of the /T/ in the block is implicitly encoded in the block type field. A valid end of packet occurs when a block containing a /T/ is followed by a control block that does not contain a /T/ or an /E/.

82.2.3.9 ordered_set (/O/)

Ordered sets are used to extend the ability to send control and status information over the link such as remote fault and local fault status. Ordered sets consist of a control character followed by three data characters followed by four zero data characters on the XLGMII/CGMII. Ordered sets always begin on the first octet of the XLGMII/CGMII. 40 Gigabit and 100 Gigabit Ethernet use one kind of ordered_set: the sequence ordered_set (see 81.3.4). An additional ordered_set, the signal ordered_set, has been reserved and it begins with another control code. The ordered_set control characters (/Q/ and /Fsig/) indicate the start of an ordered_set. The block type field plus the O code encode the specific control character for the ordered_set. See Table 82–1 for the mappings.

Sequence ordered_sets may be deleted by the PCS to adapt between clock rates. Such deletion shall only occur when two consecutive sequence ordered sets have been received and only one of the two ordered sets may be deleted. Only idle control characters may be inserted for clock compensation. Signal ordered_sets are not deleted for clock compensation.

82.2.3.10 Error (/E/)

In both the 64B/66B encoder and decoder, the /E/ is generated whenever an /E/ is detected. The /E/ is also generated when invalid blocks are detected. The /E/ allows the PCS to propagate detected errors. See R_TYPE and T_TYPE function definitions in 82.2.18.2.3 for further information.

82.2.4 Transmit process

The transmit process generates blocks based upon the TXD<63:0> and TXC<7:0> signals received from the XLGMII/CGMII. One XLGMII/CGMII data transfer is encoded into one 66-bit block. It takes 66 *inst:IS_UNITDATA_i* transfers to send a 66-bit block of data on each of the PCS lanes. The transmit process must delete idle control characters or sequence ordered sets to accommodate the transmission of alignment markers. If the PCS transmit process spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle control characters.

There are sufficient idle control characters to delete in order to make room for alignment markers, in addition to handling clock compensation. Idle control characters or sequence ordered sets are removed, if necessary, to accommodate the insertion of the 66-bit alignment markers. See 82.2.3.6 for more details.

The transmit process generates blocks as specified in the transmit process state diagram. The contents of each block are contained in a vector tx_coded<65:0>, which is passed to the scrambler. tx_coded<1:0> contains the sync header and the remainder of the bits contain the block payload.

82.2.5 Scrambler

The payload, tx_coded<65:2>, is scrambled with a self-synchronizing scrambler. The scrambler is identical to the scrambler used in Clause 49, see 49.2.6 for the definition of the scrambler. The sync bits, tx_coded<1:0>, are not scrambled.

82.2.6 Block distribution

Once the data is encoded and scrambled, it is distributed to multiple PCS lanes, 66-bit blocks at a time in a round robin distribution from the lowest to the highest numbered PCS lanes. This allows the PCS to support multiple physical lanes in the PMD and XLAUI or CAUI interfaces (see Annex 83A and Annex 83B). The 40GBASE-R PCS distributes the 66-bit blocks to 4 PCS lanes, and the 100GBASE-R PCS distributes the blocks to 20 PCS lanes. The distribution process is shown in Figure 82–6.

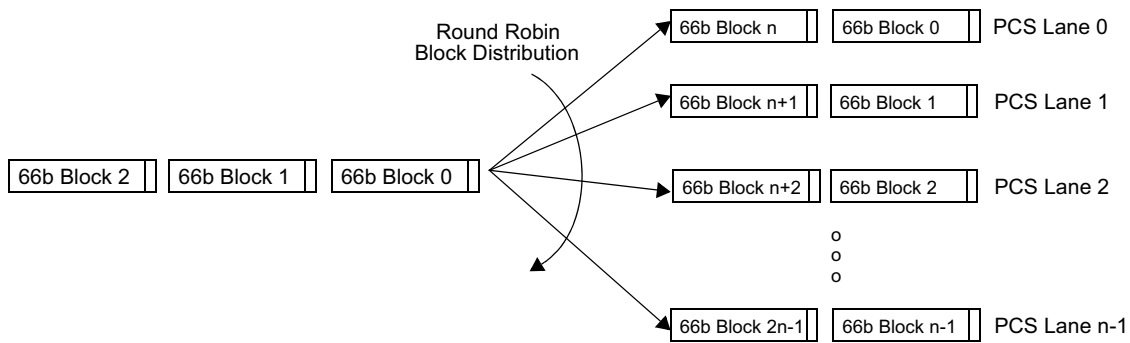


Figure 82–6—PCS Block distribution

82.2.7 Alignment marker insertion

In order to support deskew and reordering of individual PCS lanes at the receive PCS, alignment markers are added periodically to each PCS lane. The alignment marker has the form of a specially defined 66-bit block with a control block sync header. These markers interrupt any data transfer that is already in progress. This allows alignment markers to be inserted into all PCS lanes at the same time. Room for the alignment markers is created by periodically deleting IPG from the XLGMII/CGMII data stream. Other special properties of the alignment markers are that they are not scrambled and do not conform to the encoding rules as outlined in Figure 82–5. This is possible because the alignment markers are added after encoding is performed in the transmit PCS and the alignment markers are removed before 64B/66B decoding is performed in the receive PCS. The alignment markers are not scrambled in order to allow the receiver to find the alignment markers, deskew the PCS lanes, and reassemble the aggregate stream before descrambling is performed. The alignment markers themselves are formed from a known pattern that is defined to be balanced and with many transitions and therefore scrambling is not necessary for the alignment markers. The alignment markers shall be inserted after every 16383 66-bit blocks on each PCS lane. Alignment marker insertion is shown in Figure 82–7 and Figure 82–8.

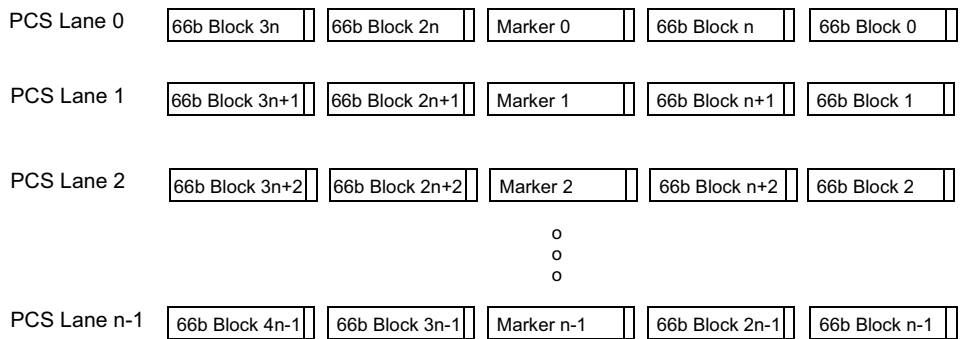


Figure 82–7—Alignment marker insertion




Figure 82–8—Alignment marker insertion period

The format of the alignment markers is shown in Figure 82–9.




Figure 82–9—Alignment marker format

The content of the alignment markers shall be as shown in Table 82–2 for 100GBASE-R and in Table 82–3 for 40GBASE-R. The contents depend on the PCS lane number and the octet number. Note that M₄ through M₆ are the bit-wise inversion of M₀ through M₂, respectively. Also BIP₇ is the bit-wise inversion of BIP₃. This property allows the alignment markers to be DC balanced. Lane markers 0 to 19 from Table 82–2 are used for the 100GBASE-R PCS and lane markers 0 to 3 from Table 82–3 are used for the 40GBASE-R PCS. As an example, the lane marker for 100GBASE-R lane number 0 is sent as (left most bit sent first):

10 1000011 00010110 10000100 BIP₃ 01111100 11101001 01111011 BIP₇

After the alignment markers are inserted, data is sent to the PMA or FEC sublayer adjacent to the PCS.

Table 82–2—100GBASE-R Alignment marker encodings

PCS lane number	Encoding ^a {M ₀ , M ₁ , M ₂ , BIP ₃ , M ₄ , M ₅ , M ₆ , BIP ₇ }	PCS lane number	Encoding ^a {M ₀ , M ₁ , M ₂ , BIP ₃ , M ₄ , M ₅ , M ₆ , BIP ₇ }
0	0xC1, 0x68, 0x21, BIP ₃ , 0x3E, 0x97, 0xDE, BIP ₇	10	0xFD, 0x6C, 0x99, BIP ₃ , 0x02, 0x93, 0x66, BIP ₇
1	0x9D, 0x71, 0x8E, BIP ₃ , 0x62, 0x8E, 0x71, BIP ₇	11	0xB9, 0x91, 0x55, BIP ₃ , 0x46, 0x6E, 0xAA, BIP ₇
2	0x59, 0x4B, 0xE8, BIP ₃ , 0xA6, 0xB4, 0x17, BIP ₇	12	0x5C, 0xB9, 0xB2, BIP ₃ , 0xA3, 0x46, 0x4D, BIP ₇
3	0x4D, 0x95, 0x7B, BIP ₃ , 0xB2, 0x6A, 0x84, BIP ₇	13	0x1A, 0xF8, 0xBD, BIP ₃ , 0xE5, 0x07, 0x42, BIP ₇
4	0xF5, 0x07, 0x09, BIP ₃ , 0x0A, 0xF8, 0xF6, BIP ₇	14	0x83, 0xC7, 0xCA, BIP ₃ , 0x7C, 0x38, 0x35, BIP ₇
5	0xDD, 0x14, 0xC2, BIP ₃ , 0x22, 0xEB, 0x3D, BIP ₇	15	0x35, 0x36, 0xCD, BIP ₃ , 0xCA, 0xC9, 0x32, BIP ₇
6	0x9A, 0x4A, 0x26, BIP ₃ , 0x65, 0xB5, 0xD9, BIP ₇	16	0xC4, 0x31, 0x4C, BIP ₃ , 0x3B, 0xCE, 0xB3, BIP ₇
7	0x7B, 0x45, 0x66, BIP ₃ , 0x84, 0xBA, 0x99, BIP ₇	17	0xAD, 0xD6, 0xB7, BIP ₃ , 0x52, 0x29, 0x48, BIP ₇
8	0xA0, 0x24, 0x76, BIP ₃ , 0x5F, 0xDB, 0x89, BIP ₇	18	0x5F, 0x66, 0x2A, BIP ₃ , 0xA0, 0x99, 0xD5, BIP ₇
9	0x68, 0xC9, 0xFB, BIP ₃ , 0x97, 0x36, 0x04, BIP ₇	19	0xC0, 0xF0, 0xE5, BIP ₃ , 0x3F, 0x0F, 0x1A, BIP ₇

^aEach octet is transmitted LSB to MSB.

Table 82–3—40GBASE-R Alignment marker encodings

PCS lane number	Encoding ^a {M ₀ , M ₁ , M ₂ , BIP ₃ , M ₄ , M ₅ , M ₆ , BIP ₇ }
0	0x90, 0x76, 0x47, BIP ₃ , 0x6F, 0x89, 0xB8, BIP ₇
1	0xF0, 0xC4, 0xE6, BIP ₃ , 0x0F, 0x3B, 0x19, BIP ₇
2	0xC5, 0x65, 0x9B, BIP ₃ , 0x3A, 0x9A, 0x64, BIP ₇
3	0xA2, 0x79, 0x3D, BIP ₃ , 0x5D, 0x86, 0xC2, BIP ₇

^aEach octet is transmitted LSB to MSB.

82.2.8 BIP calculations

A PCS lane BIP field is carried in each PCS Lane alignment marker. This allows an accurate and fast measure of the bit error ratio of a given PCS Lane. This information is used to update error counters; no state machines use this information.

Each alignment marker has two Bit Interleaved Parity fields, BIP₃ and BIP₇. BIP₇ is a bit-wise inversion of BIP₃ in order to keep the alignment marker DC balanced. The BIP₃ field contains the result of a bit

interleaved parity calculation. Each bit in the BIP field is an even parity calculation over all of the previous specified bits of a given PCS Lane, from and including the previous alignment marker, but not including the current alignment marker. Using the bit definitions as shown in Figure 82–9, Table 82–4 has the bit assignments for each BIP₃ bit. As an example, BIP₃ bit 0 contains the result of XORing 131072 bits from 16384 66-bit words. BIP₃ bit 3 and bit 4 also include one sync header bit from each 66-bit word. Bit 3 and bit 4 each contain the result of XORing 147456 bits.

Table 82–4—BIP₃ bit assignments

BIP ₃ bit number	Assigned 66-bit word bits
0	2, 10, 18, 26, 34, 42, 50, 58
1	3, 11, 19, 27, 35, 43, 51, 59
2	4, 12, 20, 28, 36, 44, 52, 60
3	0, 5, 13, 21, 29, 37, 45, 53, 61
4	1, 6, 14, 22, 30, 38, 46, 54, 62
5	7, 15, 23, 31, 39, 47, 55, 63
6	8, 16, 24, 32, 40, 48, 56, 64
7	9, 17, 25, 33, 41, 49, 57, 65

BIP₃ and BIP₇ are transmitted LSB to MSB. As an example, with BIP₃ = 0x0F, the PCS lane marker for 100GBASE-R lane number 0 is sent as (left most bit sent first):

```
10 10000011 00010110 10000100 11110000 01111100 11101001 01111011 00001111
```

82.2.9 PMA or FEC Interface

When the transmit channel is operating in normal mode, the 40GBASE-R PCS sends four data streams via *inst:IS_UNITDATA_i.request* primitives and the 100GBASE-R PCS sends twenty data streams via *inst:IS_UNITDATA_i.request* primitives.

The *inst:IS_UNITDATA_i.request* primitives are separate serial streams of bits. Since 66-bit blocks were distributed to each lane, that means for the 40GBASE-R PCS: bits 0 to 65 are sent on PCSL 0, bits 66 to 131 are sent on PCSL 1; bits 132 to 197 are sent on PCSL 2, bits 198 to 263 are sent on PCSL 3, then bits 264 to 329 are sent on PCSL 0 etc.

For 100GBASE-R it is: bits 0 to 65 on PCSL 0, bits 66 to 131 on PCSL 1, bits 132 to 197 on PCSL 2, bits 198 to 263 on PCSL 3, bits 264 to 329 on PCSL 4, bits 330 to 395 on PCSL 5, bits 396 to 461 on PCSL 6, bits 462 to 527 on PCSL 7, bits 528 to 593 on PCSL 8, bits 594 to 659 on PCSL 9, bits 660 to 725 on PCSL 10, bits 726 to 791 on PCSL 11, bits 792 to 857 on PCSL 12, bits 858 to 923 on PCSL 13, bits 924 to 989 on PCSL 14, bits 990 to 1055 on PCSL 15, bits 1056 to 1121 on PCSL 16, bits 1122 to 1187 on PCSL 17, bits 1188 to 1253 on PCSL 18, bits 1254 to 1319 on PCSL 19, then bits 1320 to 1385 on PCSL 0, etc.

82.2.10 Test-pattern generators

The PCS shall have the ability to generate and detect a scrambled idle test pattern. This test-pattern mode is suitable for receiver tests and for certain transmitter tests.

When a scrambled idle pattern is enabled, the test pattern is generated by the scrambler. No seeding of the scrambler is required during test-pattern operation. The input to the scrambler is a control block (block type=0x1E) with all idles as defined in Figure 82–5. Note that the sync headers and alignment markers are added to the stream so that the receive PCS can align and deskew the PCS lanes.

If a Clause 45 MDIO is implemented, then control of the test-pattern generation is from the BASE-R PCS test-pattern control register (bit 3.42.3).

When the transmit channel is operating in test-pattern mode, the encoded bit stream is distributed to the PCS Lanes as in normal operation (see 82.2.4).

82.2.11 Block synchronization

When the receive channel is operating in normal mode, the block synchronization function receives data via 4 (for 40GBASE-R) or 20 (for 100GBASE-R) IS_UNITDATA_0.indication primitives. The PCS forms 4 or 20 bit streams from the primitives by concatenating the bits from the indications of each primitive in order from each *inst*:IS_UNITDATA_0.indication to *inst*:IS_UNITDATA_3.indication or *inst*:IS_UNITDATA_0.indication to *inst*:IS_UNITDATA_19.indication. It obtains lock to the 66-bit blocks in each bit stream using the sync headers and outputs 66-bit blocks. Block lock is obtained as specified in the block lock state diagram shown in Figure 82–10.

82.2.12 PCS lane deskew

Once the receiver achieves block lock on a lane, then it begins obtaining alignment marker lock as specified in the alignment marker lock state diagram shown in Figure 82–11. This process identifies the PCS lane number received on a particular lane of the service interface. After alignment marker lock is achieved on all lanes (4 or 20 lanes), then all inter-lane Skew is removed as shown in the PCS deskew state diagram in Figure 82–12. The Skew budget that the PCS receiver shall support is shown in Table 82–5.

Note that Skew is defined in 80.5.

Table 82–5—Skew tolerance requirements

PCS	Maximum Skew	Maximum Skew Variation
40GBASE-R	180 ns (~1856 bits)	4ns (~41 bits)
100GBASE-R	180 ns (~928 bits)	4ns (~21 bits)

82.2.13 PCS lane reorder

Transmit PCS lanes can be received on different lanes of the service interface from which they were originally transmitted due to Skew between lanes and multiplexing by the PMA. The receive PCS shall order the received PCS lanes according to the PCS lane number.

82.2.14 Alignment marker removal

After all PCS lanes are aligned and deskewed, the PCS lanes are multiplexed together in the proper order to reconstruct the original stream of blocks and the alignment markers are deleted from the data stream. The difference in rate from the deleted alignment markers is compensated for by inserting idle control characters by a function in the Receive process. Note that an alignment marker is always deleted when a given PCS Lane is in `am_lock=true` even if it does not match the expected alignment marker value (due to a bit error for example). Repeated alignment marker errors will result in `am_lock` being set to false for a given PCS Lane, but until that happens it is sufficient to delete the block in the alignment marker position.

As part of the alignment marker removal process, the `BIP3` field is compared to the calculated BIP value for each PCS lane. If a Clause 45 MDIO is implemented, then the appropriate BIP error counter register (registers 3.200 through 3.219) is incremented by one each time the calculated BIP value does not equal the value received in the `BIP3` field. The incoming bit error ratio can be estimated by dividing the BIP block error ratio by a factor of 1081344.

82.2.15 Descrambler

The descrambler is identical to that used in Clause 49, see 49.2.10 for the definition.

82.2.16 Receive process

The receive process decodes blocks to produce `RXD<63:0>` and `RXC<7:0>` for transmission to the XLGMII/CGMII. One XLGMII/CGMII data transfer is decoded from each block. The receive process must insert idle control characters to compensate for the removal of alignment markers. If the PCS receive process spans multiple clock domains, it may also perform clock rate compensation via the deletion of idle control characters or sequence ordered sets or the insertion of idle control characters.

The receive process decodes blocks as specified in the receive state diagram shown in Figure 82–15.

82.2.17 Test-pattern checker

When the receive channel is operating in scrambled idle test-pattern mode, the scrambled idle test-pattern checker checks the bits received via `inst:IS_UNITDATA_i.indication` primitives.

The scrambled idle test-pattern checker utilizes the block lock state diagram, the alignment marker state diagram, the PCS deskew state diagram, and the descrambler operating as they do during normal data reception. The BER monitor state diagram is disabled during receive test-pattern mode. When `align_status` is true and the scrambled idle receive test-pattern mode is active, the scrambled idle test-pattern checker observes the sync header and the output from the descrambler. When the sync header and the output of the descrambler is the all idle pattern, a match is detected. When operating in scrambled idle test pattern, the test-pattern error counter counts blocks with a mismatch. Any mismatch indicates an error and shall increment the test-pattern error counter. Because of the error multiplication characteristics of the descrambler, the incoming bit error ratio can be estimated by dividing the 66-bit block error ratio by a factor of 124.

If a Clause 45 MDIO is implemented, then control of the test-pattern reception is from the BASE-R PCS test-pattern control register (bit 3.42.2). In addition errors are counted in the BASE-R PCS test-pattern error counter (3.43.15:0).

82.2.18 Detailed functions and state diagrams

82.2.18.1 State diagram conventions

The body of this subclause is comprised of state diagrams, including the associated definitions of variables, constants, and functions. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5. State diagram timers follow the conventions of 14.2.3.2. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

82.2.18.2 State variables

82.2.18.2.1 Constants

EBLOCK_R<71:0>

72 bit vector to be sent to the XLGMII/CGMII interface containing /E/ in all the eight character locations.

EBLOCK_T<65:0>

66 bit vector to be sent to the PMA containing /E/ in all the eight character locations.

LBLOCK_R<71:0>

72 bit vector to be sent to the XLGMII/CGMII interface containing one Local Fault ordered_set. The Local Fault ordered_set is defined in 81.3.4.

LBLOCK_T<65:0>

66 bit vector to be sent to the PMA containing one Local Fault ordered_set.

82.2.18.2.2 Variables

align_status

A variable set by the PCS deskew process to reflect the status of the PCS lane-to-lane alignment. Set true when all lanes are synchronized and aligned, set false when the deskew process is not complete.

alignment_valid

Boolean variable that is set true if all PCS lanes are aligned. It is valid when each lane is in am_lock, with each PCS lane locked to a unique alignment marker from Table 82–2 or Table 82–3, and when all PCS lanes are deskewed. Otherwise, alignment_valid is false.

am_lock<x>

Boolean variable that is set true when receiver acquires alignment marker delineation for a given lane of the service interface, where $x = 0:3$ for 40GBASE-R and $x = 0:19$ for 100GBASE-R.

am_slip_done

Boolean variable that is asserted true when the AM_SLIP requested by the alignment marker lock state diagram has been completed indicating that the next candidate 66-bit block position can be tested.

am_status

A Boolean variable that is true when all lanes are in am_lock and false when at least one lane is not in am_lock.

am_counter_done

Boolean variable that indicates that the alignment marker counter is done.

am_valid

Boolean variable that is set true if received block rx_coded is a valid alignment marker. A valid alignment marker will match one of the encodings in Table 82–2 or Table 82–3, excluding the BIP₃ and BIP₇ fields, and it will be repeated every 16384 blocks.

- ber_test_sh**
Boolean variable that is set true when a new sync header is available for testing and false when BER_TEST_SH state is entered. A new sync header is available for testing when the Block Sync process has accumulated enough bits from the PMA to evaluate the header of the next block.
- block_lock<x>**
Boolean variable that is set true when receiver acquires block delineation for a given lane of the service interface, where $x = 0:3$ for 40GBASE-R and $x = 0:19$ for 100GBASE-R.
- current_am**
This variable holds the value of the current alignment marker. This is compared to the variable first_am to determine if we have alignment marker lock and is always $n \times 16384$ 66-bit blocks away from the first_am.
- enable_deskew**
A Boolean variable that indicates the enabling and disabling of the deskew process. Blocks may be discarded whenever deskew is enabled. True when deskew is enabled, false when deskew is disabled.
- first_am**
A variable that holds the value of the first alignment marker that is recognized on a given lane. This is used later to compare to future alignment markers.
- hi_ber**
Boolean variable which is asserted true when the ber_cnt equals or exceeds 97 indicating a bit error ratio $>10^{-4}$
- lane_mapping<x>**
This variable indicates which PCS lane is received on lane x of the service interface when am_lock<x> = true, where $x = 0:3$ for 40GBASE-R and $x = 0:19$ for 100GBASE-R.
- r_block_type**
This variable contains the rx_coded<65:0> vector classification results, returned by the R_TYPE function. It can assume one of the following five values {C,S,T,D,E}, as defined by the R_TYPE function.
- r_block_type_next**
This variable contains the rx_coded<65:0> vector classification results, returned by the R_TYPE_NEXT function. It can assume one of the following five values {C,S,T,D,E}, as defined by the R_TYPE_NEXT function.
- r_test_mode**
Boolean variable that is asserted true when the receiver is in test-pattern mode.
- reset**
Boolean variable that controls the resetting of the PCS. It is true whenever a reset is necessary including when reset is initiated from the MDIO, during power on, and when the MDIO has put the PCS into low-power mode.
- rx_coded<65:0>**
Vector containing the input to the 64B/66B decoder. The format for this vector is shown in Figure 82–5. The leftmost bit in the figure is rx_coded<0> and the rightmost bit is rx_coded<65>.
- rx_raw<71:0>**
Vector containing one XLGMII/CGMII transfer. RXC<0> through RXC<7> are from rx_raw<0> through rx_raw<7>, respectively. RXD<0> through RXD<63> are from rx_raw<8> through rx_raw<71>, respectively.
- sh_valid**
Boolean variable that is set true if received block rx_coded has valid sync header bits. That is, sh_valid is asserted if rx_coded<0> \neq rx_coded<1> and de-asserted otherwise.
- signal_ok**
Boolean variable that is set based on the most recently received value of *inst:IS_UNITDATA.indication(SIGNAL_OK)*. It is true if the value was OK and false if the value was FAIL.

slip_done

Boolean variable that is asserted true when the SLIP requested by the Block Lock state diagram has been completed indicating that the next candidate block sync position can be tested.

t_block_type

This variable contains the tx_raw<71:0> vector classification results, returned by the T_TYPE function. It can assume one of the following five values {C,S,T,D,E}, as defined by the T_TYPE function.

test_am

Boolean variable that is set true when a new block is available for testing and false when FIND_1ST state is entered. A new block is available for testing when the Block Sync process has accumulated enough bits from the PMA to evaluate the next block

test_sh

Boolean variable that is set true when a new sync header is available for testing and false when TEST_SH and TEST_SH2 state are entered. A new sync header is available for testing when the Block Sync process has accumulated enough bits from the PMA to evaluate the header of the next block

tx_coded<65:0>

Vector containing the output from the 64B/66B encoder. The format for this vector is shown in Figure 82–5. The leftmost bit in the figure is tx_coded<0> and the rightmost bit is tx_coded<65>.

tx_raw<71:0>

Vector containing one XLGMII/CGMII transfer. TXC<0> through TXC<7> are placed in tx_raw<0> through tx_raw<7>, respectively. TXD<0> through TXD<63> are placed in tx_raw<8> through tx_raw<71>, respectively.

82.2.18.2.3 Functions**AM_SLIP**

Causes the next candidate block position to be tested. The precise method for determining the next candidate block position is not specified and is implementation dependent. However, an implementation shall ensure that all possible blocks are evaluated.

DECODE(rx_coded<65:0>)

Decodes the 66-bit vector returning rx_raw<71:0>, which is sent to the XLGMII/CGMII. The DECODE function shall decode the block as specified in 82.2.3.

ENCODE(tx_raw<71:0>)

Encodes the 72-bit vector returning tx_coded<65:0> of which tx_coded<65:2> is sent to the scrambler. The two bits of the sync header bypass the scrambler. The ENCODE function shall encode the block as specified in 82.2.3.

R_TYPE(rx_coded<65:0>)

This function classifies the current rx_coded<65:0> vector as belonging to one of the following five types, depending on its contents. The classification results are returned via the r_block_type variable.

Values: C; The vector contains a sync header of 10 and one of the following:

- a) A block type field of 0x1E and eight valid control characters other than /E/;
- b) A block type field of 0x4B.

S; The vector contains a sync header of 10 and the following:

- a) A block type field of 0x78.

T; The vector contains a sync header of 10, a block type field of 0x87, 0x99, 0xAA, 0xB4, 0xCC, 0xD2, 0xE1 or 0xFF and all control characters are valid.

D; The vector contains a sync header of 01.

E; The vector does not meet the criteria for any other value.

Valid control characters are specified in Table 82–1.

R_TYPE_NEXT

This function classifies the 66-bit rx_coded vector that immediately follows the current rx_coded<65:0> vector as belonging to one of the five types defined in R_TYPE, depending on its contents. It is intended to perform a prescient end of packet check. The classification results are returned via the r_block_type_next variable.

SLIP

Causes the next candidate block sync position to be tested. The precise method for determining the next candidate block sync position is not specified and is implementation dependent. However, an implementation shall ensure that all possible bit positions are evaluated.

T_TYPE = (tx_raw<71:0>)

This function classifies each 72-bit tx_raw vector as belonging to one of the five types depending on its contents. The classification results are returned via the t_block_type variable.

Values: C; The vector contains one of the following:

- a) Eight valid control characters other than /O/, /S/, /T/ and /E/;
- b) One valid ordered set.

S; The vector contains an /S/ in its first character, and all characters following the /S/ are data characters.

T; The vector contains a /T/ in one of its characters, all characters before the /T/ are data characters, and all characters following the /T/ are valid control characters other than /O/, /S/ and /T/.

D; The vector contains eight data characters.

E; The vector does not meet the criteria for any other value.

A tx_raw character is a control character if its associated TXC bit is asserted. A valid control character is one containing an XLGMII/CGMII control code specified in Table 82–1. A valid ordered_set consists of a valid /O/ character in the first character and data characters in the seven characters following the /O/. A valid /O/ is any character with a value for O code in Table 82–1.

82.2.18.2.4 Counters**am_counter**

This counter counts 16383 66-bit blocks that separate two consecutive alignment markers.

am_invl_d_cnt

Count of the number of invalid alignment markers seen in a row. This counter is always reset to all zeros when a valid marker is detected.

ber_count

A 22-bit counter that counts each time BER_BAD_SH state is entered. This counter is reflected in MDIO register bits 3.33.13:8 and 3.44.15:0.

ber_cnt

Count up to a maximum of 97 of the number of invalid sync headers within the current 1.25 ms (40GBASE-R) or 500 μs (100GBASE-R) period.

errored_block_count:

When the receiver is in normal mode, this 22-bit counter counts once for each time RX_E state is entered. This counter is reflected in MDIO register bits 3.33.7:0 and 3.45.13:0.

sh_cnt

Count of the number of sync headers checked within the current 64 or 1024 block window.

sh_invl_d_cnt

Count of the number of invalid sync headers within the current 64 or 1024 block window.

test_pattern_error_count:

When the receiver is in test-pattern mode, this 16-bit counter counts errors as described in 82.2.17. This counter is reflected in MDIO register bits 3.43.15:0.

82.2.18.2.5 Timers

State diagram timers follow the conventions of 14.2.3.2.

xus_timer

Timer that is triggered every 1.25 ms +1%, –25% for 40GBASE-R or 500 μ s +1%, –25% for 100GBASE-R.

82.2.18.3 State diagrams

The 40GBASE-R PCS shall implement four block lock processes as depicted in Figure 82–10. The 100GBASE-R PCS shall implement twenty block lock processes as depicted in Figure 82–10. A block lock process operates independently on each lane. Each block lock process looks for 64 valid sync headers in a row to declare lock. A valid sync header is either a 01 or a 10. Once in lock, the lock process looks for 65 invalid sync headers within a 1024 sync window to declare out of lock. An invalid sync header is a 11 or 00. Once block lock is achieved on a lane, then the alignment marker process starts.

The 40GBASE-R PCS shall implement four alignment marker lock processes as depicted in Figure 82–11. The 100GBASE-R PCS shall implement twenty alignment marker lock processes as depicted in Figure 82–11. An alignment marker lock process operates independently on each lane. The alignment marker lock state diagram shown in Figure 82–11 determines when the PCS has obtained alignment marker lock to the received bit stream for a given lane of the service interface. Each alignment marker lock process looks for two valid alignment markers 16384 66-bit blocks apart to gain alignment marker lock. On a given lane of the service interface, lane markers must match each other and an entry from Table 82–2 for 100GBASE-R or Table 82–3 for 40GBASE-R. Note that the BIP₃ and BIP₇ fields are excluded from the markers when making a match to each other or the tables. Once in lock, a lane will go out of alignment marker lock if four markers are received in a row that do not match the alignment marker that the lane is currently locked to. When the alignment marker lock process achieves lock for a lane, and if a Clause 45 MDIO is implemented, the PCS shall record the number of the PCS lane received on that lane of the service interface in the appropriate lane mapping register (3.400 to 3.419).

The PCS shall run the deskew process as depicted in Figure 82–12. The PCS deskew process is responsible for determining if the PCS is capable of presenting coherent data to the XLGMII/CGMII. The deskew process ensures that all PCS lanes have alignment marker lock, are locked to different alignment markers, and that the Skew is within the boundaries of what the PCS can deskew.

The BER Monitor state diagram shown in Figure 82–13 monitors the received aggregate signal for high bit error ratio. The high BER state shall be entered if 97 invalid 66-bit sync headers are detected within a 500 μ s window for 100GBASE-R, or a 1.25 ms window for 40GBASE-R. The high BER state is exited once there are less than 97 invalid sync headers in the same window.

The Transmit state diagram shown in Figure 82–14 controls the encoding of transmitted blocks. It makes exactly one transition for each transmit block processed. Though the Transmit state diagram sends Local Fault ordered sets when reset is asserted, the scrambler may not be operational during reset. Thus, the Local Fault ordered sets may not appear on the PMA/FEC service interface.

The Receive state diagram shown in Figure 82–15 controls the decoding of received blocks. It makes exactly one transition for each receive block processed.

The PCS shall perform the functions of block lock, alignment marker lock, PCS deskew, BER Monitor, Transmit, and Receive as specified in the respective state diagrams.

82.3 PCS Management

The following objects apply to PCS management. If an MDIO Interface is provided (see Clause 45), they are accessed via that interface. If not, it is recommended that an equivalent access be provided.

82.3.1 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PCS. Mapping of MDIO control variables to PCS control variables is shown in Table 82–6. Mapping of MDIO status variables to PMD status variables is shown in Table 82–7.

Table 82–6—MDIO/PMD control variable mapping

MDIO control variable	PCS register name	Register/ bit number	PCS control variable
Reset	PCS control 1 register	3.0.15	reset
Loopback	PCS control 1 register	3.0.14	Loopback
Transmit test-pattern enable	BASE-R PCS test-pattern control register	3.42.3	tx_test_mode
Receive test-pattern enable	BASE-R PCS test-pattern control register	3.42.2	rx_test_mode

Table 82–7—MDIO/PMD status variable mapping

MDIO status variable	PCS register name	Register/ bit number	PCS status variable
BASE-R and 10GBASE-T receive link status	BASE-R and 10GBASE-T PCS status 1 register	3.32.12	PCS_status
BASE-R and 10GBASE-T PCS high BER	BASE-R and 10GBASE-T PCS status 1 register	3.32.1	hi_ber
Block <i>x</i> lock	Multi-lane BASE-R PCS alignment status 1 and 2 registers	3.50.7:0 3.51.11:0	block_lock< <i>x</i> >
Lane <i>x</i> aligned	Multi-lane BASE-R PCS alignment status 3 and 4 registers	3.52.7:0 3.53.11:0	am_lock< <i>x</i> >
PCS lane alignment status	Multi-lane BASE-R PCS alignment status 1 register	3.50.12	align_status
BER	BASE-R and 10GBASE-T PCS status 2 register BER high order counter register	3.33.13:8 3.44.15:0	ber_count
Errored blocks	BASE-R and 10GBASE-T PCS status 2 register Errored blocks high order counter register	3.33.7:0 3.45.13:0	errored_block_count
Test-pattern error counter	BASE-R PCS test-pattern error counter register	3.43.15:0	test_pattern_error_count
BIP error counter, lane <i>x</i>	BIP error counter, lane <i>x</i> register	3.200 through 3.219	bip_counter
Lane <i>x</i> mapping	Lane <i>x</i> mapping register	3.400 through 3.419	lane_mapping

82.4 Loopback

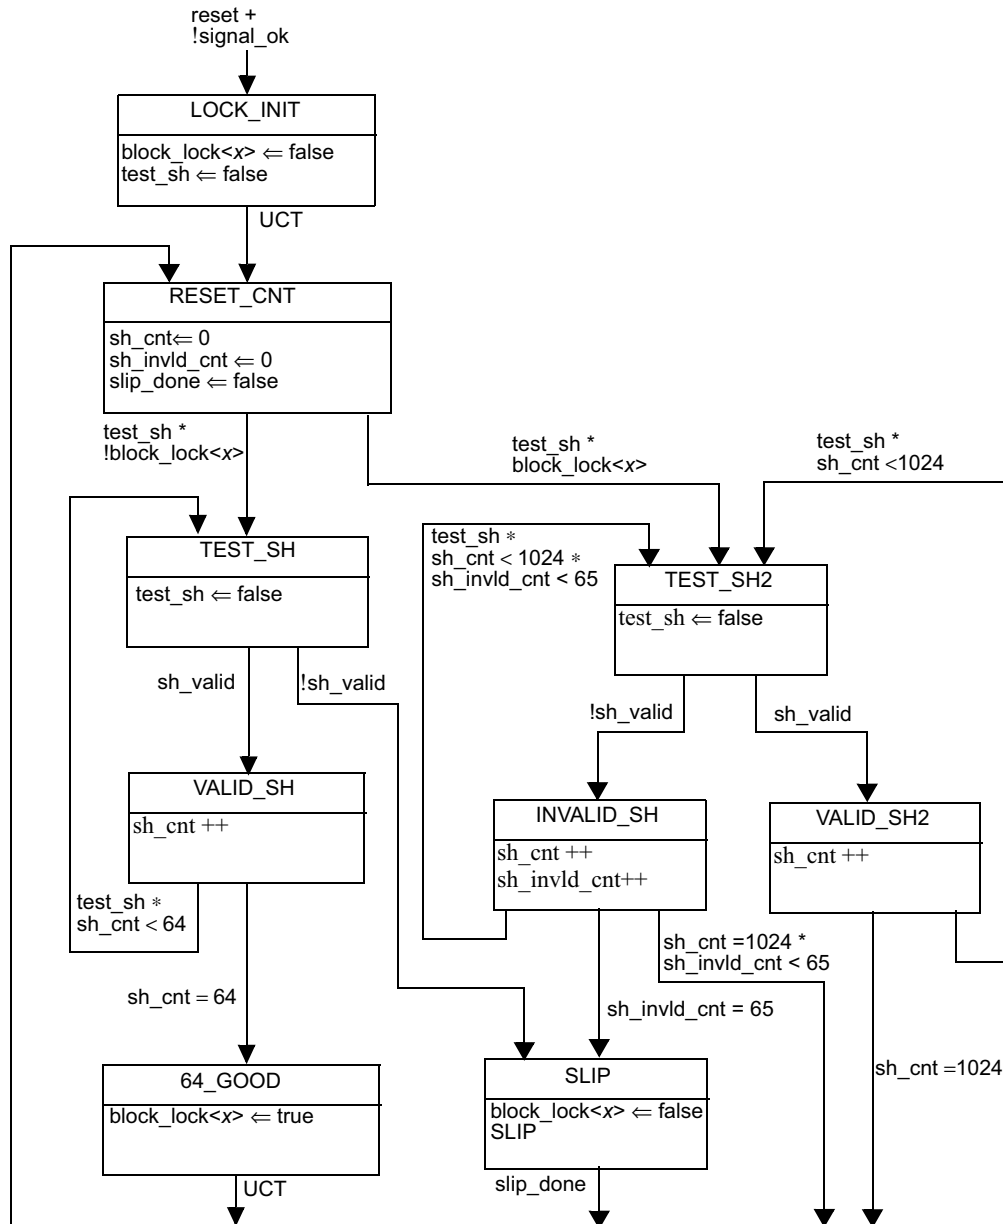
If a Clause 45 MDIO is implemented, then the PCS shall be placed in the loopback mode when the loopback bit from the PCS control 1 register (bit 3.0.14) is set to a one. In this mode, the PCS shall accept data on the transmit path from the XLGMII/CGMII and return it on the receive path to the XLGMII/CGMII. In addition, the PCS shall transmit what it receives from the XLGMII/CGMII to the PMA/FEC sublayer, and shall ignore all data presented to it by the PMA/FEC sublayer.

82.5 Delay constraints

The maximum delay contributed by the 40GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 11264 BT (22 pause_quanta or 281.6 ns). The maximum delay contributed by the 100GBASE-R PCS (sum of transmit and receive delays at one end of the link) shall be no more than 35328 BT (69 pause_quanta or 353.28 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

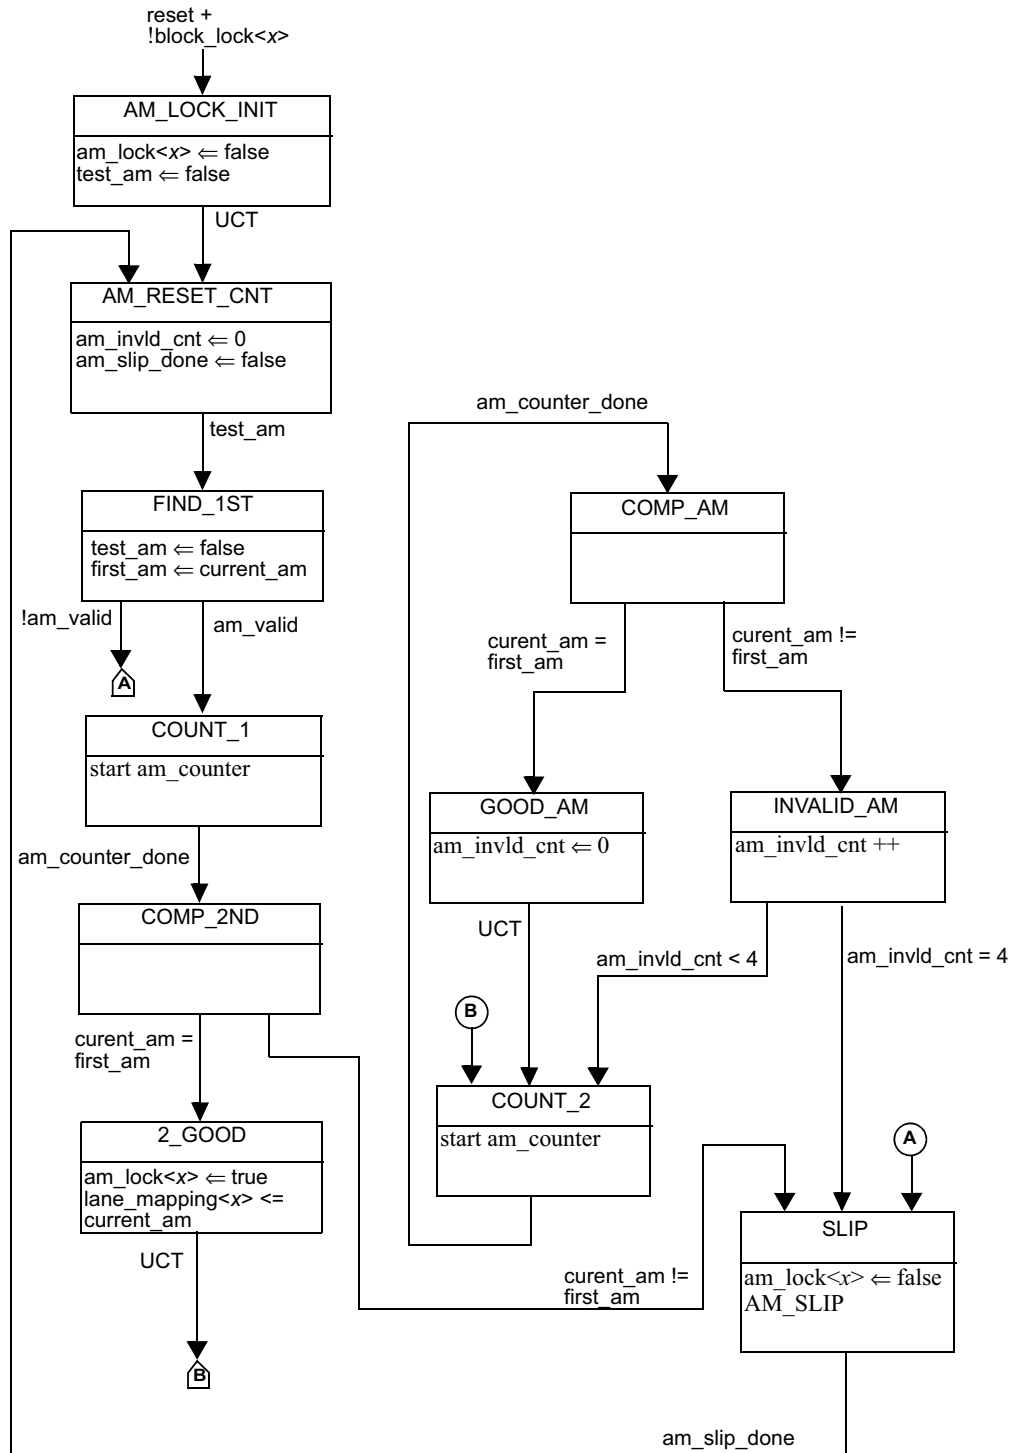
82.6 Auto-Negotiation

The following requirements apply to a PCS used with a 40GBASE-KR4 PMD, 40GBASE-CR4 PMD, or 100GBASE-CR10 PMD where support for the Auto-Negotiation process defined in Clause 73 is mandatory. The PCS shall support the primitive AN_LINK.indication(link_status) (see 73.9). The parameter link_status shall take the value FAIL when PCS_status=false and the value OK when PCS_status=true. The primitive shall be generated when the value of link_status changes.



NOTE— block_lock<x> refers to the received lane x of the service interface, where x = 0:3 (for 40GBASE-R) or 0:19 (for 100GBASE-R)

Figure 82–10—Block lock state diagram



NOTE— am_lock<x> refers to the received lane x of the service interface, where x = 0:3 (for 40GBASE-R) or 0:19 (for 100GBASE-R)

Figure 82–11—Alignment marker lock state diagram

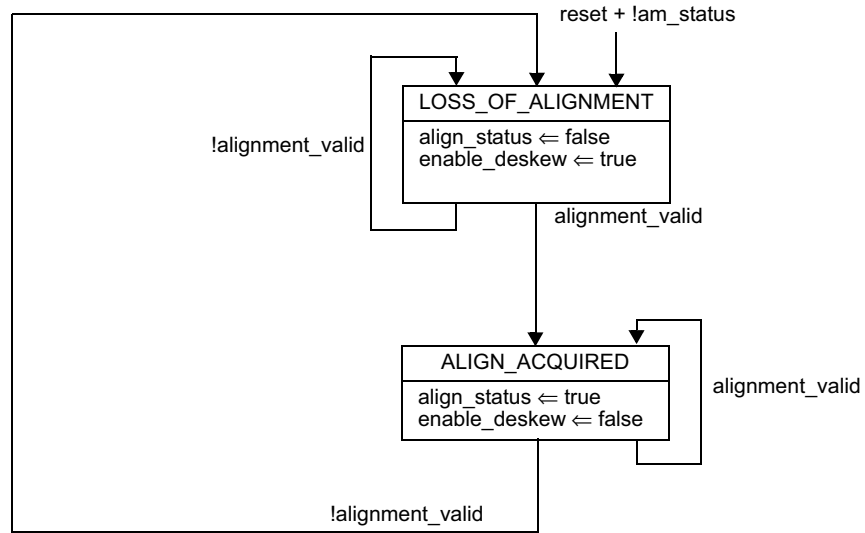


Figure 82–12—PCS deskew state diagram

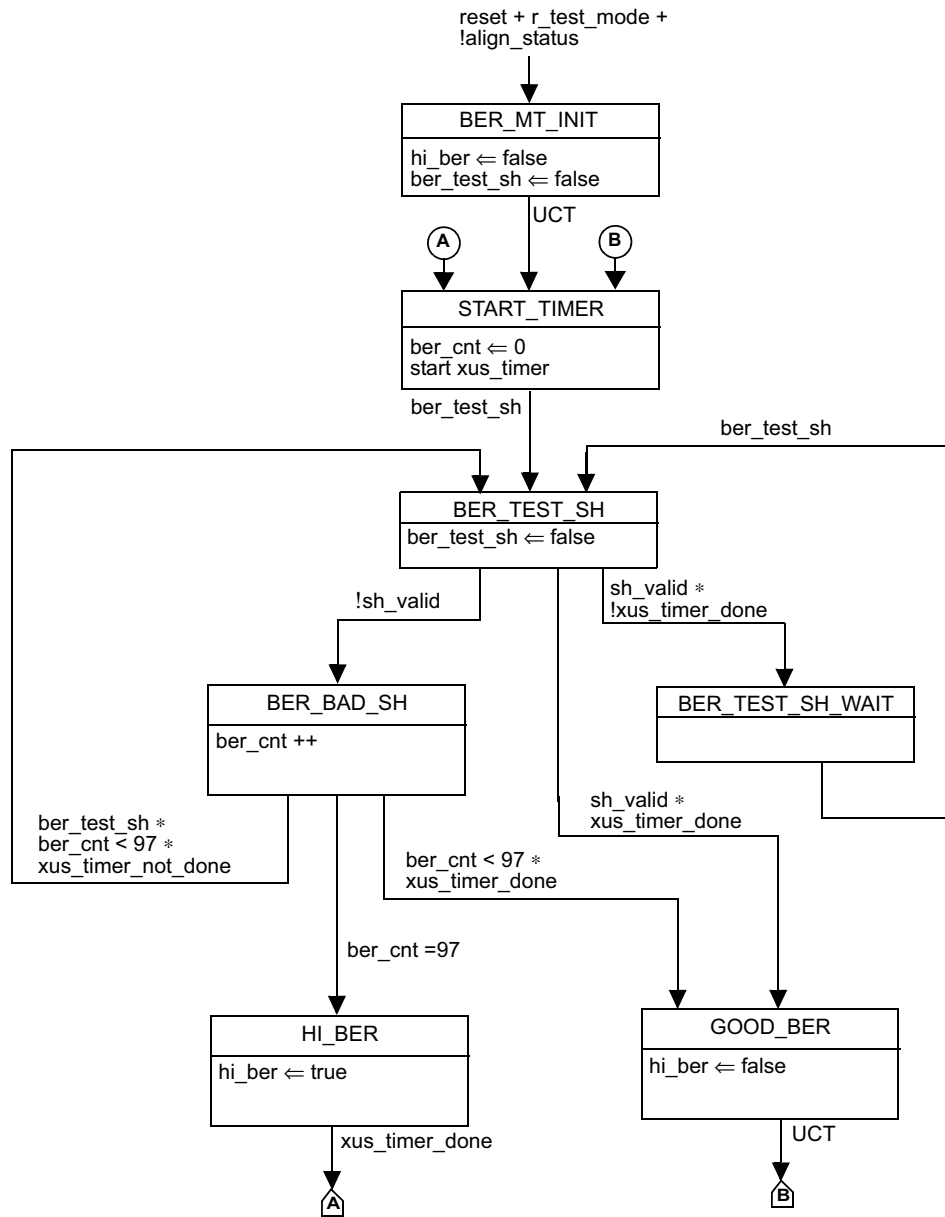


Figure 82-13—BER monitor state diagram

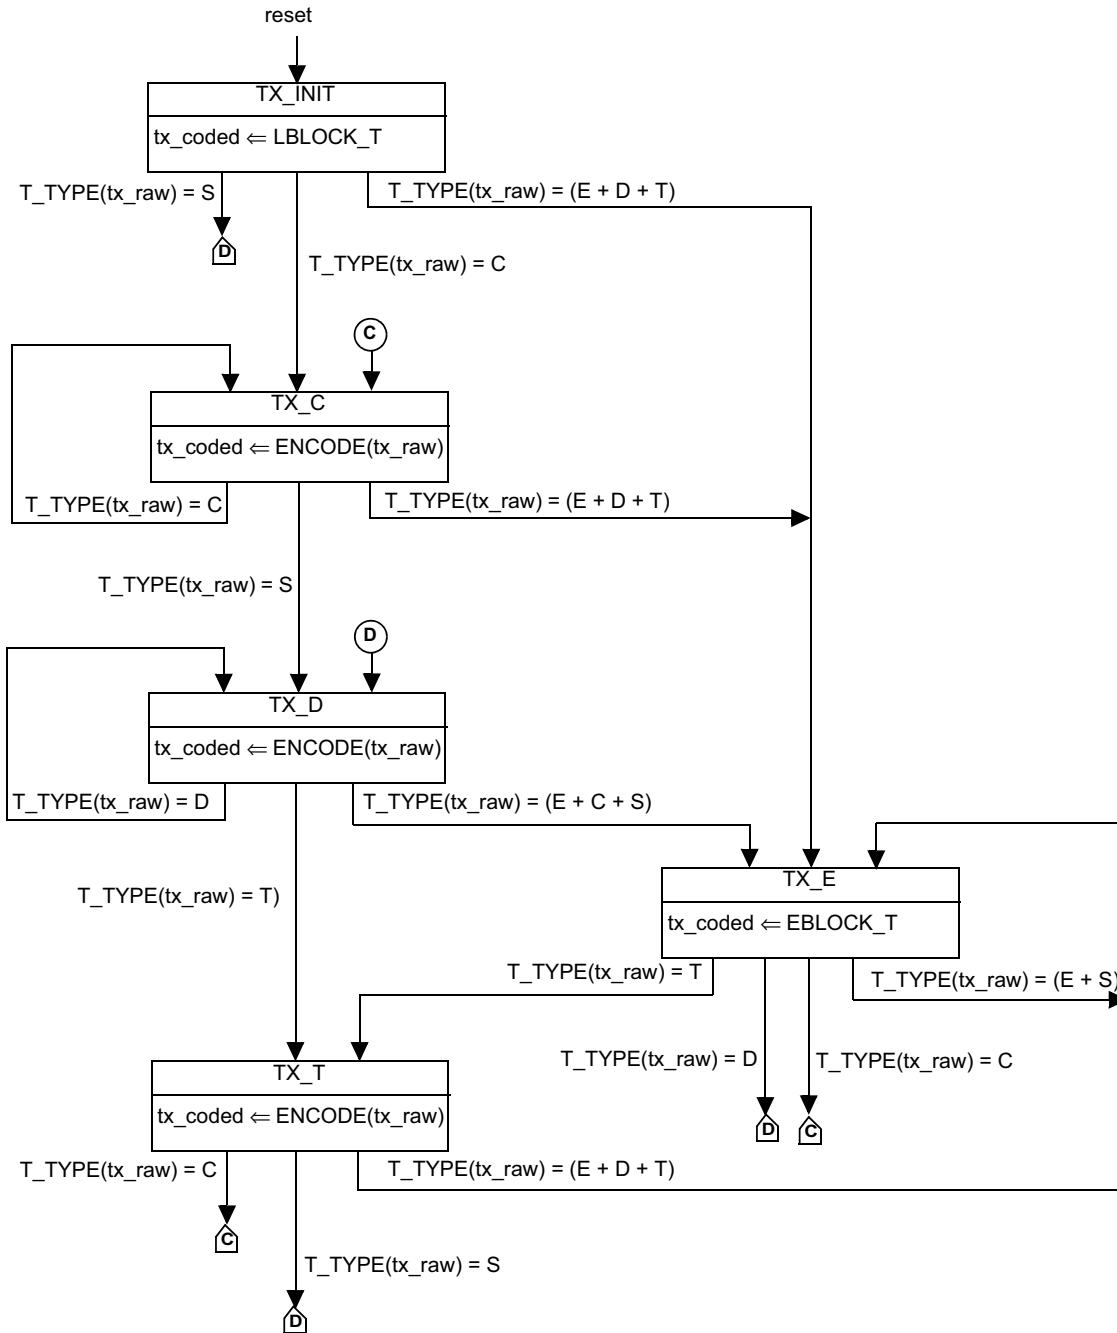


Figure 82–14—Transmit state diagram

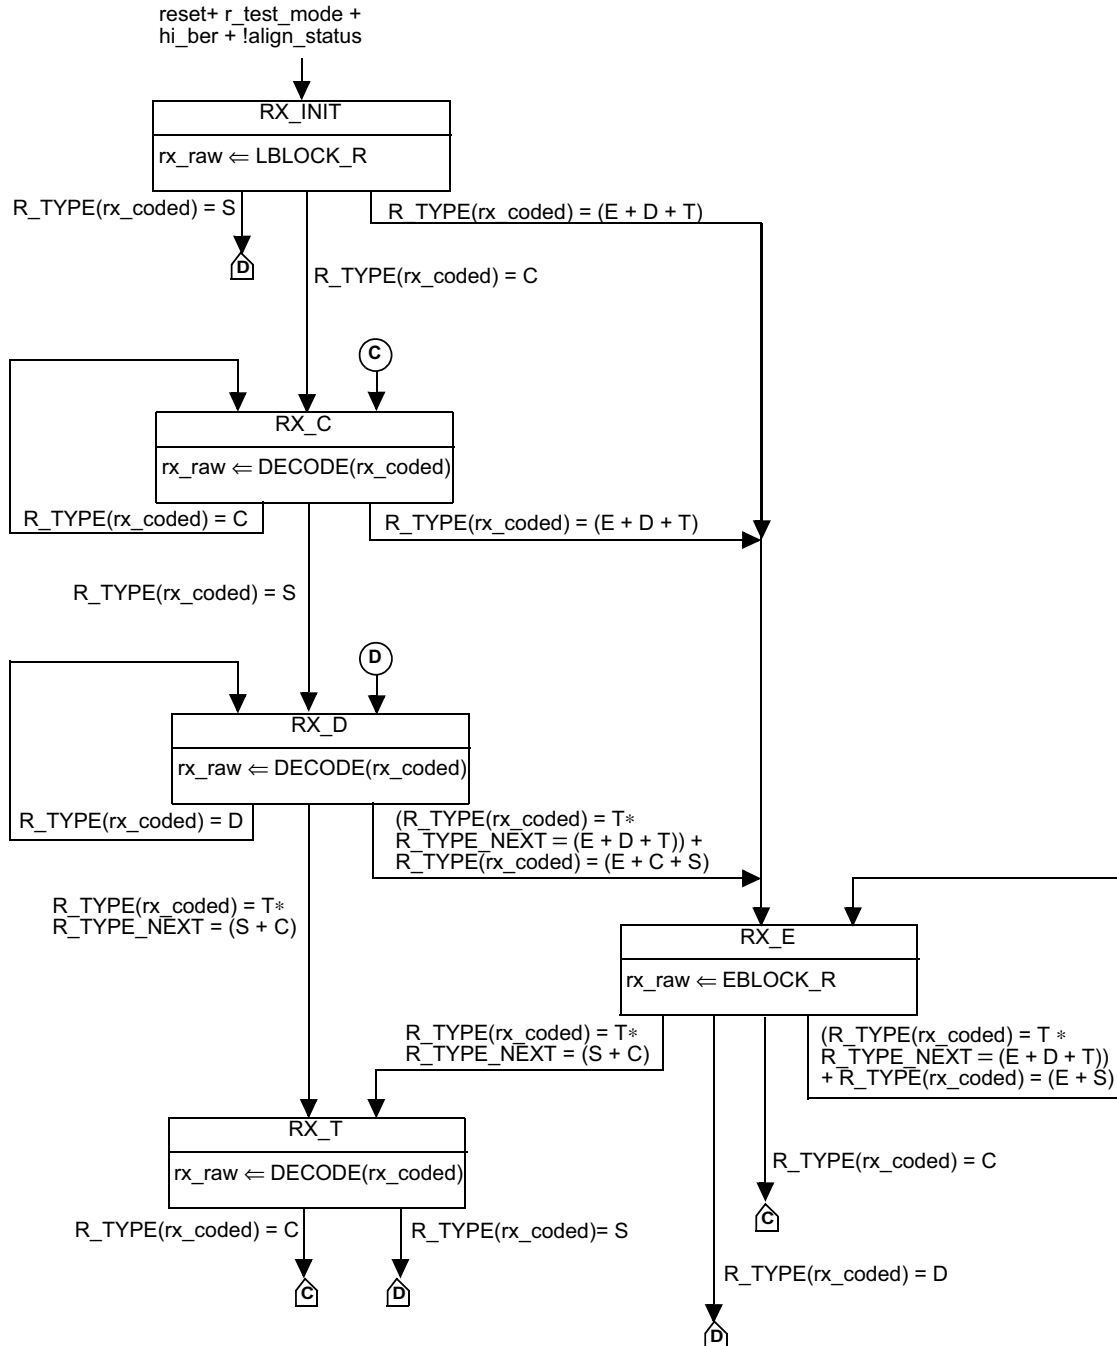


Figure 82–15—Receive state diagram

82.7 Protocol implementation conformance statement (PICS) proforma for Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R¹²

82.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

82.7.2 Identification

82.7.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1— Required for all implementations. NOTE 2— May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

82.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 82, Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

¹²Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

82.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PCS40	PCS for 40GBASE-R	82.1.1		O	Yes [] No []
*PCS100	PCS for 100GBASE-R	82.1.1		O	Yes [] No []
XGE40	XLGMII logical interface	81, 82.1.4	Logical interface is supported	O	Yes [] No []
XGE100	XLGMII logical interface	81, 82.1.4	Logical interface is supported	O	Yes [] No []
MD	MDIO	45, 82.3	Registers and interface supported	O	Yes [] No []
PMA	Supports operation directly connected to a PMA	82.1.4.2		O.1	Yes [] No []
FEC	Supports operation directly connected to a FEC sublayer	82.1.4.2		O.1	Yes [] No []
*JTM	Supports test-pattern mode	82.2.1		PMA:M	Yes [] No [] N/A []

82.7.4 PICS Proforma Tables for PCS, type 40GBASE-R and 100GBASE-R**82.7.4.1 Coding rules**

Item	Feature	Subclause	Value/Comment	Status	Support
C1	Encoder (and ENCODE function) implements the code as specified	82.2.3 and 82.2.18.2.3		M	Yes [] No []
C2	Decoder (and DECODE function) implements the code as specified	82.2.3 and 82.2.18.2.3		M	Yes [] No []
C3	Only valid block types are transmitted	82.2.3.3		M	Yes [] No []
C4	Invalid block types are treated as an error	82.2.3.3		M	Yes [] No []
C5	Only valid control characters are transmitted	82.2.3.4		M	Yes [] No []
C6	Invalid control characters are treated as an error	82.2.3.4		M	Yes [] No []
C7	Idles do not interrupt data	82.2.3.6		M	Yes [] No []
C8	IDLE control code insertion and deletion	82.2.3.6	Insertion or Deletion in groups of 8 /I/s	M	Yes [] No []
C9	Sequence ordered_set deletion	82.2.3.9	Only one whole ordered_set of two consecutive sequence ordered sets may be deleted	M	Yes [] No []

82.7.4.2 Scrambler and Descrambler

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Scrambler	82.2.5	Performs as shown in Figure 49-8	M	Yes [] No []
S2	Descrambler	82.2.15	Performs as shown in Figure 49-10	M	Yes [] No []

82.7.4.3 Deskew and Reordering

Item	Feature	Subclause	Value/Comment	Status	Support
DR1	Deskew	82.2.12	Able to deskew up to the value in Table 82–5	M	Yes [] No []
DR2	Reordering	82.2.13	Performs reordering.	M	Yes [] No []

82.7.4.4 Alignment Markers

Item	Feature	Subclause	Value/Comment	Status	Support
AM1	Alignment marker insertion	82.2.7	Alignment markers are inserted periodically as described in section 82.2.7	M	Yes [] No []
AM2	Alignment marker form	82.2.7	Alignment markers are formed as described in section 82.2.7	M	Yes [] No []
AM3	Lane mapping	82.2.18.3	PCS lane number is captured	MD:M	Yes [] No []

82.7.5 Test-pattern modes

Item	Feature	Subclause	Value/Comment	Status	Support
JT1	Scrambled idle transmit test-pattern generator is implemented	82.2.10	Performs as in 82.2.10	JTM:M	Yes [] No [] N/A []
JT2	Scrambled idle receive test-pattern checker is implemented	82.2.17	Performs as in 82.2.17	JTM:M	Yes [] No [] N/A []
JT3	Transmit and receive test-pattern modes can operate simultaneously	82.2.1		JTM:M	Yes [] No [] N/A []

82.7.5.1 Bit order

Item	Feature	Subclause	Value/Comment	Status	Support
B1	Transmit bit order	82.2.3.2	Placement of bits into the PCS lanes as shown in Figure 82–3	M	Yes [] No []
B2	Receive bit order	82.2.3.2	Placement of bits into the XLGMII/CGMII as shown in Figure 82–4	M	Yes [] No []

82.7.6 Management

Item	Feature	Subclause	Value/Comment	Status	Support
M1	Alternate access to PCS Management objects is provided	82.3		O	Yes [] No []

82.7.6.1 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SM1	40GBASE-R Block Lock	82.2.18.3	Implements 4 block lock processes as depicted in Figure 82–10	PCS40:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM2	100GBASE-R Block Lock	82.2.18.3	Implements 20 block lock processes as depicted in Figure 82–10	PCS100:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM3	The SLIP function evaluates all possible bit positions	82.2.18.2.3		M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM4	40GBASE-R Alignment Marker Lock	82.2.18.3	Implements 4 alignment marker lock processes as depicted in Figure 82–11	PCS40:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM5	100GBASE-R Alignment Marker Lock	82.2.18.3	Implements 20 alignment marker lock processes as depicted in Figure 82–11	PCS100:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM6	The AM_SLIP functions evaluates all possible blocks	82.2.18.2.3		M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM7	40GBASE-R PCS deskew state diagram	82.2.18.3	Meets the requirements of Figure 82–12	PCS40:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM8	100GBASE-R PCS deskew state diagram	82.2.18.3	Meets the requirements of Figure 82–12	PCS100:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM9	40GBASE-R BER Monitor	82.2.18.3	Meets the requirements of Figure 82–13 with xus_timer_done set to 1.25 ms	PCS40:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM10	100GBASE-R BER Monitor	82.2.18.3	Meets the requirements of Figure 82–13 with xus_timer_done set to 500 μ s	PCS100:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM11	40GBASE-R Transmit process	82.2.18.3	Meets the requirements of Figure 82–14	PCS40:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM12	100GBASE-R Transmit process	82.2.18.3	Meets the requirements of Figure 82–14	PCS100:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM13	40GBASE-R and 100GBASE-R Receive process	82.2.18.3	Meets the requirements of Figure 82–15	PCS40:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>
SM14	100GBASE-R Receive process	82.2.18.3	Meets the requirements of Figure 82–15	PCS100:M	Yes [<input type="checkbox"/> No [<input type="checkbox"/>

82.7.6.2 Loopback

Item	Feature	Subclause	Value/Comment	Status	Support
L1	Supports loopback	82.4	Performs as in 82.4	M	Yes [] No [] N/A []
L2	When in loopback, transmits what it receives from the XLGMII/CGMII	82.4	Performs as in 82.4	M	Yes [] No []

82.7.6.3 Delay constraints

Item	Feature	Subclause	Value/Comment	Status	Support
TIM1	PCS Delay Constraint	82.5	No more than 11264 BT for sum of transmit and receive path delays for 40GBASE-R and 35328 BT for 100GBASE-R	M	Yes [] No []

82.7.6.5 Auto-Negotiation for Backplane Ethernet functions

Item	Feature	Subclause	Value/Comment	Status	Support
AN1*	Support for use with a 40GBASE-KR4, 40GBASE-CR4, or 100GBASE-CR10 PMD	82.6	AN technology dependent interface described in Clause 73	O	Yes []
AN2	AN_LINK.indication primitive	82.6	Support of the primitive AN_LINK.indication(link_status), when the PCS is used with 10GBASE-KR PMD	AN1:M	Yes []
AN3	link_status parameter	82.6	Takes the value OK or FAIL, as described in 82.6	AN1:M	Yes []
AN4	Generation of AN_LINK.indication primitive	82.6	Generated when the value of link_status changes	AN1:M	Yes []

83. Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R

83.1 Overview

83.1.1 Scope

This clause specifies the Physical Medium Attachment sublayer (PMA) that is common to two families of (40 Gb/s and 100 Gb/s) Physical Layer implementations, known as 40GBASE-R and 100GBASE-R. The PMA allows the PCS (specified in Clause 82) to connect in a media-independent way with a range of physical media. The 40GBASE-R PMA(s) can support any of the 40 Gb/s PMDs in Table 80–2. The 100GBASE-R PMA(s) can support any of the 100 Gb/s PMDs in Table 80–2. The terms 40GBASE-R and 100GBASE-R are used when referring generally to Physical Layers using the PMA defined in this clause.

40GBASE-R and 100GBASE-R can be extended to support any full duplex medium requiring only that the PMD be compliant with the appropriate PMA interface.

The interfaces for the inputs of the 40GBASE-R and 100GBASE-R PMAs are defined in an abstract manner and do not imply any particular implementation. The optional physical instantiation of the PMD service interfaces for 40GBASE-SR4, 40GBASE-LR4, and 100GBASE-SR10 PMDs, known as XLPPi and CPPI, are defined in Annex 86A. The PMD service interfaces for other PMDs are defined in an abstract manner according to 80.3.1. For 40GBASE-R PMAs, electrical interfaces connecting PMA sublayers, known as XLAUI, are defined in Annex 83A and Annex 83B. For 100GBASE-R PMAs, electrical interfaces connecting PMA sublayers, known as CAUI, are defined in Annex 83A and Annex 83B.

83.1.2 Position of the PMA in the 40GBASE-R or 100GBASE-R sublayers

Figure 83–1 shows the relationship of the PMA sublayer (shown shaded) with other sublayers to the ISO Open System Interconnection (OSI) reference model.

83.1.3 Summary of functions

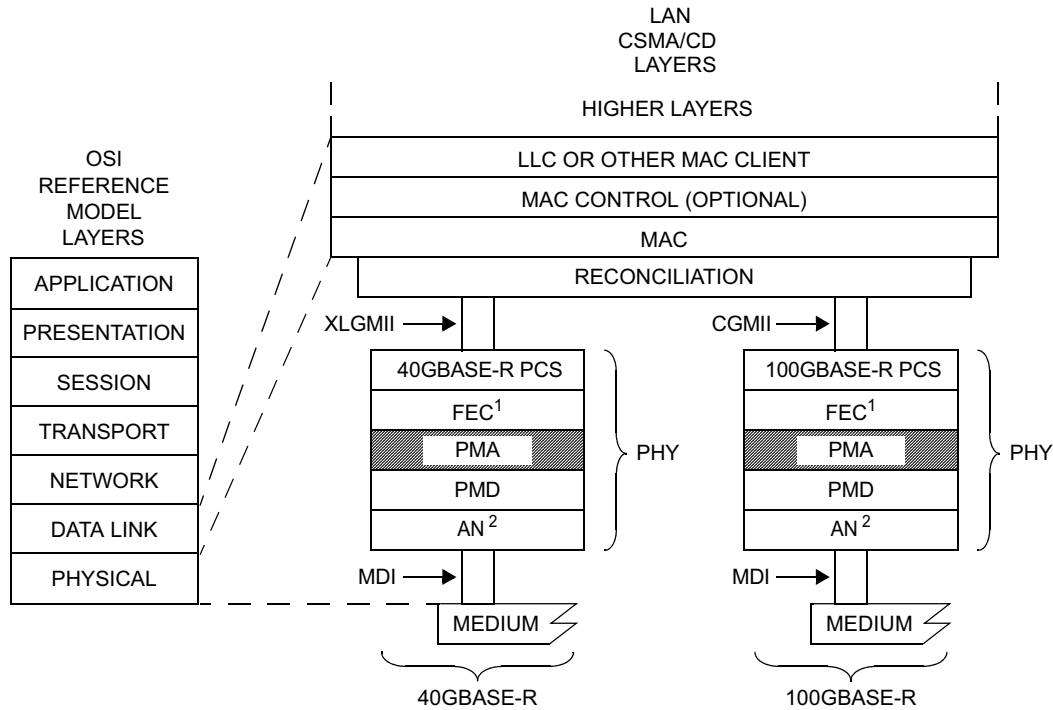
The following is a summary of the principal functions implemented (when required) by the PMA in both the transmit and receive directions:

- a) Adapt the PCSL formatted signal to the appropriate number of abstract or physical lanes.
- b) Provide per input-lane clock and data recovery.
- c) Provide bit-level multiplexing.
- d) Provide clock generation.
- e) Provide signal drivers.
- f) Optionally provide local loopback to/from the PMA service interface.
- g) Optionally provide remote loopback to/from the PMD service interface.
- h) Optionally provide test-pattern generation and detection.
- i) Tolerate Skew Variation.

In addition, the PMA provides receive link status information in the receive direction.

83.1.4 PMA sublayer positioning

An implementation may use one or more PMA sublayers to adapt the number and rate of the PCS lanes to the number and rate of the PMD lanes. The number of PMA sublayers required depends on the partitioning of functionality for a particular implementation. An example is illustrated in Figure 83–2. This example illustrates the partitioning that might arise from use of a FEC device that is separate from the PCS.



AN = AUTO-NEGOTIATION
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

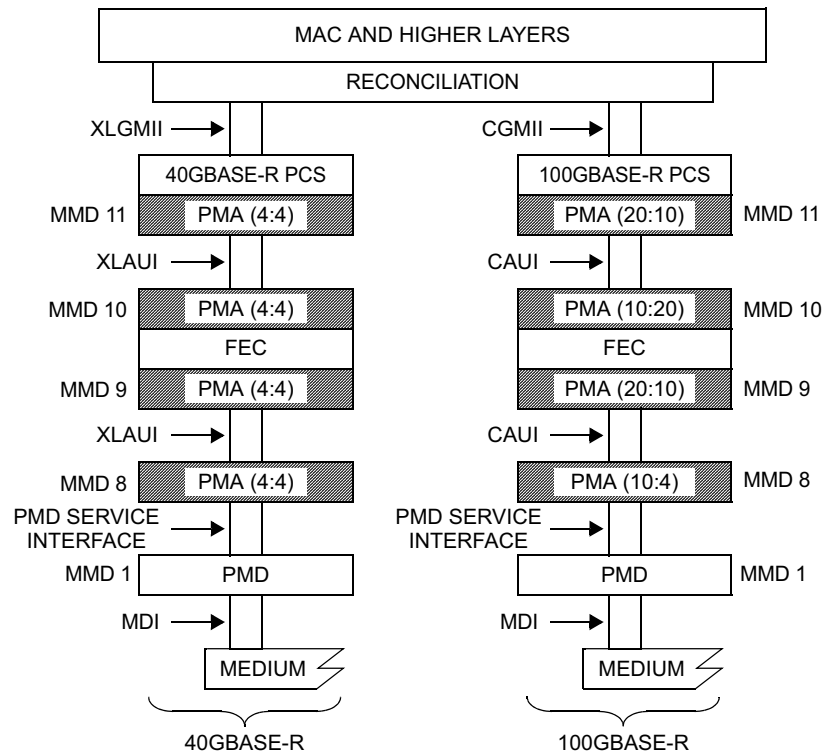
NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
 NOTE 2—CONDITIONAL BASED ON PHY TYPE

Figure 83-1—40GBASE-R and 100GBASE-R PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 CSMA/CD LAN model

Additional examples are illustrated in Annex 83C. Each PMA maps the PCSs from p PMA input lanes to q PMA output lanes in the Tx direction, and from q PMA input lanes to p PMA output lanes in the Rx direction.

Management Data Input/Output (MDIO) Manageable Device (MMD) addresses 1, 8, 9, 10, and 11 are available for addressing multiple instances of PMA sublayers (see Table 45-1 for MMD device addresses). If the PMA sublayer that is closest to the PMD is packaged with the PMD, it shares MMD 1 with the PMD. If the PMD service interface is physically instantiated as nPPI (see Annex 86A), the PMA sublayer that is closest to the PMD will be addressed as MMD 8. More addressable instances of PMA sublayers, each one separated from lower addressable instances by chip-to-chip interfaces, may be implemented and addressed allocating MMD addresses to PMAs in increasing numerical order going from the PMD toward the PCS. The example shown in Figure 83-2 could be implemented with four addressable instances: MMD 8 addressing the lowest PMA sublayer (note that this cannot share MMD 1 with the PMD as they are not packaged together in this example), MMD 9 addressing the PMA sublayer above the XLAUI/CAUI below the FEC, MMD 10 addressing the PMA sublayer below the XLAUI/CAUI above the FEC, and MMD 11 addressing the PMA sublayer closest to the PCS.

The number of input lanes and the number of output lanes for a PMA are always divisors of the number of PCSs. For PMA sublayers supporting 40GBASE-R PMDs, the number of PCSs is 4, and for PMA sublayers supporting 100GBASE-R PMDs, the number of PCSs is 20.



CAUI = 100 Gb/s ATTACHMENT UNIT INTERFACE
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 MMD = MDIO MANAGEABLE DEVICE

PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
 XLMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83-2—Example 40GBASE-R and 100GBASE-R PMA layering

The following guidelines apply to the partitioning of PMAs:

- a) The inter-sublayer service interface, defined in 80.3, is used for the PMA, FEC, and PMD service interfaces supporting a flexible architecture with optional FEC and multiple PMA sublayers.
 - 1) An instance of this interface can only connect service interfaces with the same number of lanes, where the lanes operate at the same rate.
- b) XLAUI and CAUI are physical instantiations of the connection between two adjacent PMA sublayers.
 - 1) As a physical instantiation, it defines electrical and timing specification as well as requiring a receive re-timing function
 - 2) XLAUI is a 10.3125 GBd by 4 lane physical instantiation of the respective 40 Gb/s connection
 - 3) CAUI is a 10.3125 GBd by 10 lane physical instantiation of the respective 100 Gb/s connection
- c) The abstract inter-sublayer service interface can be physically instantiated as a XLAUI or CAUI, using associated PMAs to map to the appropriate number of lanes.
- d) Opportunities for optional test-pattern generation, optional test-pattern detection, optional local loopback and optional remote loopback are dependent upon the location of the PMA sublayer in the implementation. See Figure 83-5.
- e) A minimum of one PMA sublayer is required in a PHY.
- f) A maximum of four PMA sublayers are addressable as MDIO MMDs.

83.2 PMA interfaces

All PMA variants for 40GBASE-R and 100GBASE-R signals are based on a generic specification of a bit mux function that applies to all input/output lane counts and each direction of transmission. Each direction of transmission employs one or more such bit muxes to adapt from the appropriate number of input lanes to the appropriate number of output lanes as illustrated in Figure 83–3.




Figure 83–3—PMA bit mux used in both Tx and Rx directions

Conceptually, the PMA bit mux operates in one direction of transmission by demultiplexing PCSs from m PMA input lanes and remultiplexing them into n PMA output lanes. The mapping of PCSs from input to output lanes is not specified. See 83.5.2 and Figure 83–4 for details.

Figure 83–5 provides the functional block diagram of a PMA. The parameters of a PMA include the following:

- The aggregate rate supported (40GBASE-R or 100GBASE-R).
- The numbers of input and output lanes in each direction.
- Whether the PMA is adjacent to a physically instantiated interface (XLAUI/CAUI above or below, or PMD service interface below).
- Whether the PMA is adjacent to the PCS (or adjacent to FEC when FEC is adjacent to the PCS).
- Whether the PMA is adjacent to the PMD.

83.3 PMA service interface

The PMA service interface for 40GBASE-R and 100GBASE-R is an instance of the inter-sublayer service interface defined in 80.3. The PMA service interface primitives are summarized as follows:

```
PMA:IS_UNITDATA_i.request(tx_bit)
PMA:IS_UNITDATA_i.indication(rx_bit)
PMA:IS_SIGNAL.indication(SIGNAL_OK)
```

For a PMA with p lanes at the PMA service interface, the primitives are defined for $i = 0$ to $p - 1$.

If the PMA client is the PCS or an FEC sublayer, the PMA (or PMA client) continuously sends four (for 40GBASE-R) or twenty (for 100GBASE-R) parallel bit streams to the PMA client (or PMA), each at the nominal signaling rate of the PCSL.

If the PMA client is another PMA, for a PMA supporting a 40GBASE-R PMD, the number of PCSs $z = 4$ and for a PMA supporting a 100GBASE-R PMD, the number of PCSs $z = 20$. A PMA with p input lanes receives bits on each of its input lanes at z/p times the PCSL rate. Skew may exist between the bits received

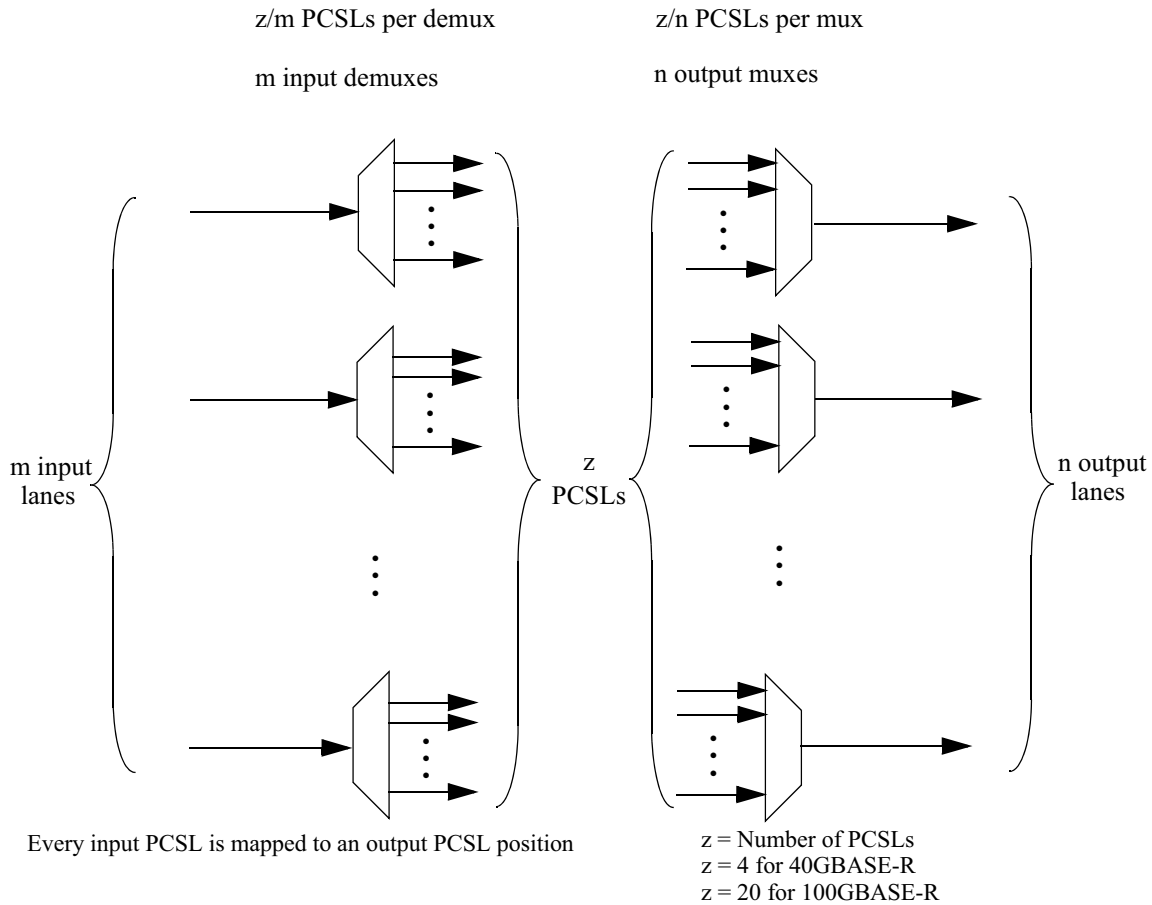


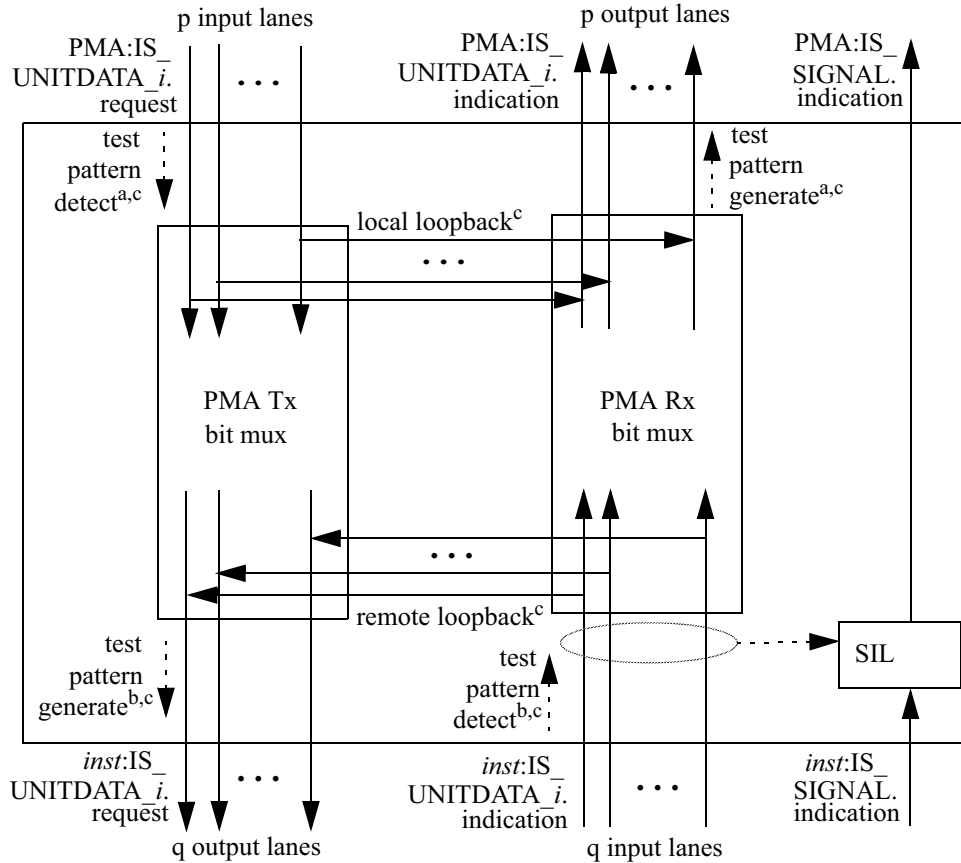
Figure 83-4—PMA bit mux operation used in both Tx and Rx directions

on each lane even though all lanes originate from the same synchronous source, so there is independence of arrival of bits on each lane.

In the Tx direction, if the bit from a PMA:IS_UNITDATA_ i .request primitive is received over a physically instantiated interface (XLAUI/CAUI), clock and data are recovered on the lane receiving the bit. The bit is routed through the PMA to an output lane through a process that may demultiplex PCSLs from the input, perform any necessary buffering to tolerate Skew Variation across input lanes, and multiplex PCSLs to output lanes. The bit is sent on an output lane to the sublayer below using the *inst*:IS_UNITDATA_ k .request (k not necessarily equal to i) primitive (see 83.4).

In the Rx direction, when data is being received from every input lane from the sublayer below the PMA that has a PCSL that is routed to a particular output lane at the PMA service interface, and (if necessary), buffers are filled to allow tolerating the Skew Variation that may appear between the input lanes, PCSLs are demultiplexed from the input lanes, remultiplexed to the output lanes, and bits are transferred over each output lane to the PMA client via the PMA:IS_UNITDATA_ i .indication primitive.

The PMA:IS_SIGNAL.indication primitive is generated through a set of Signal Indication Logic (SIL) that reports signal health based on receipt of the *inst*:IS_SIGNAL.indication from the sublayer below, data being received on all of the input lanes from the sublayer below, buffers filled (if necessary) to accommodate Skew Variation, and bits being sent to the PMA client on all of the output lanes. When these conditions are met, the SIGNAL_OK parameter sent to the PMA client via the PMA:IS_SIGNAL.indication primitive will have the value OK. Otherwise, the SIGNAL_OK primitive will have the value FAIL.



inst PMD, PMA, or FEC, depending on which sublayer is below this PMA
 SIL Signal Indication Logic

- ^a If physically instantiated interface (XLAUI/CAUI) immediately above this PMA.
- ^b If physically instantiated interface (XLAUI/CAUI or PMD service interface) immediately below this PMA, or if this is the closest PMA to the PMD.
- ^c Optional.

Figure 83-5—PMA Functional Block Diagram

83.4 Service interface below PMA

Since the architecture supports multiple PMA sublayers for various PMD lane counts and device partitioning, there are several different sublayers that may appear below a PMA, including FEC, the PMD, or another PMA. The variable *inst* represents whichever sublayer appears below the PMA (e.g., another PMA, FEC, or PMD).

The sublayer below the PMA utilizes the inter-sublayer service interface defined in 80.3. The service interface primitives provided to the PMA are summarized as follows:

```

inst:IS_UNITDATA_i.request(tx_bit)
inst:IS_UNITDATA_i.indication(rx_bit)
inst:IS_SIGNAL.indication(SIGNAL_OK)
    
```

The number of lanes *q* for the service interface matches the number of lanes expected by the PMA. The *inst:IS_UNITDATA_i* primitives are defined for each lane *i* = 0 to *q* – 1 of the service interface below the

PMA. Note that electrical and timing specifications of the service interface are defined if the interface is physically instantiated (e.g., XLAUI/CAUI or nPPI), otherwise the service interface is specified only abstractly. The interface between the PMA and the sublayer below consists of q lanes for data transfer and a status indicating a good signal sent by the sublayer below the PMA (see Figure 83–5).

In the Tx direction, when data is being received via the PMA:IS_UNITDATA_ i .request primitive from every input lane from the PMA client at the PMA service interface (see 83.3) that has a PCSL that is routed to this output lane, and (if necessary), buffers are filled to allow tolerating the Skew Variation that may appear between the input lanes from the PMA client, PCSLs are demultiplexed from the input lanes, remultiplexed to the output lanes, and bits are transferred over each output lane to the sublayer below the PMA.

In the Rx direction, if the bit is received over a physically instantiated interface (XLAUI/CAUI or nPPI), clock and data are recovered on the lane receiving the bit. The bit is routed through the PMA to an output lane toward the PMA client through a process that may demultiplex PCSLs from the input, perform any necessary buffering to tolerate Skew Variation across input lanes, and multiplex PCSLs to output lanes, and finally sending the bit on an output lane to the PMA client using the PMA:IS_UNITDATA_ k .indication (k not necessarily equal to i) primitive at the PMA service interface.

83.5 Functions within the PMA

The purpose of the PMA is to adapt the PCSL formatted signal to an appropriate number of abstract or physical lanes, to recover clock from the received signal (if appropriate), and optionally to provide test signals and loopback. Each input (Tx direction) or output (Rx direction) lane between the PMA and the PMA client carries one or more PCSLs that are bit-multiplexed. All input and output lanes between the PMA and the PMA client carry the same number of PCSLs and operate at the same nominal signaling rate. Likewise, each input (Rx direction) or output (Tx direction) lane between the PMA and the sublayer below the PMA carries one or more PCSLs that are bit-multiplexed. All input and output lanes between the PMA and the sublayer below the PMA carry the same number of PCSLs and operate at the same nominal signaling rate. As described in 83.1.4, the number of input lanes and the number of output lanes for a given PMA are divisors of the number of PCSLs for the interface type supported.

83.5.1 Per input-lane clock and data recovery

If the interface between the PMA client and the PMA is physically instantiated as XLAUI/CAUI, the PMA shall meet the electrical and timing specifications in Annex 83A or Annex 83B as appropriate. If the interface between the sublayer below the PMA and the PMA is physically instantiated as XLAUI/CAUI or nPPI, the PMA shall meet the electrical and timing specifications at the service interface as specified in Annex 83A, Annex 83B, or Annex 86A as appropriate.

83.5.2 Bit-level multiplexing

The PMA provides bit-level multiplexing in both the Tx and Rx directions. In the Tx direction, the function is performed among the bits received from the PMA client via the PMA:IS_UNITDATA_ i .request primitives (for PMA client lanes $i = 0$ to $p - 1$) with the result sent to the service interface below the PMA using the *inst*:IS_UNITDATA_ i .request primitives (for service interface lanes $i = 0$ to $q - 1$), referencing the functional block diagram shown in Figure 83–5. The bit multiplexing behavior is illustrated in Figure 83–4.

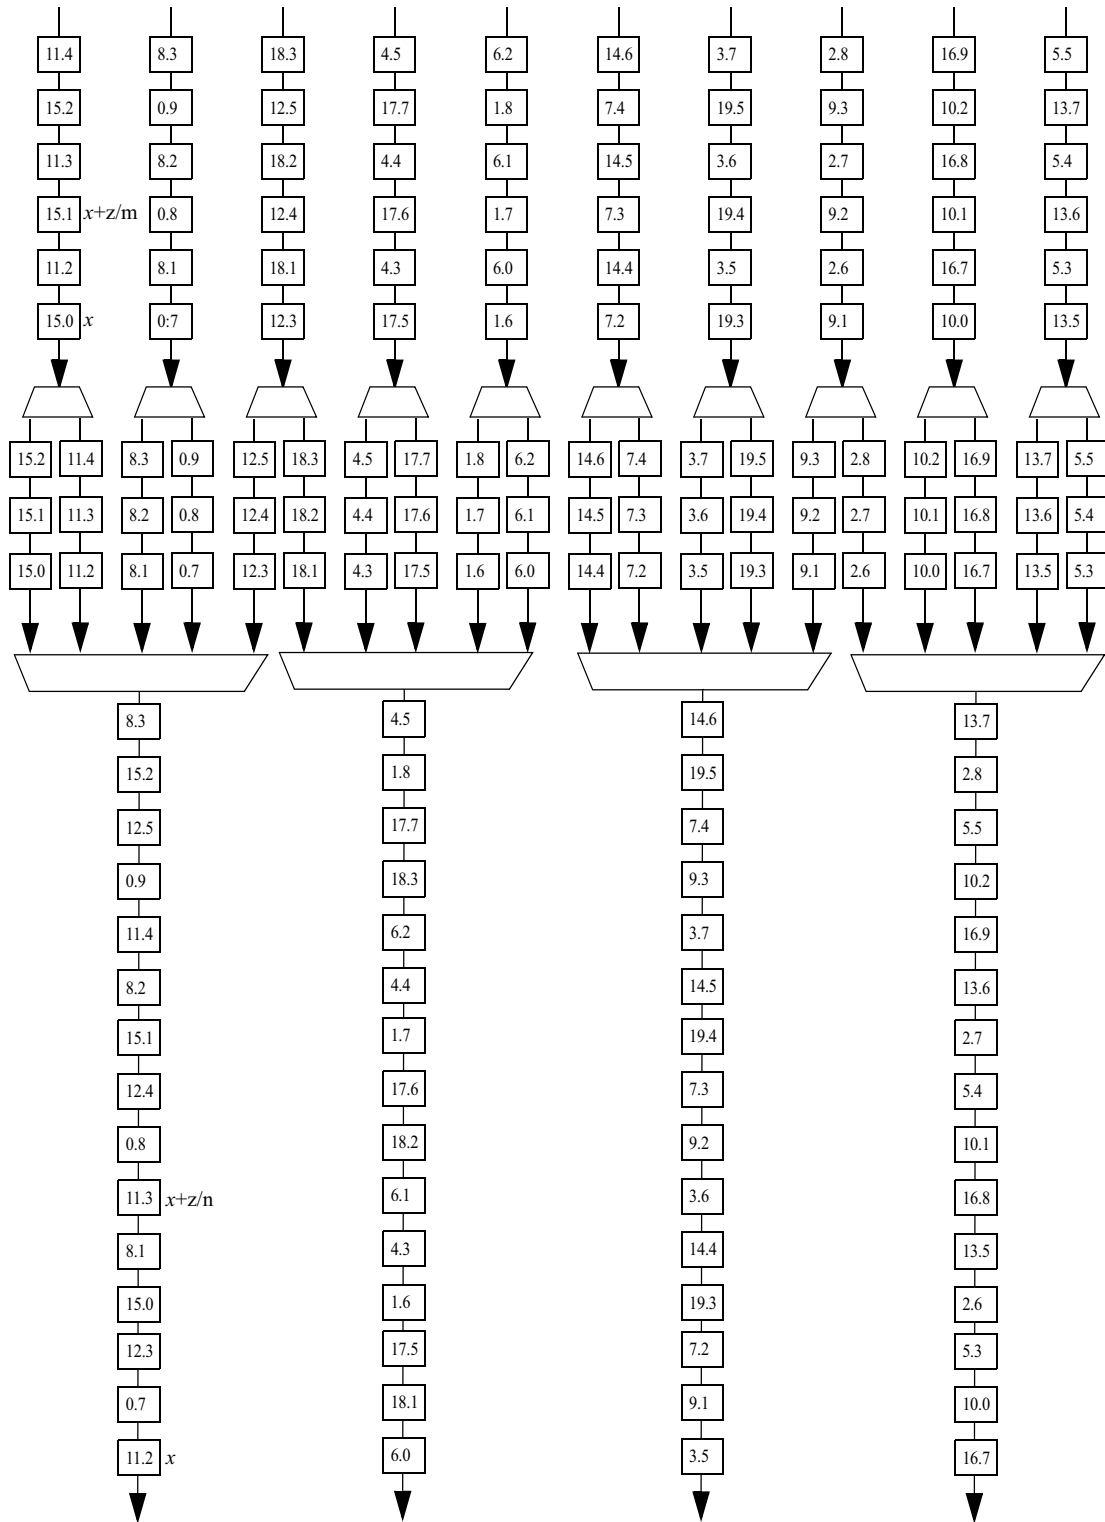
The aggregate signal carried by the group of input lanes or the group of output lanes is arranged as a set of PCSLs. For PMA sublayers supporting 40GBASE-R interfaces, the number of PCSLs z is 4, and the nominal signaling rate R of each PCSL is 10.3125 GBd. For PMA sublayers supporting 100GBASE-R interfaces, the number of PCSLs z is 20, and the nominal signaling rate R of each PCSL is 5.15625 GBd.

For a PMA with m input lanes (Tx or Rx direction), each input lane carries, bit multiplexed, z/m PCSLs. Each input lane has a nominal signaling rate of $R \times z/m$. If bit x received on an input lane belongs to a particular PCSL, the next bit of that same PCSL is received on the same input lane at bit position $x+(z/m)$. The z/m PCSLs may arrive in any sequence on a given input lane.

For a PMA with n output lanes (Tx or Rx direction), each output lane carries, bit multiplexed, z/n PCSLs. Each output lane has a nominal signaling rate of $R \times z/n$. Each PCSL is mapped from a position in the sequence on one of the z/m input lanes to a position in the sequence on one of the z/n output lanes. If bit x sent on an output lane belongs to a particular PCSL, the next bit of that same PCSL is sent on the same output lane at bit position $x + (z/n)$. The PMA shall maintain the chosen sequence of PCSLs on all output lanes while it is receiving a valid stream of bits on all input lanes.

Each PCSL received in any temporal position on an input lane is transferred into a temporal position on an output lane. As the PCS (see Clause 82) has fully flexible receive logic, an implementation is free to perform the mapping of PCSLs from input lanes to output lanes without constraint. Figure 83–6 illustrates one possible bit ordering for a 10:4 PMA bit mux. Other bit orderings are also valid.

10 Lane PMA Input



NOTE—i.k indicates bit k on PCSL i. Skew may exist between PCSLs.

Figure 83–6—Example 10: 4 PMA bit mux

83.5.3 Skew and Skew Variation

The Skew (relative delay) between the PCSs must be kept within limits so that the information on the lanes can be reassembled by the PCS.

Any PMA that combines PCSs from different input lanes onto the same output lane must tolerate Skew Variation between the input lanes without changing the PCSL positions on the output. Skew and Skew Variation are defined in 80.5. The limits for Skew and Skew Variation at physically instantiated interfaces are specified at Skew points SP1 and SP2 in the transmit direction and SP5 and SP6 in the receive direction as defined in 80.5 and illustrated in Figure 80–4 and Figure 80–5.

83.5.3.1 Skew generation toward SP1

In an implementation with one or more physically instantiated XLAUI/CAUI interfaces, the PMA that sends data in the transmit direction toward the XLAUI/CAUI that is closest to the PMD (SP1 in Figure 80–4 and Figure 80–5) shall produce no more than 29 ns of Skew between PCSs toward the XLAUI/CAUI, and no more than 200 ps of Skew Variation.

83.5.3.2 Skew tolerance at SP1

In an implementation with one or more physically instantiated XLAUI/CAUI interfaces, the PMA service interface that receives data in the transmit direction from the XLAUI/CAUI (SP1 in Figure 80–4 and Figure 80–5) shall tolerate the maximum amount of Skew Variation allowed at SP1 (200 ps) between input lanes while maintaining the bit ordering and position of each PCSL on each PMA lane in the transmit direction (toward the PMD).

83.5.3.3 Skew generation toward SP2

In an implementation with a physically instantiated PMD service interface, the PMA adjacent to the PMD service interface shall generate no more than 43 ns of Skew, and no more than 400 ps of Skew Variation between output lanes toward the PMD service interface (SP2 in Figure 80–4 and Figure 80–5). If there is a physically instantiated XLAUI/CAUI as well, then the Skew measured at SP1 is limited to no more than 29 ns of Skew and no more than 200 ps of Skew Variation.

83.5.3.4 Skew tolerance at SP5

In an implementation with a physically instantiated PMD service interface, the PMA adjacent to the PMD service interface (SP5) shall tolerate the maximum amount of Skew Variation allowed at SP5 (3.6 ns) between output lanes from the PMD service interface while maintaining the bit ordering and position of each PCSL on each PMA lane in the receive direction (toward the PCS).

83.5.3.5 Skew generation at SP6

In an implementation with one or more physically instantiated XLAUI/CAUI interfaces, at SP6 (the receive direction of the XLAUI/CAUI closest to the PCS), the PMA or group of PMAs between the PMD and the XLAUI/CAUI closest to the PCS shall deliver no more than 160 ns of Skew, and no more than 3.8 ns of Skew Variation between output lanes toward the XLAUI/CAUI in the Rx direction. If there is a physically instantiated PMD service interface as well, the Skew measured at SP5 is limited to no more than 145 ns of Skew and no more than 3.6 ns of Skew Variation. If there is no physically instantiated PMD service interface, the Skew measured at SP4 is limited to no more than 134 ns of Skew, and no more than 3.4 ns of Skew Variation.

83.5.3.6 Skew tolerance at SP6

In an implementation with one or more physically instantiated XLAUI/CAUI interfaces, the PMA between the XLAUI/CAUI closest to the PCS and the PCS shall tolerate the maximum amount of Skew Variation allowed at SP6 (3.8 ns) between input lanes while maintaining the bit order and position of PCSs on lanes sent in the receive direction towards the PCS.

83.5.4 Delay constraints

The maximum cumulative delay contributed by up to four PMA stages in a PHY (sum of transmit and receive delays at one end of the link) shall meet the values specified in Table 83–1. A description of overall system delay constraints and the definitions for bit-times and pause_quanta can be found in 80.4 and its references.

Table 83–1—Delay constraints

Sublayer	Maximum (bit time)	Maximum (pause_quanta)	Maximum (ns)
40GBASE-R PMA	4096	8	102.4
100GBASE-R PMA	9216	18	92.16

83.5.5 Clocking architecture

A PMA with m input lanes and n output lanes shall clock the output lanes at m/n times the rate of the input lanes. This applies in both the Tx and Rx directions of transmission. In the case where the interfaces between the PMA client and the PMA and/or the PMA and the sublayer below the PMA are physically instantiated, the PMA may derive its input clock(s) from the electrical interface on one or more of the input lanes, and generate the output clock(s) with an appropriate PLL multiplier/divider circuit.

There is no requirement that the PMA clock all output lanes in unison. Examples of independent clocking of output lanes include the following:

- The case where the number of input and output lanes are equal (the PMA is provided for retiming and regeneration of the signal). This may be implemented without any rearrangement of PCSs between input lanes and output lanes (although rearrangements are allowed), and such a PMA may be implemented by driving each output lane using the clock recovered from the corresponding input lane.
- If the number of input and output lanes have a common factor, the PMA may be partitioned such that PCSs from a subset of the input lanes are mapped only to a subset of the output lanes (for example, a 10:4 PMA could be implemented as two 5:2 PMAs). The output clock for one subset may be independent of the output clock for other subset(s).

83.5.6 Signal drivers

For cases where the interface between the PMA client and the PMA, or between the PMA and the sublayer below the PMA represent a physically instantiated interface, the PMA provides electrical signal drivers for that interface. The electrical and jitter/timing specifications for these interfaces appear in

- Annex 83A, which specifies the XLAUI/CAUI interface for chip-to-chip applications.
- Annex 83B, which specifies the XLAUI/CAUI interface for chip-to-module applications.

- 86.2, which specifies the PMD service interface for 40GBASE-SR4 and 100GBASE-SR10 PMDs.
- 87.2, which specifies the PMD service interface for 40GBASE-LR4 PMDs.
- Annex 86A, which specifies the Parallel Physical Interface (XLPPi and CPPI), an optional physical instantiation of the PMD service interface for 40GBASE-SR4, 40GBASE-LR4, and 100GBASE-SR10 PMDs.

83.5.7 Link status

The PMA shall provide link status information to the PMA client using the PMA:IS_SIGNAL.indication primitive. The PMA continuously monitors the link status reported by the service interface below from the *inst*:IS_SIGNAL.indication primitive, and uses this as input to Signal Indication Logic (SIL) to determine the link status to report to the layer above. Other inputs to the SIL may include the status of clock and data recovery on the lanes from the service interface below the PMA and whether buffers/FIFOs have reached the required fill level to accommodate Skew Variation so that data is being sent on the output lanes.

83.5.8 PMA local loopback mode

PMA local loopback shall be provided by the PMA adjacent to the PMD for 40GBASE-KR4, 40GBASE-CR4, and 100BASE-CR10 PMDs. PMA local loopback mode is optional for other PMDs or for PMAs not adjacent to the PMD. If it is implemented, it shall be as described in this subclause (83.5.8).

The PMA sublayer may provide a local loopback function. The function involves looping back each input lane to the corresponding output lane. Each bit received from the PMA:IS_UNITDATA_*i*.request(tx_bit) primitive is looped back in the direction of the PCS using the PMA:IS_UNITDATA_*i*.indication(rx_bit) primitive.

During local loopback, the PMA performs normal bit muxing of PCSs per 83.5.2 onto the lanes in the Tx direction toward the service interface below the PMA.

Ability to perform this function is indicated by the Local_loopback_ability status variable. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.8.0 (45.2.1.7.15). A device is placed in local loopback mode when the Local_loopback_enable control variable is set to one, and removed from local loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD control 1 register (bit 1.0.0, see 45.2.1.1.4).

83.5.9 PMA remote loopback mode (optional)

PMA remote loopback mode is optional. If implemented, it shall be as described in this subclause (83.5.9).

Remote loopback, if provided, should be implemented in a PMA sublayer close enough to the PMD to maintain the bit sequence on each individual PMD lane. When remote loopback is enabled, each bit received over a lane of the service interface below the PMA via *inst*:IS_UNITDATA_*i*.indication is looped back to the corresponding output lane toward the PMD via *inst*:IS_UNITDATA_*i*.request. Note that the service interface below the PMA can be provided by the FEC, PMD, or another PMA sublayer.

During remote loopback, the PMA performs normal bit muxing of PCSs per 83.5.2 onto the lanes in the Rx direction towards the PMA client.

The ability to perform this function is indicated by the Remote_loopback_ability status variable. If a Clause 45 MDIO is implemented, this variable is accessible through bit 1.13.15 (45.2.1.11a.1). A device is placed in remote loopback mode when the Remote_loopback_enable control variable is set to one, and removed from remote loopback mode when this variable is set to zero. If a Clause 45 MDIO is implemented, this variable is accessible through PMA/PMD Control register 1 (bit 1.0.1, see 45.2.1.1.3a).

83.5.10 PMA test patterns (optional)

Where the output lanes of the PMA appear on a physically instantiated interface XLAUI/CAUI or the PMD service interface (whether or not it is physically instantiated), the PMA may optionally generate and detect test patterns. These test patterns are used to test adjacent layer interfaces for an individual PMA sublayer or to perform testing between a physically instantiated interface of a PMA sublayer and external testing equipment.

The ability to generate each of the respective test patterns in each direction of transmission are indicated by the `PRBS9_Tx_generator_ability`, `PRBS9_Rx_generator_ability`, `PRBS31_Tx_generator_ability`, and `PRBS31_Rx_generator_ability` status variables, which if a Clause 45 MDIO is implemented are accessible through bits 1.1500.5, 1.1500.4, 1.1500.3, and 1.1500.1, respectively (see 45.2.1.95).

The ability to check PRBS31 test patterns in each direction of transmission are indicated by the `PRBS31_Tx_checker_ability` and `PRBS31_Rx_checker_ability` status variables, which, if a Clause 45 MDIO is implemented, are accessible through bits 1.1500.2 and 1.1500.0, respectively (see 45.2.1.95).

If supported, when send Tx PRBS31 test pattern is enabled by the `PRBS31_enable` and `PRBS_Tx_gen_enable` control variables, the PMA shall generate a PRBS31 pattern (as defined in 49.2.8) on each of the lanes toward the service interface below the PMA via the `inst:IS_UNITDATA_i.request` primitive. To avoid correlated crosstalk, it is highly recommended that the PRBS31 patterns generated on each lane be generated from independent, random seeds or at a minimum offset of 20 000 UI between the PRBS31 sequence on any lane and any other lane. When send Tx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2. If a Clause 45 MDIO is implemented, the `PRBS31_enable` and `PRBS_Tx_gen_enable` control variables are accessible through bits 1.1501.7 and 1.1501.3 (see 45.2.1.96).

If supported, when send Rx PRBS31 test pattern is enabled by the `PRBS31_enable` and `PRBS_Rx_gen_enable` control variables, the PMA shall generate a PRBS31 pattern on each of the lanes toward the PMA client via the `PMA:IS_UNITDATA_i.indication` primitive. While this test pattern is enabled, the PMA also generates `PMA:IS_SIGNAL.indication(SIGNAL_OK)` toward the PMA client independent of the link status at the service interface below the PMA. When send Rx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2. If a Clause 45 MDIO is implemented, the `PRBS31_enable` and `PRBS_Rx_gen_enable` control variables are accessible through bits 1.1501.7 and 1.1501.1 (see 45.2.1.96).

If supported, when check Tx PRBS31 test pattern mode is enabled by the `PRBS31_enable` and `PRBS_Tx_check_enable` control variables, the PMA shall check for the PRBS31 pattern on each of the lanes received from the PMA client via the `PMA:IS_UNITDATA_i.request` primitive. The checker shall increment the test-pattern error counter by one for each incoming bit error in the PRBS31 pattern (see 49.2.8) for isolated single bit errors. Implementations should be capable of counting at least one error whenever one or more errors occur in a sliding 1000-bit window. If a Clause 45 MDIO is implemented, the `PRBS31_enable` and `PRBS_Tx_check_enable` control variables are accessible through bits 1.1501.7 and 1.1501.2 (see 45.2.1.96). The Tx test-pattern error counters `Ln0_PRBS_Tx_test_err_counter` through `Ln9_PRBS_Tx_test_err_counter` count, per lane, errors in detecting the PRBS31 pattern on the lanes from the PMA client. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1600 through 1.1609 (see 45.2.1.98). When check Tx PRBS31 test pattern is disabled, the PMA expects normal traffic and test-pattern error counting does not continue. While in check Tx PRBS31 test-pattern mode, bit multiplexing continues as described in 83.5.2. Note that bit multiplexing of per-lane PRBS31 may produce a signal which is not meaningful for downstream sublayers.

If supported, when check Rx PRBS31 test-pattern mode is enabled by the `PRBS31_enable` and `PRBS_Rx_check_enable` control variables, the PMA shall check for the PRBS31 pattern on each of the lanes received from the service interface below the PMA via the `inst:IS_UNITDATA_i.indication` primitive.

If a Clause 45 MDIO is implemented, the PRBS31_enable and PRBS_Rx_check_enable control variables are accessible through bits 1.1501.7 and 1.1501.0 (see 45.2.1.96). The Rx test-pattern error counters Ln0_PRBS_Rx_test_err_counter through Ln9_PRBS_Rx_test_err_counter count, per lane, errors in detecting the PRBS31 pattern on the lanes from the service interface below the PMA. If a Clause 45 MDIO is implemented, these counters are accessible through registers 1.1700 through 1.1709 (see 45.2.1.99). While in check Rx PRBS31 test-pattern mode, the PMA:IS_SIGNAL.indication primitive does not indicate a valid signal. When check Rx PRBS31 test pattern is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2.

If supported, when send Tx PRBS9 test-pattern mode (see 68.6.1) is enabled by the PRBS9_enable and PRBS_Tx_gen_enable control variables, the PMA shall generate a PRBS9 pattern on each lane toward the service interface below the PMA via the *inst:IS_UNITDATA_i.request* primitive. If a Clause 45 MDIO is implemented, the PRBS9_enable and PRBS_Tx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.3 (see 45.2.1.96). When send Tx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2.

If supported, when send Rx PRBS9 test-pattern mode (see 68.6.1) is enabled by the PRBS9_enable and PRBS_Rx_gen_enable control variables, the PMA shall generate a PRBS9 pattern on each lane toward the PMA client via the PMA:IS_UNITDATA_i.indication primitive. The PMA will also generate PMA:IS_SIGNAL.indication(SIGNAL_OK) toward the PMA client independent of the link status at the service interface below the PMA. If a Clause 45 MDIO is implemented, the PRBS9_enable and PRBS_Rx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.1 (see 45.2.1.96). When send Rx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation performing bit multiplexing as described in 83.5.2.

Note that PRBS9 is intended to be checked by external test gear, and no PRBS9 checking function is provided within the PMA.

Transmit square wave test-pattern mode optionally applies to each lane of the Tx direction PMA towards a physically instantiated XLAUI/CAUI or towards the PMD service interface whether or not it is physically instantiated. The ability to perform this function is indicated by the Square_wave_ability status variable. If a Clause 45 MDIO is implemented, the Square_wave_ability status variable is accessible through the Square wave test ability bit 1.1500.12 (see 45.2.1.95). If implemented, the transmit square wave test-pattern mode is enabled by control variables Square_wave_enable_0 through Square_wave_enable_9. If a Clause 45 MDIO is implemented, these control variables are accessible through the square wave testing control and status register bits 1.1510.0 through 1.1510.9 (limited to the number of lanes of the service interface below the PMA, see 45.2.1.97). When enabled, the PMA shall generate a square wave test pattern (8 ones followed by 8 zeros) on the square wave enabled lanes toward the service interface below the PMA via the *inst:IS_UNITDATA_i.request* primitive. Lanes for which square wave is not enabled will transmit normal data resulting from the bit multiplexing operations described in 83.5.2 or test patterns as determined by other registers. When transmit square wave test pattern is disabled for all lanes, the PMA will perform normal operation performing bit multiplexing as described in 83.5.2 or transmit test patterns as determined by other registers.

83.6 PMA MDIO function mapping

The optional MDIO capability described in Clause 45 describes several variables that provide control and status information for and about the PMA. Since a given implementation may employ more than one PMA sublayer, the PMA control and status information is organized into multiple addressable instances, one for each possible PMA sublayer. See 45.2.1 and 83.1.4 for the allocation of MMD addresses to PMA sublayers. Control and status registers for MMD 8, 9, 10, and 11 will use the Extended PMA control and status registers at identical locations to those for MMD 1.

Mapping of MDIO control variables to PMA control variables is shown in Table 83–2. Mapping of MDIO status variables to PMA status variables is shown in Table 83–3. Mapping of MDIO counter to PMA counters is shown in Table 83–4. These tables provide the register and bit numbers for the PMA addressed as MMD 1. For implementations with multiple PMA sublayers, additional PMA sublayers use the corresponding register and bit numbers in MMDs 8, 9, 10, and 11 as necessary.

Table 83–2—MDIO/PMA control variable mapping

MDIO variable	PMA/PMD register name	Register/ bit number	PMA control variable
PMA remote loopback	PMA/PMD control 1	1.0.1	Remote_loopback_enable
PMA local loopback	PMA/PMD control 1	1.0.0	Local_loopback_enable
PRBS31 pattern enable	PRBS pattern testing control	1.1501.7	PRBS31_enable
PRBS9 pattern enable	PRBS pattern testing control	1.1501.6	PRBS9_enable
Tx generator enable	PRBS pattern testing control	1.1501.3	PRBS_Tx_gen_enable
Tx checker enable	PRBS pattern testing control	1.1501.2	PRBS_Tx_check_enable
Rx generator enable	PRBS pattern testing control	1.1501.1	PRBS_Rx_gen_enable
Rx checker enable	PRBS pattern testing control	1.1501.0	PRBS_Rx_check_enable
Lane 0 SW enable	Square wave testing control	1.1510.0	Square_wave_enable_0
Lane 1 SW enable	Square wave testing control	1.1510.1	Square_wave_enable_1
Lane 2 SW enable	Square wave testing control	1.1510.2	Square_wave_enable_2
Lane 3 SW enable	Square wave testing control	1.1510.3	Square_wave_enable_3
Lane 4 SW enable	Square wave testing control	1.1510.4	Square_wave_enable_4
Lane 5 SW enable	Square wave testing control	1.1510.5	Square_wave_enable_5
Lane 6 SW enable	Square wave testing control	1.1510.6	Square_wave_enable_6
Lane 7 SW enable	Square wave testing control	1.1510.7	Square_wave_enable_7
Lane 8 SW enable	Square wave testing control	1.1510.8	Square_wave_enable_8
Lane 9 SW enable	Square wave testing control	1.1510.9	Square_wave_enable_9

Table 83–3—MDIO/PMA status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMA status variable
PMA remote loopback ability	40G/100G PMA/PMD extended ability register	1.13.15	Remote_loopback_ability
PMA local loopback ability	PMA/PMD status 2 register	1.8.0	Local_loopback_ability
PRBS9 Tx generator ability	Test pattern ability register	1.1500.5	PRBS9_Tx_generator_ability
PRBS9 Rx generator ability	Test pattern ability register	1.1500.4	PRBS9_Rx_generator_ability
PRBS31Tx generator ability	Test pattern ability register	1.1500.3	PRBS31_Tx_generator_ability
PRBS31Tx checker ability	Test pattern ability register	1.1500.2	PRBS31_Tx_checker_ability
PRBS31 Rx generator ability	Test pattern ability register	1.1500.1	PRBS31_Rx_generator_ability
PRBS31 Rx checker ability	Test pattern ability register	1.1500.0	PRBS31_Rx_checker_ability
Square wave test ability	Test pattern ability register	1.1500.12	Square_wave_ability

Table 83–4—MDIO/PMA counters mapping

MDIO variable	PMA/PMD register name	Register/bit number	PMA status variable
Error counter Tx, lane 0	PRBS Tx pattern testing error counter, lane 0	1.1600	Ln0_PRBS_Tx_test_err_counter
Error counter Tx, lane 1	PRBS Tx pattern testing error counter, lane 1	1.1601	Ln1_PRBS_Tx_test_err_counter
Error counter Tx, lane 2	PRBS Tx pattern testing error counter, lane 2	1.1602	Ln2_PRBS_Tx_test_err_counter
Error counter Tx, lane 3	PRBS Tx pattern testing error counter, lane 3	1.1603	Ln3_PRBS_Tx_test_err_counter
Error counter Tx, lane 4	PRBS Tx pattern testing error counter, lane 4	1.1604	Ln4_PRBS_Tx_test_err_counter
Error counter Tx, lane 5	PRBS Tx pattern testing error counter, lane 5	1.1605	Ln5_PRBS_Tx_test_err_counter
Error counter Tx, lane 6	PRBS Tx pattern testing error counter, lane 6	1.1606	Ln6_PRBS_Tx_test_err_counter
Error counter Tx, lane 7	PRBS Tx pattern testing error counter, lane 7	1.1607	Ln7_PRBS_Tx_test_err_counter
Error counter Tx, lane 8	PRBS Tx pattern testing error counter, lane 8	1.1608	Ln8_PRBS_Tx_test_err_counter
Error counter Tx, lane 9	PRBS Tx pattern testing error counter, lane 9	1.1609	Ln9_PRBS_Tx_test_err_counter

Table 83–4—MDIO/PMA counters mapping (continued)

MDIO variable	PMA/PMD register name	Register/bit number	PMA status variable
Error counter Rx, lane 0	PRBS Rx pattern testing error counter, lane 0	1.1700	Ln0_PRBS_Rx_test_err_counter
Error counter Rx, lane 1	PRBS Rx pattern testing error counter, lane 1	1.1701	Ln1_PRBS_Rx_test_err_counter
Error counter Rx, lane 2	PRBS Rx pattern testing error counter, lane 2	1.1702	Ln2_PRBS_Rx_test_err_counter
Error counter Rx, lane 3	PRBS Rx pattern testing error counter, lane 3	1.1703	Ln3_PRBS_Rx_test_err_counter
Error counter Rx, lane 4	PRBS Rx pattern testing error counter, lane 4	1.1704	Ln4_PRBS_Rx_test_err_counter
Error counter Rx, lane 5	PRBS Rx pattern testing error counter, lane 5	1.1705	Ln5_PRBS_Rx_test_err_counter
Error counter Rx, lane 6	PRBS Rx pattern testing error counter, lane 6	1.1706	Ln6_PRBS_Rx_test_err_counter
Error counter Rx, lane 7	PRBS Rx pattern testing error counter, lane 7	1.1707	Ln7_PRBS_Rx_test_err_counter
Error counter Rx, lane 8	PRBS Rx pattern testing error counter, lane 8	1.1708	Ln8_PRBS_Rx_test_err_counter
Error counter Rx, lane 9	PRBS Rx pattern testing error counter, lane 9	1.1709	Ln9_PRBS_Rx_test_err_counter

83.7 Protocol implementation conformance statement (PICS) proforma for Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R¹³

83.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the same, can be found in Clause 21.

83.7.2 Identification

83.7.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

83.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010 Clause 83, Physical Medium Attachment (PMA) sublayer, type 40GBASE-R and 100GBASE-R
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	
Date of Statement	

¹³Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

83.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*PMA40	PMA for 40GBASE-R	83.1.1		O.1	Yes [] No []
*PMA100	PMA for 100GBASE-R	83.1.1		O.1	Yes [] No []
LANES_UPSTREAM	Number of lanes in direction of PCS	83.1.4	Divisor of number of PCS lanes	PMA40:M	4 []
				PMA100:M	4 [] 10 [] 20 []
LANES_DOWNSTREAM	Number of lanes in direction of PMD	83.1.4	Divisor of number of PCS lanes	PMA40:M	4 []
				PMA100:M	4 [] 10 [] 20 []
RX_CLOCK	Signaling rate of output lanes in Rx direction	83.5.5	LANES_DOWNSTREAM / LANES_UPSTREAM times signaling rate of input lanes in Rx direction	M	Yes [] No []
TX_CLOCK	Signaling rate of output lanes in Tx direction	83.5.5	LANES_UPSTREAM / LANES_DOWNSTREAM times signaling rate of input lanes in Tx direction	M	Yes [] No []
LANE_MAPPING	Maintain lane mapping while link is in operation	83.5.2	Maintain sequence of PCSLs on all output lanes	M	Yes [] No []
LNKS	PMA link status	83.5.7	Meets the requirements of 83.5.7	M	Yes [] No []
JTP	Supports test-pattern mode	83.5.10		O	Yes [] No [] N/A []
*KRCCR	PMA adjacent to the PMD for 40GBASE-KR4, 40GBASE-CR4, or 100GBASE-CR10	83.5.8		O	Yes [] No []
*LBL	PMA local loopback	83.5.8	Supports local loopback	KRCCR:M !KRCCR:O	Yes [] No [] N/A []
*LBR	PMA remote loopback	83.5.9	Supports remote loopback	O	Yes [] No [] N/A []
MD	MDIO	83.6	Registers and interface supported	O	Yes [] No []

Item	Feature	Subclause	Value/Comment	Status	Support
*USP1SP6	Physically instantiated XLAUI or CAUI above (toward PCS)	83.5.3		O	Yes [] No []
*DSP1SP6	Physically instantiated XLAUI or CAUI below (toward PMD)	83.5.3		O	Yes [] No []
*SP2SP5	Physically instantiated PMD service interface	83.5.3		O	Yes [] No []
*PPI	PMD service interface instantiated as nPPI	83.5.1, 83.5.5		O	Yes [] No []
UNAU1	Electrical and timing requirements of Annex 83A or Annex 83B as appropriate met by upstream XLAUI/CAUI	83.5.1, 83.5.5		USP1SP6: M	Yes [] No []
DNAUI	Electrical and timing requirements of Annex 83A or Annex 83B as appropriate met by downstream XLAUI/CAUI	83.5.1, 83.5.5		DSP1SP6: M	Yes [] No []
PPIET	Electrical and timing requirements of Annex 86A met by PMD service interface	83.5.1, 83.5.5		PPI:M	Yes [] No []
DELAY40	Roundtrip delay limit for 40GBASE-R	83.5.4	No more than 4096 BT or 8 pause quanta	PMA40:M	Yes [] No []
DELAY100	Roundtrip delay limit for 100GBASE-R	83.5.4	No more than 9216 BT or 18 pause quanta	PMA100: M	Yes [] No []

83.7.4 Skew generation and tolerance

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Skew generation toward SP1 in Tx direction	83.5.3.1	≤ 29 ns	DSP1SP6:M	Yes [] No []
S2	Skew variation generation toward SP1 in Tx direction	83.5.3.1	≤ 200 ps	DSP1SP6:M	Yes [] No []
S3	Skew variation tolerance at SP1	83.5.3.2	Minimum 200 ps	USP1SP6:M	Yes [] No []
S4	Skew generation toward SP2 in Tx direction	83.5.3.3	≤ 43 ns	SP2SP5:M	Yes [] No []
S5	Skew variation generation toward SP2 in Tx direction	83.5.3.3	≤ 400 ps	SP2SP5:M	Yes [] No []
S6	Skew variation tolerance at SP5	83.5.3.4	Minimum 3.6 ns	SP2SP5:M	Yes [] No []
S7	Skew generation toward SP6 in Rx direction	83.5.3.5	≤ 160 ns	USP1SP6:M	Yes [] No []
S8	Skew variation generation toward SP6 in Rx direction	83.5.3.5	≤ 3.8 ns	USP1SP6:M	Yes [] No []
S9	Skewvariation tolerance at SP6	83.5.3.6	Minimum 3.8 ns	DSP1SP6:M	Yes [] No []

83.7.5 Test patterns

Item	Feature	Subclause	Value/Comment	Status	Support
*JTP1	Physically Instantiated XLAUI/CAUI between PMA and PMA client	83.5.10		O	Yes [] No []
*JTP2	Physically Instantiated XLAUI/CAUI between PMA and sublayer below the PMA, or adjacent to PMD whether or not PMD service interface is physically instantiated as nPPI	83.5.10		O	Yes [] No []
J1	Send PRBS31 Tx	83.5.10		JTP2:O	Yes [] No []
J2	Send PRBS31 Rx	83.5.10		JTP1:O	Yes [] No []
J3	Check PRBS31 Tx	83.5.10		JTP1:O	Yes [] No []
J4	Check PRBS31 Rx	83.5.10		JTP2:O	Yes [] No []
J5	Send PRBS9 Tx	83.5.10		JTP2:O	Yes [] No []

Item	Feature	Subclause	Value/Comment	Status	Support
J6	Send PRBS9 Rx	83.5.10		JTP1:O	Yes [] No []
J7	Send square wave Tx	83.5.10		JTP2:O	Yes [] No []

83.7.6 Loopback modes

Item	Feature	Subclause	Value/Comment	Status	Support
LB1	PMA local loopback implemented	83.5.8	Meets the requirements of 83.5.8	LBL:M	Yes [] No []
LB2	PMA remote loopback implemented	83.5.9	Meets the requirements of 83.5.9	LBR:M	Yes [] No []

84. Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4

84.1 Overview

This clause specifies the 40GBASE-KR4 PMD. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 84–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 84–1—Physical Layer clauses associated with the 40GBASE-KR4 PMD

Associated clause	40GBASE-KR4
81—RS	Required
81—XLGMII ^a	Optional
82—PCS for 40GBASE-R	Required
74—FEC for BASE-R	Optional
83—PMA for 40GBASE-R	Required
83A—XLAUI	Optional
73—Auto-Negotiation	Required

^aThe XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.

Figure 84–1 shows the relationship of the 40GBASE-KR4 PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model.

84.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-KR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data. The PMD translates the encoded data to and from signals suitable for the medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_ i.request
PMD:IS_UNITDATA_ i.indication
PMD:IS_SIGNAL.indication
```

The 40GBASE-KR4 PMD has four parallel bit streams, hence $i = 0$ to 3.

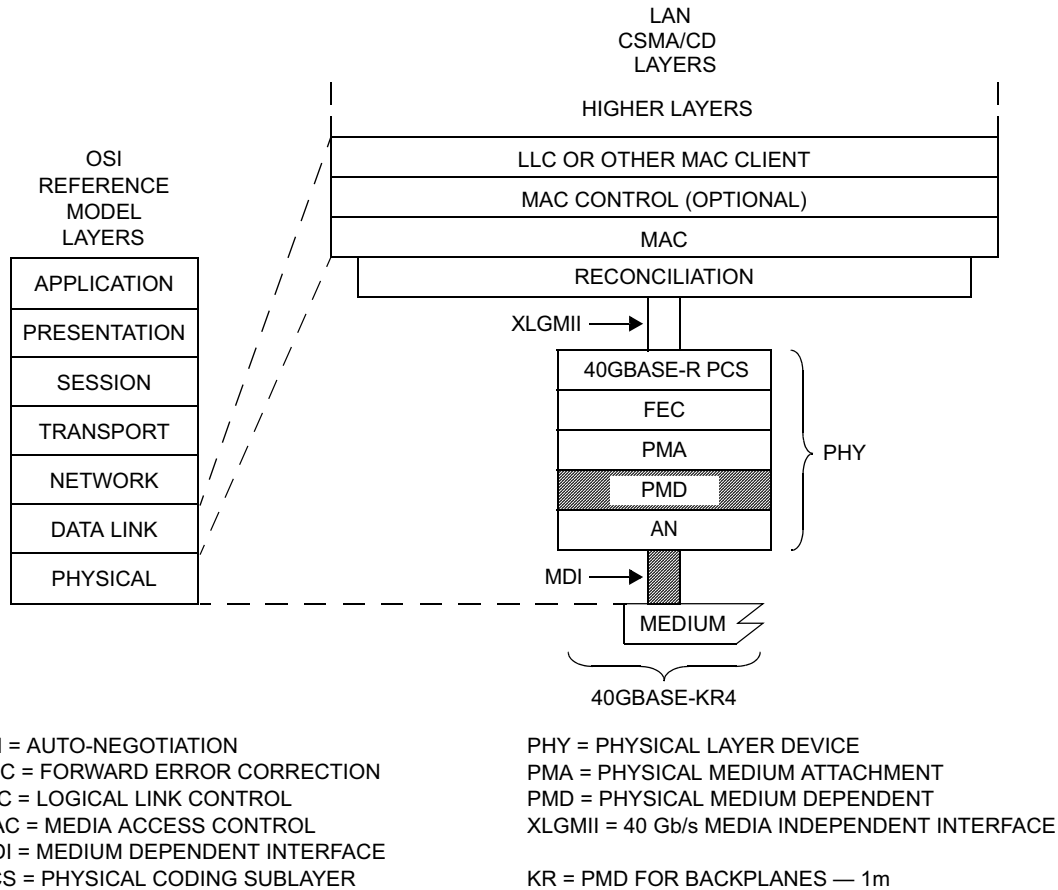


Figure 84–1—40GBASE-KR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 CSMA/CD LAN model

The PMA (or the PMD) continuously sends four parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 10.3125 GBd.

SIGNAL_DETECT in 40GBASE-KR4 indicates the successful completion of the start-up protocol on all four lanes.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the PMD:IS_UNITDATA_*i*.indication parameters are undefined.

The SIGNAL_DETECT parameter maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication primitive.

84.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

84.4 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-KR4 PMD, AN, and the medium in one direction shall be no more than 2048 bit times (4 pause_quanta or 51.2 ns). It is assumed that the one way delay through the medium is no more than 320 bit times (8 ns).

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

84.5 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–4 and Figure 80–5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5.

84.6 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control variables to PMD control variables as shown in Table 84–2, and MDIO status variables to PMD status variables as shown in Table 84–3.

84.7 PMD functional specifications

84.7.1 Link block diagram

The 40GBASE-KR4 PMD uses the same Link block diagram as 10GBASE-KX4, as defined in 71.6.1.

84.7.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate electrical streams. A positive output voltage of SL<p> minus SL<n> (differential voltage) shall correspond to tx_bit = one.

Table 84–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0
Restart training	BASE-R PMD control register	1.150.0	mr_restart_training
Training enable	BASE-R PMD control register	1.150.1	mr_training_enable

Table 84–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0
Receiver status 3	BASE-R PMD status register	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status register	1.151.13	frame_lock_3
Start-up protocol status 3	BASE-R PMD status register	1.151.14	training_3
Training failure 3	BASE-R PMD status register	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status register	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status register	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status register	1.151.10	training_2
Training failure 2	BASE-R PMD status register	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status register	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status register	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status register	1.151.6	training_1

Table 84–3—MDIO/PMD status variable mapping (continued)

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Training failure 1	BASE-R PMD status register	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status register	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status register	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status register	1.151.2	training_0
Training failure 0	BASE-R PMD status register	1.151.3	training_failure_0

84.7.3 PMD receive function

The PMD Receive function shall convert the four electrical streams from the MDI into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication. A positive input voltage of DL<p> minus DL<n> (differential voltage) shall correspond to rx_bit = one.

84.7.4 Global PMD signal detect function

The Global PMD signal detect function shall continuously report the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) to the PMD service interface. SIGNAL_DETECT, while normally intended to be an indicator of signal presence, is used by 40GBASE-KR4 to indicate the successful completion of the start-up protocol on all lanes.

SIGNAL_DETECT shall be set to FAIL following system reset or the manual reset of the training state diagram. Upon successful completion of training on all lanes, SIGNAL_DETECT shall be set to OK.

If training is disabled by management, SIGNAL_DETECT shall be set to OK.

If the MDIO interface is implemented, then Global_PMD_signal_detect (1.10.0) shall be continuously set to the value of SIGNAL_DETECT as described in 45.2.1.9.5.

84.7.5 PMD lane-by-lane signal detect function

If the MDIO interface is implemented, then PMD_signal_detect_0 (1.10.1), PMD_signal_detect_1 (1.10.2), PMD_signal_detect_2 (1.10.3), and PMD_signal_detect_3 (1.10.4) shall be set to one or zero depending on whether a particular lane's signal_detect, as defined by the training state diagram in Figure 72-5, returns true or false (see 45.2.1.9).

84.7.6 Global PMD transmit disable function

The Global_PMD_transmit_disable function is optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

- a) When the Global_PMD_transmit_disable variable is set to one, this function shall turn off all of the transmitters such that each transmitter drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage with TX disabled as specified in Table 72-6.
- b) If a PMD_fault (84.7.9) is detected, then the PMD may turn off the electrical transmitter in all lanes.
- c) Loopback, as defined in 84.7.8, shall not be affected by Global_PMD_transmit_disable.

If the MDIO interface is implemented, then this function shall map to the Global_PMD_transmit_disable bit as specified in 45.2.1.8.5.

84.7.7 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_*i* function (where *i* represents the lane number in the range 0:3) is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

- a) When a PMD_transmit_disable_*i* variable is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage with TX disabled as specified in Table 72-6.
- b) If a PMD_fault (84.7.9) is detected, then the PMD may set each PMD_transmit_disable_*i* to one, turning off the electrical transmitter in each lane.
- c) Loopback, as defined in 84.7.8, shall not be affected by PMD_transmit_disable_*i*.

84.7.8 Loopback mode

Local loopback shall be provided by the adjacent PMA (see 83.5.8) for the 40GBASE-KR4 PMD as a test function to the device. When loopback mode is enabled, transmission requests passed to each transmitter are sent directly to the corresponding receiver, overriding any signal detected by each receiver on its attached link. Note that loopback mode does not affect the state of the transmitter that continues to send data (unless disabled). The method of implementing loopback mode is not defined by this standard.

Control of the loopback function is specified in 45.2.1.1.4.

NOTE 1—The signal paths that are exercised in the loopback mode are implementation specific, but it is recommended that these signal paths encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

84.7.9 PMD_fault function

If the MDIO is implemented, PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault.

84.7.10 PMD transmit fault function

The PMD_transmit_fault function is optional. The faults detected by this function are implementation specific, but should not include the assertion of the Global_PMD_transmit_disable function.

If a PMD_transmit_fault (optional) is detected, then the Global_PMD_transmit_disable function should also be asserted.

If the MDIO interface is implemented, then this function shall be mapped to the Transmit fault bit as specified in 45.2.1.7.4.

84.7.11 PMD receive fault function

The PMD_receive_fault function is optional. The faults detected by this function are implementation specific.

If the MDIO interface is implemented, then this function shall contribute to PMA/PMD Receive fault bit as specified in 45.2.1.7.5.

84.7.12 PMD control function

Each lane of the 40GBASE-KR4 PMD shall use the same control function as 10GBASE-KR, as defined in 72.6.10.

The random seed for the training pattern described in 72.6.10.2.6 shall be different for each of the lanes.

The variables $rx_trained_i$, $frame_lock_i$, $training_i$, and $training_failure_i$ (where i goes from 0 to 3) report status for each lane and are equivalent to $rx_trained$, $frame_lock$, $training$, and $training_failure$ as defined in 72.6.10.3.1.

If the MDIO interface is implemented, then this function shall map these variables to the appropriate bits in the BASE-R PMD status register (Register 1.151) as specified in 45.2.1.78.

84.8 40GBASE-KR4 electrical characteristics

84.8.1 Transmitter characteristics

Transmitter electrical characteristics at TP1 for 40GBASE-KR4 shall be the same as 10GBASE-KR, as detailed in 72.7.1.1 through 72.7.1.11.

84.8.1.1 Test fixture

The test fixture defined for 10GBASE-KR in 72.7.1.1 shall be used on all lanes.

84.8.2 Receiver characteristics

Receiver electrical characteristics at TP4 for 40GBASE-KR4 shall be the same as 10GBASE-KR, as detailed in 72.7.2.1 through 72.7.2.5.

84.8.2.1 Receiver interference tolerance

The receiver interference tolerance tests shall be the same as those described for 10GBASE-KR in 72.7.2.1 and Annex 69A.

For 40GBASE-KR4, each lane shall be tested individually using the methodology defined in Annex 69A with the transmitters of the unused lanes active and terminated by the reference impedance.

84.9 Interconnect characteristics

Informative interconnect characteristics for 40GBASE-KR4 are provided in Annex 69B.

84.10 Environmental specifications

84.10.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

84.10.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

84.10.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

84.10.4 Electromagnetic compatibility

A system integrating the 40GBASE-KR4 PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

84.10.5 Temperature and humidity

A system integrating the 40GBASE-KR4 PHY is 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 standard.

84.11 Protocol implementation conformance statement (PICS) proforma for Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4¹⁴

84.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

84.11.2 Identification

84.11.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

84.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 84, Physical Medium Dependent sublayer and baseband medium, type 40GBASE-KR4
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

¹⁴Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

84.11.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XGE	XLGMII	84.1	Interface is supported	O	Yes [] No []
PCS	Support of 40GBASE-R PCS	84.1		M	Yes []
PMA	Support of 40GBASE-R PMA	84.1		M	Yes []
XLAUI	XLAUI	84.1	Interface is supported	O	Yes [] No []
AN	Auto-Negotiation	84.1	Device implements Auto-Negotiation	M	Yes []
FEC	Forward Error Correction	84.1	Device implements BASE-R Forward Error Correction	O	Yes []
DC	Delay constraints	84.4	Device conforms to delay constraints	M	Yes []
DSC	Skew constraints	84.5	Device conforms to Skew and Skew Variation constraints	M	Yes []
*MD	MDIO interface	84.6	Device implements MDIO	O	Yes [] No []
*TD	Global_PMD_transmit_disable	84.7.6		O	Yes [] No []

84.11.4 PICS proforma tables for Clause 84, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-KR4

84.11.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	Transmit function	84.7.2	Conveys bits from PMD service interface to MDI	M	Yes []
FS2	Transmitter signal	84.7.2	A positive differential voltage corresponds to tx_bit = one	M	Yes []
FS3	Receive function	84.7.3	Conveys bits from MDI to PMD service interface	M	Yes []
FS4	Receiver signal	84.7.3	A positive differential voltage corresponds to rx_bit = one	M	Yes []
FS5	Signal detect	84.7.4	Report to PMD service interface	M	Yes []
FS6	SIGNAL_DETECT value	84.7.4	Set to FAIL on reset	M	Yes []
FS7	SIGNAL_DETECT value	84.7.4	Set to OK when training completes successfully	M	Yes []
FS8	SIGNAL_DETECT value	84.7.4	Set to OK when training disabled	M	Yes []
FS9	Transmit disable requirements	84.7.6	Requirements of 84.7.6 and 84.7.7	TD:M	Yes [] N/A []
FS10	Loopback support	84.7.8	Provided for 40GBASE-KR4 PMD by transmitter and receiver	M	Yes []
FS11	Control function	84.7.12	The same control function as 10GBASE-KR is used	M	Yes []
FS12	Random seed	84.7.12	The random seed shall be different for each lane	M	Yes []

84.11.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	MDIO Variable Mapping	84.6	Per Table 84–2, and Table 84–3	MD:M	Yes [] N/A []
MF2	Global signal detect	84.7.4	Value described in 45.2.1.9.5	MD:M	Yes [] N/A []
MF3	Lane-by-lane signal detect	84.7.5	Value as described in 84.7.5	MD:M	Yes [] N/A []
MF4	PMD_transmit_fault function	84.7.10	Sets PMD_transmit_fault as specified in 45.2.1.7.4	MD:M	Yes [] N/A []
MF5	PMD_receive_fault function	84.7.11	Sets PMD_receive_fault as specified in 45.2.1.7.5	MD:M	Yes [] N/A []
MF6	PMD training status	84.7.12	Sets training status as specified in 45.2.1.78	MD:M	Yes [] N/A []

84.11.4.3 Transmitter electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Transmitter characteristics	84.8.1	The same transmitter characteristics as 10GBASE-KR are used	M	Yes []
TC2	Test fixture	84.8.1.1	The same test fixture as 10GBASE-KR is used	M	Yes []

84.11.4.4 Receiver electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Receiver characteristics	84.8.2	The same receiver characteristics as 10GBASE-KR are used	M	Yes []
RC2	Receiver interference tolerance	84.8.2.1	The same receiver interference tolerance test as 10GBASE-KR is used	M	Yes []
RC3	Receiver interference tolerance testing	84.8.2.1	Each lane shall be tested individually using the methodology defined in Annex 69A with the transmitters of the unused lanes active and terminated by the reference impedance	M	Yes []

84.11.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	84.10.1	Complies with applicable section of IEC 60950-1	M	Yes []
ES2	Electromagnetic interference	84.10.4	Complies with applicable local and national codes	M	Yes []

85. Physical Medium Dependent sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10

85.1 Overview

This clause specifies the 40GBASE-CR4 PMD and the 100GBASE-CR10 PMD (including MDI) and the baseband medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 85–1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

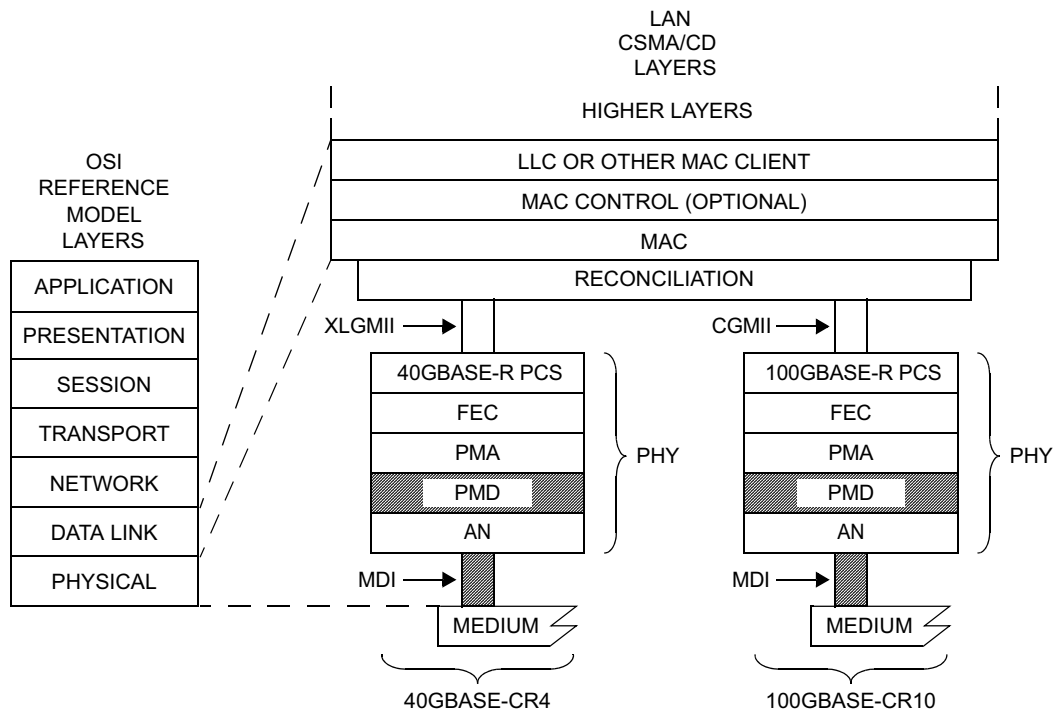
Table 85–1—Physical Layer clauses associated with the 40GBASE-CR4 and 100GBASE-CR10 PMDs

Associated clause	40GBASE-CR4	100GBASE-CR10
81—RS	Required	Required
81—XLGMII ^a	Optional	Not applicable
81—CGMII ^b	Not applicable	Optional
82—PCS for 40GBASE-R	Required	Not applicable
82—PCS for 100GBASE-R	Not applicable	Required
74—FEC	Optional	Optional
83—PMA for 40GBASE-R	Required	Not applicable
83—PMA for 100GBASE-R	Not applicable	Required
83A—XLAUI	Optional	Not applicable
83A—CAUI	Not applicable	Optional
73—Auto-Negotiation	Required	Required

^aThe XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.

^bThe CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.

Figure 85–1 shows the relationship of the 40GBASE-CR4 and 100GBASE-CR10 PMD sublayers and MDI to the ISO/IEC Open System Interconnection (OSI) reference model.



AN = AUTO-NEGOTIATION

CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE

FEC = FORWARD ERROR CORRECTION

LLC = LOGICAL LINK CONTROL

MAC = MEDIA ACCESS CONTROL

MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE

PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT

XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

CR = PMD FOR COPPER — 7 m

Figure 85–1—40GBASE-CR4 and 100GBASE-CR10 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 CSMA/CD LAN model

85.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-CR4 and 100GBASE-CR10 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA and PMD entities. The PMD translates the encoded data to and from signals suitable for the specified medium. The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_i.request
PMD:IS_UNITDATA_i.indication
PMD:IS_SIGNAL.indication
```

The 40GBASE-CR4 PMD has four parallel bit streams, hence $i = 0$ to 3 for 40GBASE-CR4, and the 100GBASE-CR10 PMD has ten parallel bit streams, hence $i = 0$ to 9 for 100GBASE-CR10. The PMA (or the PMD) continuously sends four or ten parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 10.3125 GBd.

For 40GBASE-CR4, SIGNAL_DETECT indicates the successful completion of the start-up protocol on all four lanes. For 100GBASE-CR10, SIGNAL_DETECT indicates the successful completion of the start-up protocol on all ten lanes.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the PMD:IS_UNITDATA_*i*.indication parameters are undefined.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication primitive.

85.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD is required to support the AN service interface primitive AN_LINK.indication defined in 73.9. (See 82.6.)

85.4 Delay constraints

A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-CR4 PMD, AN, and the medium in one direction shall be no more than 6144 bit times (12 pause_quanta or 153.6 ns). It is assumed that the one way delay through the medium is no more than 2072 bit times (51.8 ns).

The sum of the transmit and the receive delays at one end of the link contributed by the 100GBASE-CR10 PMD, AN, and the medium in one direction shall be no more than 14848 bit times (29 pause_quanta or 148.48 ns). It is assumed that the one way delay through the medium is no more than 5180 bit times (51.8 ns).

85.5 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and are specified at the points SP1 to SP6 shown in Figure 80–4 and Figure 80–5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns. For more information on Skew and Skew Variation, see 80.5.

85.6 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMD. Mapping of MDIO control variables to PMD control variables is illustrated in Table 85–2 for 40GBASE-CR4 and 100GBASE-CR10. Mapping of MDIO status variables to PMD status variables is illustrated in Table 85–3 for 40GBASE-CR4 and 100GBASE-CR10.

Table 85–2—40GBASE-CR4 and 100GBASE-CR10 MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	Global_PMD_transmit_disable
PMD transmit disable 9 to PMD transmit disable 0	PMD transmit disable register	1.9.10 to 1.9.1	PMD_transmit_disable_9 to PMD_transmit_disable_0
Restart training	BASE-R PMD control register	1.150.0	mr_restart_training
Training enable	BASE-R PMD control register	1.150.1	mr_training_enable

Table 85–3—40GBASE-CR4 and 100GBASE-CR10 MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	Global_PMD_signal_detect
PMD receive signal detect 9 to PMD receive signal detect 0	PMD receive signal detect register	1.10.10 to 1.10.1	PMD_signal_detect_9 to PMD_signal_detect_0
Receiver status 9	BASE-R PMD status 3 register	1.157.4	rx_trained_9
Frame lock 9	BASE-R PMD status 3 register	1.157.5	frame_lock_9
Start-up protocol status 9	BASE-R PMD status 3 register	1.157.6	training_9
Training failure 9	BASE-R PMD status 3 register	1.157.7	training_failure_9
Receiver status 8	BASE-R PMD status 3 register	1.157.0	rx_trained_8
Frame lock 8	BASE-R PMD status 3 register	1.157.1	frame_lock_8
Start-up protocol status 8	BASE-R PMD status 3 register	1.157.2	training_8
Training failure 8	BASE-R PMD status 3 register	1.157.3	training_failure_8
Receiver status 7	BASE-R PMD status 2 register	1.156.12	rx_trained_7
Frame lock 7	BASE-R PMD status 2 register	1.156.13	frame_lock_7

Table 85–3—40GBASE-CR4 and 100GBASE-CR10 MDIO/PMD status variable mapping (continued)

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Start-up protocol status 7	BASE-R PMD status 2 register	1.156.14	training_7
Training failure 7	BASE-R PMD status 2 register	1.156.15	training_failure_7
Receiver status 6	BASE-R PMD status 2 register	1.156.8	rx_trained_6
Frame lock 6	BASE-R PMD status 2 register	1.156.9	frame_lock_6
Start-up protocol status 6	BASE-R PMD status 2 register	1.156.10	training_6
Training failure 6	BASE-R PMD status 2 register	1.156.11	training_failure_6
Receiver status 5	BASE-R PMD status 2 register	1.156.4	rx_trained_5
Frame lock 5	BASE-R PMD status 2 register	1.156.5	frame_lock_5
Start-up protocol status 5	BASE-R PMD status 2 register	1.156.6	training_5
Training failure 5	BASE-R PMD status 2 register	1.156.7	training_failure_5
Receiver status 4	BASE-R PMD status 2 register	1.156.0	rx_trained_4
Frame lock 4	BASE-R PMD status 2 register	1.156.1	frame_lock_4
Start-up protocol status 4	BASE-R PMD status 2 register	1.156.2	training_4
Training failure 4	BASE-R PMD status 2 register	1.156.3	training_failure_4
Receiver status 3	BASE-R PMD status register	1.151.12	rx_trained_3
Frame lock 3	BASE-R PMD status register	1.151.13	frame_lock_3
Start-up protocol status 3	BASE-R PMD status register	1.151.14	training_3
Training failure 3	BASE-R PMD status register	1.151.15	training_failure_3
Receiver status 2	BASE-R PMD status register	1.151.8	rx_trained_2
Frame lock 2	BASE-R PMD status register	1.151.9	frame_lock_2
Start-up protocol status 2	BASE-R PMD status register	1.151.10	training_2
Training failure 2	BASE-R PMD status register	1.151.11	training_failure_2
Receiver status 1	BASE-R PMD status register	1.151.4	rx_trained_1
Frame lock 1	BASE-R PMD status register	1.151.5	frame_lock_1
Start-up protocol status 1	BASE-R PMD status register	1.151.6	training_1
Training failure 1	BASE-R PMD status register	1.151.7	training_failure_1
Receiver status 0	BASE-R PMD status register	1.151.0	rx_trained_0
Frame lock 0	BASE-R PMD status register	1.151.1	frame_lock_0
Start-up protocol status 0	BASE-R PMD status register	1.151.2	training_0
Training failure 0	BASE-R PMD status register	1.151.3	training_failure_0

85.7 PMD functional specifications

85.7.1 Link block diagram

A 40GBASE-CR4 or 100GBASE-CR10 link in one direction is illustrated in Figure 85–2. For purposes of system conformance, the PMD sublayer is standardized at the test points described in this subclause. The electrical transmit signal is defined at TP2. Unless specified otherwise, all transmitter measurements and tests defined in Table 85–5 are made at TP2 utilizing the test fixture specified in 85.8.3.5. Unless specified otherwise, all receiver measurements and tests defined in 85.8.4 are made at TP3 utilizing the test fixture specified in 85.8.3.5. A mated connector pair has been included in both the transmitter and receiver specifications defined in 85.8.3 and 85.8.4. The maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is specified in 85.8.3.4.

The 40GBASE-CR4 and 100GBASE-CR10 channels are defined between the transmitter (TP0) and receiver blocks (TP5) to include the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss as illustrated in Figure 85–2. Annex 85A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system. All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 85–2. The cable assembly test fixture of Figure 85–13 or its functional equivalent, is required for measuring the cable assembly specifications in 85.10 at TP1 and TP4. Two mated connector pairs and the cable assembly test fixture have been included in the cable assembly specifications defined in 85.10. Transmitter and receiver differential controlled impedance printed circuit board insertion losses defined between TP0–TP1 and TP4–TP5 respectively are provided informatively in Annex 85A.

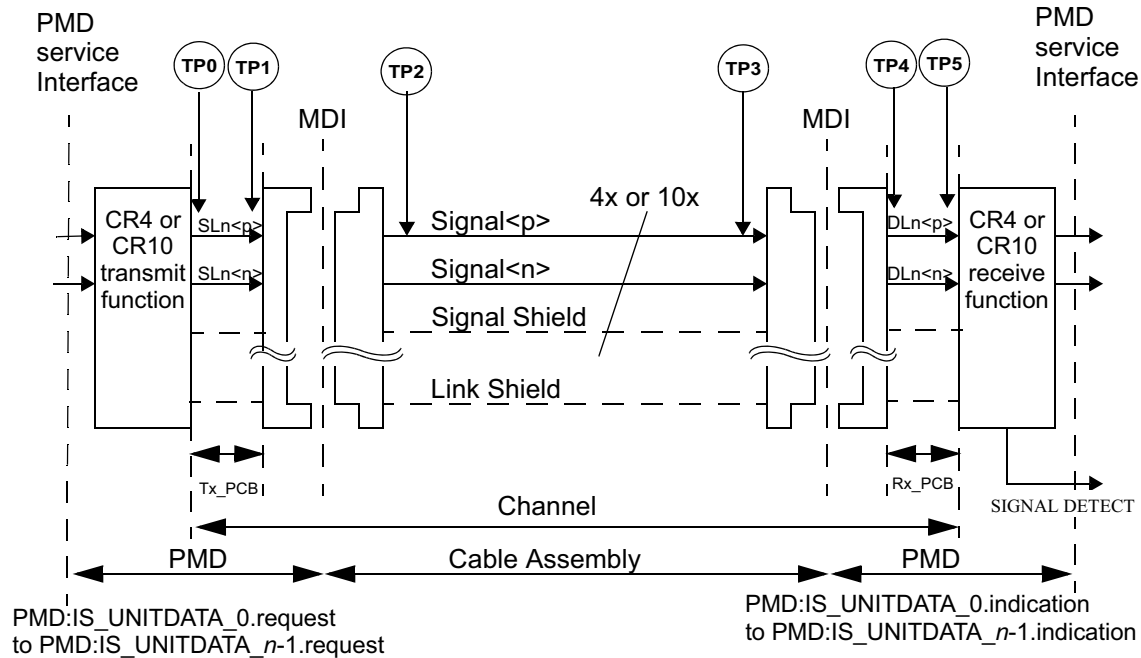


Figure 85–2—40GBASE-CR4 or 100GBASE-CR10 link (half link is illustrated)

Note that the source lanes (SL), signals $SL_n<p>$, and $SL_n<n>$ are the positive and negative sides of the transmitters differential signal pairs and the destination lanes (DL) signals, $DL_n<p>$, and $DL_n<n>$ are the positive and negative sides of the receivers differential signal pairs for lane n ($n = 0, 1, 2, 3$ or $n = 0, 1, 2, 3, 4, 5, 6, 7, 8, 9$).

Table 85–4 describes the defined test points illustrated in Figure 85–2.

Table 85–4—40GBASE-CR4 and 100GBASE-CR10 test points

Test points	Description
TP0 to TP5	The 40GBASE-CR4 and 100GBASE-CR10 channels including the transmitter and receiver differential controlled impedance printed circuit board insertion loss and the cable assembly insertion loss.
TP1 to TP4	All cable assembly measurements are to be made between TP1 and TP4 as illustrated in Figure 85–2. The cable assembly test fixture of Figure 85–13 or its functional equivalent, is required for measuring the cable assembly specifications in 85.10 at TP1 and TP4.
TP0 to TP1 TP4 to TP5	A mated connector pair has been included in both the transmitter and receiver specifications defined in 85.8.3 and 85.8.4. The maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is specified in 85.8.3.4.
TP2	Unless specified otherwise, all transmitter measurements and tests defined in Table 85–5 are made at TP2 utilizing the test fixture specified in 85.8.3.5.
TP3	Unless specified otherwise, all receiver measurements and tests defined in 85.8.4 are made at TP3 utilizing the test fixture specified in 85.8.3.5.

85.7.2 PMD Transmit function

The 40GBASE-CR4 PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate electrical streams. The four electrical signal streams shall then be delivered to the MDI, all according to the transmit electrical specifications in 85.8.3. A positive output voltage of $SL_{<p>} - SL_{<n>}$ (differential voltage) shall correspond to $tx_bit = one$. The 100GBASE-CR10 PMD Transmit function shall convert the ten bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_9.request. The ten electrical signal streams shall then be delivered to the MDI, all according to the transmit electrical specifications in 85.8.3. A positive output voltage of $SL_{<p>} - SL_{<n>}$ (differential voltage) shall correspond to $tx_bit = one$.

85.7.3 PMD Receive function

The 40GBASE-CR4 PMD Receive function shall convert the four electrical streams from the MDI into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication. A positive input voltage of $DL_{<p>} - DL_{<n>}$ (differential voltage) shall correspond to $rx_bit = one$. The 100GBASE-CR10 PMD Receive function shall convert the ten electrical streams from the MDI into ten bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_9.indication. A positive input voltage of $DL_{<p>} - DL_{<n>}$ (differential voltage) shall correspond to $rx_bit = one$.

85.7.4 Global PMD signal detect function

The Global PMD signal detect function shall continuously report the message PMD:IS_SIGNAL.indication (SIGNAL_DETECT) to the PMD service interface. SIGNAL_DETECT, while normally intended to be an indicator of signal presence, is used by 40GBASE-CR4 and 100GBASE-CR10 to indicate the successful completion of the start-up protocol on all lanes. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication primitive.

SIGNAL_DETECT shall be set to FAIL following system reset or the manual reset of the training state diagram. Upon successful completion of training on all lanes, SIGNAL_DETECT shall be set to OK.

If training is disabled by management, SIGNAL_DETECT shall be set to OK.

If the MDIO interface is implemented, then Global_PMD_signal_detect (1.10.0) shall be continuously set to the value of SIGNAL_DETECT as described in 45.2.1.9.5.

85.7.5 PMD lane-by-lane signal detect function

When the MDIO is implemented, each PMD_signal_detect_{*i*} value, where *i* represents the lane number in the range 0:3 for 40GBASE-CR4 and 0:9 for 100GBASE-CR10, shall be continuously updated as described in the following two paragraphs.

For 40GBASE-CR4 PMD_signal_detect_0 (1.10.1), PMD_signal_detect_1 (1.10.2), PMD_signal_detect_2 (1.10.3) and PMD_signal_detect_3 (1.10.4) shall be set to one or zero depending on whether a particular lane's signal_detect, as defined by the training state diagram in Figure 72-5, returns true or false.

For 100GBASE-CR10 PMD_signal_detect_0 (1.10.1), PMD_signal_detect_1 (1.10.2), PMD_signal_detect_2 (1.10.3), PMD_signal_detect_3 (1.10.4), PMD_signal_detect_4 (1.10.5), PMD_signal_detect_5 (1.10.6), PMD_signal_detect_6 (1.10.7), PMD_signal_detect_7 (1.10.8), PMD_signal_detect_8 (1.10.9), and PMD_signal_detect_9 (1.10.10) shall be set to one or zero depending on whether a particular lane's signal_detect, as defined by the training state diagram in Figure 72-5, returns true or false.

85.7.6 Global PMD transmit disable function

The Global_PMD_transmit_disable function is optional. When implemented, it allows all of the transmitters to be disabled with a single variable.

- a) When Global_PMD_transmit_disable variable is set to one, this function shall turn off all of the transmitters such that each transmitter drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage in Table 85-5.
- b) If a PMD_fault (85.7.9) is detected, then the PMD may turn off the electrical transmitter in all lanes.
- c) Loopback, as defined in 85.7.8, shall not be affected by Global_PMD_transmit_disable.

85.7.7 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_{*i*} function (where *i* represents the lane number in the range 0:3 or 0:9) is optional and allows the electrical transmitter in each lane to be selectively disabled. When this function is supported, it shall meet the following requirements:

- a) When a PMD_transmit_disable_{*i*} variable is set to one, this function shall turn off the transmitter associated with that variable such that it drives a constant level (i.e., no transitions) and does not exceed the maximum differential peak-to-peak output voltage specified in Table 85-5.
- b) If a PMD_fault (85.7.9) is detected, then the PMD may set each PMD_transmit_disable_{*i*} to one, turning off the electrical transmitter in each lane.
- c) Loopback, as defined in 85.7.8, shall not be affected by PMD_transmit_disable_{*i*}.

85.7.8 Loopback mode

Local loopback mode shall be provided by the adjacent PMA (see 83.5.8) for the 40GBASE-CR4 and 100GBASE-CR10 PMDs as a test function to the device. When loopback mode is enabled, transmission requests passed to each transmitter are sent directly to the corresponding receiver, overriding any signal

detected by each receiver on its attached link. Note that loopback mode does not affect the state of the transmitter, which continues to send data (unless disabled).

Control of the loopback function is specified in 45.2.1.1.4.

NOTE 1—The signal path that is exercised in the loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

85.7.9 PMD_fault function

If the MDIO is implemented, PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault. The PMD fault function shall be mapped to bit 1.1.7 as listed in Table 85–3.

85.7.10 PMD transmit fault function

The PMD_transmit_fault function is optional. The faults detected by this function are implementation specific, but should not include the assertion of the Global_PMD_transmit_disable function.

If a PMD_transmit_fault (optional) is detected, then the Global_PMD_transmit_disable function should also be asserted.

If the MDIO interface is implemented, then this function shall be mapped to the Transmit fault bit as specified in 45.2.1.7.4.

85.7.11 PMD receive fault function

The PMD_receive_fault function is optional. The faults detected by this function are implementation specific.

If the MDIO interface is implemented, then this function shall contribute to the Receive fault bit as specified in 45.2.1.7.5.

85.7.12 PMD control function

Each lane of the 40GBASE-CR4 or 100GBASE-CR10 PMD shall use the same control function as 10GBASE-KR, as defined in 72.6.10.

The random seed for the training pattern described in 72.6.10.2.6 shall be different for each of the lanes.

85.8 MDI Electrical specifications for 40GBASE-CR4 and 100GBASE-CR10

85.8.1 Signal levels

The 40GBASE-CR4 and 100GBASE-CR10 MDI is a low-swing AC coupled differential interface. AC coupling at the receiver, as defined in 85.8.4.5, allows for interoperability between components operating from different supply voltages. Low-swing differential signaling provides noise immunity and improved electromagnetic interference (EMI).

85.8.2 Signal paths

The 40GBASE-CR4 and 100GBASE-CR10 MDI signal paths are point-to-point connections. Each path corresponds to a 40GBASE-CR4 or 100GBASE-CR10 MDI lane and comprises two complementary signals, which form a balanced differential pair. For 40GBASE-CR4, there are four differential paths in each direction for a total of eight pairs, or sixteen connections. For 100GBASE-CR10, there are ten differential paths in each direction for a total of 20 pairs, or forty connections. The signal paths are intended to operate on twinaxial cable assemblies ranging from 0.5 m to 7 m in length, as described in 85.10.

85.8.3 Transmitter characteristics

Transmitter characteristics shall meet specifications summarized in Table 85–5 at TP2 unless otherwise noted. The transmitter specifications at TP0 are provided informatively in Annex 85A, Table 85A–1.

Table 85–5—Transmitter characteristics at TP2 summary

Parameter	Subclause reference	Value	Units
Signaling rate, per lane	85.8.3.8	10.3125 ± 100 ppm	GBd
Unit interval nominal	85.8.3.8	96.969697	ps
Differential peak-to-peak output voltage (max) with Tx disabled	85.8.3.3	30	mV
Common-mode voltage limits	72.7.1.4	0 to 1.9	V
Differential output return loss (min)	85.8.3.1	See Equation (85–1)	dB
Common-mode output return loss (min)		See Equation (72–6) and Equation (72–7)	dB
Common-mode AC output voltage (max., RMS)		30	mV
Amplitude peak-to-peak (max)	72.7.1.4	1200 ^a	mV
Transmitter DC amplitude ^b	85.8.3.3	0.34 min, 0.6 max	V
Linear fit pulse (min) ^c	85.8.3.3	$0.63 \times$ Transmitter DC amplitude	V
Transmitted waveform max normalized error(linear fit), “e”	85.8.3.3	0.037	
abs coefficient step size	85.8.3.3.2	0.0083 min, 0.05 max	
minimum precursor fullscale range	85.8.3.3.3	1.54	
minimum post cursor fullscale range	85.8.3.3.3	4	
Far-end transmit output noise (max) Low insertion loss channel High insertion loss channel	85.8.3.2	2 See Equation (85–2) 1 See Equation (85–3)	mV
Max output jitter (peak-to-peak) Random jitter ^d Duty Cycle Distortion ^e Total jitter excluding data dependent jitter ^f		0.15 0.035 0.25	UI UI UI

^aThe 40GBASE-CR4 Style-1 connector may support 40GBASE-CR4 or XLPPPI interfaces. For implementations that support both interfaces, the transmitter should not exceed the XLPPPI voltage maximum until a 40GBASE-CR4 cable assembly has been identified. The 100GBASE-CR10 connector may support 100GBASE-CR10 or CPPI interfaces. For implementations that support both interfaces, the transmitter should not exceed the CPPI voltage maximum until a 100GBASE-CR10 cable assembly has been identified.

^bThe transmitter DC amplitude is the sum of linear fit pulse response $p(k)$ from step 3) divided by M from step 3).

^cThe peak of the linear fit pulse response $p(k)$ from step 3).

^dRandom jitter is specified at a BER of 10^{-12} .

^eSee 72.7.1.8 for duty cycle distortion definition.

^fTotal jitter at a BER of 10^{-12} measured per 83A.5.1 excluding data dependent jitter (DDJ). DDJ is a jitter component where jitter that is not correlated to the data pattern has been removed. DDJ is measured with PRBS9 as specified in 83.5.10.

85.8.3.1 Transmitter differential output return loss

The differential output return loss, in dB, of the transmitter shall meet Equation (85–1). This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements shall be 100Ω .

$$Return_loss(f) \geq \begin{cases} 12 - 2\sqrt{f} & 0.05 \leq f < 4.11 \\ 6.3 - 13\log_{10}(f/5.5) & 4.11 \leq f \leq 10 \end{cases} \quad (\text{dB}) \quad (85-1)$$

where

f is the frequency in GHz
 $Return_loss(f)$ is the return loss at frequency f

85.8.3.2 Transmitter noise parameter measurements

The far-end transmitter output noise is an additional source of noise to the cable assembly's integrated crosstalk noise (ICN) specified in 85.10.7. The far-end transmitter output noise parameter is characterized using two reference channels; a "low-loss" cable assembly with insertion loss on the reference pair of $6 \text{ dB} \pm 1.5 \text{ dB}$ at 5.15625 GHz and cable assembly integrated crosstalk noise (ICN) meeting the requirements of 85.10.7 and a "high-loss" cable assembly with insertion loss on the reference pair of $15.7 \text{ dB} \pm 1.5 \text{ dB}$ at 5.15625 GHz and cable assembly ICN meeting the requirements of 85.10.7. The far-end transmitter output noise is characterized as a deviation from the cable assembly ICN using the following procedure:

- 1) Compute the far-end integrated crosstalk noise σ_{fx} into the reference lane of the cable assembly using the methodology of 85.10.7 and the parameters in Table 85–11.
- 2) Denote σ_l as the far-end ICN for the low-loss cable assembly.
- 3) Denote σ_h as the far-end ICN for the high-loss cable assembly.
- 4) The transmitter under test is connected to one end of the reference cable assembly and the other end is connected to the cable assembly test fixture specified in 85.10.8.
- 5) All lanes of the cable assembly test fixture are terminated in the reference impedance with the reference lane connected to the measuring instrument.
- 6) The reference lane of the transmitter under test sends a square wave test pattern as specified in 83.5.10 while all other adjacent transmitter lanes send either scrambled idle or PRBS31.
- 7) A fixed point on the square wave test pattern is chosen and the RMS deviation from the mean voltage at this observation point is measured. The histogram for RMS noise measurement is 1 UI wide.
- 8) The measurement should not include the measurement system noise.

For the low-loss cable assembly, the measured RMS deviation from the cable assembly ICN due to the far-end transmitter output noise shall meet the values determined using Equation (85–2).

$$RMSI_{dev} \leq \sqrt{\sigma_l^2 + 2^2} \quad (\text{mV}) \quad (85-2)$$

For the high-loss cable assembly, the measured RMS deviation from the cable assembly ICN due to the far-end transmitter output noise shall meet the values determined using Equation (85–3).

$$RMS h_{dev} \leq \sqrt{\sigma_h^2 + 1^2} \text{ (mV)} \quad (85-3)$$

85.8.3.3 Transmitter output waveform

The 40GBASE-CR4 and 100GBASE-CR10 transmit function includes programmable equalization to compensate for the frequency-dependent loss of the channel and facilitate data recovery at the receiver. The functional model for the transmit equalizer is the three tap transversal filter shown in Figure 85–3.

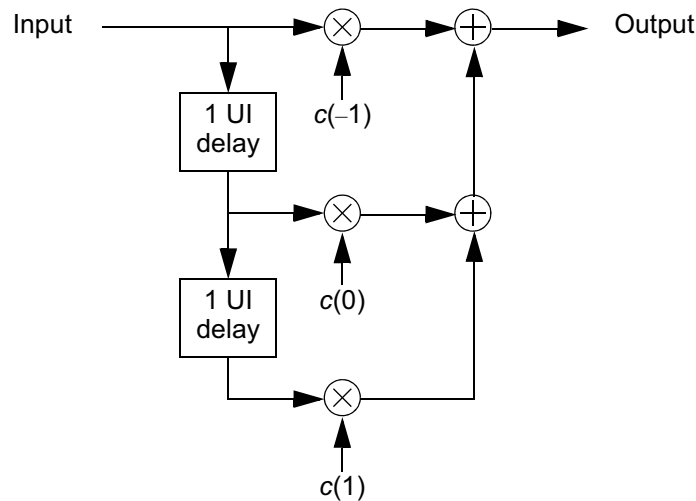


Figure 85–3— Transmit equalizer functional model

The state of the transmit equalizer and hence the transmitted output waveform may be manipulated via the PMD control function defined in 85.7.12 or via the management interface. The transmit function responds to a set of commands issued by the link partner's receive function and conveyed by a back-channel communications path.

This command set includes instructions to

- a) Increment coefficient $c(i)$.
- b) Decrement coefficient $c(i)$.
- c) Hold coefficient $c(i)$ at its current value.
- d) Set the coefficients to a pre-defined value (preset or initialize).

In response, the transmit function relays status information to the link partner's receive function. The status messages indicate that

- a1) The requested update to coefficient $c(i)$ has completed (updated).
- b1) Coefficient $c(i)$ is at its minimum value.
- c1) Coefficient $c(i)$ is at its maximum value.
- d1) Coefficient $c(i)$ is ready for the next update request (not_updated).

The requirements for the 40GBASE-CR4 and 100GBASE-CR10 transmit equalizer are intended to be similar to the requirements for 10GBASE-KR specified in 72.7.1.10. However, the signal path from the transmit function to TP2 introduces frequency-dependent loss and phase shift that distorts the signal and

makes it difficult to accurately characterize equalizer performance at TP2 using the methodology specified for 10GBASE-KR. The following process enables accurate characterization of the equalizer performance at TP2 by determining and correcting for the frequency-dependent loss and phase shift of the signal path from the transmit function to TP2.

- 1) The transmitter under test is preset as specified in 72.6.10.2.3.1 such that $c(-1)$ and $c(1)$ are zero and $c(0)$ is its maximum value.
- 2) Capture at least one complete cycle of the test pattern PRBS9 as specified in 83.5.10 at TP2 per 85.8.3.3.4.
- 3) Compute the linear fit to the captured waveform and the linear fit pulse response $p(k)$ per 85.8.3.3.5.
- 4) Define t_x to be the time where the rising edge of the linear fit pulse, p , from step 3) crosses 50% of its peak amplitude.
- 5) Sample the linear fit pulse, p , at symbol-spaced intervals relative to the time $t_0 = t_x + 0.5 \text{ UI}$, interpolating as necessary to yield the sampled pulse p_i .
- 6) Use p_i to compute the vector of coefficients, w , of a N_w -tap symbol-spaced transversal filter that equalizes for the transfer function from the transmit function to TP2 per 85.8.3.3.6.

The parameters of the pulse fit and the equalizing filter are given in Table 85–6. The DC amplitude, the sum of linear fit pulse response, $p(k)$, from step 3) divided by M from step 3), shall be greater than 0.34 V and less than or equal to 0.6 V. The peak of the linear fit pulse response from step 3) shall be greater than $0.63 \times \text{DC amplitude}$. The RMS value of the error between the linear fit from step 3) and the measured waveform, e , normalized to the peak value of the pulse, p , must be no greater than 0.037.

Table 85–6—Normalized transmit pulse template

Description	Symbol	Value	Units
Linear fit pulse length	N_p	7	UI
Linear fit pulse delay	D_p	1	UI
Equalizer length	N_w	7	UI
Equalizer delay	D_w	1	UI

For each configuration of the transmit equalizer

- 7) Configure the transmitter under test as required by the test.
- 8) Capture at least one complete cycle of the test pattern PRBS9 as specified in 83.5.10 at TP2 per 85.8.3.3.4.
- 9) Compute the linear fit to the captured waveform and the linear fit pulse response $p(k)$ per 85.8.3.3.5.
- 10) Define t_x to be the time where the rising edge of the linear fit pulse response, $p(k)$, from step 9 crosses 50% of its peak amplitude.
- 11) Sample the linear fit pulse response, $p(k)$, at symbol-spaced intervals relative to the time $t_0 = t_x + 0.5 \text{ UI}$, interpolating as necessary to yield the sampled pulse p_i .
- 12) Equalize the sampled pulse p_i using the coefficient vector, w , computed in step 6) per 85.8.3.3.6 to yield the equalized pulse q_i .

The RMS value of the error between the linear fit from step 9) and the measured waveform, e , normalized to the peak value of the pulse, p , must be no greater than 0.037.

The normalized amplitude of coefficient $c(-1)$ is the value of q_i at time $t_0 + (D_p - 1)$ UI. The normalized amplitude of coefficient $c(0)$ is the value of q_i at time $t_0 + D_p$ UI. The normalized amplitude of coefficient $c(1)$ is the value of q_i at time $t_0 + (D_p + 1)$ UI.

85.8.3.3.1 Coefficient initialization

When the PMD enters the INITIALIZE state of the Training state diagram (Figure 72-5) or receives a valid request to “initialize” from the link partner, the coefficients of the transmit equalizer shall be configured such that the ratio $(c(0)+c(1)-c(-1))/(c(0)+c(1)+c(-1))$ is $1.29 \pm 10\%$ and the ratio $(c(0)-c(1)+c(-1))/(c(0)+c(1)+c(-1))$ is $2.57 \pm 10\%$. These requirements apply upon the assertion a coefficient status report of “updated” for all coefficients.

85.8.3.3.2 Coefficient step size

The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to “increment” that coefficient shall be between 0.0083 and 0.05. The change in the normalized amplitude of coefficient $c(i)$ corresponding to a request to “decrement” that coefficient shall be between -0.05 and -0.0083 .

The change in the normalized amplitude of the coefficient is defined to be the difference in the value measured prior to the assertion of the “increment” or “decrement” request (e.g., the coefficient update request for all coefficients is “hold”) and the value upon the assertion of a coefficient status report of “updated” for that coefficient.

85.8.3.3.3 Coefficient range

When sufficient “increment” or “decrement” requests have been received for a given coefficient, the coefficient will reach a lower or upper bound based on the coefficient range or restrictions placed on the minimum steady state differential output voltage or the maximum peak-to-peak differential output voltage.

With $c(-1)$ set to zero and both $c(0)$ and $c(1)$ having received sufficient “decrement” requests so that they are at their respective minimum values, the ratio $(c(0) - c(1))/(c(0) + c(1))$ shall be greater than or equal to 4.

With $c(1)$ set to zero and both $c(-1)$ and $c(0)$ having received sufficient “decrement” requests so that they are at their respective minimum values, the ratio $(c(0) - c(-1))/(c(0) + c(-1))$ shall be greater than or equal to 1.54.

Note that a coefficient may be set to zero by first asserting a coefficient preset request and then manipulating the other coefficients as required by the test.

85.8.3.3.4 Waveform acquisition

The transmitter under test repetitively transmits the specified test pattern. The waveform shall be captured with an effective sample rate that is M times the signaling rate of the transmitter under test. The value of M shall be an integer not less than 7. Averaging multiple waveform captures is recommended.

The captured waveform shall represent an integer number of repetitions of the test pattern totaling N bits. Hence the length of the captured waveform should be MN samples. The waveform should be aligned such that the first M samples of waveform correspond to the first bit of the test pattern, the second M samples to the second bit, and so on.

85.8.3.3.5 Linear fit to the waveform measurement at TP2

Given the captured waveform $y(k)$ and corresponding aligned symbols $x(n)$ derived from the procedure defined in 85.8.3.3.4, define the M -by- N waveform matrix Y as shown in Equation (85-4).

$$Y = \begin{bmatrix} y(1) & y(M+1) & \dots & y(M(N-1)+1) \\ y(2) & y(M+2) & \dots & y(M(N-1)+2) \\ \dots & \dots & \dots & \dots \\ y(M) & y(2M) & \dots & y(MN) \end{bmatrix} \quad (85-4)$$

Rotate the symbols vector x by the specified pulse delay D_p to yield x_r as shown in Equation (85-5).

$$x_r = \begin{bmatrix} x(D_p+1) & x(D_p+2) & \dots & x(N) & x(1) & \dots & x(N-D_p) \end{bmatrix} \quad (85-5)$$

Define the matrix X to be an N -by- N matrix derived from x_r as shown in Equation (85-6).

$$X = \begin{bmatrix} x_r(1) & x_r(2) & \dots & x_r(N) \\ x_r(N) & x_r(1) & \dots & x_r(N-1) \\ \dots & \dots & \dots & \dots \\ x_r(2) & x_r(3) & \dots & x_r(1) \end{bmatrix} \quad (85-6)$$

Define the matrix X_1 to be the first N_p rows of X concatenated with a row vector of ones of length N . The M -by- (N_p+1) coefficient matrix, P , corresponding to the linear fit is then defined by Equation (85-7). The superscript “ T ” denotes the matrix transpose operator.

$$P = YX_1^T(X_1X_1^T)^{-1} \quad (85-7)$$

The error waveform, $e(k)$, is then read column-wise from the elements of E as shown in Equation (85-8).

$$E = PX_1 - Y = \begin{bmatrix} e(1) & e(M+1) & \dots & e(M(N-1)+1) \\ e(2) & e(M+2) & \dots & e(M(N-1)+2) \\ \dots & \dots & \dots & \dots \\ e(M) & e(2M) & \dots & e(MN) \end{bmatrix} \quad (85-8)$$

Define P_1 to be a matrix consisting of the first N_p columns of the matrix P as shown in Equation (85-9). The linear fit pulse response, $p(k)$, is then read column-wise from the elements of P_1 .

$$P_1 = \begin{bmatrix} p(1) & p(M+1) & \dots & p(M(N_p-1)+1) \\ p(2) & p(M+2) & \dots & p(M(N_p-1)+2) \\ \dots & \dots & \dots & \dots \\ p(M) & p(2M) & \dots & p(MN_p) \end{bmatrix} \quad (85-9)$$

85.8.3.3.6 Transfer function between the transmit function and TP2

Rotate the sampled pulse response p_i by the specified equalizer delay D_w to yield p_r as shown in Equation (85-10).

$$p_r = \begin{bmatrix} p_i(D_w+1) & p_i(D_w+2) & \dots & p_i(N_p) & p_i(1) & \dots & p_i(N_p-D_w) \end{bmatrix} \quad (85-10)$$

Define the matrix P_2 to be an N_p -by- N_p matrix derived from p_r as shown in Equation (85-11).

$$P_2 = \begin{bmatrix} p_r(1) & p_r(N_p) & \dots & p_r(2) \\ p_r(2) & p_r(1) & \dots & p_r(3) \\ \dots & \dots & \dots & \dots \\ p_r(N_p) & p_r(N_p - 1) & \dots & p_r(1) \end{bmatrix} \quad (85-11)$$

Define the matrix P_3 to be the first N_w columns of P_2 . Define a unit pulse column vector x_p of length N_p . The value of element $x_p(D_p + 1)$ is 1 and all other elements have a value of 0. The vector of filter coefficients w that equalizes p_i is then defined by Equation (85-12).

$$w = (P_3^T P_3)^{-1} P_3^T x_p \quad (85-12)$$

Given the column vector of equalizer coefficients, w , the equalized pulse response q_i is determined by Equation (85-13).

$$q_i = P_3 w \quad (85-13)$$

85.8.3.4 Insertion loss TP0 to TP2 or TP3 to TP5

Transmitter measurements and tests defined in Table 85-5 are made at TP2 or TP3 using the test fixture of Figure 85-5, or its functional equivalent. The maximum insertion loss from TP0 to TP2 or TP3 to TP5 including the test fixture is determined using Equation (85-14).

$$Insertion_loss(f) \leq \left\{ \begin{array}{ll} 0.114 + 0.8914\sqrt{f} + 0.846f & 0.01 \leq f < 7 \\ -35.91 + 6.3291f & 7 \leq f < 8 \\ 14.72 & 8 \leq f \leq 10 \end{array} \right\} \text{ (dB)} \quad (85-14)$$

where

f is the frequency in GHz
 $Insertion_loss(f)$ is the insertion loss at frequency f

The maximum insertion loss of TP0 to TP2 or TP3 to TP5 is illustrated in Figure 85-4.




Figure 85-4—Maximum insertion loss TP0 to TP2 or TP3 to TP5

85.8.3.5 Test fixture

The test fixture of Figure 85-5, or its functional equivalent, is required for measuring the transmitter specifications in 85.8.3 at TP2 and the receiver return loss at TP3. TP2 and TP3 are illustrated in Figure 85-2. Figure 85-5 illustrates the test fixture attached to TP2 or TP3.

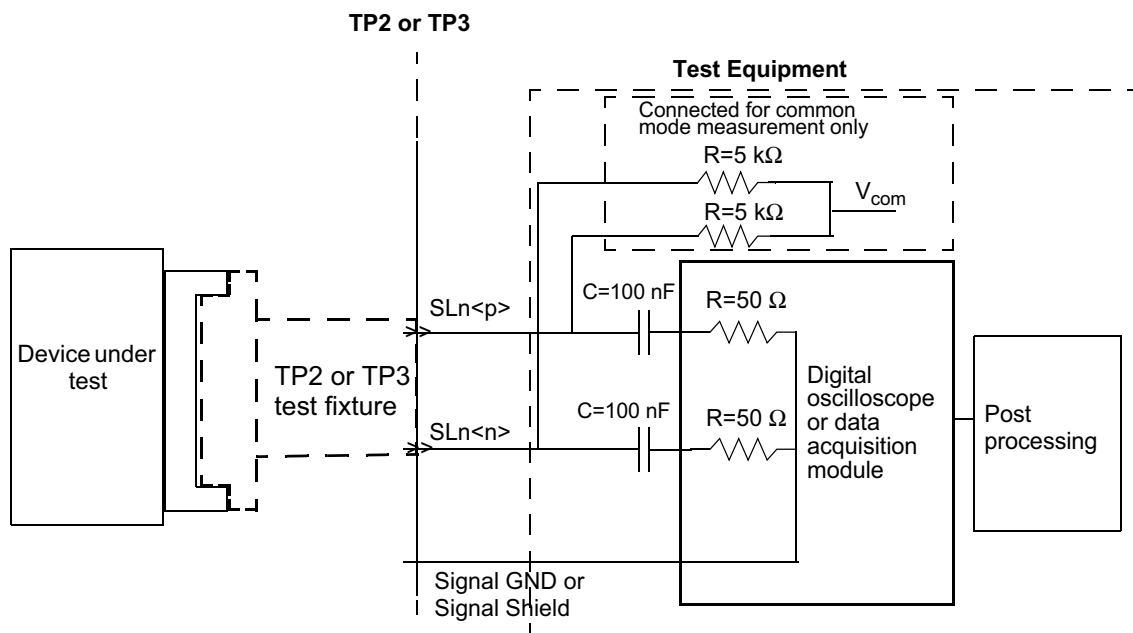


Figure 85-5—Transmitter test fixture

85.8.3.6 Test fixture impedance

The differential return loss, in dB, of the test fixture is specified in a mated state and shall meet the requirements of 85.10.9.2.

85.8.3.7 Test fixture insertion loss

The reference test fixture printed circuit board insertion loss shall meet the values determined using Equation (85–15). The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

$$IL_{tref}(f) = 0.01 + 0.3\sqrt{f} + 0.11f \text{ (dB)} \quad (85-15)$$

for $0.01 \text{ GHz} \leq f \leq 10 \text{ GHz}$

where

f is the frequency in GHz
 $IL_{tref}(f)$ is the reference test fixture PCB insertion loss at frequency f

85.8.3.8 Data dependent jitter (DDJ)

An oscilloscope, time interval analyzer, or other instrument with equivalent capability may be used to measure DDJ. A repeating PRBS9 pseudo-random test pattern, 511 bits long, is used. For DDJ measurements, the measurement bandwidth should be at least 12 GHz. If the measurement bandwidth affects the result, it can be corrected for by post-processing.

Establish a crossing level equal to the average value of the entire waveform being measured. Synchronize the instrument to the pattern repetition frequency and average the waveforms or the crossing times sufficiently to remove the effects of random jitter and noise in the system. The PRBS9 pattern has 128 positive-going transitions and 128 negative-going transitions. The mean time of each crossing is then compared to the expected time of the crossing, and a set of 256 timing variations is determined. Crossings earlier than expected give a negative variation. Crossings later than expected give a positive variation. DDJ is the range (maximum minus minimum) of the timing variations. Note that it may be convenient to align the expected time of one of the crossings with the measured mean crossing.

Figure 85–6 illustrates the method. The vertical axis is in arbitrary units, and the horizontal axis is plotted in UI. The waveform is AC coupled to an average value of 0, therefore 0 is the appropriate crossing level. The rectangular waveform shows the ideal crossing times, and the other is the waveform with jitter that is being measured. Only 16 UI are shown (out of 511). The waveforms have been arbitrarily aligned with $(\Delta t_2 = 0)$ at 5 UI.

Data dependent jitter is defined as in Equation (85–16).

$$DDJ = \max(\Delta t_1, \Delta t_2, \dots, \Delta t_{256}) - \min(\Delta t_1, \Delta t_2, \dots, \Delta t_{256}) \quad (85-16)$$

85.8.3.9 Signaling rate range

The 40GBASE-CR4 and 100GBASE-CR10 MDI signaling rate shall be 10.3125 GBd \pm 100 ppm per lane. The corresponding unit interval is nominally 96.969697 ps.




Figure 85-6—Data dependent jitter test method

85.8.4 Receiver characteristics at TP3 summary

The receiver characteristics shall meet the specifications summarized in Table 85-7 at TP3 unless otherwise noted.

Table 85-7—Receiver characteristics at TP3 summary

Parameter	Subclause reference	Value	Units
Bit error ratio	85.8.4.3	10^{-12} or better	
Signaling rate, per lane	85.8.4.4	10.3125 ± 100 ppm	GBd
Unit interval (UI) nominal	85.8.4.4	96.969697	ps
Differential peak-to-peak input amplitude tolerance (max)	72.7.2.4	1200	mV
Differential input return loss (min) ^a	85.8.4.1	Equation (85-17)	dB
Differential to common-mode input return loss		10 min from 10 MHz to 10 GHz	dB

^aRelative to 100 Ω differential.

85.8.4.1 Receiver differential input return loss

The differential input return loss, in dB, of the receiver shall meet Equation (85-17). This return loss requirement applies at all valid input levels. The reference impedance for differential return loss measurements is 100 Ω.

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.01 \leq f < 4.11 \\ 6.3 - 13 \log_{10}(f/5.5) & 4.11 \leq f \leq 10 \end{array} \right\} \quad (\text{dB}) \quad (85-17)$$

where

f is the frequency in GHz
 $Return_loss(f)$ is the return loss at frequency f

85.8.4.2 Receiver interference tolerance test

The receiver interference tolerance of each lane shall comply with both test 1 and test 2 using the parameters of Table 85–8 when measured according to the requirements of 85.8.4.2.1 to 85.8.4.2.5.

Table 85–8—10GBASE-CR4 and 10GBASE-CR10 interference tolerance parameters

Parameter	Test 1 values	Test 2 values	Units
Target BER	10^{-12}	10^{-12}	
Maximum fitted insertion loss coefficients	$a_1 = 2.15$ $a_2 = 0.78$ $a_4 = 0.03$	$a_1 = 6.04$ $a_2 = 0.94$ $a_4 = 0.08$	dB/ $\sqrt{\text{GHz}}$ dB/GHz dB/GHz ²
Applied SJ ^a (min peak-to-peak)	0.115	0.115	UI
Applied RJ ^b (min peak-to-peak)	0.13	0.13	UI
Applied DCD (min peak-to-peak)	0.035	0.035	UI
Calibrated far-end crosstalk (min RMS)	6.3	2.2	mV
Calibrated ICN (min, RMS) – σ_{nx}	3.7	3.7	mV

^aApplied SJ frequency >15 MHz, specified at TP0.

^bApplied random jitter at TP0 is specified at 10^{-12} .

85.8.4.2.1 Test setup

The interference tolerance test is performed with the setup shown in Figure 85–7. The requirements of this subclause are verified at the pattern generator connection (PGC) or test references in Figure 85–7 and Figure 85–8. The lanes under test (LUT) are illustrated in Figure 85–7 and Figure 85–8. The cable assembly receive lanes are terminated in 100 Ω differentially.

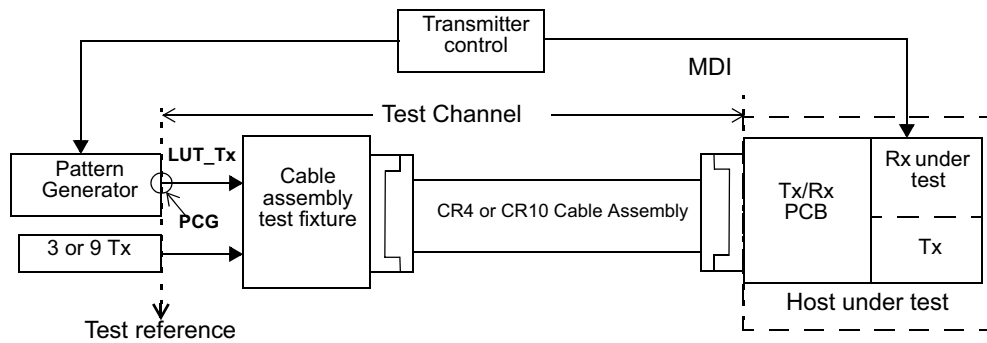


Figure 85-7—Interference tolerance test setup

85.8.4.2.2 Test channel

The test channel consists of the following:

- a) A cable assembly
- b) A cable assembly test fixture
- c) A connecting path from the pattern generator to the cable assembly test fixture

85.8.4.2.3 Test channel calibration

The insertion loss, near-end integrated crosstalk noise, and far-end crosstalk of the test channels are characterized at the test references as illustrated in Figure 85-7 using the cable assembly test fixtures specified in 85.10.8.

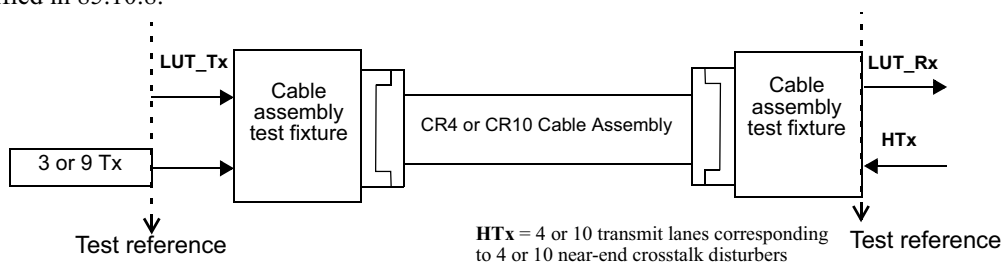


Figure 85-8—Test channel calibration

The minimum fitted insertion loss coefficients of the lane under test (LUT), derived using the fitting procedure in 85.10.2, shall meet the test values in Table 85-8. It is recommended that the deviation between the insertion loss and the fitted insertion loss be as small as practical and that the fitting parameters be as close as practical to the values given in Table 85-8.

The MDNEXT is measured from points HTx to adjacent point LUT_Rx in Figure 85-8. HTx is the set of 4 or 10 transmit lanes of the device under test corresponding to the 4 or 10 near-end crosstalk disturbers with the parameters given in Table 85-11. The RMS value of the integrated MDNEXT crosstalk noise, determined using Equation (85-30) through Equation (85-34), shall meet the test values in Table 85-8.

The far-end crosstalk disturbers consist of 40GBASE-CR4 or 100GBASE-CR10 transmitters. It is recommended that the transition time, equalization setting, and path from the far-end crosstalk disturbers to the cable assembly test fixture emulate the pattern generator as much as practical. For 40GBASE-CR4 test channels, the crosstalk that is coupled into a receive lane is from three transmitters. For 100GBASE-CR10 test channels, the crosstalk that is coupled into a receiver lane is from nine transmitters. The disturber transmitters send either scrambled idle codes or PRBS31. The amplitudes of each of the disturbers should not deviate more than 3 dB from the mean of the disturber amplitudes. The amplitudes of the disturbers should be such that the calibrated far-end crosstalk in Table 85–8 is met in the calibration setup at the LUT point with no signal applied at the PGC, and HTx and PGC terminated in 100 Ω differentially.

85.8.4.2.4 Pattern generator

The pattern generator transmits data to the device under test. Its output amplitude shall be no more than 800 mV peak-to-peak differential when measured on an alternating one zero pattern. The rise and fall times of the pattern generator, as defined in 72.7.1.7, are 47 ps. If the rise and fall times of the pattern generator, T_r , are less than 47 ps the value of da_4 in Table 85–8 is increased by da_4 from Equation (85–18).

$$da_4 = 60.51 \times 10^{-6} (47^2 - T_r^2) \quad (85-18)$$

where T_r is the rise time in ps.

The pattern generator shall meet the jitter specification in Table 85–8. The output waveform of the pattern generator shall comply to 72.7.1.11.

85.8.4.2.5 Test procedure

For 40GBASE-CR4 or 100GBASE-CR10 testing, the pattern generator is first configured to transmit the training pattern defined in 72.6.10.2. During this initialization period, the device under test (DUT) configures the pattern generator equalizer, via transmitter control, to the coefficient settings it would select using the protocol described in 72.6.10 and the receiver will be tuned using its optimization method.

After the pattern generator equalizer has been configured and the receiver tuned, the pattern generator is set to transmit test pattern 3 as defined in 86.8.2. The receiver under test shall meet the target BER listed in Table 85–8. During the tests, the disturbers transmit at their calibrated level and all of the transmitters in the device under test transmit either scrambled idle characters or PRBS31, with the maximum compliant amplitude and equalization turned off (preset condition).

85.8.4.3 Bit error ratio

The receiver shall operate with a BER 10^{-12} or better when receiving a compliant transmit signal, as defined in 85.8.3, through a compliant cable assembly as defined in 85.10.

85.8.4.4 Signaling rate range

A 40GBASE-CR4 and 100GBASE-CR10 receiver shall comply with the requirements of 85.8.4.3 for any signaling rate in the range 10.3125 GBd \pm 100 ppm. The corresponding unit interval is nominally 96.969697 ps.

85.8.4.5 AC coupling

The 40GBASE-CR4 and 100GBASE-CR10 receivers are AC coupled. AC coupling shall be part of the receive function for Style-2 40GBASE-CR4 connectors. For Style-1 40GBASE-CR4 and 100GBASE-CR10 plug connectors, the receive lanes are AC coupled; the coupling capacitors shall be within the plug

connectors. It should be noted that there may be various methods for AC coupling in actual implementations. The low frequency 3 dB cutoff of the AC coupling shall be less than 50 kHz.

It is recommended that the value of the coupling capacitors be 100 nF. This will limit the inrush currents and baseline wander.

85.9 Channel characteristics

The 40GBASE-CR4 and 100GBASE-CR10 channel is defined between TP0 and TP5 to include the transmitter and receiver differential controlled impedance printed circuit board and the cable assembly as illustrated in Figure 85–2. The channel parameters insertion loss, insertion loss deviation (ILD), insertion loss to crosstalk ratio, and the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane are provided informatively in 85A.4 through 85A.7.

85.10 Cable assembly characteristics

The 40GBASE-CR4 and 100GBASE-CR10 cable assembly contains insulated conductors terminated in a connector at each end for use as a link segment between MDIs. This cable assembly is primarily intended as a point-to-point interface of up to 7 m between network ports using controlled impedance cables. All cable assembly measurements are to be made between TP1 and TP4 with cable assembly test fixtures as specified in 85.10.8 and illustrated in Figure 85–13. These cable assembly specifications are based upon twinaxial cable characteristics, but other cable types are acceptable if the specifications of 85.10 are met.

Table 85–9 provides a summary of the cable assembly differential characteristics at 5.15625 GHz and references to the subclauses addressing each parameter.

Table 85–9—Cable assembly differential characteristics summary

Description	Reference	Value	Unit
Maximum insertion loss at 5.15625 GHz	85.10.2	17.04	dB
Minimum insertion loss at 5.15625 GHz		3	dB
Insertion loss deviation at 5.15625 GHz	85.10.3	max = 1.73 min = –1.73	dB
Minimum return loss at 5.15625 GHz	85.10.4	6.66	dB
MDNEXT loss	85.10.5	Equation (85–26)	dB
MDFEXT loss	85.10.6	Equation (85–27)	dB
Maximum integrated crosstalk noise	85.10.7	Equation (85–33)	mV

85.10.1 Characteristic impedance and reference impedance

The nominal differential characteristic impedance of the cable assembly is 100 Ω . The differential reference impedance for cable assembly specifications shall be 100 Ω .

85.10.2 Cable assembly insertion loss

The fitted cable assembly insertion loss $IL_{fitted}(f)$ as a function of frequency f is defined in Equation (85–19).

$$IL_{\text{fitted}}(f) = a_1\sqrt{f} + a_2f + a_4f^2 \quad (\text{dB}) \quad (85-19)$$

where

f is the frequency in GHz
 $IL_{\text{fitted}}(f)$ is the fitted cable assembly insertion loss at frequency f

Given the cable assembly insertion loss measured between TP1 and TP4 is at N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 7500 MHz with a maximum frequency spacing of 10 MHz, the coefficients of the fitted insertion loss are determined using Equation (85-20) and Equation (85-21).

Define the frequency matrix F as shown in Equation (85-20).

$$F = \begin{bmatrix} \sqrt{f_1} & f_1 & f_1^2 \\ \sqrt{f_2} & f_2 & f_2^2 \\ \dots & \dots & \dots \\ \sqrt{f_N} & f_N & f_N^2 \end{bmatrix} \quad (85-20)$$

The polynomial coefficients a_1 , a_2 , and a_4 are determined using Equation (85-21). In Equation (85-21), T denotes the matrix transpose operator and IL is a column vector of the measured insertion loss values, IL_n at each frequency f_n .

$$\begin{bmatrix} a_1 \\ a_2 \\ a_4 \end{bmatrix} = (F^T F)^{-1} F^T IL \quad (85-21)$$

The maximum allowed values of the polynomial coefficients a_1 , a_2 , and a_4 of the fitted cable assembly insertion loss of each pair of the 40GBASE-CR4 and 100GBASE-CR10 in Equation (85-19) and the maximum insertion loss at 5.15625 GHz shall meet the specifications summarized in Table 85-10 unless otherwise noted. The fitted insertion loss corresponding to one example of the maximum insertion loss at 5.15625 GHz and the maximum allowed values of a_1 , a_2 , and a_4 is illustrated in Figure 85-9.

Table 85-10—Maximum cable assembly insertion loss characteristics

Description	Value	Unit
Maximum insertion loss at 5.15625 GHz	17.04 ^a	dB
Maximum fitted insertion loss coefficient a_1	6	dB/ $\sqrt{\text{GHz}}$
Maximum fitted insertion loss coefficient a_2	1	dB/GHz
Maximum fitted insertion loss coefficient a_4	0.08	dB/GHz ²

^aThe limit on the maximum insertion loss at 5.15625 GHz precludes the coefficients a_1 , a_2 , and a_4 from simultaneous maximum values.




Figure 85-9—Example maximum cable assembly insertion loss

85.10.3 Cable assembly insertion loss deviation (ILD)

The cable assembly insertion loss deviation is the difference between the cable assembly insertion loss and the fitted cable assembly insertion loss determined using Equation (85-22).

$$ILD(f) = IL(f) - IL_{fitted}(f) \tag{85-22}$$

where

- f is the frequency in MHz
- $ILD(f)$ is the cable assembly insertion loss deviation at frequency f

The ILD shall be within the region defined by Equation (85-23) and Equation (85-24). This includes the insertion loss of the differential cabling pairs and the cable assembly connectors.

$$ILD(f) \geq ILD_{min}(f) = -0.7 - 0.2 \times 10^{-3} f \text{ (dB)} \tag{85-23}$$

$$ILD(f) \leq ILD_{max}(f) = 0.7 + 0.2 \times 10^{-3} f \text{ (dB)} \tag{85-24}$$

for $50 \text{ MHz} \leq f \leq 7500 \text{ MHz}$

where

- f is the frequency in MHz

The insertion loss deviation limits are illustrated in Figure 85–10.




Figure 85–10—Maximum cable assembly insertion loss deviation

85.10.4 Cable assembly return loss

The return loss of each pair of the 40GBASE-CR4 and 100GBASE-CR10 cable assembly shall meet the values determined using Equation (85–25).

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.05 \leq f < 4.1 \\ 6.3 - 13 \log_{10}(f/5.5) & 4.1 \leq f \leq 10 \end{array} \right\} \text{ (dB)} \tag{85-25}$$

where

- f is the frequency in GHz
- $Return_loss(f)$ is the return loss at frequency f

The minimum cable assembly return loss is illustrated in Figure 85–11.




Figure 85-11—Minimum cable assembly return loss

85.10.5 Cable assembly multiple disturber near-end crosstalk (MDNEXT) loss

Since four lanes or ten lanes are used to transfer data between PMDs, the NEXT that is coupled into a receive lane will be from the four or ten transmit lanes. Multiple Disturber Near-End Crosstalk (MDNEXT) loss is determined using the individual NEXT losses.

MDNEXT loss is determined from the four or ten individual pair-to-pair differential NEXT loss values using Equation (85-26).

$$MDNEXT_loss(f) = -10\log_{10}\left(\sum_{i=0}^{i=3\text{ or }9} 10^{-NL_i(f)/10}\right) \text{ (dB)} \tag{85-26}$$

for $50 \text{ MHz} \leq f \leq 10000 \text{ MHz}$

where

- MDNEXT_loss(f)* is the MDNEXT loss at frequency *f*
- NL_i(f)* is the NEXT loss at frequency *f* of pair combination *i*, in dB
- f* is the frequency in MHz
- i* is the 0 to 3 (pair-to-pair combination) or 0 to 9 (pair-to-pair combination)

85.10.6 Cable assembly multiple disturber far-end crosstalk (MDFEXT) loss

Since four lanes or ten lanes are used to transfer data between PMDs, the FEXT that is coupled into a data carrying lane will be from the three other lanes or nine other lanes in the same direction. MDFEXT loss is specified using the individual FEXT losses. MDFEXT loss is determined from the three or nine individual pair-to-pair differential FEXT loss values using Equation (85-27).

$$MDFEXT_loss(f) = -10\log_{10}\left(\sum_{i=0}^{i=2\text{ or }8} 10^{-NL_i(f)/10}\right) \text{ (dB)} \tag{85-27}$$

for $50 \text{ MHz} \leq f \leq 10000 \text{ MHz}$

where

$MDFEXT_loss(f)$	is the MDFEXT loss at frequency f
$NL_i(f)$	is the FEXT loss at frequency f of pair combination i , in dB
f	is the frequency in MHz
i	is the 0 to 2 (pair-to-pair combination) or 0 to 8 (pair-to-pair combination)

85.10.7 Cable assembly integrated crosstalk noise (ICN)

In order to limit multiple disturber crosstalk noise at a receiver, the cable assembly integrated crosstalk noise (ICN) is specified in relationship to the measured insertion loss. ICN is calculated from the MDFEXT and MDNEXT. Given the multiple disturber near-end crosstalk loss $MDNEXT_loss(f)$ and multiple disturber far-end crosstalk loss $MDFEXT_loss(f)$ measured over N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 10000 MHz with a maximum frequency spacing of 10 MHz, the RMS value of the integrated crosstalk noise shall be determined using Equation (85–28) through Equation (85–32). The RMS crosstalk noise is characterized at the output of a specified receive filter utilizing a specified transmitter waveform and the measured multiple disturber crosstalk transfer functions. The transmitter and receiver filters are defined in Equation (85–28) and Equation (85–29) as weighting functions to the multiple disturber crosstalk in Equation (85–30) and Equation (85–31). The sinc function is defined by $\text{sinc}(x) = \sin(\pi x)/(\pi x)$.

Define the weight at each frequency f_n using Equation (85–28) and Equation (85–29).

$$W_{nt}(f_n) = (A_{nt}^2/f_b) \text{sinc}(f_n/f_b)^2 \left[\frac{1}{1 + (f_n/f_{nt})^4} \right] \left[\frac{1}{1 + (f_n/f_r)^8} \right] \quad (85-28)$$

$$W_{ft}(f_n) = (A_{ft}^2/f_b) \text{sinc}(f_n/f_b)^2 \left[\frac{1}{1 + (f_n/f_{ft})^4} \right] \left[\frac{1}{1 + (f_n/f_r)^8} \right] \quad (85-29)$$

where the equation parameters are given in Table 85–11.

Note that the 3 dB transmit filter bandwidths f_{nt} and f_{ft} are inversely proportional to the 20% to 80% rise and fall times T_{nt} and T_{ft} respectively. The constant of proportionality is 0.2365 (e.g., $T_{nt}f_{nt} = 0.2365$; with f_{nt} in hertz and T_{nt} in seconds). In addition, f_r is the 3 dB reference receiver bandwidth, which is set to 7.5 GHz.

The near-end integrated crosstalk noise σ_{nx} is calculated using Equation (85–30).

$$\sigma_{nx} = \left[2\Delta f \sum_n W_{nt}(f_n) 10^{-MDNEXT_{loss}(f_n)/10} \right]^{1/2} \quad (85-30)$$

The far-end integrated crosstalk noise σ_{fx} is calculated using Equation (85–31).

$$\sigma_{fx} = \left[2\Delta f \sum_n W_{ft}(f_n) 10^{-MDFEXT_{loss}(f_n)/10} \right]^{1/2} \quad (85-31)$$

where Δf is the uniform frequency step of f_n .

The total integrated crosstalk noise σ_x is calculated using Equation (85–32).

$$\sigma_x = \sqrt{\sigma_{nx}^2 + \sigma_{fx}^2} \tag{85-32}$$

The total integrated crosstalk noise for the cable assembly shall be computed using the parameters shown in Table 85-11.

Table 85-11—Cable assembly integrated crosstalk parameters

Description	Symbol	Value	Units
Symbol rate	f_b	10.3125	GBd
Near-end disturber peak differential output amplitude	A_{nt}	600	mV
Far-end disturber peak differential output amplitude	A_{ft}	600	mV
Near-end disturber 20% to 80% rise and fall times	T_{nt}	24	ps
Far-end disturber 20% to 80% rise and fall times	T_{ft}	24	ps

The total integrated crosstalk RMS noise voltage shall meet the values determined by Equation (85-33) illustrated in Figure 85-12.

$$\sigma_{x, ca} \leq \left\{ \begin{array}{ll} 10 & 3 \leq IL \leq 5.3 \\ 12.4 - 0.45IL & 5.3 < IL \leq 17.04 \end{array} \right\} \text{ (mV)} \tag{85-33}$$

where IL is the value of the cable assembly insertion loss in dB at 5.15625 GHz.




Figure 85-12—Integrated crosstalk noise limits

85.10.8 Cable assembly test fixture

The test fixture of Figure 85–13 or its functional equivalent, is required for measuring the cable assembly specifications in 85.10 at TP1 and TP4. TP1 and TP4 are illustrated in Figure 85–2 and Figure 85–13. The test fixture return loss is equivalent to the test fixture return loss specified in 85.8.3.6. The reference test fixture printed circuit board insertion loss is given in Equation (85–34). The effects of differences between the insertion loss of an actual test fixture and the reference insertion loss are to be accounted for in the measurements.

$$IL_{\text{catf}}(f) = 0.0006 + 0.16\sqrt{f} + 0.0587f \text{ (dB)} \tag{85-34}$$

for $0.01 \text{ GHz} \leq f \leq 10 \text{ GHz}$

where

f is the frequency in GHz
 $IL_{\text{catf}}(f)$ is the reference test fixture printed circuit board insertion loss at frequency f




Figure 85–13—Cable assembly test fixtures

85.10.9 Mated test fixtures

The test fixtures of Figure 85–5 and Figure 85–13 are specified in a mated state illustrated in Figure 85–14.




Figure 85–14—Mated test fixtures

85.10.9.1 Mated test fixtures insertion loss

The insertion loss of the mated test fixtures shall meet the values determined using Equation (85–35) and Equation (85–36).

$$IL(f) \geq IL_{MTFmin}(f) = -0.11 + 0.46\sqrt{f} + 0.16f \quad 0.01 \leq f \leq 10 \quad (\text{dB}) \tag{85-35}$$

$$IL(f) \leq IL_{MTEmax}(f) = \begin{cases} 0.029 + 0.861\sqrt{f} + 0.158f & 0.01 \leq f < 5.5 \\ 0.2 + 0.65f & 5.5 \leq f \leq 10 \end{cases} \quad (\text{dB}) \tag{85-36}$$

where

f is the frequency in GHz
 $IL(f)$ is the mated test fixture insertion loss at frequency f

The mated test fixtures insertion loss limits are illustrated in Figure 85–15.




Figure 85–15—Mated test fixtures Insertion loss

85.10.9.2 Mated test fixtures return loss

The return loss of the mated test fixtures measured at either test fixture test interface shall meet the values determined using Equation (85–37).

$$Return_loss(f) \geq \begin{cases} 20 - 2f & 0.01 \leq f < 2.5 \\ 15 & 2.5 \leq f < 5 \\ 13.8 - 28.85\log_{10}(f/5.5) & 5 \leq f \leq 10 \end{cases} \quad (\text{dB}) \tag{85-37}$$

where

f is the frequency in GHz
 $Return_loss(f)$ is the return loss at frequency f

The mated test fixtures return loss is illustrated in Figure 85–16.




Figure 85–16—Mated test fixtures return loss

85.10.9.3 Mated test fixtures common-mode return loss

The common-mode return loss of the mated test fixtures measured at either test fixture test interface shall meet the values determined using Equation (85–38).

$$Return_loss(f) \geq \begin{cases} 12 - 2.8f & 0.01 \leq f < 2.5 \\ 5.2 - 0.08f & 2.5 \leq f \leq 10 \end{cases} \text{ (dB)} \quad (85-38)$$

where

f is the frequency in GHz
 $Return_loss(f)$ is the common-mode return loss at frequency f

The mated test fixtures common-mode return loss is illustrated in Figure 85–17.




Figure 85-17—Common-mode return loss

85.10.9.4 Mated test fixtures common-mode conversion loss

The common-mode conversion loss of the mated test fixtures measured at either test fixture test interface shall meet the values determined using Equation (85-39).

$$Conversion_loss(f) \geq \left\{ \begin{array}{ll} 30 - 2.91f & 0.01 \leq f < 5.5 \\ 14 & 5.5 \leq f \leq 10 \end{array} \right\} \text{ (dB)} \tag{85-39}$$

where

f is the frequency in GHz
 $Conversion_loss(f)$ is the return loss at frequency f

The mated test fixtures common-mode conversion loss is illustrated in Figure 85-18.




Figure 85-18—Common-mode conversion loss

85.10.9.5 Mated test fixtures integrated crosstalk noise

The mated test fixtures integrated crosstalk RMS noise voltages for the single-disturber near-end crosstalk loss and the single-disturber far-end crosstalk loss are determined using Equation (85-28) through Equation (85-32) by substituting the single disturber near-end for the multiple disturber near-end crosstalk loss and the single disturber far-end crosstalk loss for the multiple disturber far-end crosstalk loss. The values of the mated test fixtures integrated crosstalk RMS noise voltages determined using Equation (85-28) through Equation (85-32) for the single-disturber near-end crosstalk loss, the single-disturber far-end crosstalk loss, the multiple disturber near-end crosstalk loss, and the multiple disturber far-end crosstalk loss shall meet the specifications in Table 85-12.

Table 85-12—Mated test fixtures integrated crosstalk noise

Parameter	Value	Units
Near-end integrated crosstalk noise voltage (RMS)	0.7	mV
Far-end integrated crosstalk noise voltage (RMS)	2.5	mV
MDNEXT integrated crosstalk noise voltage (RMS)	1	mV
MDFEXT integrated crosstalk noise voltage (RMS)	3.5	mV

85.10.10 Shielding

The cable assembly shall provide Class 2 or better shielding in accordance with IEC 61196-1.

85.10.11 Crossover function

The cable assembly shall be wired in a crossover fashion as illustrated in Figure 85-19, with each of the four or ten pairs being attached to the transmitter contacts at one end and the receiver contacts at the other end.




Figure 85-19—Cable assembly wiring

Note that the source lanes (SL) signals SLn<p> and SLn<n> are the positive and negative sides of the transmitters differential signal pairs and the destination lanes (DL) signals DLn<p> and DLn<n> are the positive and negative sides of the receivers differential signal pairs for lane n (n = 0, 1, 2, 3 or n = 0, 1, 2, 3, 4, 5, 6, 7, 8, 9). Signal_Shield_n is the signal shield of the differential signal pair for Lane n.

85.11 MDI specification

This subclause defines the Media Dependent Interface (MDI). The 40GBASE-CR4 and 100GBASE-CR10 PMD, as per 85.7, is coupled to the cable assembly, as per 85.10, by the MDI.

85.11.1 40GBASE-CR4 MDI connectors

Connectors meeting the requirements of 85.11.1.1 (Style-1) or 85.11.1.2 (Style-2) shall be used as the mechanical interface between the PMD of 85.7 and the cable assembly of 85.10. The plug connector shall be used on the cable assembly and the receptacle on the PHY. Style-1 or Style-2 connectors may be used as the MDI interface.

85.11.1.1 Style-1 40GBASE-CR4 MDI connectors

The connector for each end of the cable assembly shall be the quad small form factor pluggable (QSFP+) with the mechanical mating interface defined by SFF-8436 Rev 3.4 and illustrated in Figure 85-20. The MDI connector shall be the quad small form factor pluggable (QSFP+) receptacle with the mechanical mating interface defined by SFF-8436 Rev 3.4 and illustrated in Figure 85-21. These connectors have contact assignments matching that in Table 85-13, and electrical performance consistent with the signal quality and electrical requirements of 85.8 and 85.9.




Figure 85-20—Example Style-1 cable assembly plug




Figure 85–21—Example Style-1 MDI board receptacle

The Style-1 MDI connector of the 40GBASE-CR4 PMD comprises 38 signal connections. The Style-1 40GBASE-CR4 MDI connector contact assignments shall be as defined in Table 85–13.

Table 85–13—Style-1 40GBASE-CR4 lane to MDI connector contact mapping

Tx lane	MDI connector contact	Rx lane	MDI connector contact
signal gnd	S1	signal gnd	S13
SL1<n>	S2	DL2<p>	S14
SL1<p>	S3	DL2<n>	S15
signal gnd	S4	signal gnd	S16
SL3<n>	S5	DL0<p>	S17
SL3<p>	S6	DL0<n>	S18
signal gnd	S7	signal gnd	S19
SL2<p>	S33	DL1<n>	S21
SL2<n>	S34	DL1<p>	S22
signal gnd	S35	signal gnd	S23
SL0<p>	S36	DL3<n>	S24
SL0<n>	S37	DL3<p>	S25
signal gnd	S38	signal gnd	S26

NOTE—Although the 40GBASE-CR4 Style-1 MDI supports 38 connections only the transmitter and receiver contact assignments are specified.

85.11.1.1.1 Style-1 AC coupling

For Style-1 40GBASE-CR4 plug connectors the receive lanes are AC coupled; the coupling capacitors are contained within the plug connectors as specified in 85.8.4.5.

85.11.1.2 Style-2 40GBASE-CR4 MDI connectors

The connector for each end of the cable assembly shall be the latch-type plug with the mechanical mating interface defined by IEC 61076-3-113 and illustrated in Figure 85–22. The MDI connector shall be the latch-type receptacle with the mechanical mating interface defined by IEC 61076-3-113 and illustrated in Figure 85–23. These connectors have a pinout matching that in Table 85–14, and electrical performance consistent with the signal quality and electrical requirements of 85.8 and 85.9. Note that support of compatibility with 10GBASE-CX4 is at the Style-2 40GBASE-CR4 MDI.




Figure 85–22—Example Style-2 cable assembly plug




Figure 85–23—Example Style-2 MDI board receptacle

85.11.1.2.1 Style-2 40GBASE-CR4 Connector pin assignments

The MDI connector of the PMD comprises sixteen signal connections, eight signal shield connections, and one link shield connection. The 40GBASE-CR4 MDI connector pin assignments shall be as defined in Table 85–14.

Table 85–14—Style-2 40GBASE-CR4 lane to MDI connector pin mapping

Rx lane	MDI connector pin	Tx lane	MDI connector pin
DL0<p>	S1	SL0<p>	S16
DL0<n>	S2	SL0<n>	S15
DL1<p>	S3	SL1<p>	S14
DL1<n>	S4	SL1<n>	S13
DL2<p>	S5	SL2<p>	S12
DL2<n>	S6	SL2<n>	S11
DL3<p>	S7	SL3<p>	S10
DL3<n>	S8	SL3<n>	S9
Signal Shield	G1	Signal Shield	G5
Signal Shield	G2	Signal Shield	G6
Signal Shield	G3	Signal Shield	G7
Signal Shield	G4	Signal Shield	G8
—	—	Link Shield	G9

85.11.2 100GBASE-CR10 MDI connectors

The connector for each end of the cable assembly shall be the plug with the mechanical mating interface defined in SFF-8642 Rev 2.4 and illustrated in Figure 85–24. The MDI connector shall be the receptacle with the mechanical mating interface defined by SFF-8642 Rev 2.4 and illustrated in Figure 85–25. These connectors have contact assignments matching that in Table 85–15, and electrical performance consistent with the signal quality and electrical requirements of 85.8 and 85.9.




Figure 85–24—Example cable assembly plug




Figure 85–25—Example MDI board receptacle

The MDI connector of the PMD comprises 84 connections. The 100GBASE-CR10 PMD MDI connector contact assignments shall be as defined in Table 85–15.

Table 85–15—100GBASE-CR10 lane to MDI connector contact mapping

Tx lane	MDI connector contact	Tx lane	MDI connector contact	Rx lane	MDI connector contact	Rx lane	MDI connector contact
signal gnd	A1	signal gnd	B1	signal gnd	C1	signal gnd	D1
SL0<p>	A2	—	B2	DL0<p>	C2	—	D2
SL0<n>	A3	—	B3	DL0<n>	C3	—	D3
signal gnd	A4	signal gnd	B4	signal gnd	C4	signal gnd	D4
SL2<p>	A5	SL1<p>	B5	DL2<p>	C5	DL1<p>	D5
SL2<n>	A6	SL1<n>	B6	DL2<n>	C6	DL1<n>	D6
signal gnd	A7	signal gnd	B7	signal gnd	C7	signal gnd	D7
SL4<p>	A8	SL3<p>	B8	DL4<p>	C8	DL3<p>	D8
SL4<n>	A9	SL3<n>	B9	DL4<n>	C9	DL3<n>	D9
signal gnd	A10	signal gnd	B10	signal gnd	C10	signal gnd	D10
SL6<p>	A11	SL5<p>	B11	DL6<p>	C11	DL5<p>	D11
SL6<n>	A12	SL5<n>	B12	DL6<n>	C12	DL5<n>	D12
signal gnd	A13	signal gnd	B13	signal gnd	C13	signal gnd	D13
SL8<p>	A14	SL7<p>	B14	DL8<p>	C14	DL7<p>	D14
SL8<n>	A15	SL7<n>	B15	DL8<n>	C15	DL7<n>	D15
signal gnd	A16	signal gnd	B16	signal gnd	C16	signal gnd	D16
—	A17	SL9<p>	B17	—	C17	DL9<p>	D17
—	A18	SL9<n>	B18	—	C18	DL9<n>	D18
signal gnd	A19	signal gnd	B19	signal gnd	C19	signal gnd	D19

NOTE—Although the 100GBASE-CR10 MDI supports 84 connections only the transmitter and receiver contact assignments are specified.

85.11.2.1 100GBASE-CR10 MDI AC coupling

For 100GBASE-CR10 plug connectors, the receive lanes are AC coupled; the coupling capacitors are contained within the plug connectors as specified in 85.8.4.5.

85.11.3 Electronic keying

Electronic keying can be used to enable the detection of Style-1 40GBASE-CR4 MDI connectors or 100GBASE-CR10 MDI cable assembly plugs versus fiber modules or no modules present. Specifications of electronic keying are beyond the scope of this standard.

85.12 Environmental specifications

All equipment subject to this clause shall conform to the applicable requirements of 14.7.

85.13 Protocol implementation conformance statement (PICS) proforma for Clause 85, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10¹⁵

85.13.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 85 Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

85.13.2 Identification

85.13.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1— Required for all implementations. NOTE 2— May be completed as appropriate in meeting the requirements for the identification. NOTE 3— The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

85.13.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 85, Physical Medium Dependent (PMD) sublayer and baseband medium, type 40GBASE-CR4 and 100GBASE-CR10
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

¹⁵Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

85.13.3 PICS proforma tables for 40GBASE-CR4 and 100GBASE-CR10 PMDs and baseband medium

85.13.4 Major capabilities/options

Item ^a	Feature	Subclause	Value/Comment	Status	Support
XLGMII	XLGMII	85.1	Interface is supported	O	Yes [] No []
CGMII	CGMII	85.1	Interface is supported	O	Yes [] No []
XLAUI	XLAUI	85.1		O	Yes [] No []
CAUI	CAUI	85.1		O	Yes [] No []
CR4	40GBASE-CR4 PMD	85.1	Can operate as 40GBASE-CR4 PMD	O.1	Yes []
CR10	100GBASE-CR10 PMD	85.1	Can operate as 100GBASE-CR10 PMD	O.1	Yes []
PCS	Support of 40GBASE-R PCS	85.1		CR4:M	Yes []
PCS	Support of 100GBASE-R PCS	85.1		CR10:M	Yes []
PMA	Support of 40GBASE-R PMA	85.1		CR4:M	Yes []
PMA	Support of 100GBASE-R PMA	85.1		CR10:M	Yes []
FEC	Forward error correction	85.1	Device implements BASE-R Forward Error Correction	O	Yes [] No []

Item ^a	Feature	Subclause	Value/Comment	Status	Support
AN	Auto-negotiation	85.1	Device implements Auto-Negotiation	M	Yes []
DC	Delay constraints	85.4	Device conforms to delay constraints specified in 85.4	M	Yes []
DSC	Skew constraints	85.5	Device conforms to Skew and Skew Variation constraints specified in 85.5	M	Yes []
*MD	MDIO capability	85.6	Registers and interface supported	O	Yes [] No []
*CBL	Cable assembly	85.10	Items marked with CBL include cable assembly specifications not applicable to a PHY manufacturer	O	Yes [] No []
CA401	40GBASE-CR4 Style-1 cable assembly	85.10	Cable assembly supports 40GBASE-CR4 Style-1	CBL:O.3	Yes [] No []
CA402	40GBASE-CR4 Style-2 cable assembly	85.10	Cable assembly supports 40GBASE-CR4 Style-2	CBL:O.3	Yes [] No []
CA100	100GBASE-CR10 cable assembly	85.10	Cable assembly supports 100GBASE-CR10	CBL:O.3	Yes [] No []
MDIST1	Style-1 MDI connector	85.11.1.1	40GBASE-CR4 device uses Style-1 MDI	O:2	Yes [] N/A []
MDIST2	Style-2 MDI connector	85.11.1.2	40GBASE-CR4 device uses Style-2 MDI	O:2	Yes [] N/A []

^aA “*” preceding an “Item” identifier indicates there are other PICS that depend on whether or not this item is supported.

85.13.4.1 PMD functional specifications

Item	Feature	Sub clause	Value/Comment	Status	Support
PF1	Transmit function	85.7.2	Converts four or ten logical bit streams into four or ten separate electrical streams	M	Yes []
PF2	Transmit function	85.7.2	Conveys four or ten bits from PMD service interface to MDI lanes	M	Yes []
PF3	Transmitter signal	85.7.2	A positive differential voltage corresponds to tx_bit = one	M	Yes []
PF4	Receive function	85.7.3	Converts four or ten electrical signal streams from the MDI into four or ten logical bit streams	M	Yes []
PF5	Receive function	85.7.3	Conveys four or ten logical bit streams from the MDI lanes to the PMD service interface	M	Yes []
PF6	Receiver signal	85.7.3	A positive differential voltage corresponds to rx_bit = one	M	Yes []
PF7	Global PMD Signal Detect function	85.7.4	Report state via PMD:IS_SIGNAL.indication (SIGNAL_DETECT)	M	Yes []
PF8	Signal_Detect value	85.7.4	Set to FAIL following reset	M	Yes []
PF9	Signal_Detect value	85.7.4	Set to OK on completion of training	M	Yes []
PF10	Signal_Detect value	85.7.4	Set to OK if training disabled by management	M	Yes []
PF11	Global_PMD_transmit_disable	85.7.6	Disables all transmitters by forcing a constant output state	O	Yes [] No []
PF12	Global_PMD_transmit_disable	85.7.6	Loopback not affected	O	Yes [] No []
PF13	Lane-by-lane_PMD_transmit_disable	85.7.7	Allows each lane transmitter to be selectively disabled	O	Yes [] No []
PF14	Lane-by-lane_PMD_transmit_disable	85.7.7	Disables transmitters by forcing a constant output state	O	Yes [] No []
PF15	Lane-by-lane_PMD_transmit_disable	85.7.7	Loopback not affected	O	Yes [] No []
PF16	Loopback	85.7.8	Loopback function provided	M	Yes []
PF18	PMD fault function	85.7.9	Mapped to bit 1.1.7 as listed in Table 85-3	M	Yes []
PF19	PMD control function	85.7.12	Pattern described in 72.6.10 different for each of the lanes	M	Yes []

85.13.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	Global_PMD_signal_detect	85.7.4	Set to the value described in 45.2.1.9.5	O	Yes []
MF2	Global_PMD_signal_detect	85.7.7	Set defined by the training state diagram in Figure 72-5	O	Yes [] No []
MF3	Lane-by-Lane Signal Detect function	85.7.5	Sets PMD_signal_detect_n values on a lane-by-lane basis per requirements of 85.7.4	MD:M	Yes [] N/A []
MF4	PMD_transmit_fault function	85.7.10	Mapped to the PMD_transmit_fault bit as specified in 45.2.1.7.4	MD:M	Yes [] N/A []
MF5	PMD_receive_fault function	85.7.11	Contributes to the PMA/PMD receive fault bit as specified in 45.2.1.7.5	MD:M	Yes [] N/A []

85.13.4.3 Transmitter specifications

Item	Feature	Subclause	Value/Comment	Status	Support
DS1	Meets specifications at TP2	85.8.3	Unless otherwise noted per Table 85-5	M	Yes []
DS2	Test load	85.8.3.6	100 Ω differential load with return loss using Equation (85-1)	M	Yes []
DS3	Test fixture return loss	85.8.3.6	Equation (85-37)	M	Yes []
DS4	Test fixture insertion loss	85.8.3.7	Per Equation (85-15)	M	Yes []
DS5	Signaling rate, per lane	85.8.3.8	10.3125 GBd \pm 100 ppm	M	Yes []

85.13.4.4 Receiver specifications

Item	Feature	Subclause	Value/Comment	Status	Support
RS1	Receiver tolerance	85.8.4.2	BER of better than 10^{-12}	M	Yes []
RS2	Receiver tolerance	85.8.4.2	Maximum fitted insertion loss coefficients	M	Yes []
RS3	Receiver tolerance	85.8.4.2	MDNEXT crosstalk noise	M	Yes []
RS4	Receiver tolerance	85.8.4.2	Pattern generator output amplitude	M	Yes []
RS5	Receiver tolerance	85.8.4.2	Pattern generator jitter specification	M	Yes []
RS6	Bit Error Ratio	85.8.4.2.5	BER of better than 10^{-12}	M	Yes []
RS7	Meets specifications at TP3	85.8.4	Unless otherwise noted per Table 85-7	M	Yes []
RS8	Signaling rate, per lane	85.8.4.4	10.3125 GBd \pm 100 ppm	M	Yes []
RS9	AC coupling	85.8.4.5	3 dB cutoff	M	Yes []

85.13.4.5 Cable assembly specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CA1	Differential reference impedance	85.10.1	100 Ω	CBL:M	Yes [] N/A []
CA2	Insertion loss	85.10.2	Per Equation (85–19) and Table 85–10	CBL:M	Yes [] N/A []
CA3	Insertion loss deviation	85.10.3	Per Equation (85–23) and Equation (85–24)	CBL:M	Yes [] N/A []
CA4	Return loss	85.10.4	Per Equation (85–25)	CBL:M	Yes [] N/A []
CA5	Integrated crosstalk noise	85.10.7	Integrated crosstalk noise voltage per Equation (85–28) through Equation (85–33) and Table 85–11	CBL:M	Yes [] N/A []
CA6	Cable assembly test fixture insertion loss	85.10.8	Per Equation (85–34)	CBL:M	Yes [] N/A []
CA7	Mated test fixture insertion loss	85.10.9.1	Per Equation (85–35) and Equation (85–36)	CBL:M	Yes [] N/A []
CA8	Mated test fixture return loss	85.10.9.2	Per Equation (85–39)	CBL:M	Yes [] N/A []
CA9	Mated test fixtures integrated crosstalk noise	85.10.9.5	Per Equation (85–28), through Equation (85–32) and Table 85–12	CBL:M	Yes [] N/A []
CA10	Shielding	85.10.10	Class 2 or better in accordance with IEC 61196-1	CBL:M	Yes [] N/A []
CA11	Crossover function	85.10.11	Per Figure 85–19	CBL:M	Yes [] N/A []
CA12	Cable assembly connector type	85.11.1	40GBASE-CR4 Style-1 plug (SFF-8436 plug)	CA401:M	Yes [] N/A []
CA13	Pin assignments	85.11.1.1	Per Table 85–13		Yes [] N/A []
CA14	Cable assembly connector type	85.11.1	IEC 61076-3-113 latch-type plug	CA402:M	Yes [] N/A []
CA15	Pin assignments	85.11.1.2.1	Per Table 85–14		Yes [] N/A []
CA16	Cable assembly connector type	85.11.2	100GBASE-CR10 plug (SFF-8642 plug)	CA100:M	Yes [] N/A []
CA17	Pin assignments	85.11.2	Per Table 85–15		Yes [] N/A []
CA18	AC coupling	85.8.4.5	3 dB cutoff	M	Yes []

85.13.4.6 MDI connector specifications

Item	Feature	Subclause	Value/Comment	Status	Support
MDC1	MDI connector type	85.11.1.1	40GBASE-CR4 Style-1 receptacle (SFF-8436 receptacle)	CR4*MDIST1:M	Yes [] N/A []
MDC2	MDI connector type	85.11.1.2	IEC 61076-3-113 latch-type receptacle	CR4*MDIST2:M	Yes []
MDC3	MDI connector type	85.11.2.1	100GBASE-CR10 receptacle (SFF-8642 receptacle)	CR10:M	Yes [] N/A []

85.13.4.7 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	Environmental specifications	85.12		M	Yes []

86. Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10

86.1 Overview

This clause specifies the 40GBASE–SR4 PMD and the 100GBASE–SR10 PMD together with the multimode fiber medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 86–1, to the medium through the MDI, and optionally to the management functions that are accessible through the management interface defined in Clause 45, or equivalent.

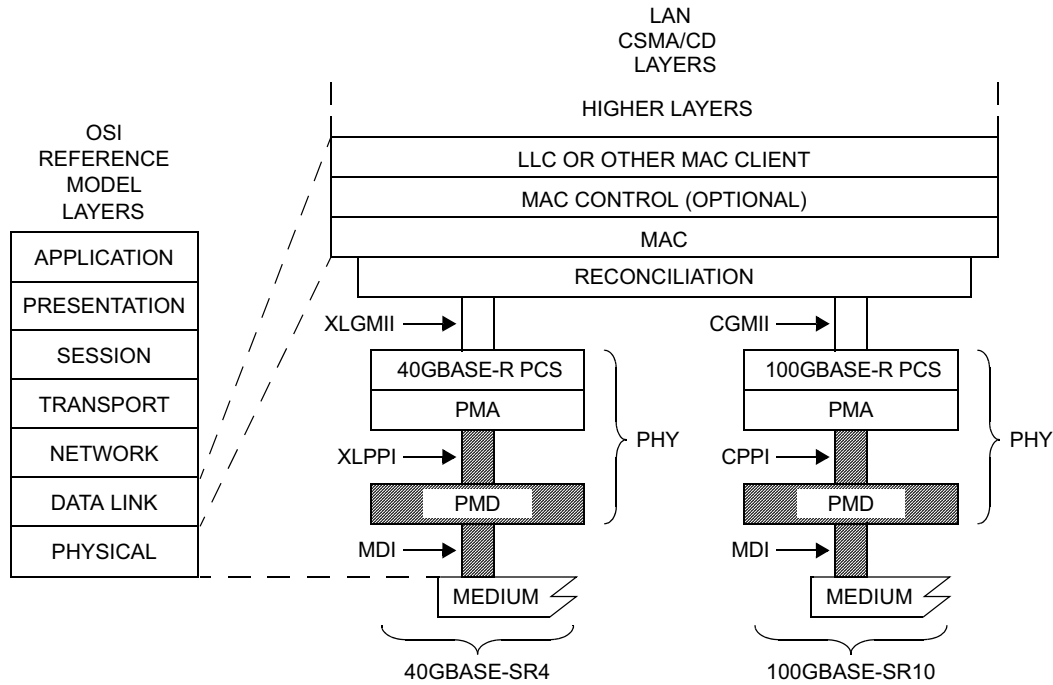
Table 86–1—Physical Layer clauses associated with the 40GBASE-SR4 and 100GBASE-SR10 PMDs

Associated clause	40GBASE–SR4	100GBASE–SR10
81—RS	Required	Required
81—XLGMII ^a	Optional	Not applicable
81—CGMII ^a	Not applicable	Optional
82—PCS for 40GBASE-R	Required	Not applicable
82—PCS for 100GBASE-R	Not applicable	Required
83—PMA for 40GBASE-R4	Required	Not applicable
83—PMA for 100GBASE-R10	Not applicable	Required
83A—XLAUI ^b	Optional	Not applicable
83A—CAUI ^b	Not applicable	Optional
83B—Chip to module XLAUI ^b	Optional	Not applicable
83B—Chip to module CAUI ^b	Not applicable	Optional
86A—XLPPi	Optional	Not applicable
86A—CPPI	Not applicable	Optional

^a XLGMII and CGMII are optional interfaces. However, if the appropriate interface is not implemented, a conforming implementation must behave functionally as though the RS, and XLGMII or CGMII, were present.

^b If XLAUI or CAUI is present, there is at least a PMA between the XLAUI or CAUI and the PMD.

Figure 86–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2. Further relevant information may be found in Clause 1 (terminology and conventions, references, definitions and abbreviations) and Annex A (bibliography, referenced as [B1], [B2], etc.).



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 CPPI = 100 Gb/s PARALLEL PHYSICAL INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
 XLPPI = 40 Gb/s PARALLEL PHYSICAL INTERFACE
 SR = PMD FOR MULTIMODE FIBER

Figure 86–1—40GBASE-SR4 and 100GBASE-SR10 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 CSMA/CD LAN model

The 40GBASE-SR4 and 100GBASE-SR10 PMD sublayers provide point-to-point 40 Gb/s and 100 Gb/s Ethernet links over four or ten pairs of multimode fiber, up to at least 100 m. Table 86–2 shows the primary attributes of each PMD type.

Table 86–2—Summary of 40GBASE-SR4 and 100GBASE-SR10

PMD Type	40GBASE-SR4	100GBASE-SR10	Unit
Fiber type	50/125 μm multimode, type A1a.2 ^a (OM3) or OM4 ^b		
Number of fiber pairs	4	10	
Nominal wavelength	850		nm
Required operating range	0.5 to 100 for OM3		m
	0.5 to 150 for OM4 ^c		
Signaling rate, each lane	10.3125 ±100 ppm		GBd

^a Type A1a.2 (OM3) specified in IEC 60793-2-10. See 86.10.2.1.

^b OM4 specified in TIA-492AAAD. See 86.10.2.1.

^c This is an engineered link with maximum 1 dB connection and splice loss.

40GBASE-SR4 uses four identical lanes, while 100GBASE-SR10 uses ten of the same lanes. In this clause, where there are four or ten items (depending on PMD type) such as lanes, the items are numbered from 0 to $n - 1$, and an example item is numbered i . Thus n is 4 or 10.

The connection to the PMA may use the optional physical instantiation of the PMD service interface called XLPPi (four lanes, for 40GBASE-SR4) or CPPI (ten lanes, for 100GBASE-SR10). The term “nPPI” is used to denote either XLPPi or CPPI, or both.

This clause is arranged as follows: following the overview and an abstract description of the PMD service interface, delay and Skew specifications, control and status variables and registers, a block diagram and high-level specification of the PMD functions, and lane assignments, 86.7 (four parts) contains the optical specifications for 40GBASE-SR4 and 100GBASE-SR10. 86.8 defines test points and optical and dual-use parameters. 86.9 addresses safety, installation, environment and labeling, 86.10 defines the optical channel, and 86.11 contains the PICS. Annex 86A contains the electrical specifications for XLPPi and CPPI, electrical compliance boards, definitions of electrical parameters, and a recommended PCB response for the host (PMA).

86.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-SR4 and 100GBASE-SR10 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation, although an optional implementation of the PMD service interface, the Parallel Physical Interface (nPPI), is specified in 86A.4.1 and 86A.4.2. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_ i.request
PMD:IS_UNITDATA_ i.indication
PMD:IS_SIGNAL.indication
```

The 40GBASE-SR4 PMD has four parallel bit streams, hence $i = 0$ to 3 for 40GBASE-SR4 and the 100GBASE-SR10 PMD has ten parallel bit streams, hence $i = 0$ to 9 for 100GBASE-SR10.

The PMA (or the PMD) continuously sends four or ten parallel bit streams to the PMD (or the PMA), one per lane, each at a nominal signaling rate of 10.3125 GBd (see 83.4).

Upon receipt of the PMD:IS_UNITDATA_ i .request primitive, the PMD converts the specified streams of bits into the appropriate signals on the MDI.

The PMD:IS_UNITDATA_ i .indication primitive corresponds to one of the signals received from the MDI. This primitive is received by the client (the PMA), as described in 83.4.

The PMD:IS_SIGNAL.indication(SIGNAL_DETECT) primitive is generated by the PMD to report the parameter SIGNAL_DETECT, which indicates the status of the signals being received from the MDI (see 86.5.4). There is one parameter and primitive, reporting on all the lanes as a group. This primitive is received by the client (the PMA) as PMD:IS_SIGNAL.indication(SIGNAL_OK), as described in 83.4.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the 10^{-12} BER objective.

86.3 Delay and Skew

86.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 40GBASE-SR4 PMD including 2 m of fiber in one direction shall be no more than 1024 bit times (2 pause_quanta or 25.6 ns). The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-SR10 PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

86.3.2 Skew and Skew Variation constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–4 and Figure 80–5. Skew points as they relate to the nPPI are shown in Figure 86–3.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 86.8.3.1.

86.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If MDIO is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 86–3, and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 86–4.

Table 86–3—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 9 ^a	PMD transmit disable register	1.9.10	PMD_transmit_disable_9
PMD transmit disable 8 ^a to PMD transmit disable 0	PMD transmit disable register	1.9.9 to 1.9.1	PMD_transmit_disable_8 to PMD_transmit_disable_0

^a For 40GBASE-SR4, the highest-numbered six of the ten lane-by-lane transmit disables do not apply.

Table 86–4—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 9 ^a	PMD receive signal detect register	1.10.10	PMD_signal_detect_9
PMD receive signal detect 8 ^a to PMD receive signal detect 0	PMD receive signal detect register	1.10.9 to 1.10.1	PMD_signal_detect_8 to PMD_signal_detect_0

^a For 40GBASE-SR4, the highest-numbered six of the ten lane-by-lane signal detects do not apply.

86.5 PMD functional specifications

The 40GBASE-SR4 and 100GBASE-SR10 PMDs perform the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

86.5.1 PMD block diagram

The PMD block diagram is shown in Figure 86–2. Figure 86–3 shows the test points. It is not required that the PMD service interface be exposed or measurable (nPPI as defined in Annex 86A with compliance points TP1, TP1a, TP4, TP4a). However, if it is not, a conforming implementation must behave as though the interface were compliant.

For purposes of system conformance, the PMD sublayer is standardized at the test points described in 86.8.1. The transmit side electrical signal (PMA output and PMD electrical input) is defined at TP1 and TP1a; see Annex 86A. The optical transmit signal is defined at the output end of a 50 μm multimode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all optical transmitter measurements and tests defined in 86.8 are made at TP2. The optical receive signal is defined at the output

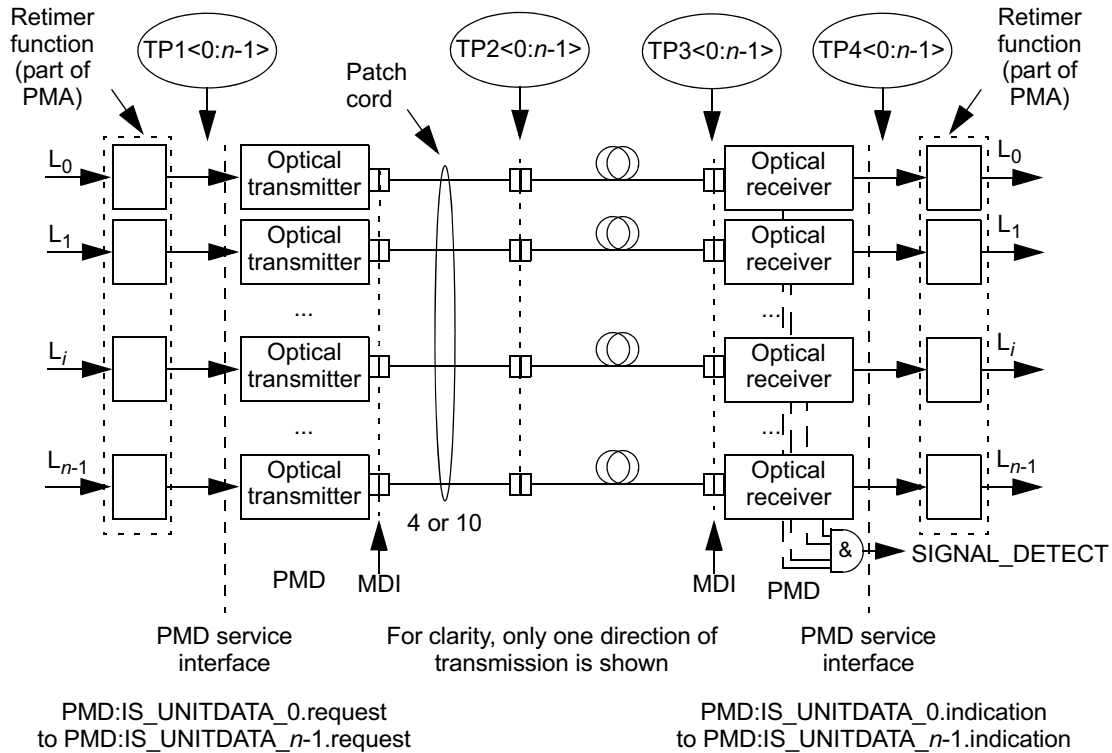


Figure 86–2—Block diagram for 40GBASE–SR4 and 100GBASE–SR10 transmit/receive paths

of the fiber optic cabling (TP3) at the MDI (see 86.10.3). Unless specified otherwise, all optical receiver measurements and tests defined in 86.8 are made at TP3. The receive side electrical signal (PMD electrical output and PMA input) is defined at TP4 and TP4a; see Annex 86A.

86.5.2 PMD transmit function

The PMD Transmit function shall convert the four or ten electronic bit streams requested by the PMD service interface messages `PMD:IS_UNITDATA_0.request` to `PMD:IS_UNITDATA_n-1.request` into the same number of optical signal streams. The optical signal streams are delivered to the MDI, which contains four or ten parallel light paths for transmit, according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to `tx_bit = one`.

86.5.3 PMD receive function

The PMD Receive function shall convert the four or ten parallel optical signal streams received from the MDI into separate electronic bit streams for delivery to the PMD service interface using the messages `PMD:IS_UNITDATA_0.indication` to `PMD:IS_UNITDATA_n-1.indication`, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to `rx_bit = one`.

86.5.4 PMD global signal detect function

The PMD global signal detect function shall report the state of `SIGNAL_DETECT` via the PMD service interface. The `SIGNAL_DETECT` parameter is signaled continuously, while the `PMD:IS_SIGNAL.indication` message is generated when a change in the value of `SIGNAL_DETECT` occurs.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 86–5. The PMD receiver is not required to verify whether a compliant 40GBASE-SR4 or 100GBASE-SR10 signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 86–5—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 \leq -30 dBm	FAIL
For all lanes; [(Optical power at TP3 \geq Minimum OMA, each lane, in Table 86–7) and (compliant 40GBASE-SR4 or 100GBASE-SR10 signal input as appropriate)]	OK
All other conditions	Unspecified

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

86.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect_{*i*}, where *i* represents the lane number in the range 0:*n*-1, shall be continuously set in response to the optical signal on its associated lane, according to the requirements of Table 86–5.

86.5.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

86.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

- a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 86–6.
- b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

86.5.8 PMD lane-by-lane transmit disable function (optional)

The PMD_transmit_disable_ *i* function (where *i* represents the lane number in the range 0:*n*-1) is optional and allows the optical transmitter in each lane to be selectively disabled.

- a) When a PMD_transmit_disable_ *i* variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 86–6.
- b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_ *i* to one, turning off the optical transmitter in each lane.

If the optional PMD_transmit_disable_ *i* function is not implemented in MDIO, an alternative method may be provided to independently disable each transmit lane.

86.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set PMD_fault to one. If the MDIO interface is implemented, PMD_fault shall be mapped to the PMA/PMD fault bit as specified in 45.2.1.2.1.

86.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the PMD_transmit_fault variable to one. If the MDIO interface is implemented, PMD_transmit_fault shall be mapped to the PMA/PMD transmit fault bit as specified in 45.2.1.7.4.

86.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the PMD_receive_fault variable to one. If the MDIO interface is implemented, PMD_receive_fault shall be mapped to the PMA/PMD receive fault bit as specified in 45.2.1.7.5.

86.6 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for 40GBASE-SR4 or 100GBASE-SR10. While it is expected that a PMD will map electrical lane *i* to optical lane *i* and vice versa, there is no need to define the physical ordering of the lanes, as the PCS is capable of receiving the lanes in any arrangement. The positioning of transmit and receive lanes at the MDI is specified in 86.10.3.

86.7 PMD to MDI specifications for 40GBASE-SR4 or 100GBASE-SR10

The required operating range for the 40GBASE-SR4 and 100GBASE-SR10 PMD is defined in Table 86–2. A compliant PMD operates on 50/125 μm multimode fibers according to the specifications of Table 86–14. A PMD which exceeds the operating range requirement while meeting all other optical specifications is considered compliant (e.g., operating at 125 m meets the operating range requirement of 0.5 m to 100 m). The signaling rate for a lane of a 40GBASE-SR4 or 100GBASE-SR10 PMD shall be as defined in Table 86–2. The optical signal at the transmit and receive side of the MDI is specified in 86.7.1 and 86.7.3. The range of optical signals within the optical medium is defined in 86.7.2, and an illustrative link power budget is provided in 86.7.4. Test points are defined in 86.8.1.

86.7.1 Transmitter optical specifications

Each lane of a 40GBASE-SR4 or 100GBASE-SR10 optical transmitter shall meet the specifications of Table 86-6 per the definitions in 86.8.

Table 86-6—40GBASE-SR4 or 100GBASE-SR10 optical transmit characteristics

Description	Type	Value	Unit
Center wavelength	Range	840 to 860	nm
RMS spectral width ^a	Max	0.65	nm
Average launch power, each lane	Max	2.4	dBm
Average launch power, each lane	Min	-7.6	dBm
Optical Modulation Amplitude (OMA), each lane	Max	3	dBm
Optical Modulation Amplitude (OMA), each lane	Min	-5.6 ^b	dBm
Difference in launch power between any two lanes (OMA)	Max	4	dB
Peak power, each lane	Max	4	dBm
Launch power in OMA minus TDP, each lane	Min	-6.5	dBm
Transmitter and dispersion penalty (TDP), each lane	Max	3.5	dB
Extinction ratio	Min	3	dB
Optical return loss tolerance	Max	12	dB
Encircled flux ^c		≥ 86% at 19 μm, ≤ 30% at 4.5 μm	
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3} Hit ratio 5×10^{-5} hits per sample	Spec values	0.23, 0.34, 0.43, 0.27, 0.35, 0.4	
Average launch power of OFF transmitter, each lane	Max	-30	dBm

^a RMS spectral width is the standard deviation of the spectrum.

^b Even if the TDP < 0.9 dB, the OMA (min) must exceed this value.

^c If measured into type A1a.2 50 μm fiber in accordance with IEC 61280-1-4.

86.7.2 Characteristics of signal within, and at the receiving end of, a compliant optical channel

Table 86–7 gives the characteristics of a signal within, and at the receiving end of, a lane of a compliant 40GBASE–SR4 or 100GBASE–SR10 optical channel, and the aggregate signal. A signal with power in OMA or average power not within the ranges given cannot be compliant. However, a signal with power values within the ranges is not necessarily compliant.

Table 86–7—Characteristics of signal within, and at the receiving end of, a compliant optical channel

Description	Minimum		Maximum	Unit
	OM3	OM4		
Fiber type	OM3	OM4		
Total average power for 40GBASE–SR4	–3.5	–3.1	+8.4	dBm
Total average power for 100GBASE–SR10	+0.5	+0.9	+12.4	dBm
Average power, each lane	–9.5	–9.1	+2.4	dBm
Optical Modulation Amplitude (OMA), each lane	–7.5	–7.1	+3	dBm

NOTE—Table 86–7 provides information for diagnostic purposes that is needed by network operators in maintenance. There is no need to assure compliance to it in normal circumstances.

86.7.3 40GBASE–SR4 or 100GBASE–SR10 receiver optical specifications

Each lane of a 40GBASE–SR4 or 100GBASE–SR10 optical receiver shall meet the specifications defined in Table 86–8 per the definitions in 86.8.

Table 86–8—40GBASE–SR4 or 100GBASE–SR10 optical receiver characteristics

Description		Type	Value	Unit
Center wavelength, each lane		Range	840 to 860	nm
Damage threshold ^a		Min	+3.4	dBm
Average power at receiver input, each lane		Max	+2.4	dBm
		Min	–9.5	dBm
Receiver reflectance		Max	–12	dB
Optical Modulation Amplitude (OMA), each lane		Max	3	dBm
Stressed receiver sensitivity in OMA, each lane ^b		Max	–5.4	dBm
Peak power, each lane		Max	4	dBm
Conditions of stressed receiver sensitivity test:				
	Vertical eye closure penalty (VECP) ^c , each lane	—	1.9	dB
	Stressed eye J2 Jitter ^c , each lane	—	0.3	UI
	Stressed eye J9 Jitter ^c , each lane	—	0.47	UI
	OMA of each aggressor lane	—	–0.4	dBm
Receiver jitter tolerance in OMA, each lane ^d		Max	–5.4	dBm
Conditions of receiver jitter tolerance test:				
	Jitter frequency and peak-to-peak amplitude	—	(75, 5)	(kHz, UI)
	Jitter frequency and peak-to-peak amplitude	—	(375, 1)	(kHz, UI)
	OMA of each aggressor lane	—	–0.4	dBm

^a The receiver shall be able to tolerate, without damage, continuous exposure to a modulated optical input signal having this power level on one lane. The receiver does not have to operate correctly at this input power.

^b Measured with conformance test signal at TP3 (see 86.8.4.7).

^c Vertical eye closure penalty and stressed eye jitter are test conditions for measuring stressed receiver sensitivity. They are not characteristics of the receiver. The apparent discrepancy between VECP and TDP is because VECP is defined at eye center while TDP is defined with ± 0.15 UI offsets of the sampling instant.

^d This is a test of the optical receiver's ability to track low-frequency jitter and is inappropriate for any subsystem that does not include a CRU.

86.7.4 40GBASE–SR4 or 100GBASE–SR10 illustrative link power budget

Illustrative power budgets and penalties for 40GBASE–SR4 or 100GBASE–SR10 optical channels are shown in Table 86–9.

Table 86–9—40GBASE–SR4 or 100GBASE–SR10 illustrative link power budgets

Parameter	OM3	OM4	Unit
Effective modal bandwidth at 850 nm	2000 ^a	4700 ^b	MHz•km
Power budget (for maximum TDP)	8.3		dB
Operating distance	0.5 to 100	0.5 to 150	m
Channel insertion loss ^c	1.9	1.5	dB
Allocation for penalties (for maximum TDP) ^d	6.4	6.5	dB
Unallocated margin	0	0.3 ^e	dB
Additional insertion loss allowed	0		dB

^a Per IEC 60793-2-10.

^b Per TIA-492AAAD.

^c The channel insertion loss is calculated using the maximum distances specified in Table 86–2 and cabled optical fiber attenuation of 3.5 dB/km at 850 nm plus an allocation for connection and splice loss given in 86.10.2.2.1.

^d Link penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

^e This unallocated margin is not available for use.

86.8 Definitions of optical and dual-use parameters and measurement methods

In this section, the test points and the parameters applicable to optical signals and for dual use (both optical and electrical) are defined. Test points are defined in 86.8.1, test patterns in 86.8.2, and parameters in 86.8.3 and 86.8.4. Further parameters for electrical use, and the use of electrical compliance boards, are defined in 86A.5.

86.8.1 Test points and compliance boards

Figure 86–3 shows the six test points for 40GBASE-SR4 and 100GBASE-SR10. These are TP1, TP1a, TP2, TP3, TP4, and TP4a; four of these are Skew points SP2, SP3, SP4, and SP5 as shown. Figure 86–3 also shows the substitution of compliance boards for PMD or PMA. This is explained in Annex 86A. Table 86–10 shows the parameters or signals measured at each point, including the electrical compliance points.

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified. A patch cord that connects the MDI transmit side to 4 or 10 individual connectors may be suitable.

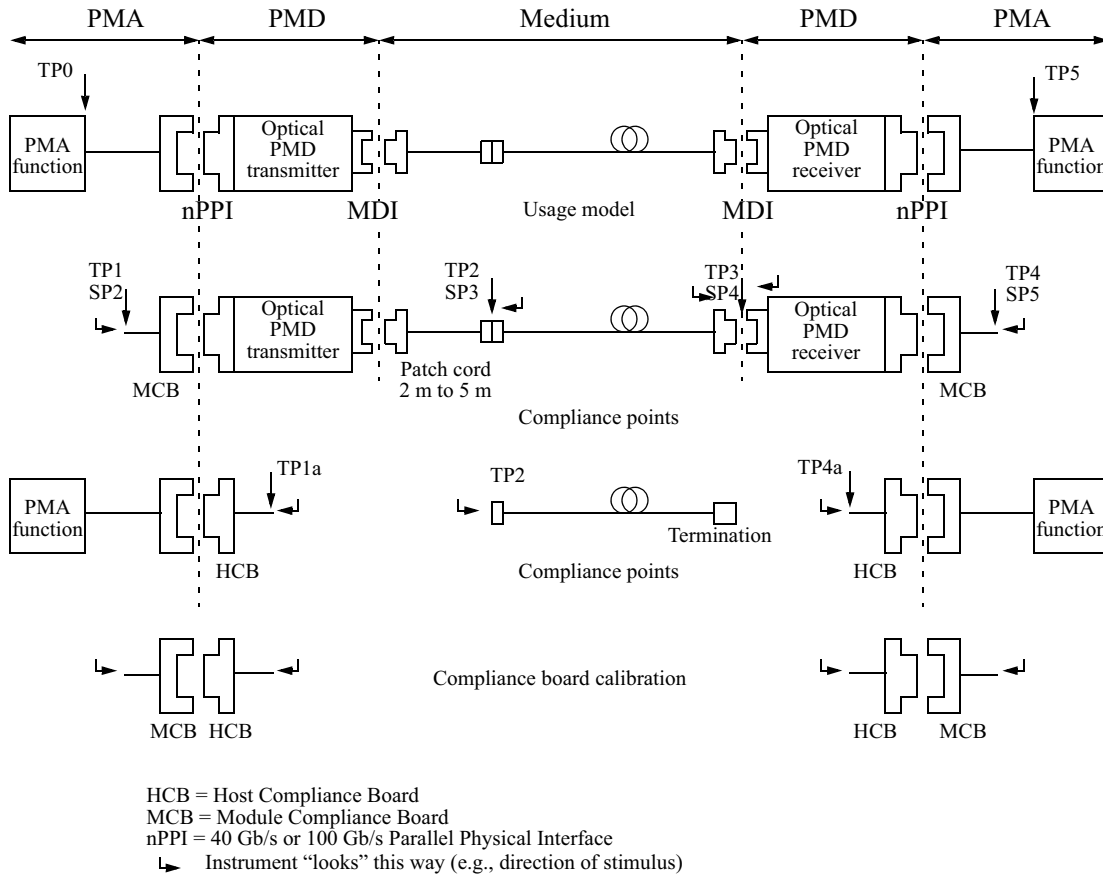


Figure 86-3—Test points for 40GBASE-SR4 and 100GBASE-SR10

Table 86-10—Parameters defined at each test point

Test point	Direction	Parameter
TP1	Looking downstream into PMD transmitter input	PMD transmitter input return loss
TP1a	Looking upstream into PMA transmitter output	PMA transmitter output signal and output return loss, PMD transmitter compliance signal calibration, PMA receiver compliance crosstalk signal calibration
TP2	Looking upstream into optical transmitter output	PMD transmitted signal, PMD transmitter reflectance
	Looking downstream into fiber	Optical return loss, connector reflections
TP3	Looking upstream into fiber	PMD receiver compliance signal
	Looking downstream into optical receiver input	PMD receiver reflectance
TP4	Looking upstream into PMD receiver output	PMD receiver output signal and output return loss, PMA receiver compliance signal calibration
TP4a	Looking downstream into PMA receiver input	PMA receiver input return loss

86.8.2 Test patterns and related subclauses

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 86–11 lists the defined test patterns, and

Table 86–11—Test patterns

Pattern no.	Pattern	Pattern defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.10

Table 86–12 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. As Pattern 3 is more demanding than Pattern 5 (which itself is the same or more demanding than other 40GBASE-R or 100GBASE-R bit streams) an item that is compliant using Pattern 5 is considered compliant even if it does not meet the required limit using Pattern 3. Test patterns for further electrical parameters are given in Table 86A–6.

Table 86–12—Test patterns and related subclauses

Parameter	Pattern	Related subclause
Wavelength, spectral width	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.4.1
Average optical power	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.4.2
Transmitter OMA (modulated optical power)	Square wave or 4	86.8.4.3
Extinction ratio	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.4.5
Transmitted optical waveform (eye mask)	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.2, 86.8.4.6
TDP (transmitter and dispersion penalty)	3 or 5	86.8.4.4
Stressed receiver sensitivity	3 or 5	86.8.4.7
Receiver jitter tolerance	3 or 5	86.8.4.8
Calibration of OMA for receiver tests	Square wave or 4	52.9.9
Vertical eye closure penalty calibration	3 or 5	52.9.9
J2 Jitter	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.3.1
J9 Jitter	3 or 5	86.8.3.3.2

Table 86–12—Test patterns and related subclauses (continued)

Parameter	Pattern	Related subclause
Data Dependent Pulse Width Shrinkage (DDPWS)	4	86A.5.3.4
AC common-mode voltage	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86A.5.3.1
Transition time	Square wave or 4	86A.5.3.3
Electrical waveform (eye mask)	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.2, 86A.5.3.6

86.8.2.1 Multi-lane testing considerations

TDP is defined for each lane, at a BER of 10^{-12} on that lane. Stressed receiver sensitivity, receiver jitter tolerance and host input signal tolerance (in Annex 86A) are defined for an interface BER of 10^{-12} . The interface BER is the average of the four or ten BERs of the receive lanes when they are stressed.

Measurements with Pattern 3 (PRBS31) allow lane-by-lane BER measurements. Measurements with Pattern 5 (scrambled idle) give the interface BER if all lanes are stressed at the same time. If each lane is stressed in turn, the BER is diluted by the three or nine unstressed lanes, and the BER for that stressed lane alone must be found, e.g., by multiplying by 4 or 10 if the unstressed lanes have low BER. To allow TDP measurement with Pattern 5, unstressed lanes for the error detector may be created by setting the power at the reference receivers well above their sensitivities, or by copying the contents of the transmit lanes not under BER test to the error detector by other means. In stressed receiver sensitivity and receiver jitter tolerance measurements, unstressed lanes may be created by setting the power at the receiver under test well above its sensitivity and/or not stressing those lanes with ISI and jitter, or by other means. Each receive lane is stressed in turn while all are operated. All aggressor lanes are operated as specified. To find the interface BER, the BERs of all the lanes when stressed are averaged.

Where relevant, parameters are defined with all co-propagating and counter-propagating lanes operational so that crosstalk effects are included. Where not otherwise specified, the maximum amplitude (OMA or VMA) for a particular situation is used, and for counter-propagating lanes, the minimum transition time is used. Alternative test methods that generate equivalent results may be used. While the lanes in a particular direction may share a common clock, the Tx and Rx directions are not synchronous to each other. If Pattern 3 is used for the lanes not under test using a common clock, there is at least 31 UI delay between the PRBS31 patterns on one lane and any other lane.

86.8.3 Parameters applicable to both electrical and optical signals

86.8.3.1 Skew and Skew Variation

Skew and Skew Variation are defined in 80.5 and are required to remain within the limits given in 86.3.2 over the time that the link is in operation. Skew points as they relate to the nPPI are shown in Figure 86–3. The measurement of Skew and Skew Variation is made by acquiring the data on each lane using a clock and data recovery unit with a high-frequency corner bandwidth and slope as specified in 86.8.3.2. The arrival times of the one to zero transition of the alignment marker sync bits on each lane are then compared. This arrangement ensures that any high-frequency jitter that is present on the signals is not included in the Skew measurement.

86.8.3.2 Eye diagrams

Eye diagrams can be used to assess both electrical and optical signals. The measurement of the electrical eye is defined in 86A.5.3.6 and the measurement of the optical transmitter waveform is defined in 86.8.4.6. Whether electrical or optical eye diagrams, all co-propagating and counter-propagating lanes are active as crosstalk sources, using one of patterns 3, 5, or a valid 40GBASE-R or 100GBASE-R signal. The input lanes of the item under test are receiving signals that are asynchronous to those being output.

Normalized times of 0 and 1 on the unit interval scale are determined by the eye crossing means measured at the average value of the eye pattern. A clock recovery unit (CRU) is used to trigger the oscilloscope for mask measurements, as shown in Figure 52-9. It has a high-frequency corner bandwidth of 4 MHz and a slope of -20 dB/decade. The CRU tracks acceptable levels of low-frequency jitter and wander.

Consideration should be given as to whether a correction is needed for actual instrument properties.

86.8.3.2.1 Eye mask acceptable hit count examples

An example calculation relating hit count to hit ratio for an eye mask measurement using an oscilloscope is detailed below.

If an oscilloscope records 1350 samples/screen, and the time-base is set to 0.2 UI per division with 10 divisions across the screen, and the measurement is continued for 200 waveforms, then a signal that averages less than 6.75 hits is compliant to a hit ratio specification of 5×10^{-5} , i.e.,

$$\frac{5 \times 10^{-5} \times 200 \times 1350}{0.2 \times 10} = 6.75 \quad (86-1)$$

Likewise, if a measurement is continued for 1000 waveforms, then an average of less than 33.75 hits is compliant. An extended measurement is expected to give a more accurate result, and a single reading of 6 hits in 200 waveforms would not give a statistically significant pass or fail.

The hit ratio limit has been chosen to avoid misleading results due to signal and oscilloscope noise.

86.8.3.3 Jitter

J2 Jitter and J9 Jitter are specified with all co-propagating and counter-propagating lanes active as crosstalk sources, using one of patterns 3, 5, or a valid 40GBASE-R or 100GBASE-R signal. The input lanes of the item under test are receiving signals that are asynchronous to those being output.

86.8.3.3.1 J2 Jitter

J2 Jitter is defined as the time interval that includes all but 10^{-2} of the jitter distribution, which is the time interval from the 0.5th to the 99.5th percentile of the jitter histogram. This may be measured using an oscilloscope, or if measured by plotting BER vs. decision time, J2 is the time interval between the two points with a BER of 2.5×10^{-3} . Oscilloscope histograms should include at least 10 000 hits, and should be taken over about 1% of the signal amplitude. Test patterns are given in Table 86-12.

86.8.3.3.2 J9 Jitter

J9 Jitter is defined as the time interval that includes all but 10^{-9} of the jitter distribution. If measured by plotting BER vs. decision time, it is the time interval between the two points with a BER of 2.5×10^{-10} . Test patterns are given in Table 86-12.

86.8.4 Optical parameter definitions

86.8.4.1 Wavelength and spectral width

The wavelength of each optical lane shall be within the range given in Table 86–6 if measured using the method given in TIA–455–127-A. The lane under test is modulated using the test pattern defined in Table 86–12.

86.8.4.2 Average optical power

Average optical power is defined by the methods given in IEC 61280-1-1.

86.8.4.3 Optical Modulation Amplitude (OMA)

OMA shall be as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each optical lane is tested individually. See 86.8.2 for test pattern information.

86.8.4.4 Transmitter and dispersion penalty (TDP)

Transmitter and dispersion penalty (TDP) shall be as defined for 10GBASE-S in 52.9.10 with the following exceptions:

- a) Each optical lane is tested individually with all other lanes in operation.
- b) The test pattern is as defined in Table 86–12.
- c) The transmitter is tested using an optical channel with an optical return loss of 12 dB.
- d) The reference receiver (including the effect of the decision circuit) has a fourth order Bessel-Thomson filter response with a bandwidth of 6.1 GHz. The transversal filter of 52.9.10.3 is not used.
- e) The reference sensitivity S and the measurement P_{DUT} are both measured with the sampling instant displaced from the eye center by ± 0.15 UI. For each of the two cases (early and late), if $P_{DUT}(i)$ is larger than $S(i)$, the $TDP(i)$ for the transmitter under test is the difference between $P_{DUT}(i)$ and $S(i)$, $TDP(i) = P_{DUT}(i) - S(i)$. Otherwise, $TDP(i) = 0$. The TDP is the larger of the two $TDP(i)$.
- f) The test setup illustrated in Figure 52-12 shows the reference method. Other measurement implementations may be used with suitable calibration.
- g) The BER of 10^{-12} is for the lane under test on its own. See 86.8.2.1 for multi-lane pattern considerations.

NOTE—Because practical receivers and decision circuits have noise and timing impairments, the sampling instant offsets have to be calibrated. One method of doing this is via a jitter bathtub method using a known low-jitter signal.

86.8.4.5 Extinction ratio

Extinction ratio is defined by the methods of IEC 61280–2–2 using the test pattern defined in Table 86–12.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 86–12).

86.8.4.6 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram specified in Table 86–6 and shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 86–12 shall meet specifications according to the methods specified in 86.8.4.6.1 with the filter nominal reference frequency f_r of 7.5 GHz and filter tolerances as specified for STM-64 in ITU-T G.691.

86.8.4.6.1 Optical transmitter eye mask

The eye is measured with respect to the mask specified in Table 86–6 and shown in Figure 86–4 using a receiver with the fourth-order Bessel-Thomson response having a transfer function given by Equation (86–2) and Equation (86–3):

$$H(y) = \frac{105}{105 + 105y + 45y^2 + 10y^3 + y^4} \tag{86-2}$$

where:

$$y = 2.114p; \quad p = \frac{j\omega}{\omega_r}; \quad \omega_r = 2\pi f_r; \quad f_r = \text{Reference frequency in GHz} \tag{86-3}$$




Figure 86–4—Transmitter eye mask definition

The Bessel-Thomson receiver is not intended to represent the noise filter used within a compliant optical receiver, but is intended to provide uniform measurement conditions at the transmitter. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

Normalized levels of 0 and 1 represent logic zero and one respectively. These are defined by the means of the lower and upper halves of the central 0.2 UI of the eye.

Further requirements are given in 86.8.3.2.

The transmitter shall achieve a hit ratio lower than the limit of hits per sample specified in the appropriate table or 5×10^{-5} hits per sample if not otherwise specified. “Hits” are the number of samples within the grey areas of Figure 86–4, and the sample count is the total number of samples from 0 UI to 1 UI. Some illustrative examples are provided in 86.8.3.2.1.

Further information on optical eye pattern measurement procedures may be found in IEC 61280-2-2.

86.8.4.7 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 86–8 if measured using the method defined by 52.9.9 with the conformance test signal at TP3 and with the following exceptions:

- a) The reference test procedure for a single lane is defined in 52.9.9. See 86.8.2.1 and below for multi-lane considerations.
- b) The sinusoidal amplitude interferer is replaced by a Gaussian noise generator.
- c) The fourth-order Bessel-Thomson filter is replaced by a low-pass filter followed by a limiter and a fourth-order Bessel-Thomson filter.
- d) The sinusoidal jitter is at a fixed 80 MHz frequency and between 0 and 0.05 UI peak-to-peak amplitude.
- e) The Gaussian noise generator, the amplitude of the sinusoidal jitter, and the Bessel-Thomson filter are adjusted so that the VECP, J2 Jitter and J9 Jitter specifications given in Table 86–8 are simultaneously met (the random noise effects such as RIN, random clock jitter do not need to be minimized).
- f) The pattern for the received compliance signal is specified in Table 86–12.
- g) The interface BER of the PMD receiver is the average of the BER of all receive lanes while stressed and at the specified receive OMA.
- h) Where nPPI or XLAUI/CAUI is exposed, a PMD receiver is considered compliant if it meets the module electrical output specifications at TP4 given in Table 86A–3 for nPPI, or the requirements in Table 83B–3 for XLAUI/CAUI.

Stressed receiver sensitivity is defined with all transmit and receive lanes in operation. Each receive lane is tested in turn while all aggressor receive lanes are operated as specified in Table 86–8. Pattern 3 or Pattern 5, or a valid 40GBASE–R4 or 100GBASE–R10 signal is sent from the transmit section of the receiver under test. The signal being transmitted is asynchronous to the received signal. If Pattern 3 is used for the transmit and receive lanes not under test, there is at least 31 UI delay between the PRBS31 patterns generated on one lane and any other lane.

For 40GBASE-SR4 and 100GBASE-SR10, the relevant BER is the interface BER. The interface BER is the average of the four or ten BERs of the receive lanes when stressed: see 86.8.2.1.

86.8.4.8 Receiver jitter tolerance

Receiver jitter tolerance shall be as defined as in 68.6.11, with the following differences:

- a) The reference test procedure for a single lane is defined in 68.6.11. See 86.8.2.1 for multi-lane considerations.
- b) The pattern to be received is specified in Table 86–12.
- c) The parameters of the signal are specified in Table 86–8 and the power in OMA at the receiver is set to the maximum for receiver jitter tolerance in OMA given in Table 86–8.
- d) Each receive lane is tested in turn while all are operated. All aggressor lanes are operated as specified in Table 86–8.
- e) The receive lanes not being tested are receiving Pattern 3, Pattern 5, or a valid 40GBASE-R4 or 100GBASE-R10 signal.
- f) The transmitter is transmitting one of these signals using all lanes.
- g) The transmitter and the receiver are not synchronous.
- h) The interface BER of the PMD receiver is the average of the BER of all receive lanes when stressed.
- i) The mode-conditioning patch cord suitable for 62.5/125 μm fiber is not used.

86.9 Safety, installation, environment, and labeling

86.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

86.9.2 Laser safety

40GBASE-SR4 and 100GBASE-SR10 optical transceivers shall conform to Class 1M laser requirements as defined in IEC 60825-1 and IEC 60825-2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.¹⁶

86.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

86.9.4 Environment

The 40GBASE-SR4 and 100GBASE-SR10 operating environment specifications are as defined in 52.11, as defined in 52.11.1 for electromagnetic emission, and as defined in 52.11.2 for temperature, humidity, and handling.

86.9.5 PMD labeling

The 40GBASE-SR4 and 100GBASE-SR10 labeling recommendations and requirements are as defined in 52.12.

86.10 Optical channel

The fiber optic cabling (channel) contains 4 or 10 optical fibers for each direction to support 40GBASE-SR4 or 100GBASE-SR10, respectively. The fiber optic cabling interconnects the transmitters at the MDI on one end of the channel to the receivers at the MDI on the other end of the channel. As defined in 86.10.3, the optical lanes appear in defined locations at the MDI but the locations are not assigned specific lane numbers within this standard because any transmitter lane may be connected to any receiver lane.

86.10.1 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 86-5.

The channel shall comply with the specifications in Table 86-13.

¹⁶A host system that fails to meet the manufacturers requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.




Figure 86-5—Fiber optic cabling model

Table 86-13—Fiber optic cabling (channel) characteristics at 850 nm

Description	Type	OM3	OM4	Unit
Operating distance	Max	100	150	m
Cabling Skew	Max	79		ns
Cabling Skew Variation ^a	Max	2.5		ns
Channel insertion loss	Min	0		dB
Channel insertion loss ^b	Max	1.9 ^c	1.5 ^d	dB

^a An additional 300 ps of Skew Variation could be caused by wavelength changes, which are attributable to the transmitter not the channel.

^b These channel insertion loss values include cable, connectors, and splices.

^c 1.5 dB allocated for connection and splice loss.

^d 1 dB allocated for connection and splice loss.

A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, modal dispersion, reflections and losses of all connectors and splices meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with IEC 61280-4-1:2009. As OM4 optical fiber meets the requirements for OM3, a channel compliant to the “OM3” column may use OM4 optical fiber, or a combination of OM3 and OM4.

86.10.2 Characteristics of the fiber optic cabling (channel)

86.10.2.1 Optical fiber cable

The fiber contained within the 40GBASE-SR4 or 100GBASE-SR10 fiber optic cabling shall comply with the specifications and parameters of Table 86-14. A variety of multimode cable types may satisfy these requirements, provided the resulting channel also meets the specifications of Table 86-13. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

86.10.2.2 Optical fiber connection

An optical fiber connection, as shown in Figure 86-5, consists of a mated pair of optical connectors.

Table 86–14—Optical fiber and cable characteristics

Description	OM3 ^a	OM4 ^b	Unit
Nominal core diameter	50		μm
Nominal fiber specification wavelength	850		nm
Effective modal bandwidth (min) ^c	2000	4700	MHz•km
Cabled optical fiber attenuation (max)	3.5		dB/km
Zero dispersion wavelength (λ_0)	$1295 \leq \lambda_0 \leq 1340$		nm
Chromatic dispersion slope (max) (S_0)	0.105 for $1295 \leq \lambda_0 \leq 1310$ and 0.000375×(1590 – λ_0) for $1310 \leq \lambda_0 \leq 1340$		ps/nm ² km

^a IEC 60793-2-10 type A1a.2

^b See TIA-492AAAD

^c When measured with the launch conditions specified in Table 86–6.

86.10.2.2.1 Connection insertion loss

The maximum operating distances are based on an allocation of 1.5 dB or 1 dB total connection and splice loss. For example, these allocations support two connections, each with an insertion loss of 0.75 dB or 0.5 dB respectively. Connections with lower loss characteristics may be used provided the requirements of Table 86–13 are met. However, the loss of a single connection shall not exceed 0.75 dB.

86.10.2.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –20 dB.

86.10.3 Medium Dependent Interface (MDI)

The 40GBASE–SR4 or 100GBASE–SR10 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 86–5). The 40GBASE–SR4 PMD is coupled to the fiber optic cabling through one connector plug into the MDI optical receptacle as shown in Figure 86–6. The 100GBASE–SR10 PMD is coupled to the fiber optic cabling through one or two connector plugs into the MDI optical receptacle(s), depending on choice of implementation, as shown in Figure 86–7. Example constructions of the MDI include the following:

- a) PMD with a connectorized fiber pigtail plugged into an adapter;
- b) PMD receptacle.

86.10.3.1 Optical lane assignments for 40GBASE-SR4

The four transmit and four receive optical lanes of 40GBASE-SR4 shall occupy the positions depicted in Figure 86–6 when looking into the MDI receptacle with the connector keyway feature on top. The interface contains eight active lanes within twelve total positions. The transmit optical lanes occupy the leftmost four positions. The receive optical lanes occupy the rightmost four positions. The four center positions are unused.




Figure 86–6—40GBASE-SR4 optical lane assignments

86.10.3.2 Optical lane assignments for 100GBASE-SR10

The ten transmit and ten receive optical lanes of 100GBASE-SR10 shall occupy the positions depicted in Figure 86–7 when looking into the MDI optical receptacle(s) with the connector keyway feature(s) on top. The single-receptacle Option A is recommended, the two-receptacle Option B and Option C are alternatives. The interface contains 20 active lanes within up to 24 total positions arranged in two rows of 10 or 12 positions. One row is dedicated to transmit optical lanes and the other row to receive optical lanes. For the depicted 12-position rows, the optical lanes occupy the center ten positions of each row with the outermost positions unused.

86.10.3.3 Medium Dependent Interface (MDI) requirements

The MDI adapter or receptacle shall meet the dimensional specifications of IEC 61754-7 interface 7-3, the MPO adapter interface. The plug terminating the optical fiber cabling shall meet the dimensional specifications of IEC 61754-7 interface 7-4, MPO female plug connector flat interface. The MDI shall optically mate with the plug on the optical fiber cabling. Figure 86–8 shows an MPO female plug connector with flat interface, and an MDI as a PMD receptacle using an MPO adapter interface.




Figure 86–8—MPO female plug and MDI as a PMD receptacle using MPO adapter

The MDI shall meet the interface performance specifications of IEC 61753-1-1 and IEC 61753-022-2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 86.5.1, not at the MDI.

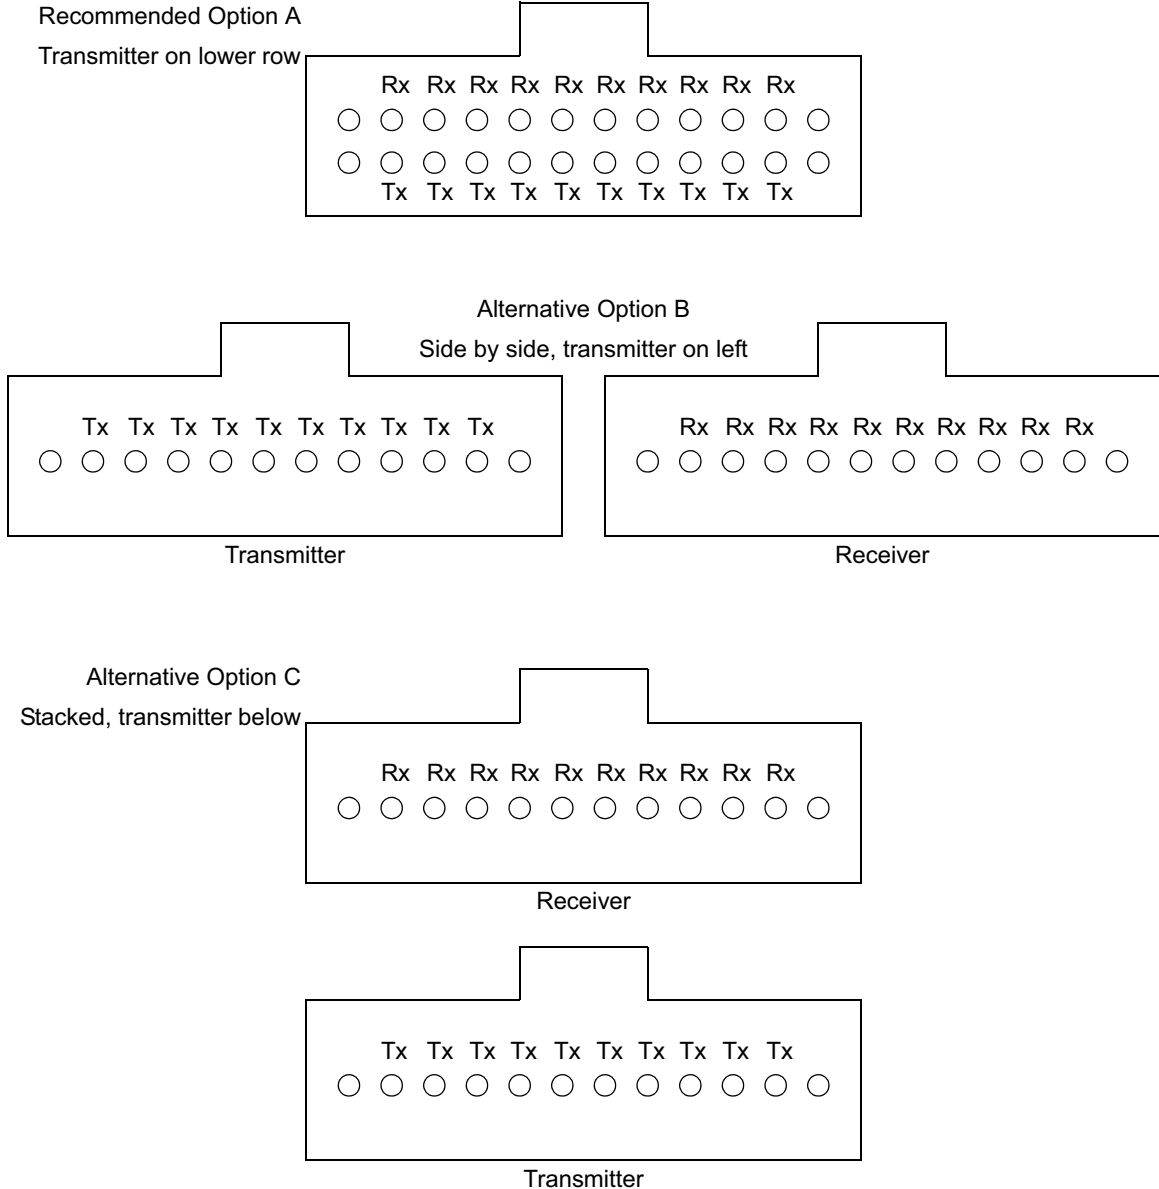


Figure 86-7—100GBASE-SR10 optical lane assignments

86.11 Protocol implementation conformance statement (PICS) proforma for Clause 86, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10¹⁷

86.11.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 86, Physical Medium Dependent sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

86.11.2 Identification

86.11.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
<p>NOTE 1—Required for all implementations.</p> <p>NOTE 2—May be completed as appropriate in meeting the requirements for the identification.</p> <p>NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).</p>	

86.11.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 86, Physical Medium Dependent sublayer and medium, type 40GBASE–SR4 and 100GBASE–SR10
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
<p>Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)</p>	

Date of Statement	
-------------------	--

¹⁷Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

86.11.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*SR4	40GBASE-SR4 PMD	86.7	Can operate as 40GBASE-SR4 PMD	O.1	Yes [] No []
*SR10	100GBASE-SR10 PMD	86.7	Can operate as 100GBASE-SR10 PMD	O.1	Yes [] No []
*INS	Installation / cable	86.10.1	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [] No []
*TP1	Compliance point TP1 exposed and available for testing	86.5.1	See text	O	Yes [] No []
*TP4	Compliance point TP4 exposed and available for testing	86.5.1	See text	O	Yes [] No []
*PIT	nPPI Tx interface	86.1	Uses XLPPi or CPPI host to module (see Annex 86A)	TP1:O	Yes [] No []
*PIR	nPPI Rx interface	86.1	Uses XLPPi or CPPI module to host (see Annex 86A)	TP4:O	Yes [] No []
*MD	MDIO capability	86.4	Registers and interface supported	O	Yes [] No []

86.11.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, types 40GBASE-SR4 and 100GBASE-SR10

86.11.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SF1	Compatible with 40GBASE-R or 100GBASE-R PCS and PMA	86.1	See text	M	Yes []
SF2	Integration with management functions	86.1	See text	O	Yes [] No []
D	Delay constraints	86.3.1	For SR4, max 25.6 ns. For SR10, max 20.48 ns.	M	Yes []
SF3	Skew and Skew Variation at SP3 and SP4	86.3.2	At SP3, less than 54 ns, 600 ps. At SP4, less than 134 ns, 3.4 ns.	M	Yes []
SF4	Skew and Skew Variation at SP5 (TP4)	86.3.2	If measurable, less than 145 ns, 3.6 ns.	TP4: M	Yes [] N/A []
SF5	Transmit function	86.5.2	Conveys bits from PMD service interface to MDI	M	Yes []
SF6	Delivery to the MDI	86.5.2	4 or 10 optical signal streams for delivery to the MDI	M	Yes []
SF7	Mapping between optical signal and logical signal for transmitter	86.5.2	Higher optical power is a one	M	Yes []
SF8	Receive function	86.5.3	Conveys bits from MDI to PMD service interface	M	Yes []
SF9	Conversion of optical signals to electrical signals	86.5.3	For delivery to the PMD service interface	M	Yes []
SF10	Mapping between optical signal and logical signal for receiver	86.5.3	Higher optical power is a one	M	Yes []
SF11	Global Signal Detect function	86.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes []
SF12	Global Signal Detect behavior	86.5.4	Global indicator of the presence of optical signals on all lanes per Table 86-5	M	Yes []
SF13	Lane-by-lane Signal Detect function	86.5.5	Sets PMD_signal_detect _i values on a lane-by-lane basis per Table 86-5	MD:O	Yes [] No [] N/A []
SF14	PMD_reset function	86.5.6	Resets the PMD sublayer	MD:O	Yes [] No [] N/A []

86.11.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
SM1	Management register set	86.4	See 86.4	MD:M	Yes [] N/A []
SM2	Global transmit disable function	86.5.7	Disables all of the optical transmitters with Global_PMD_transmit_disable	MD:O	Yes [] No [] N/A []
SM3	PMD_transmit_disable_ <i>i</i> function	86.5.8	Disables the optical transmitter on the lane associated with PMD_transmit_disable_ <i>i</i>	MD:O	Yes [] No [] N/A []
SM4	PMD lane-by-lane transmit disable	86.5.8	Disables each optical transmitter independently if MD = No	!MD:O	Yes [] No [] N/A []
SM5	PMD_fault function	86.5.9	Sets PMD_fault to one if any local fault is detected	O	Yes [] No []
SM6	PMD_fault mapping to MDIO	86.5.9	Mapped to PMA/PMD fault bit	SM5*MD :M	Yes [] N/A []
SM7	PMD_transmit_fault function	86.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	O	Yes [] No []
SM8	PMD_transmit_fault mapping to MDIO	86.5.10	Mapped to PMA/PMD transmit fault bit	SM7*MD :M	Yes [] N/A []
SM9	PMD_receive_fault function	86.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	O	Yes [] No []
SM10	PMD_receive_fault mapping to MDIO	86.5.11	Mapped to PMA/PMD receive fault bit	SM9*MD :M	Yes [] N/A []

86.11.4.3 Optical specifications for 40GBASE-SR4 or 100GBASE-SR10

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Signaling rate per Table 86-2	86.7	10.3125 ±100 ppm	M	Yes [] N/A []
S2	Optical transmitter meets specifications in Table 86-6	86.7.1	Per definitions in 86.8	M	Yes [] N/A []
S3	Optical receiver meets specifications in Table 86-8	86.7.3	Per definitions in 86.8	M	Yes [] N/A []

86.11.4.4 Definitions of parameters and measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
SOM1	Measurement cable	86.8.1	2 m to 5 m in length	M	Yes []
SOM2	Center wavelength	86.8.4.1	Per TIA-455-127-A	M	Yes []
SOM3	Average optical power	86.8.4.2	Per IEC 61280-1-1	M	Yes []
SOM4	OMA measurements	86.8.4.3	Each lane, per methods of 52.9.5 or 68.6.2	M	Yes []
SOM5	TDP	86.8.4.4	Each lane	M	Yes []
SOM6	Extinction ratio	86.8.4.5	Per IEC 61280-2-2	M	Yes []
SOM7	Transmit eye	86.8.4.6	Per 86.8.4.6.1 and 86.8.3.2, hit ratio lower than specified limit	M	Yes []
SOM8	Stressed receiver conformance	86.8.4.7	See 86.8.4.7	M	Yes []
SOM9	Receiver jitter tolerance	86.8.4.8	Per 68.6.11 as modified	M	Yes []

86.11.4.5 Environmental and safety specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SES1	General safety	86.9.1	Conforms to IEC 60950-1	M	Yes []
SES2	Laser safety	86.9.2	Conforms to Class 1M laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes []
SES3	Electromagnetic interference	86.9.4	As 52.11. Complies with codes for the limitation of electromagnetic interference	M	Yes []

86.11.4.6 Optical channel and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
SOC1	Fiber optic cabling	86.10.1	Specified in Table 86–13	INS:M	Yes [] N/A []
SOC2	Optical fiber characteristics	86.10.2.1	Per Table 86–14	INS:M	Yes [] N/A []
SOC3	Maximum loss per connection	86.10.2.2.1	0.75 dB	INS:M	Yes [] N/A []
SOC4	Maximum discrete reflectance	86.10.2.2.2	Less than –20 dB	INS:M	Yes [] N/A []
SOC5	40G MDI layout	86.10.3.1	Optical lane assignments per Figure 86–6	SR4:M	Yes [] N/A []
SOC6	100G MDI layout	86.10.3.2	Optical lane assignments per Figure 86–7	SR10:M	Yes [] N/A []
SOC7	MDI dimensions	86.10.3.3	Per IEC 61754-7 interface 7-3	M	Yes []
SOC8	Cabling connector dimensions	86.10.3.3	Per IEC 61754-7 interface 7-4	INS:M	Yes [] N/A []
SOC9	MDI mating	86.10.3.3	MDI optically mates with plug on the cabling	M	Yes []
SOC10	MDI requirements	86.10.3.3	Meets IEC 61753-1-1 and IEC 61753-022-2	M INS:M	Yes [] N/A []

87. Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4

87.1 Overview

This clause specifies the 40GBASE-LR4 PMD together with the single-mode fiber medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 87-1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 87-1—Physical Layer clauses associated with the 40GBASE-LR4 PMD

Associated clause	40GBASE-LR4
81—RS	Required
81—XLGMI ^a	Optional
82—PCS for 40GBASE-R	Required
83—PMA for 40GBASE-R	Required
83A—XLAUI	Optional
83B—Chip to module XLAUI	Optional
86A—XLPPi	Optional

^aThe XLGMII is an optional interface. However, if the XLGMII is not implemented, a conforming implementation must behave functionally as though the RS and XLGMII were present.

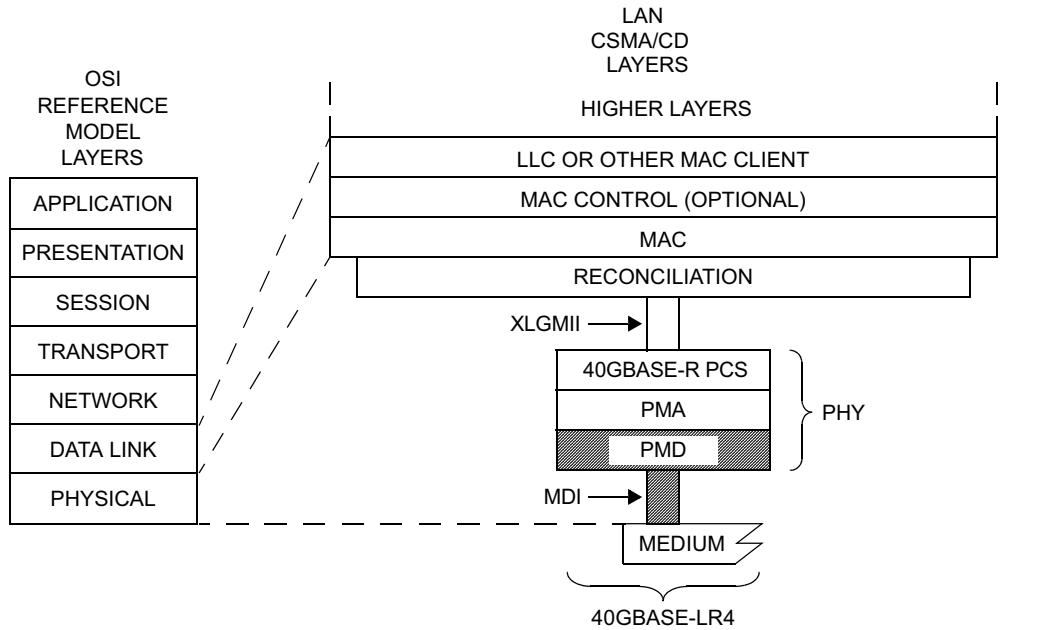
Figure 87-1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

87.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 40GBASE-LR4 PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation, although an optional implementation of the PMD service interface, the 40 Gb/s Parallel Physical Interface (XLPPi), is specified in Annex 86A. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_i.request
PMD:IS_UNITDATA_i.indication
PMD:IS_SIGNAL.indication
```



LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE

PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

LR = PMD FOR SINGLE-MODE FIBER — 10 km

Figure 87–1—40GBASE-LR4 PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 CSMA/CD LAN model

The 40GBASE-LR4 PMD has four parallel bit streams, hence $i = 0$ to 3.

In the transmit direction, the PMA continuously sends four parallel bit streams to the PMD, one per lane, each at a nominal signaling rate of 10.3125 GBd. The PMD converts these streams of bits into appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel bit streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 10.3125 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service primitive defined in 80.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the 10^{-12} BER objective.

87.3 Delay and Skew

87.3.1 Delay constraints

The sum of the transmit and the receive delays at one end of the link contributed by the 40GBASE-LR4 PMD including 2 m of fiber in one direction shall be no more than 1024 bit times (2 pause_quanta or 25.6 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

87.3.2 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80–4 and Figure 80–5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 87.8.2.

87.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 87–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 87–3.

87.5 PMD functional specifications

The 40GBASE–LR4 PMD performs the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

87.5.1 PMD block diagram

The PMD block diagram is shown in Figure 87–2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 87.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 87.11.3). Unless specified otherwise, all receiver measurements and tests defined in 87.8 are made at TP3.

Table 87–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0

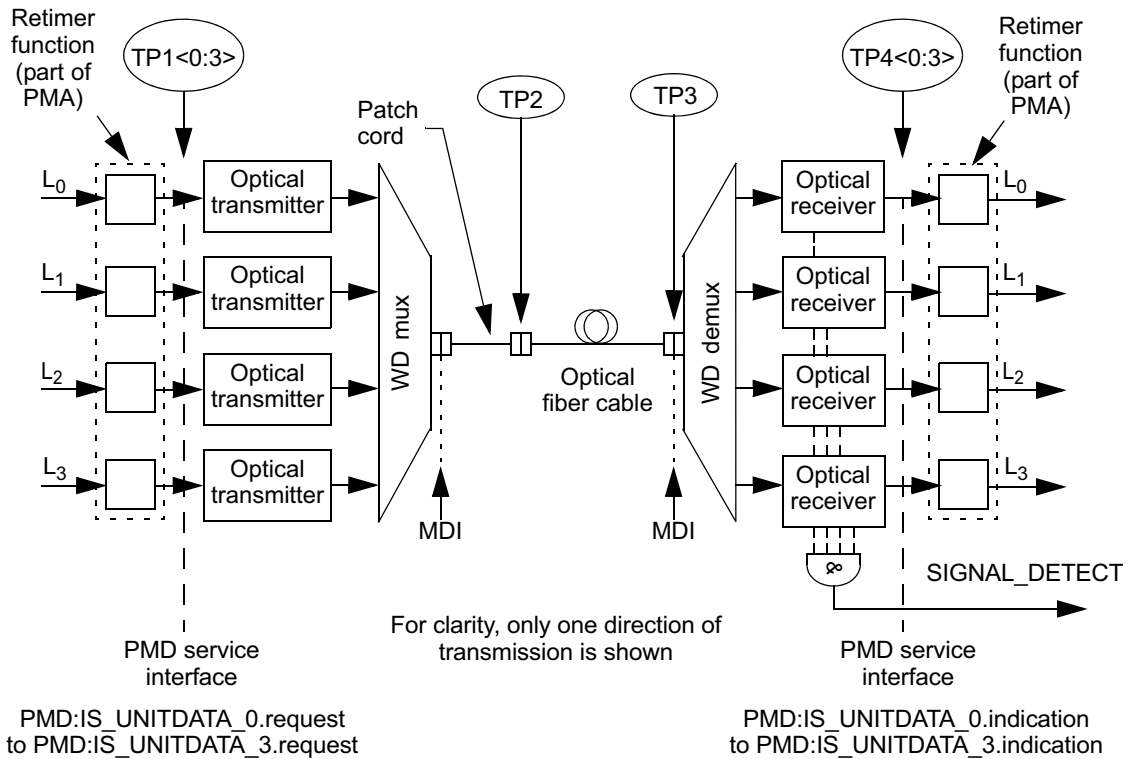
Table 87–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0

TP1<0:3> and TP4<0:3> are informative reference points that may be useful to implementors for testing components (these test points will not typically be accessible in an implemented system).

87.5.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate optical signal streams. The four optical signal streams shall then be wavelength division multiplexed and delivered to the MDI, all according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to tx_bit = one.



WD = Wavelength division

NOTE—Specification of the retimer function is beyond the scope of this standard.

Figure 87–2—Block diagram for 40GBASE-LR4 transmit/receive paths

87.5.3 PMD receive function

The PMD Receive function shall demultiplex the composite optical signal stream received from the MDI into four separate optical signal streams. The four optical signal streams shall then be converted into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_bit = one.

87.5.4 PMD global signal detect function

The PMD global signal detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 87–4. The PMD receiver is not required to verify whether a compliant 40GBASE-R signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

Table 87–4—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 \leq -30 dBm	FAIL
For all lanes; [(Optical power at TP3 \geq receiver sensitivity (max) in OMA in Table 87–8) AND (compliant 40GBASE-R signal input)]	OK
All other conditions	Unspecified

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

87.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect_{*i*}, where *i* represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 87–4.

87.5.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

87.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

- a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 87–7.
- b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

87.5.8 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_{*i*} (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitters in each lane to be selectively disabled.

- a) When a PMD_transmit_disable_{*i*} variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 87–7.
- b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_{*i*} to one, turning off the optical transmitter in each lane.

If the optional `PMD_transmit_disable_i` function is not implemented in MDIO, an alternative method shall be provided to independently disable each transmit lane for testing purposes.

87.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set `PMD_fault` to one. If the MDIO interface is implemented, `PMD_fault` shall be mapped to the fault bit as specified in 45.2.1.2.1.

87.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the `PMD_transmit_fault` variable to one. If the MDIO interface is implemented, `PMD_transmit_fault` shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

87.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the `PMD_receive_fault` variable to one. If the MDIO interface is implemented, `PMD_receive_fault` shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

87.6 Wavelength-division-multiplexed lane assignments

The wavelength range for each lane of the 40GBASE-LR4 PMD is defined in Table 87-5. The center wavelengths are members of the CWDM wavelength grid defined in ITU-T G.694.2 and are spaced at 20 nm.

Table 87-5—Wavelength-division-multiplexed lane assignments

Lane	Center wavelength	Wavelength range
L ₀	1271 nm	1264.5 to 1277.5 nm
L ₁	1291 nm	1284.5 to 1297.5 nm
L ₂	1311 nm	1304.5 to 1317.5 nm
L ₃	1331 nm	1324.5 to 1337.5 nm

NOTE—There is no requirement to associate a particular electrical lane with a particular optical lane, as the PCS is capable of receiving lanes in any arrangement.

87.7 PMD to MDI optical specifications for 40GBASE-LR4

The operating range for the 40GBASE-LR4 PMD is defined in Table 87-6. A 40GBASE-LR4 compliant PMD operates on type B1.1, B1.3 or B6_A single-mode fibers according to the specifications defined in Table 87-14. A PMD that exceeds the required operating range while meeting all other optical specifications is considered compliant (e.g., operating at 12.5 km meets the operating range requirement of 2 m to 10 km).

Table 87–6—40GBASE–LR4 operating range

PMD type	Required operating range
40GBASE–LR4	2 m to 10 km

87.7.1 40GBASE–LR4 transmitter optical specifications

The 40GBASE–LR4 transmitter shall meet the specifications defined in Table 87–7 per the definitions in 87.8.

Table 87–7—40GBASE-LR4 transmit characteristics

Description	Value	Unit
Signaling rate, each lane (range)	10.3125 ± 100 ppm	GBd
Lane wavelengths (range)	1264.5 to 1277.5 1284.5 to 1297.5 1304.5 to 1317.5 1324.5 to 1337.5	nm
Side-mode suppression ratio (SMSR), (min)	30	dB
Total average launch power (max)	8.3	dBm
Average launch power, each lane (max)	2.3	dBm
Average launch power, each lane ^a (min)	–7	dBm
Optical Modulation Amplitude (OMA), each lane (max)	3.5	dBm
Optical Modulation Amplitude (OMA), each lane (min) ^b	–4	dBm
Difference in launch power between any two lanes (OMA) (max)	6.5	dB
Launch power in OMA minus TDP, each lane (min)	–4.8	dBm
Transmitter and dispersion penalty (TDP), each lane (max)	2.6	dB
Average launch power of OFF transmitter, each lane (max)	–30	dBm
Extinction ratio (min)	3.5	dB
RIN ₂₀ OMA (max)	–128	dB/Hz
Optical return loss tolerance (max)	20	dB
Transmitter reflectance ^c (max)	–12	dB
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.4, 0.45, 0.25, 0.28, 0.4}	

^aAverage launch power, each lane (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.

^bEven if the TDP < 0.8dB, the OMA (min) must exceed this value.

^cTransmitter reflectance is defined looking into the transmitter.

87.7.2 40GBASE–LR4 receive optical specifications

The 40GBASE–LR4 receiver shall meet the specifications defined in Table 87–8 per the definitions in 87.8.

Table 87–8—40GBASE–LR4 receive characteristics

Description	Value	Unit
Signaling rate, each lane (range)	10.3125 ± 100 ppm	GBd
Lane wavelengths (range)	1264.5 to 1277.5 1284.5 to 1297.5 1304.5 to 1317.5 1324.5 to 1337.5	nm
Damage threshold ^a (min)	3.3	dBm
Average receive power, each lane (max)	2.3	dBm
Average receive power, each lane ^b (min)	–13.7	dBm
Receive power, each lane (OMA) (max)	3.5	dBm
Difference in receive power between any two lanes (OMA) (max)	7.5	dB
Receiver reflectance (max)	–26	dB
Receiver sensitivity (OMA), each lane ^c (max)	–11.5	dBm
Receiver 3 dB electrical upper cutoff frequency, each lane (max)	12.3	GHz
Stressed receiver sensitivity (OMA), each lane ^d (max)	–9.6	dBm
Conditions of stressed receiver sensitivity test:		
Vertical eye closure penalty, ^e each lane	1.9	dB
Stressed eye J2 Jitter, ^e each lane	0.3	UI
Stressed eye J9 Jitter, ^e each lane	0.47	UI

^aThe receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level

^bAverage receive power, each lane (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.

^cReceiver sensitivity (OMA), each lane (max) is informative.

^dMeasured with conformance test signal at TP3 (see 87.8.11) for BER = 10^{–12}.

^eVertical eye closure penalty, stressed eye J2 Jitter, and stressed eye J9 Jitter are test conditions for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

87.7.3 40GBASE–LR4 illustrative link power budget

An illustrative power budget and penalties for 40GBASE–LR4 channels are shown in Table 87–9.

Table 87–9—40GBASE–LR4 illustrative link power budget

Parameter	Value	Unit
Power budget (for max TDP)	9.3	dB
Operating distance	10	km
Channel insertion loss ^a	6.7	dB
Maximum discrete reflectance	–26	dB
Allocation for penalties ^b (for max TDP)	2.6	dB
Additional insertion loss allowed	0	dB

^aThe channel insertion loss is calculated using the maximum distance specified in Table 87–6 and cabled optical fiber attenuation of 0.47 dB/km at 1264.5 nm plus an allocation for connection and splice loss given in 87.11.2.1.

^bLink penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

87.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

87.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 87–11 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 87–11 may be used to perform that test. The test patterns used in this clause are shown in Table 87–10.

Table 87–10—Test patterns

Pattern	Pattern description	Defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.10

87.8.2 Skew and Skew Variation

Skew and Skew Variation are defined in 80.5 and are required to remain within the limits given in 87.3.2 over the time that the link is in operation. The measurement of Skew and Skew Variation is made by separating optical lanes with an optical demultiplexer and then acquiring the data on each lane using clock and data recovery units with high-frequency corner bandwidths as specified in 86.8.3.2 and a slope of –20 dB/decade. The arrival times of the one to zero transition of the alignment marker sync bits on each lane are

Table 87–11—Test-pattern definitions and related subclauses

Parameter	Pattern	Related subclause
Wavelength	3, 5 or valid 40GBASE-LR signal	87.8.3
Side mode suppression ratio	3, 5 or valid 40GBASE-LR signal	—
Average optical power	3, 5 or valid 40GBASE-LR signal	87.8.4
Optical modulation amplitude (OMA)	Square wave or 4	87.8.5
Transmitter and dispersion penalty (TDP)	3 or 5	87.8.6
Extinction ratio	3, 5 or valid 40GBASE-LR signal	87.8.7
RIN ₂₀ OMA	Square wave or 4	87.8.8
Transmitter optical waveform	3, 5 or valid 40GBASE-LR signal	87.8.9
Stressed receiver sensitivity	3 or 5	87.8.11
Calibration of OMA for receiver tests	Square wave or 4	87.8.11
Vertical eye closure penalty calibration	3 or 5	87.8.11
Receiver 3 dB electrical upper cutoff frequency	3, 5 or valid 40GBASE-LR signal	87.8.12

then compared. This arrangement ensures that any high frequency jitter that is present on the signals is not included in the Skew measurement.

87.8.3 Wavelength

The wavelength of each optical lane shall be within the ranges given in Table 87–5 if measured per TIA/EIA–455–127–A or IEC 61280–1–3. The lane under test is modulated using the test pattern defined in Table 87–11.

87.8.4 Average optical power

The average optical power of each lane shall be within the limits given in Table 87–7 for 40GBASE–LR4 if measured using the methods given in IEC 61280–1–1, with the sum of the optical power from all of the lanes not under test below –30 dBm, per the test setup in Figure 53–6.

87.8.5 Optical Modulation Amplitude (OMA)

OMA shall be as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.

87.8.6 Transmitter and dispersion penalty

Transmitter and dispersion penalty (TDP) shall be as defined in 52.9.10 with the exception that each optical lane is tested individually using an optical filter to separate the lane under test from the others. The measurement procedure for 40GBASE-LR4 is detailed in 87.8.6.1 to 87.8.6.4.

The optical filter passband ripple shall be limited to 0.5 dB peak-to-peak and the isolation is chosen such that the ratio of the power in the lane being measured to the sum of the powers of all of the other lanes is greater than 20 dB (see ITU-T G.959.1 Annex B). The lanes not under test shall be operating with PRBS31 or valid 40GBASE-LR4 bit streams.

87.8.6.1 Reference transmitter requirements

The reference transmitter is a high-quality instrument-grade device that can be implemented by a CW laser modulated by a high-performance modulator. The basic requirements are as follows:

- a) Rise/fall times of less than 25 ps at 20% to 80%.
- b) The output optical eye is symmetric and passes the transmitter optical waveform test of 87.8.9.
- c) In the center 20% region of the eye, the worst-case vertical eye closure penalty as defined in 52.9.9.2 is less than 0.5 dB.
- d) Total Jitter less than 0.2 UI peak-to-peak.
- e) RIN of less than -136 dB/Hz.

87.8.6.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 87–12.

Table 87–12—Transmitter compliance channel specifications

PMD type	Dispersion ^a (ps/nm)		Insertion loss ^b	Optical return loss ^c
	Minimum	Maximum		
40GBASE-LR4	$0.2325 \times \lambda \times [1 - (1324 / \lambda)^4]$	$0.2325 \times \lambda \times [1 - (1300 / \lambda)^4]$	Minimum	20 dB

^aThe dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes 10 km for 40GBASE-LR4.

^bThere is no intent to stress the sensitivity of the BERT's optical receiver.

^cThe optical return loss is applied at TP2.

A 40GBASE-LR4 transmitter is to be compliant with a total dispersion at least as negative as the “minimum dispersion” and at least as positive as the “maximum dispersion” columns specified in Table 87–12 for the wavelength of the device under test. This may be achieved with channels consisting of fibers with lengths chosen to meet the dispersion requirements.

To verify that the fiber has the correct amount of dispersion, the measurement method defined in IEC 60793-1-42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 87–12. The state of polarization of the back reflection is adjusted to create the greatest RIN.

87.8.6.3 Reference receiver requirements

The reference receiver is required to have the bandwidth given in 87.8.9. The sensitivity of the reference receiver is limited by Gaussian noise. The receiver has minimal threshold offset, deadband, hysteresis, baseline wander, deterministic jitter, or other distortions. Decision sampling has minimal uncertainty and setup/hold times.

The nominal sensitivity of the reference receiver, S , is measured in OMA using the setup of Figure 52–12 without the test fiber and with the transversal filter removed. The sensitivity S must be corrected for any significant reference transmitter impairments including any vertical eye closure. It is measured while sampling at the eye center or corrected for off-center sampling. It is calibrated at the wavelength of the transmitter under test.

For all transmitter and dispersion penalty measurements, determination of the center of the eye is required. Center of the eye is defined as the time halfway between the left and right sampling points within the eye where the measured BER is greater than or equal to 1×10^{-3} .

The clock recovery unit (CRU) used in the TDP measurement has a corner frequency of 4 MHz and a slope of 20 dB/decade. When using a clock recovery unit as a clock for BER measurements, passing of low-frequency jitter from the data to the clock removes this low-frequency jitter from the measurement.

87.8.6.4 Test procedure

The test procedure is as defined in 52.9.10.4 with the exception that all lanes are operational in both directions (transmit and receive), each lane is tested individually using an optical filter to separate the lane under test from the others, and the BER of 1×10^{-12} is for the lane under test on its own.

87.8.7 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 87–7 for 40GBASE–LR4 if measured using the methods specified in IEC 61280–2–2, with the sum of the optical power from all of the lanes not under test below –30 dBm. The extinction ratio is measured using the test pattern defined in Table 87–11.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 87–11).

87.8.8 Relative Intensity Noise (RIN_{20OMA})

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

- a) The optical return loss is 20 dB.
- b) Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.
- c) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 10.3 GHz).

87.8.9 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 87–11 shall meet specifications according to the methods specified in 86.8.4.6.1 with the filter nominal reference frequency f_r of 7.5 GHz and filter tolerances as specified for

STM-64 in ITU-T G.691. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

87.8.10 Receiver sensitivity

Receiver sensitivity, which is defined for an ideal input signal, is informative and compliance is not required. If measured, the test signal should have negligible impairments such as intersymbol interference (ISI), rise/fall times, jitter, and RIN. Instead, the normative requirement for receivers is stressed receiver sensitivity.

87.8.11 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 87–8 for 40GBASE–LR4 if measured using the method described in 87.8.11.1 and 87.8.11.5 with the conformance test signal at TP3 as described in 87.8.11.2. The BER is required to be met for the lane under test on its own.

For each lane, the stressed receiver sensitivity is defined with the transmit section in operation on all four lanes and with the receive lanes not under test also in operation. Pattern 3 or Pattern 5, or a valid 40GBASE–R signal, is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal.

87.8.11.1 Stressed receiver conformance test block diagram

A block diagram for the receiver conformance test is shown in Figure 87–3. The patterns used for testing the receiver are specified in Table 87–11. The optical test signal is conditioned (stressed) using the stressed receiver methodology defined in 87.8.11.2, and has sinusoidal jitter applied as specified in 87.8.11.4. A suitable test set is needed to characterize and verify that the signal used to test the receiver has the appropriate characteristics.

The low-pass filter is used to create ISI-induced vertical eye closure penalty (VECP). The low-pass filter, when combined with the E/O converter, should have a frequency response that results in the appropriate level of initial vertical eye closure before the sinusoidal terms are added.

The sinusoidal amplitude interferer 1 causes jitter that is intended to emulate instantaneous bit shrinkage that can occur with DDJ. This type of jitter cannot be created by simple phase modulation. The sinusoidal amplitude interferer 2 causes additional eye closure, but in conjunction with the finite edge rates from the limiter, also causes some jitter. The sinusoidally jittered clock represents other forms of jitter and also verifies that the receiver under test can track low-frequency jitter. The sinusoidal amplitude interferers may be set at any frequency between 100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferers, the sinusoidal jitter, the signaling rate, and the pattern repetition rate. The Gaussian noise generator, the amplitude of the sinusoidal interferers, and the low-pass filter are adjusted so that the VECP, stressed eye J2 Jitter, and stressed eye J9 Jitter specifications are met simultaneously.

For improved visibility for calibration, all elements in the signal path (cables, DC blocks, E/O converter, etc.) should have wide and smooth frequency response, and linear phase response, throughout the spectrum of interest. Baseline wander and overshoot and undershoot should be minimized.

The stressed receiver conformance signal verification is described in 87.8.11.3.

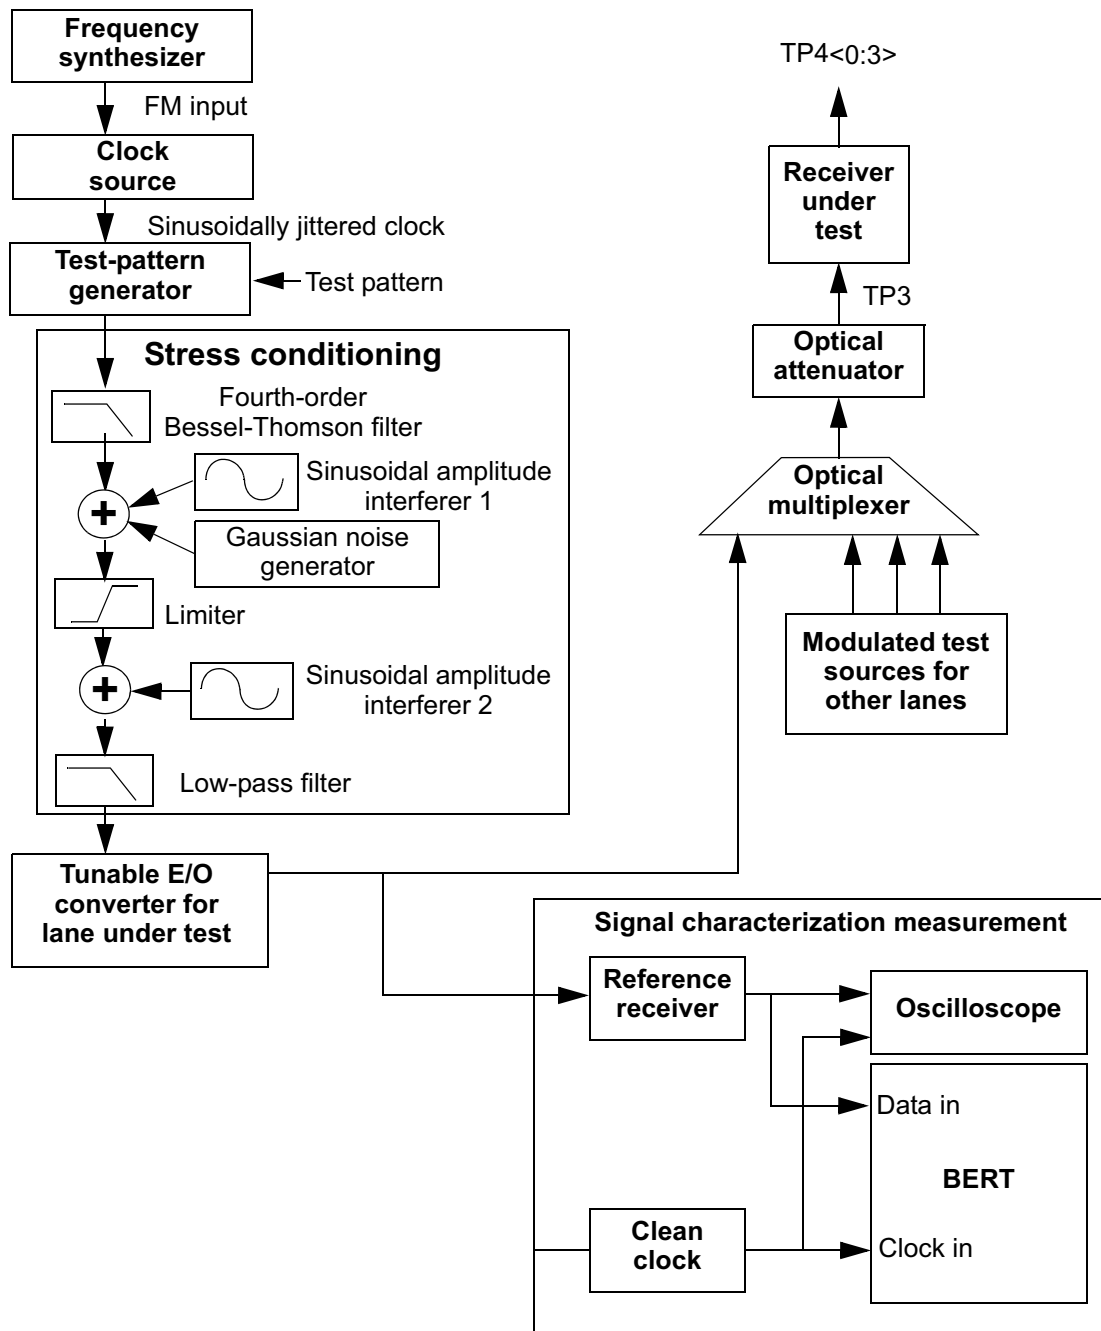


Figure 87–3—Stressed receiver conformance test block diagram

87.8.11.2 Stressed receiver conformance test signal characteristics and calibration

The conformance test signal is used to validate that the PMD receiver of the lane under test meets BER requirements with near worst-case waveforms at TP3.

The primary parameters of the conformance test signal are vertical eye closure penalty (VECP), stressed eye J2 Jitter, and stressed eye J9 Jitter. VECP is measured at the time center of the eye (halfway between 0 and 1

on the unit interval scale as defined in 52.9.7). Stressed eye J2 Jitter is defined in 86.8.3.3.1 and stressed eye J9 Jitter is defined in 86.8.3.3.2. The values of these components are defined by their histogram results, as measured at the average optical power, which can be obtained with AC coupling.

The vertical eye closure penalty is given by Equation (87–1).

$$\text{Vertical eye closure penalty} = 10\log_{10}\frac{OMA}{A_O} \quad (\text{dB}) \quad (87-1)$$

where

A_O is the amplitude of the eye opening from the 99.95th percentile of the lower histogram to the 0.05th percentile of the upper histogram

OMA is the optical modulation amplitude as defined in 87.8.5




Figure 87–4—Required characteristics of the conformance test signal at TP3

An example stressed receiver conformance test setup is shown in Figure 87–3; however, any approach that modulates or creates the appropriate levels and frequencies of the VECP and jitter components is acceptable.

Residual low-probability noise and jitter should be minimized so that the outer slopes of the final amplitude histograms are as steep as possible.

The following steps describe a possible method for setting up and calibrating a stressed eye conformance signal when using a stressed receiver conformance test setup as shown in Figure 87–3:

- 1) Set the signaling rate of the test-pattern generator to meet the requirements in Table 87–8.
- 2) With sinusoidal interferers and sinusoidal jitter turned off, set the extinction ratio of the E/O to approximately the minimum specified in Table 87–7.

- 3) The required values of vertical eye closure penalty, stressed eye J2 Jitter, and stressed eye J9 Jitter of the stressed receiver conformance signal are given in Table 87–8.

With the sinusoidal interferers and sinusoidal jitter turned off, greater than two thirds of the dB value of the VECP should be created by the selection of the appropriate bandwidth for the low-pass filter. Any remaining VECP must be created with sinusoidal interferer 2 or sinusoidal jitter.

The sinusoidal amplitude interferers may be set at any frequency between 100 MHz and 2 GHz, although care should be taken to avoid harmonic relationships between the sinusoidal interferers, the sinusoidal jitter, the signaling rate and the pattern repetition rate.

Sinusoidal jitter is added as specified in Table 87–13. When calibrating the conformance signal, the sinusoidal jitter frequency should be well within the 4 MHz to 10 times LB as defined in Table 87–13 and illustrated in Figure 87–5.

Iterate the adjustments of sinusoidal interferers and Gaussian noise generator until the values of VECP, stressed eye J2 Jitter and stressed eye J9 Jitter meet the requirements in Table 87–8, and sinusoidal jitter above 4 MHz is as specified in Table 87–13. The resulting stressed eye conformance signal is required to have at least 0.05 UI of pulse width shrinkage.

Figure 87–3 shows the stress conditioned signal being applied to a tunable E/O converter. However, any optical source may be used that can meet the OMA and wavelength requirements for the lane under test as described in 87.8.11.5. Similarly, the other test sources that supply modulated signals to the other lanes may use any tunable or fixed sources that meet the OMA and wavelength requirements described in 87.8.11.5.

Each receiver lane is conformance tested in turn. The source for the lane under test is adjusted to supply a signal at the input to the receiver under test at the stressed receiver sensitivity OMA specified in Table 87–8, and the test sources for the other lanes are set to the required OMA and wavelength as described in 87.8.11.5.

87.8.11.3 Stressed receiver conformance test signal verification

The stressed receiver conformance test signal can be verified using an optical reference receiver with an ideal fourth-order Bessel-Thomson response with a reference frequency f_r of 7.5 GHz. Use of G.691 tolerance filters may significantly degrade this calibration. The clock output from the clock source in Figure 87–3 will be modulated with the sinusoidal jitter. To use an oscilloscope to calibrate the final stressed eye J2 Jitter and stressed eye J9 Jitter that includes the sinusoidal jitter component, a separate clock source (clean clock of Figure 87–3) is required that is synchronized to the source clock, but not modulated with the jitter source.

Care should be taken when characterizing the test signal because excessive noise/jitter in the measurement system will result in an input signal that does not fully stress the receiver under test. Running the receiver tolerance test with a signal that is under-stressed may result in the deployment of non-compliant receivers. Care should be taken to minimize the noise/jitter introduced by the reference O/E, filters, and BERT and/or to correct for this noise. While the details of a BER scan measurement and test equipment are beyond the scope of this standard, it is recommended that the implementor fully characterize the test equipment and apply appropriate guard bands to ensure that the stressed receiver conformance input signal meets the stress and sinusoidal jitter specified in 87.8.11.2 and 87.8.11.4.


87.8.11.4 Sinusoidal jitter for receiver conformance test

The sinusoidal jitter is used to test receiver jitter tolerance. The amplitude of the applied sinusoidal jitter is dependent on frequency as specified in Table 87–13 and is illustrated in Figure 87–5.

Table 87–13—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak (UI)
$f < 40 \text{ kHz}$	Not specified
$40 \text{ kHz} < f \leq 4 \text{ MHz}$	$2 \times 10^5 / f$
$4 \text{ MHz} < f < 10 \text{ LB}^a$	0.05

^aLB = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

**Figure 87–5—Illustration of the mask of the sinusoidal component of jitter tolerance**

87.8.11.5 Stressed receiver conformance test procedure for WDM conformance testing

The receiver tests requiring the TP3 conformance test signal are performed on a per lane basis. For each lane, the stressed receiver sensitivity is defined with the transmit section in operation on all four lanes and with the receive lanes not under test also in operation. Pattern 3 or Pattern 5, or a valid 40GBASE-R signal is sent from the transmit section of the PMD under test. The signal being transmitted is asynchronous to the received signal. All test sources are modulated simultaneously, using valid 40GBASE-LR4 signals. The transmitter of the transceiver under test is operating with test patterns as defined in Table 87–11.

A rigorous method for testing WDM conformance using tunable test sources is described as follows:

- The OMAs of the test sources in the lanes other than the lane under test are set to the highest OMA relative to the test source in the lane under test allowed by the “difference in receive power between any two lanes (OMA) (max)” parameter in Table 87–8.
- The test sources in the lanes other than the lane under test are tuned to wavelengths within their wavelength range corresponding to the worst-case crosstalk to the lane under test.
- The test source for the lane under test is tuned to the wavelength within its wavelength range corresponding to the worst-case sensitivity for the receiver under test.

It is recognized that a test setup that uses multiple tunable sources, while allowing rigorous worst-case measurements to be made, is likely to be onerous in practice.

A more practical, alternative method for testing WDM conformance is described below. The alternative WDM conformance test avoids the need for tunable sources but may result in some over-stressing of the receiver under test:

- 1) The test sources in each lane can be at any wavelength within each lane's wavelength range.
- 2) The OMAs of the test sources in the lanes other than the lane under test are set to the highest OMA relative to the test source in the lane under test allowed by the "difference in receive power between any two lanes (OMA) (max)" parameter in Table 87–8, plus an increment corresponding to insertion loss variation within the lane under test (NOTE 1), plus an increment corresponding to the isolation variation of the lane in question (NOTE 2).

NOTE 1—The increment corresponding to the variation of insertion loss with wavelength within the lane under test is equal to the dB variation of measured receiver sensitivity as the test source wavelength is swept across the wavelength range of the lane under test.

NOTE 2—For each of the test sources in the lanes not under test, an increment corresponding to the isolation variation is applied that is equal to the dB variation in optical crosstalk measured in the lane under test as the wavelength of each test source in the lanes not under test is swept across its respective wavelength range.

There are many ways to determine the size of the increments required, two example methods are given in the paragraphs that follow. Note that each lane will have one insertion loss variation value, and three values of optical crosstalk variation (one for each of the other lanes).

Example method 1: If a measure of the received signal strength is available for each lane, scan a tunable laser across the wavelength range of each lane while simultaneously recording the received signal in all lanes. As the laser sweeps across the wavelength range of a particular lane the dB variation in observed signal in that lane is equal to the insertion loss variation. The dB variation in signal observed in each of the other lanes is equal to the optical crosstalk variation.

Example method 2: If a discrete optical demultiplexer and discrete receivers are used, scan a tunable laser across the wavelength range of each channel while simultaneously recording the optical power of all four demultiplexer outputs. As the laser sweeps across the wavelength range of a particular channel, the dB variation in observed optical power in that lane's output is equal to the insertion loss variation. The dB variation in optical power observed in each of the other lanes outputs is equal to the optical crosstalk variation.

87.8.12 Receiver 3 dB electrical upper cutoff frequency

The receiver 3dB electrical upper cutoff frequency shall be within the limits given in Table 87–8 if measured as described in 52.9.11. Each optical lane is measured using an optical signal or signals with its/their wavelength within the specified wavelength range of the lane to be tested. The test may use an electrical combiner and one optical source as in 53.9.13.

87.9 Safety, installation, environment, and labeling

87.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

87.9.2 Laser safety

40GBASE-LR4 optical transceivers shall conform to Class 1 laser requirements as defined in IEC 60825–1 and IEC 60825–2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.¹⁸

87.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

87.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 40GBASE-LR4 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

87.9.4.1 Electromagnetic emission

A system integrating a 40GBASE-LR4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

87.9.4.2 Temperature, humidity, and handling

The optical link is 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 standard.

87.9.5 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g., 40GBASE-LR4).

Labeling requirements for Class 1 lasers are given in the laser safety standards referenced in 87.9.2.

¹⁸A host system that fails to meet the manufacturers requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

87.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 87–6.




Figure 87–6—Fiber optic cabling model

The channel insertion loss is given in Table 87–14. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA–526–7/method A–1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 87–14—Fiber optic cabling (channel) characteristics for 40GBASE–LR4

Description	Value	Unit
Operating distance (max)	10	km
Channel insertion loss ^{a, b} (max)	6.7	dB
Channel insertion loss (min)	0	dB
Positive dispersion ^b (max)	33.5	ps/nm
Negative dispersion ^b (min)	–59.5	ps/nm
DGD_max ^c	10	ps
Optical return loss (min)	21	dB

^aThese channel insertion loss values include cable, connectors, and splices.

^bOver the wavelength range 1264.5 nm to 1337.5 nm.

^cDifferential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were transmitted in the two principal states of polarization of an optical signal. DGD_max is the maximum differential group delay that the system must tolerate.

87.11 Characteristics of the fiber optic cabling (channel)

The 40GBASE–LR4 fiber optic cabling shall meet the specifications defined in Table 87–14. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

87.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793–2–50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_A (bend insensitive) fibers and the requirements in Table 87–15 where they differ.

Table 87–15—Optical fiber and cable characteristics for 40GBASE–LR4

Description	Value	Unit
Nominal fiber specification wavelength	1310	nm
Cabled optical fiber attenuation (max)	0.47 ^a or 0.5 ^b	dB/km
Zero dispersion wavelength (λ_0)	$1300 \leq \lambda_0 \leq 1324$	nm
Dispersion slope (max) (S_0)	0.093	ps/nm ² km

^aThe 0.47 dB/km at 1264.5 nm attenuation for optical fiber cables is derived from Appendix I of ITU-T G.695.

^bThe 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA/EIA 568-B.3-2000. Using 0.5 dB/km may not support operation at 10 km for 40GBASE–LR4.

87.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 87–6, consists of a mated pair of optical connectors.

87.11.2.1 Connection insertion loss

The maximum link distances for single-mode fiber are calculated based on an allocation of 2 dB total connection and splice loss. For example, this allocation supports four connections with an average insertion loss per connection of 0.5 dB. Connections with different loss characteristics may be used provided the requirements of Table 87–14 are met.

87.11.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –26 dB.

87.11.3 Medium Dependent Interface (MDI) requirements

The 40GBASE–LR4 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 87–6). Examples of an MDI include the following:

- a) Connectorized fiber pigtail
- b) PMD receptacle

When the MDI is a connector plug and receptacle connection, it shall meet the interface performance specifications of IEC 61753–1–1 and IEC 61753–021–2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 87.5.1, not at the MDI.

87.12 Protocol implementation conformance statement (PICS) proforma for Clause 87, Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4¹⁹

87.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 87, Physical Medium Dependent sublayer and medium, type 40GBASE-LR4, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

87.12.2 Identification

87.12.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

87.12.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 87, Physical Medium Dependent sublayer and medium, type 40GBASE-LR4
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

¹⁹Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

87.12.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
LR4	40GBASE-LR4 PMD	87.7	Device supports requirements for 40GBASE-LR4 PHY	O	Yes [] No []
*INS	Installation/cable	87.10	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [] No []
XLTP1	Reference point TP1 exposed and available for testing	87.5.1	This point may be made available for use by implementors to certify component conformance	O	Yes [] No []
XLTP4	Reference point TP4 exposed and available for testing	87.5.1	This point may be made available for use by implementors to certify component conformance	O	Yes [] No []
XLDC	Delay constraints	87.3.1	Device conforms to delay constraints	M	Yes []
XLSC	Skew constraints	87.3.2	Device conforms to Skew and Skew Variation constraints	M	Yes []
*MD	MDIO capability	87.4	Registers and interface supported	O	Yes [] No []

87.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 40GBASE-LR4

87.12.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
XLF1	Compatible with 40GBASE-R PCS and PMA	87.1		M	Yes []
XLF2	Integration with management functions	87.1		O	Yes [] No []
XLF3	Transmit function	87.5.2	Conveys bits from PMD service interface to MDI	M	Yes []
XLF4	Optical multiplexing and delivery to the MDI	87.5.2	Optically multiplexes the four optical signal streams for delivery to the MDI	M	Yes []
XLF5	Mapping between optical signal and logical signal for transmitter	87.5.2	Higher optical power is a one	M	Yes []
XLF6	Receive function	87.5.3	Conveys bits from MDI to PMD service interface	M	Yes []
XLF7	Conversion of four optical signals to four electrical signals	87.5.3	For delivery to the PMD service interface	M	Yes []
XLF8	Mapping between optical signal and logical signal for receiver	87.5.3	Higher optical power is a one	M	Yes []
XLF9	Global Signal Detect function	87.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes []
XLF10	Global Signal Detect behavior	87.5.4	SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes	M	Yes []
XLF11	Lane-by-Lane Signal Detect function	87.5.5	Sets PMD_signal_detect_i values on a lane-by-lane basis per requirements of Table 87-4	MD:O	Yes [] No [] N/A []
XLF12	PMD_reset function	87.5.6	Resets the PMD sublayer	MD:O	Yes [] No [] N/A []

87.12.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
XLM1	Management register set	87.4		MD:M	Yes [] N/A []
XLM2	Global transmit disable function	87.5.7	Disables all of the optical transmitters with the PMD_global_transmit_disable variable	MD:O	Yes [] No [] N/A []
XLM3	PMD_lane_by_lane_transmit_disable function	87.5.8	Disables the optical transmitter on the lane associated with the PMD_transmit_disable_i variable	MD:O.2	Yes [] No [] N/A []
XLM4	PMD_lane_by_lane_transmit_disable	87.5.8	Disables each optical transmitter independently if XLM3 = No	O.2	Yes [] No []
XLM5	PMD_fault function	87.5.9	Sets PMD_fault to one if any local fault is detected	MD:O	Yes [] No [] N/A []
XLM6	PMD_transmit_fault function	87.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	MD:O	Yes [] No [] N/A []
XLM7	PMD_receive_fault function	87.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	MD:O	Yes [] No [] N/A []

87.12.4.3 PMD to MDI optical specifications for 40GBASE-LR4

Item	Feature	Subclause	Value/Comment	Status	Support
XLLR1	Transmitter meets specifications in Table 87-7	87.7.1	Per definitions in 87.8	M	Yes [] N/A []
XLLR2	Receiver meets specifications in Table 87-8	87.7.2	Per definitions in 87.8	M	Yes [] N/A []

87.12.4.4 Optical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
XLOM1	Measurement cable	87.8	2 m to 5 m in length	M	Yes []
XLOM2	Center wavelength	87.8.3	Per TIA-455-127-A or IEC 61280-1-3 under modulated conditions	M	Yes []
XLOM3	Average optical power	87.8.4	Per IEC 61280-1-1	M	Yes []
XLOM4	Transmitter and dispersion penalty	87.8.6	Each lane	M	Yes []
XLOM5	OMA	87.8.5	Each lane	M	Yes []
XLOM6	Extinction ratio	87.8.7	Per IEC 61280-2-2	M	Yes []
XLOM7	RIN ₂₀ OMA	87.8.8	Each lane	M	Yes []
XLOM8	Transmit eye	87.8.9	Each lane	M	Yes []
XLOM9	Stressed receiver sensitivity	87.8.11	Each lane	M	Yes []
XLOM10	Receiver 3 dB electrical upper cutoff frequency	87.8.12	Each lane	M	Yes []

87.12.4.5 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
XLES1	General safety	87.9.1	Conforms to IEC 60950-1	M	Yes []
XLES2	Laser safety—IEC Class 1	87.9.2	Conforms to Class 1 laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes []
XLES3	Electromagnetic interference	87.9.4.1	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes []

87.12.4.6 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
XLOC1	Fiber optic cabling	87.11	Meets the requirements specified in Table 87-14	INS:M	Yes [] N/A []
XLOC2	Maximum discrete reflectance	87.11.2.2	Less than -26 dB	INS:M	Yes [] N/A []
XLOC3	MDI requirements	87.11.3	Meets the interface performance specifications listed	INS:M	Yes [] N/A []

88. Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4

88.1 Overview

This clause specifies the 100GBASE-LR4 PMD and the 100GBASE-ER4 PMD together with the single-mode fiber medium. When forming a complete Physical Layer, a PMD shall be connected to the appropriate PMA as shown in Table 88-1, to the medium through the MDI and to the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent.

Table 88-1—Physical Layer clauses associated with the 100GBASE-LR4 and 100GBASE-ER4 PMDs

Associated clause	100GBASE-LR4	100GBASE-ER4
81—RS	Required	Required
81—CGMII ^a	Optional	Optional
82—PCS for 100GBASE-R	Required	Required
83—PMA for 100GBASE-R	Required	Required
83A—CAUI	Optional	Optional
83B—Chip to module CAUI	Optional	Optional

^aThe CGMII is an optional interface. However, if the CGMII is not implemented, a conforming implementation must behave functionally as though the RS and CGMII were present.

Figure 88-1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 40 Gb/s and 100 Gb/s Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

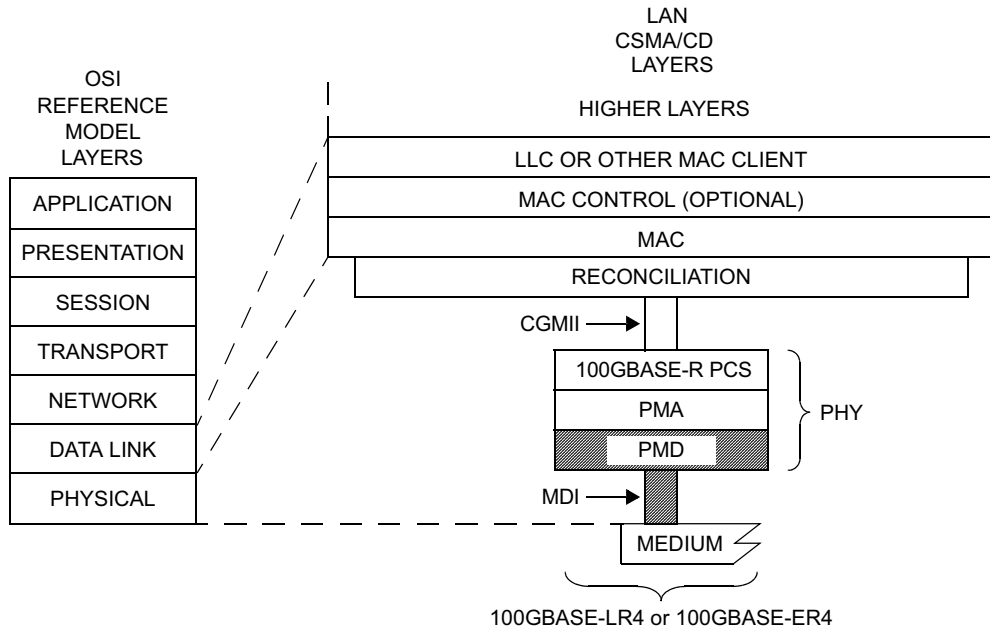
88.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-LR4 and 100GBASE-ER4 PMDs. The service interfaces for these PMDs are described in an abstract manner and do not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The PMD service interface is an instance of the inter-sublayer service interface defined in 80.3. The PMD service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_i.request
PMD:IS_UNITDATA_i.indication
PMD:IS_SIGNAL.indication
```

The 100GBASE-LR4 and 100GBASE-ER4 PMDs have four parallel bit streams, hence $i = 0$ to 3.



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE

PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 LR = PMD FOR SINGLE-MODE FIBER — 10 km
 ER = PMD FOR SINGLE-MODE FIBER — 40 km

Figure 88–1—100GBASE-LR4 and 100GBASE-ER4 PMDs relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 CSMA/CD LAN model

In the transmit direction, the PMA continuously sends four parallel bit streams to the PMD, one per lane, each at a nominal signaling rate of 25.78125 GBd. The PMD then converts these streams of bits into the appropriate signals on the MDI.

In the receive direction, the PMD continuously sends four parallel bit streams to the PMA corresponding to the signals received from the MDI, one per lane, each at a nominal signaling rate of 25.78125 GBd.

The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the PMD:IS_SIGNAL.indication(SIGNAL_OK) inter-sublayer service interface primitive defined in 80.3.

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL. When SIGNAL_DETECT = FAIL, the rx_bit parameters are undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx_bit parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the 10^{-12} BER objective.

88.3 Delay and Skew

88.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-LR4 or 100GBASE-ER4 PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

88.3.2 Skew constraints

The Skew (relative delay) between the lanes must be kept within limits so that the information on the lanes can be reassembled by the PCS. The Skew Variation must also be limited to ensure that a given PCS lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP1 to SP6 shown in Figure 80-4 and Figure 80-5.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation see 80.5. The measurements of Skew and Skew Variation are defined in 87.8.2 with the exception that the clock and data recovery units' high-frequency corner bandwidths are 10 MHz.

88.4 PMD MDIO function mapping

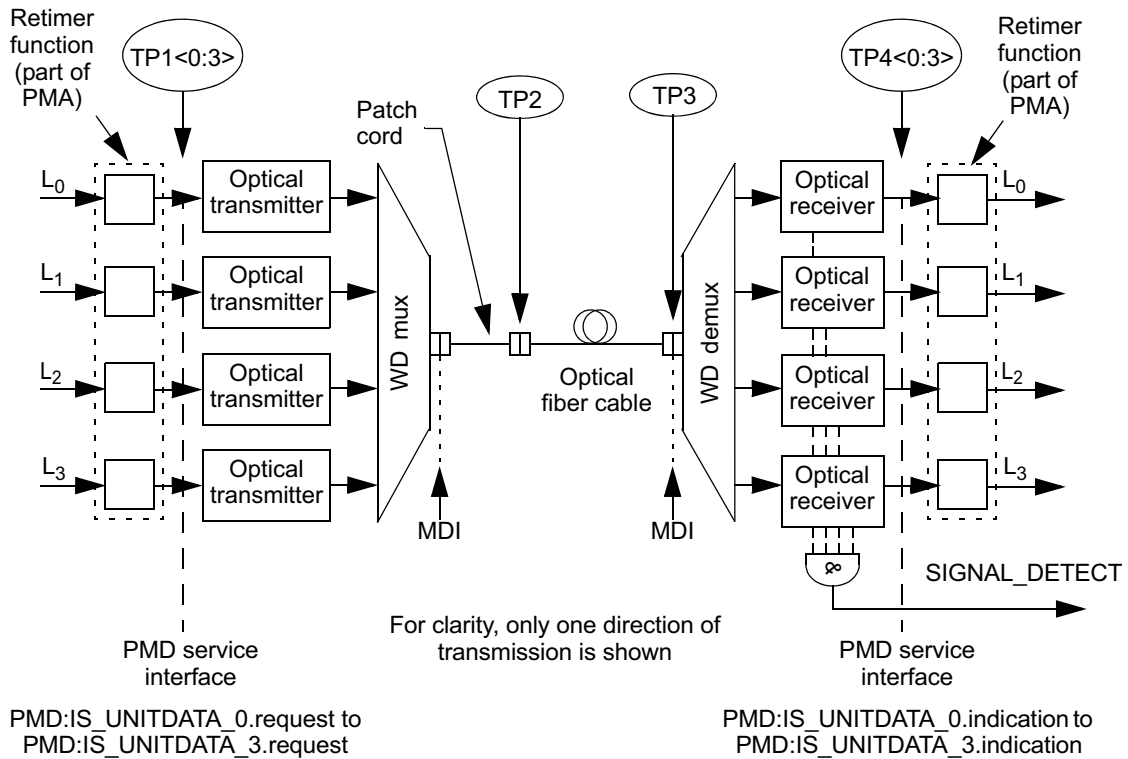
The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 88-2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 88-3.

88.5 PMD functional specifications

The 100GBASE-LR4 and 100GBASE-ER4 PMDs perform the Transmit and Receive functions, which convey data between the PMD service interface and the MDI.

88.5.1 PMD block diagram

The PMD block diagram is shown in Figure 88-2. For purposes of system conformance, the PMD sublayer is standardized at the points described in this subclause. The optical transmit signal is defined at the output end of a single-mode fiber patch cord (TP2), between 2 m and 5 m in length. Unless specified otherwise, all transmitter measurements and tests defined in 88.8 are made at TP2. The optical receive signal is defined at the output of the fiber optic cabling (TP3) at the MDI (see 88.11.3). Unless specified otherwise, all receiver measurements and tests defined in 88.8 are made at TP3.



NOTE—Specification of the retimer function and the electrical implementation of the PMD service interface is beyond the scope of this standard.

Figure 88–2—Block diagram for 100GBASE-LR4 and 100GBASE-ER4 transmit/receive paths

Table 88–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
PMD transmit disable 3	PMD transmit disable register	1.9.4	PMD_transmit_disable_3
PMD transmit disable 2	PMD transmit disable register	1.9.3	PMD_transmit_disable_2
PMD transmit disable 1	PMD transmit disable register	1.9.2	PMD_transmit_disable_1
PMD transmit disable 0	PMD transmit disable register	1.9.1	PMD_transmit_disable_0

TP1<0:3> and TP4<0:3> are informative reference points that may be useful to implementors for testing components (these test points will not typically be accessible in an implemented system).

Table 88–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
PMD receive signal detect 3	PMD receive signal detect register	1.10.4	PMD_signal_detect_3
PMD receive signal detect 2	PMD receive signal detect register	1.10.3	PMD_signal_detect_2
PMD receive signal detect 1	PMD receive signal detect register	1.10.2	PMD_signal_detect_1
PMD receive signal detect 0	PMD receive signal detect register	1.10.1	PMD_signal_detect_0

88.5.2 PMD transmit function

The PMD Transmit function shall convert the four bit streams requested by the PMD service interface messages PMD:IS_UNITDATA_0.request to PMD:IS_UNITDATA_3.request into four separate optical signal streams. The four optical signal streams shall then be wavelength division multiplexed and delivered to the MDI, all according to the transmit optical specifications in this clause. The higher optical power level in each signal stream shall correspond to tx_bit = one.

88.5.3 PMD receive function

The PMD Receive function shall demultiplex the composite optical signal stream received from the MDI into four separate optical signal streams. The four optical signal streams shall then be converted into four bit streams for delivery to the PMD service interface using the messages PMD:IS_UNITDATA_0.indication to PMD:IS_UNITDATA_3.indication, all according to the receive optical specifications in this clause. The higher optical power level in each signal stream shall correspond to rx_bit = one.

88.5.4 PMD global signal detect function

The PMD global signal detect function shall report the state of SIGNAL_DETECT via the PMD service interface. The SIGNAL_DETECT parameter is signaled continuously, while the PMD:IS_SIGNAL.indication message is generated when a change in the value of SIGNAL_DETECT occurs. The SIGNAL_DETECT parameter defined in this clause maps to the SIGNAL_OK parameter in the inter-sublayer service interface primitives defined in 80.3.

SIGNAL_DETECT shall be a global indicator of the presence of optical signals on all four lanes. The value of the SIGNAL_DETECT parameter shall be generated according to the conditions defined in Table 88–4. The PMD receiver is not required to verify whether a compliant 100GBASE-R signal is being received. This standard imposes no response time requirements on the generation of the SIGNAL_DETECT parameter.

As an unavoidable consequence of the requirements for the setting of the SIGNAL_DETECT parameter, implementations must provide adequate margin between the input optical power level at which the SIGNAL_DETECT parameter is set to OK, and the inherent noise level of the PMD including the effects of crosstalk, power supply noise, etc.

Table 88–4—SIGNAL_DETECT value definition

Receive conditions	SIGNAL_DETECT value
For any lane; Average optical power at TP3 \leq -30 dBm	FAIL
For all lanes; [(Optical power at TP3 \geq receiver sensitivity (max) in OMA in Table 88–8) AND (compliant 100GBASE-R signal input)]	OK
All other conditions	Unspecified

Various implementations of the Signal Detect function are permitted by this standard, including implementations that generate the SIGNAL_DETECT parameter values in response to the amplitude of the modulation of the optical signal and implementations that respond to the average optical power of the modulated optical signal.

88.5.5 PMD lane-by-lane signal detect function

Various implementations of the Signal Detect function are permitted by this standard. When the MDIO is implemented, each PMD_signal_detect_{*i*}, where *i* represents the lane number in the range 0:3, shall be continuously set in response to the magnitude of the optical signal on its associated lane, according to the requirements of Table 88–4.

88.5.6 PMD reset function

If the MDIO interface is implemented, and if PMD_reset is asserted, the PMD shall be reset as defined in 45.2.1.1.1.

88.5.7 PMD global transmit disable function (optional)

The PMD_global_transmit_disable function is optional and allows all of the optical transmitters to be disabled.

- a) When the PMD_global_transmit_disable variable is set to one, this function shall turn off all of the optical transmitters so that each transmitter meets the requirements of the average launch power of the OFF transmitter in Table 88–7.
- b) If a PMD_fault is detected, then the PMD may set the PMD_global_transmit_disable to one, turning off the optical transmitter in each lane.

88.5.8 PMD lane-by-lane transmit disable function

The PMD_transmit_disable_{*i*} (where *i* represents the lane number in the range 0:3) function is optional and allows the optical transmitters in each lane to be selectively disabled.

- a) When a PMD_transmit_disable_{*i*} variable is set to one, this function shall turn off the optical transmitter associated with that variable so that the transmitter meets the requirements of the average launch power of the OFF transmitter in Table 88–7.
- b) If a PMD_fault is detected, then the PMD may set each PMD_transmit_disable_{*i*} to one, turning off the optical transmitter in each lane.

If the optional `PMD_transmit_disable_i` function is not implemented in MDIO, an alternative method shall be provided to independently disable each transmit lane for testing purposes.

88.5.9 PMD fault function (optional)

If the PMD has detected a local fault on any of the transmit or receive paths, the PMD shall set `PMD_fault` to one.

If the MDIO interface is implemented, `PMD_fault` shall be mapped to the fault bit as specified in 45.2.1.2.1.

88.5.10 PMD transmit fault function (optional)

If the PMD has detected a local fault on any transmit lane, the PMD shall set the `PMD_transmit_fault` variable to one.

If the MDIO interface is implemented, `PMD_transmit_fault` shall be mapped to the transmit fault bit as specified in 45.2.1.7.4.

88.5.11 PMD receive fault function (optional)

If the PMD has detected a local fault on any receive lane, the PMD shall set the `PMD_receive_fault` variable to one.

If the MDIO interface is implemented, `PMD_receive_fault` shall be mapped to the receive fault bit as specified in 45.2.1.7.5.

88.6 Wavelength-division-multiplexed lane assignments

The wavelength range for each lane of the 100GBASE-LR4 and 100GBASE-ER4 PMDs is defined in Table 88-5. The center frequencies are members of the frequency grid for 100 GHz spacing and above defined in ITU-T G.694.1 and are spaced at 800 GHz.

Table 88-5—Wavelength-division-multiplexed lane assignments

Lane	Center frequency	Center wavelength	Wavelength range
L ₀	231.4 THz	1295.56 nm	1294.53 to 1296.59 nm
L ₁	230.6 THz	1300.05 nm	1299.02 to 1301.09 nm
L ₂	229.8 THz	1304.58 nm	1303.54 to 1305.63 nm
L ₃	229 THz	1309.14 nm	1308.09 to 1310.19 nm

NOTE—There is no requirement to associate a particular electrical lane with a particular optical lane, as the PCS is capable of receiving lanes in any arrangement.

88.7 PMD to MDI optical specifications for 100GBASE-LR4 and 100GBASE-ER4

The operating ranges for the 100GBASE-LR4 and 100GBASE-ER4 PMDs are defined in Table 88-6. A 100GBASE-LR4 or 100GBASE-ER4 compliant PMD operates on type B1.1, B1.3, or B6_A single-mode fibers according to the specifications defined in Table 88-14. A PMD that exceeds the operating range

requirement while meeting all other optical specifications is considered compliant (e.g., a 100GBASE-LR4 PMD operating at 12.5 km meets the operating range requirement of 2 m to 10 km). The 100GBASE-ER4 PMD interoperates with the 100GBASE-LR4 PMD provided that the channel requirements for 100GBASE-LR4 are met.

Table 88–6—100GBASE-LR4 and 100GBASE-ER4 operating ranges

PMD type	Required operating range
100GBASE-LR4	2 m to 10 km
100GBASE-ER4	2 m to 30 km
	2 m to 40 km ^a

^aLinks longer than 30 km for the same link power budget are considered engineered links. Attenuation for such links needs to be less than the worst case specified for B1.1, B1.3, or B6_A single-mode fiber.

88.7.1 100GBASE-LR4 and 100GBASE-ER4 transmitter optical specifications

The 100GBASE-LR4 transmitter shall meet the specifications defined in Table 88–7 per the definitions in 88.8. The 100GBASE-ER4 transmitter shall meet the specifications defined in Table 88–7 per the definitions in 88.8.

Table 88–7—100GBASE-LR4 and 100GBASE-ER4 transmit characteristics

Description	100GBASE-LR4	100GBASE-ER4	Unit
Signaling rate, each lane (range)	25.78125 ± 100 ppm		GBd
Lane wavelengths (range)	1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19		nm
Side-mode suppression ratio (SMSR), (min)	30		dB
Total average launch power (max)	10.5	8.9	dBm
Average launch power, each lane (max)	4.5	2.9	dBm
Average launch power, each lane ^a (min)	-4.3	-2.9	dBm
Optical Modulation Amplitude (OMA), each lane (max)	4.5		dBm
Optical Modulation Amplitude (OMA), each lane (min)	-1.3 ^b	0.1	dBm
Difference in launch power between any two lanes (OMA) (max)	5	—	dB
Difference in launch power between any two lanes (Average and OMA) (max)	—	3.6	
Launch power in OMA minus TDP, each lane (min)	-2.3	—	dBm
Transmitter and dispersion penalty (TDP), each lane (max)	2.2	2.5	dB
Average launch power of OFF transmitter, each lane (max)	-30		dBm
Extinction ratio (min)	4	8	dB
RIN ₂₀ OMA (max)	-130		dB/Hz
Optical return loss tolerance (max)	20		dB
Transmitter reflectance ^c (max)	-12		dB
Transmitter eye mask definition {X1, X2, X3, Y1, Y2, Y3}	{0.25, 0.4, 0.45, 0.25, 0.28, 0.4}		

^aAverage launch power, each lane (min) is informative and not the principal indicator of signal strength. A transmitter with launch power below this value cannot be compliant; however, a value above this does not ensure compliance.

^bEven if the TDP < 1 dB, the OMA (min) must exceed this value.

^cTransmitter reflectance is defined looking into the transmitter.

88.7.2 100GBASE-LR4 and 100GBASE-ER4 receive optical specifications

The 100GBASE-LR4 receiver shall meet the specifications defined in Table 88–8 per the definitions in 88.8. The 100GBASE-ER4 receiver shall meet the specifications defined in Table 88–8 per the definitions in 88.8.

Table 88–8—100GBASE-LR4 and 100GBASE-ER4 receive characteristics

Description	100GBASE-LR4	100GBASE-ER4	Unit
Signaling rate, each lane (range)	25.78125 ± 100 ppm		GBd
Lane wavelengths (range)	1294.53 to 1296.59 1299.02 to 1301.09 1303.54 to 1305.63 1308.09 to 1310.19		nm
Damage threshold ^a	5.5		dBm
Average receive power, each lane (max)	4.5 ^b		dBm
Average receive power, each lane ^c (min)	–10.6	–20.9	dBm
Receive power, each lane (OMA) (max)	4.5		dBm
Difference in receive power between any two lanes (OMA) (max)	5.5	—	dB
Difference in receive power between any two lanes (Average and OMA) (max)	—	4.5	
Receiver reflectance (max)	–26		dB
Receiver sensitivity (OMA), each lane ^d (max)	–8.6	–21.4	dBm
Receiver 3 dB electrical upper cutoff frequency, each lane (max)	31		GHz
Stressed receiver sensitivity (OMA), each lane ^e (max)	–6.8	–17.9	dBm
Conditions of stressed receiver sensitivity test			
Vertical eye closure penalty, ^f each lane	1.8	3.5	dB
Stressed eye J2 Jitter, ^f each lane	0.3		UI
Stressed eye J9 Jitter, ^f each lane	0.47		UI

^aThe receiver shall be able to tolerate, without damage, continuous exposure to an optical input signal having this average power level.

^bThe average receive power, each lane (max) for 100GBASE-ER4 is larger than the 100GBASE-ER4 transmitter value to allow compatibility with 100GBASE-LR4 units at short distances.

^cAverage receive power, each lane (min) is informative and not the principal indicator of signal strength. A received power below this value cannot be compliant; however, a value above this does not ensure compliance.

^dReceiver sensitivity (OMA), each lane (max) is informative.

^eMeasured with conformance test signal at TP3 (see 88.8.10) for BER = 10^{–12}.

^fVertical eye closure penalty, stressed eye J2 Jitter, and stressed eye J9 Jitter are test conditions for measuring stressed receiver sensitivity. They are not characteristics of the receiver.

88.7.3 100GBASE-LR4 and 100GBASE-ER4 illustrative link power budgets

Illustrative power budgets and penalties for 100GBASE-LR4 and 100GBASE-ER4 channels are shown in Table 88-9.

Table 88-9—100GBASE-LR4 and 100GBASE-ER4 illustrative link power budgets

Parameter	100GBASE-LR4	100GBASE-ER4		Unit
Power budget (for maximum TDP)	8.5	—		dB
Power budget	—	21.5		dB
Operating distance	10	30	40 ^a	km
Channel insertion loss	6.3 ^b	15	18	dB
Maximum discrete reflectance	−26	−26		dB
Allocation for penalties ^c (for maximum TDP)	2.2	—		dB
Allocation for penalties ^c	—	3.5		
Additional insertion loss allowed	0	3	0	dB

^aLinks longer than 30 km are considered engineered links. Attenuation for such links needs to be less than the worst case for B1.1, B1.3, or B6_A single-mode cabled optical fiber

^bThe channel insertion loss is calculated using the maximum distance specified in Table 88-6 for 100GBASE-LR4 and fiber attenuation of 0.43 dB/km at 1295 nm plus an allocation for connection and splice loss given in 88.11.2.1.

^cLink penalties are used for link budget calculations. They are not requirements and are not meant to be tested.

88.8 Definition of optical parameters and measurement methods

All transmitter optical measurements shall be made through a short patch cable, between 2 m and 5 m in length, unless otherwise specified.

88.8.1 Test patterns for optical parameters

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 88-11 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Any of the test patterns given for a particular test in Table 88-11 may be used to perform that test. The test patterns used in this clause are shown in Table 88-10.

88.8.2 Wavelength

The wavelength of each optical lane shall be within the ranges given in Table 88-5 if measured per TIA/EIA-455-127-A or IEC 61280-1-3. The lane under test is modulated using the test pattern defined in Table 88-11.

Table 88–10—Test patterns

Pattern	Pattern description	Defined in
Square wave	Square wave (8 ones, 8 zeros)	83.5.10
3	PRBS31	83.5.10
4	PRBS9	83.5.10
5	Scrambled idle	82.2.10

Table 88–11—Test-pattern definitions and related subclauses

Parameter	Pattern	Related subclause
Wavelength	3, 5 or valid 100GBASE-R signal	88.8.2
Side mode suppression ratio	3, 5 or valid 100GBASE-R signal	—
Average optical power	3, 5 or valid 100GBASE-R signal	88.8.3
Optical modulation amplitude (OMA)	Square wave or 4	88.8.4
Transmitter and dispersion penalty (TDP)	3 or 5	88.8.5
Extinction ratio	3, 5 or valid 100GBASE-R signal	88.8.6
RIN ₂₀ OMA	Square wave or 4	88.8.7
Transmitter optical waveform	3, 5 or valid 100GBASE-R signal	88.8.8
Stressed receiver sensitivity	3 or 5	88.8.10
Calibration of OMA for receiver tests	Square wave or 4	87.8.11
Vertical eye closure penalty calibration	3 or 5	87.8.11
Receiver 3 dB electrical upper cutoff frequency	3, 5 or valid 100GBASE-R signal	88.8.11

88.8.3 Average optical power

The average optical power of each lane shall be within the limits given in Table 88–7 if measured using the methods given in IEC 61280–1–1, with the sum of the optical power from all of the lanes not under test below –30 dBm, per the test setup in Figure 53–6.

88.8.4 Optical Modulation Amplitude (OMA)

OMA shall be as defined in 52.9.5 for measurement with a square wave (8 ones, 8 zeros) test pattern or 68.6.2 (from the variable MeasuredOMA in 68.6.6.2) for measurement with a PRBS9 test pattern, with the exception that each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.

88.8.5 Transmitter and dispersion penalty (TDP)

Transmitter and dispersion penalty (TDP) shall be as defined in 52.9.10 with the exception that each optical lane is tested individually using an optical filter to separate the lane under test from the others. The measurement procedure for 100GBASE-LR4 and 100GBASE-ER4 is detailed in 88.8.5.1 to 88.8.5.4.

The optical filter passband ripple shall be limited to 0.5 dB peak-to-peak and the isolation is chosen such that the ratio of the power in the lane being measured to the sum of the powers of all of the other lanes is greater than 20 dB (see ITU-T G.959.1 Annex B). The lanes not under test shall be operating with PRBS31 or valid 100GBASE-R bit streams.

88.8.5.1 Reference transmitter requirements

The reference transmitter is a high-quality instrument-grade device, which can be implemented by a CW laser modulated by a high-performance modulator. The basic requirements are as follows:

- Rise/fall times of less than 12 ps at 20% to 80%.
- The output optical eye is symmetric and passes the transmitter optical waveform test of 88.8.8.
- In the center 20% region of the eye, the worst-case vertical eye closure penalty as defined in 87.8.11.2 is less than 0.5 dB.
- Total Jitter less than 0.2 UI peak-to-peak.
- RIN of less than -138 dB/Hz.

88.8.5.2 Channel requirements

The transmitter is tested using an optical channel that meets the requirements listed in Table 88–12.

Table 88–12—Transmitter compliance channel specifications

PMD type	Dispersion ^a (ps/nm)		Insertion loss ^b	Optical return loss ^c	Max mean DGD
	Minimum	Maximum			
100GBASE-LR4	$0.2325 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$	$0.2325 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$	Minimum	20 dB	0.8 ps
100GBASE-ER4	$0.93 \cdot \lambda \cdot [1 - (1324 / \lambda)^4]$	$0.93 \cdot \lambda \cdot [1 - (1300 / \lambda)^4]$	Minimum	20 dB	0.8 ps

^aThe dispersion is measured for the wavelength of the device under test (λ in nm). The coefficient assumes 10 km for 100GBASE-LR4 and 40 km for 100GBASE-ER4.

^bThere is no intent to stress the sensitivity of the BERT's optical receiver.

^cThe optical return loss is applied at TP2.

A 100GBASE-LR4 or 100GBASE-ER4 transmitter is to be compliant with a total dispersion at least as negative as the “minimum dispersion” and at least as positive as the “maximum dispersion” columns specified in Table 88–12 for the wavelength of the device under test. This may be achieved with channels consisting of fibers with lengths chosen to meet the dispersion requirements.

To verify that the fiber has the correct amount of dispersion, the measurement method defined in IEC 60793–1–42 may be used. The measurement is made in the linear power regime of the fiber.

The channel provides an optical return loss specified in Table 88–12. The state of polarization of the back reflection is adjusted to create the greatest RIN.

The mean DGD of the channel is to be less than the value specified in Table 88–12.

88.8.5.3 Reference receiver requirements

The reference receiver is required to have the bandwidth given in 88.8.8. The sensitivity of the reference receiver is limited by Gaussian noise. The receiver has minimal threshold offset, deadband, hysteresis, baseline wander, deterministic jitter, or other distortions. Decision sampling has minimal uncertainty and setup/hold times.

The nominal sensitivity of the reference receiver, S , is measured in OMA using the setup of Figure 52–12 without the test fiber and with the transversal filter removed. The sensitivity S must be corrected for any significant reference transmitter impairments including any vertical eye closure. It is measured while sampling at the eye center or corrected for off-center sampling. It is calibrated at the wavelength of the transmitter under test.

For all transmitter and dispersion penalty measurements, determination of the center of the eye is required. Center of the eye is defined as the time halfway between the left and right sampling points within the eye where the measured BER is greater than or equal to 1×10^{-3} .

The clock recovery unit (CRU) used in the TDP measurement has a corner frequency of 10 MHz and a slope of 20 dB/decade. When using a clock recovery unit as a clock for BER measurements, passing of low-frequency jitter from the data to the clock removes this low-frequency jitter from the measurement.

88.8.5.4 Test procedure

The test procedure is as defined in 52.9.10.4 with the exception that all lanes are operational in both directions (transmit and receive), each lane is tested individually using an optical filter to separate the lane under test from the others and the BER of 1×10^{-12} is for the lane under test on its own.

88.8.6 Extinction ratio

The extinction ratio of each lane shall be within the limits given in Table 88–7 if measured using the methods specified in IEC 61280–2–2, with the sum of the optical power from all of the lanes not under test below –30 dBm. The extinction ratio is measured using the test pattern defined in Table 88–11.

NOTE—Extinction ratio and OMA are defined with different test patterns (see Table 88–11).

88.8.7 Relative Intensity Noise (RIN_{20OMA})

RIN shall be as defined by the measurement methodology of 52.9.6 with the following exceptions:

- a) The optical return loss is 20 dB.
- b) Each lane may be tested individually with the sum of the optical power from all of the lanes not under test being below –30 dBm, or if other lanes are operating, a suitable optical filter may be used to separate the lane under test.
- c) The upper –3 dB limit of the measurement apparatus is to be approximately equal to the signaling rate (i.e., 25.8 GHz).

88.8.8 Transmitter optical waveform (transmit eye)

The required optical transmitter pulse shape characteristics are specified in the form of a mask of the transmitter eye diagram as shown in Figure 86–4. The transmitter optical waveform of a port transmitting the test pattern specified in Table 88–11 shall meet specifications according to the methods specified in 86.8.4.6.1 with the exception that the clock recovery unit's high-frequency corner bandwidth is 10 MHz. The filter nominal reference frequency f_c is 19.34 GHz and the filter tolerances are as specified for STM-64

in ITU-T G.691. Compensation may be made for variation of the reference receiver filter response from an ideal fourth-order Bessel-Thomson response.

88.8.9 Receiver sensitivity

Receiver sensitivity, which is defined for an ideal input signal, is informative and compliance is not required. If measured, the test signal should have negligible impairments such as intersymbol interference (ISI), rise/fall times, jitter and RIN. Instead, the normative requirement for receivers is stressed receiver sensitivity.

88.8.10 Stressed receiver sensitivity

Stressed receiver sensitivity shall be within the limits given in Table 88–8 if measured using the method defined in 87.8.11 with the following exceptions:

- a) Added sinusoidal jitter is as specified in Table 88–13.
- b) The stressed eye J2 Jitter, stressed eye J9 Jitter, and vertical eye closure penalty are as given in Table 88–8.
- c) The test pattern is as given in Table 88–11.
- d) The reference receiver used to verify the conformance test signal is required to have the bandwidth given in 88.8.8.

Table 88–13—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak (UI)
$f < 100$ kHz	Not specified
100 kHz $< f \leq 10$ MHz	$5 \times 10^5 / f$
10 MHz $< f < 10$ LB ^a	0.05

^aLB = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

88.8.11 Receiver 3 dB electrical upper cutoff frequency

The receiver 3 dB electrical upper cutoff frequency shall be within the limits given in Table 88–8 if measured as described in 52.9.11. Each optical lane is measured using an optical signal or signals with its/their wavelength within the specified wavelength range of the lane to be tested. The test may use an electrical combiner and one optical source as in 53.9.13.

88.9 Safety, installation, environment, and labeling

88.9.1 General safety

All equipment subject to this clause shall conform to IEC 60950-1.

88.9.2 Laser safety

100GBASE-LR4 and 100GBASE-ER4 optical transceivers shall conform to Class 1 laser requirements as defined in IEC 60825–1 and IEC 60825–2, under any condition of operation. This includes single fault conditions whether coupled into a fiber or out of an open bore.

Conformance to additional laser safety standards may be required for operation within specific geographic regions.

Laser safety standards and regulations require that the manufacturer of a laser product provide information about the product's laser, safety features, labeling, use, maintenance, and service. This documentation explicitly defines requirements and usage restrictions on the host system necessary to meet these safety certifications.²⁰

88.9.3 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

88.9.4 Environment

Normative specifications in this clause shall be met by a system integrating a 100GBASE-LR4 or 100GBASE-ER4 PMD over the life of the product while the product operates within the manufacturer's range of environmental, power, and other specifications.

It is recommended that manufacturers indicate in the literature associated with the PHY the operating environmental conditions to facilitate selection, installation, and maintenance.

It is recommended that manufacturers indicate, in the literature associated with the components of the optical link, the distance and operating environmental conditions over which the specifications of this clause will be met.

88.9.5 Electromagnetic emission

A system integrating a 100GBASE-LR4 or 100GBASE-ER4 PMD shall comply with applicable local and national codes for the limitation of electromagnetic interference.

88.9.6 Temperature, humidity, and handling

The optical link is 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 standard.

88.9.7 PMD labeling requirements

It is recommended that each PHY (and supporting documentation) be labeled in a manner visible to the user, with at least the applicable safety warnings and the applicable port type designation (e.g., 100GBASE-LR4).

Labeling requirements for Class 1 lasers are given in the laser safety standards referenced in 88.9.2.

²⁰A host system that fails to meet the manufacturer's requirements and/or usage restrictions may emit laser radiation in excess of the safety limits of one or more safety standards. In such a case, the host manufacturer is required to obtain its own laser safety certification.

88.10 Fiber optic cabling model

The fiber optic cabling model is shown in Figure 88–3.




Figure 88–3—Fiber optic cabling model

The channel insertion loss is given in Table 88–14. A channel may contain additional connectors as long as the optical characteristics of the channel, such as attenuation, dispersion, reflections, and polarization mode dispersion meet the specifications. Insertion loss measurements of installed fiber cables are made in accordance with ANSI/TIA/EIA–526–7/method A–1. The fiber optic cabling model (channel) defined here is the same as a simplex fiber optic link segment. The term channel is used here for consistency with generic cabling standards.

Table 88–14—Fiber optic cabling (channel) characteristics

Description	100GBASE–LR4	100GBASE–ER4		Unit
Operating distance (max)	10	30	40	km
Channel insertion loss ^{a, b} (max)	6.3	18	18	dB
Channel insertion loss (min)	0	0		dB
Positive dispersion ^b (max)	9.5	28	36	ps/nm
Negative dispersion ^b (min)	–28.5	–85	–114	ps/nm
DGD_max ^c	8	10.3	10.3	ps
Optical return loss (min)	21	21	21	dB

^aThese channel insertion loss values include cable, connectors, and splices.

^bOver the wavelength range 1294.53 nm to 1310.19 nm

^cDifferential Group Delay (DGD) is the time difference at reception between the fractions of a pulse that were transmitted in the two principal states of polarization of an optical signal. DGD_max is the maximum differential group delay that the system must tolerate.

88.11 Characteristics of the fiber optic cabling (channel)

The 100GBASE–LR4 and 100GBASE–ER4 fiber optic cabling shall meet the specifications defined in Table 88–14. The fiber optic cabling consists of one or more sections of fiber optic cable and any intermediate connections required to connect sections together.

88.11.1 Optical fiber cable

The fiber optic cable requirements are satisfied by cables containing IEC 60793–2–50 type B1.1 (dispersion un-shifted single-mode), type B1.3 (low water peak single-mode), or type B6_A (bend insensitive) fibers and the requirements in Table 88–15 where they differ.

Table 88–15—Optical fiber and cable characteristics

Description	Value	Unit
Nominal fiber specification wavelength	1310	nm
Cabled optical fiber attenuation (max)	0.43 ^a or 0.5 ^b	dB/km
Zero dispersion wavelength (λ_0)	$1300 \leq \lambda_0 \leq 1324$	nm
Dispersion slope (max) (S_0)	0.093	ps/nm ² km

^aThe 0.43 dB/km at 1295 nm attenuation for optical fiber cables is derived from Appendix I of ITU-T G.695.

^bThe 0.5 dB/km attenuation is provided for Outside Plant cable as defined in ANSI/TIA/EIA 568-B.3-2000. Using 0.5 dB/km may not support operation at 10 km for 100GBASE–LR4 or 40 km for 100GBASE–ER4.

88.11.2 Optical fiber connection

An optical fiber connection, as shown in Figure 88–3, consists of a mated pair of optical connectors.

88.11.2.1 Connection insertion loss

The maximum link distance is based on an allocation of 2 dB total connection and splice loss. For example, this allocation supports four connections with an average insertion loss per connection of 0.5 dB. Connections with different loss characteristics may be used provided the requirements of Table 88–14 are met.

88.11.2.2 Maximum discrete reflectance

The maximum discrete reflectance shall be less than –26 dB.

88.11.3 Medium Dependent Interface (MDI) requirements

The 100GBASE–LR4 or 100GBASE–ER4 PMD is coupled to the fiber optic cabling at the MDI. The MDI is the interface between the PMD and the “fiber optic cabling” (as shown in Figure 88–3). Examples of an MDI include the following:

- a) Connectorized fiber pigtail
- b) PMD receptacle

When the MDI is a connector plug and receptacle connection, it shall meet the interface performance specifications of IEC 61753–1–1 and IEC 61753–021–2.

NOTE—Transmitter compliance testing is performed at TP2 as defined in 88.5.1, not at the MDI.

88.12 Protocol implementation conformance statement (PICS) proforma for Clause 88, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4²¹

88.12.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 88, Physical Medium Dependent sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in Clause 21.

88.12.2 Identification

88.12.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

88.12.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Clause 88, Physical Medium Dependent sublayer and medium, type 100GBASE-LR4 and 100GBASE-ER4
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

²¹Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

88.12.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*LR4	100GBASE-LR4 PMD	88.7	Device supports requirements for 100GBASE-LR4 PHY	O.1	Yes [] No []
*ER4	100GBASE-ER4 PMD	88.7	Device supports requirements for 100GBASE-ER4 PHY	O.1	Yes [] No []
*INS	Installation / cable	88.10	Items marked with INS include installation practices and cable specifications not applicable to a PHY manufacturer	O	Yes [] No []
CTP1	Reference point TP1 exposed and available for testing	88.5.1	This point may be made available for use by implementors to certify component conformance	O	Yes [] No []
CTP4	Reference point TP4 exposed and available for testing	88.5.1	This point may be made available for use by implementors to certify component conformance	O	Yes [] No []
CDC	Delay constraints	88.3.1	Device conforms to delay constraints	M	Yes []
CSC	Skew constraints	88.3.2	Device conforms to Skew and Skew Variation constraints	M	Yes []
*MD	MDIO capability	88.4	Registers and interface supported	O	Yes [] No []

88.12.4 PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, types 100GBASE-LR4 and 100GBASE-ER4

88.12.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CF1	Compatible with 100GBASE-R PCS and PMA	88.1		M	Yes []
CF2	Integration with management functions	88.1		O	Yes [] No []
CF3	Transmit function	88.5.2	Conveys bits from PMD service interface to MDI	M	Yes []
CF4	Optical multiplexing and delivery to the MDI	88.5.2	Optically multiplexes the four optical signal streams for delivery to the MDI	M	Yes []
CF5	Mapping between optical signal and logical signal for transmitter	88.5.2	Higher optical power is a one	M	Yes []
CF6	Receive function	88.5.3	Conveys bits from MDI to PMD service interface	M	Yes []
CF7	Conversion of four optical signals to four electrical signals	88.5.3	For delivery to the PMD service interface	M	Yes []
CF8	Mapping between optical signal and logical signal for receiver	88.5.3	Higher optical power is a one	M	Yes []
CF9	Global Signal Detect function	88.5.4	Report to the PMD service interface the message PMD:IS_SIGNAL.indication(SIGNAL_DETECT)	M	Yes []
CF10	Global Signal Detect behavior	88.5.4	SIGNAL_DETECT is a global indicator of the presence of optical signals on all four lanes	M	Yes []
CF11	Lane-by-lane Signal Detect function	88.5.5	Sets PMD_signal_detect _i values on a lane-by-lane basis per requirements of Table 88-4	MD:O	Yes [] No [] N/A []
CF12	PMD reset function	88.5.6	Resets the PMD sublayer	MD:O	Yes [] No [] N/A []

88.12.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
CM1	Management register set	88.4		MD:M	Yes [] N/A []
CM2	Global transmit disable function	88.5.7	Disables all of the optical transmitters with the PMD_global_transmit_disable variable	MD:O	Yes [] No [] N/A []
CM3	PMD_lane_by_lane_transmit_disable function	88.5.8	Disables the optical transmitter on the lane associated with the PMD_transmit_disable_i variable	MD:O.2	Yes [] No [] N/A []
CM4	PMD lane-by-lane transmit disable	88.5.8	Disables each optical transmitter independently if CM3 = No	O.2	Yes [] No []
CM5	PMD_fault function	88.5.9	Sets PMD_fault to one if any local fault is detected	MD:O	Yes [] No [] N/A []
CM6	PMD_transmit_fault function	88.5.10	Sets PMD_transmit_fault to one if a local fault is detected on any transmit lane	MD:O	Yes [] No [] N/A []
CM7	PMD_receive_fault function	88.5.11	Sets PMD_receive_fault to one if a local fault is detected on any receive lane	MD:O	Yes [] No [] N/A []

88.12.4.3 PMD to MDI optical specifications for 100GBASE-LR4

Item	Feature	Subclause	Value/Comment	Status	Support
CLR1	Transmitter meets specifications in Table 88-7	88.7.1	Per definitions in 88.8	LR4:M	Yes [] N/A []
CLR2	Receiver meets specifications in Table 88-8	88.7.2	Per definitions in 88.8	LR4:M	Yes [] N/A []

88.12.4.4 PMD to MDI optical specifications for 100GBASE-ER4

Item	Feature	Subclause	Value/Comment	Status	Support
CER1	Transmitter meets specifications in Table 88-7	88.7.1	Per definitions in 88.8	ER4:M	Yes [] N/A []
CER2	Receiver meets specifications in Table 88-8	88.7.2	Per definitions in 88.8	ER4:M	Yes [] N/A []

88.12.4.5 Optical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
COM1	Measurement cable	88.8	2 m to 5 m in length	M	Yes []
COM2	Center wavelength	88.8.2	Per TIA/EIA-455-127-A or IEC 61280-1-3 under modulated conditions	M	Yes []
COM3	Average optical power	88.8.3	Per IEC 61280-1-1	M	Yes []
COM4	OMA measurements	88.8.4	Each lane	M	Yes []
COM5	Transmitter and dispersion penalty	88.8.5	Each lane	M	Yes []
COM6	Extinction ratio	88.8.6	Per IEC 61280-2-2	M	Yes []
COM7	RIN ₂₀ OMA measurement procedure	88.8.7	Each lane	M	Yes []
COM8	Transmit eye	88.8.8	Each lane	M	Yes []
COM9	Stressed receiver sensitivity	88.8.10	Each lane	M	Yes []
COM10	Receiver 3 dB electrical upper cutoff frequency	88.8.11	Each lane	M	Yes []

88.12.4.6 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
CES1	General safety	88.9.1	Conforms to IEC 60950-1	M	Yes []
CES2	Laser safety—IEC Class 1	88.9.2	Conforms to Class 1 laser requirements defined in IEC 60825-1 and IEC 60825-2	M	Yes []
CES3	Electromagnetic interference	88.9.5	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes []

88.12.4.7 Characteristics of the fiber optic cabling and MDI

Item	Feature	Subclause	Value/Comment	Status	Support
COC1	Fiber optic cabling	88.11	Meets requirements specified in Table 88-14	INS:M	Yes [] N/A []
COC2	Maximum discrete reflectance	88.11.2.2	Less than -26 dB	INS:M	Yes [] N/A []
COC3	MDI requirements	88.11.3	Meets IEC 61753-1-1 and IEC 61753-021-2	INS:M	Yes [] N/A []

Annex A

(informative)

Bibliography

Insert the following references in alphabetical order and renumber the list:

[Bx1] ITU-T G.709—Interfaces for the Optical Transport Network (OTN).

NOTE—This annex is numbered in correspondence to its associated clause; i.e., Annex 4A corresponds to Clause 4.

Annex 4A

(normative)

Simplified full duplex media access control

This annex is based on the Clause 4 MAC, with simplifications for use in networks that do not require the half duplex operational mode. Additional functionality is included for managing Physical Layer congestion and for support of interframe spacing outside this sublayer. This annex stands alone and does not rely on information within Clause 4 to be implemented.

4A.4.2 MAC parameters

Insert the following note below Table 4A-2 (above the warning box) for 40 Gb/s and 100 Gb/s MAC data rates, and renumber notes as appropriate:

NOTE 4—For 40 Gb/s and 100 Gb/s operation, the received interPacketGap (the spacing between two packets, from the last bit of the FCS field of the first packet to the first bit of the Preamble of the second packet) can have a minimum value of 8 BT (bit times), as measured at the XLGMII or CGMII receive signals at the DTE due to clock tolerance and lane alignment requirements.

Annex 31B

(normative)

MAC Control PAUSE operation

31B.3.7 Timing considerations for PAUSE operation

Change subclause 31B.3.7 (IEEE Std 802.3-2008/Cor 1-2009) as follows:

In a full duplex mode DTE, it is possible to receive PAUSE frames asynchronously with respect to the transmission of MAC frames. For effective flow control, it is necessary to place an upper bound on the length of time that a DTE can transmit data frames after receiving a valid PAUSE frame with a non-zero `pause_time` `request_operand`.

Reception of a PAUSE frame shall not affect the transmission of a MAC frame that has been submitted by the MAC Control sublayer to the underlying MAC (i.e., the `MAC:MA_DATA.request` service primitive is synchronous, and is never interrupted).

At operating speeds of 100 Mb/s or less, a station that implements an exposed MII, shall not begin to transmit a (new) frame (assertion of `TX_EN` at the MII, see 22.2.2.3) more than `pause_quantum` bit times after the reception of a valid PAUSE frame (de-assertion of `RX_DV` at the MII, see 22.2.2.6) that contains a non-zero value of `pause_time`. Stations that do not implement an exposed MII, shall measure this time at the MDI, with the timing specification increased to $(\text{pause_quantum} + 64)$ bit times.

At an operating speed of 1000 Mb/s, a station shall not begin to transmit a (new) frame more than two `pause_quantum` bit times after the reception of a valid PAUSE frame that contains a non-zero value of `pause_time`, as measured at the MDI.

At operating speeds of 10 Gb/s, a station with a 10GBASE-T PHY shall not begin to transmit a (new) frame more than 74 `pause_quantum` bit times after the reception of a valid PAUSE frame that contains a non-zero value of `pause_time`, as measured at the MDI. A station using any other 10 Gb/s PHY shall not begin to transmit a (new) frame more than 60 `pause_quantum` bit times after the reception of a valid PAUSE frame that contains a non-zero value of `pause_time`, as measured at the MDI.

At operating speeds of 40 Gb/s, a station shall not begin to transmit a (new) frame more than 118 `pause_quantum` bit times after the reception of a valid PAUSE frame that contains a non-zero value of `pause_time`, as measured at the MDI.

At operating speeds of 100 Gb/s, a station shall not begin to transmit a (new) frame more than 394 `pause_quantum` bit times after the reception of a valid PAUSE frame that contains a non-zero value of `pause_time`, as measured at the MDI.

In addition to DTE and MAC Control delays, system designers should take into account the delay of the link segment when designing devices that implement the PAUSE operation.

31B.4 Protocol implementation conformance statement (PICS) proforma for MAC Control PAUSE operation²²

31B.4.3 Major capabilities/options

Change the table in 31B.4.3 as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
*PST	Support for transmit of PAUSE frames	31B.3.2	N/A	O	Yes [] No []
*MIIa	At operating speeds of 100 Mb/s or less, MII connection exists and is accessible for test.	31B.3.7	N/A	O	Yes [] No []
*MIIb	At operating speeds of 100 Mb/s or less, MII connection does not exist or is not accessible for test.	31B.3.7	N/A	O	Yes [] No []
*MIIc	At operating speeds above 100 of <u>1000</u> Mb/s.	31B.3.7	N/A	O	Yes [] No []
*MIId	<u>At operating speeds of 10 Gb/s with PHY types other than 10GBASE-T.</u>	<u>31B.3.7</u>	<u>N/A</u>	<u>O</u>	<u>Yes []</u> <u>No []</u>
*MIIe	<u>At operating speeds of 10 Gb/s with PHY types of 10GBASE-T.</u>	<u>31B.3.7</u>	<u>N/A</u>	<u>O</u>	<u>Yes []</u> <u>No []</u>
*MIIf	<u>At operating speeds of 40 Gb/s.</u>	<u>31B.3.7</u>	<u>N/A</u>	<u>O</u>	<u>Yes []</u> <u>No []</u>
*MIIg	<u>At operating speeds of 100 Gb/s.</u>	<u>31B.3.7</u>	<u>N/A</u>	<u>O</u>	<u>Yes []</u> <u>No []</u>

²²Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

31B.4.6 PAUSE command MAC timing considerations

Change subclause 31B.4.6 (IEEE Std 802.3-2008/Cor 1-2009) as follows:

Item	Feature	Subclause	Value/Comment	Status	Support
TIM1	Effect of PAUSE frame on a frame already submitted to underlying MAC	31B.3.7	Has no effect	M	Yes []
	Delay from receiving valid PAUSE command, with non-zero value for pause_time, to cessation of transmission	31B.3.7	Measured as described		
TIM2	Measurement point for station with MII		Delay at MII \leq pause_quantum bits	MIIa: M	N/A [] M: Yes []
TIM3	Measurement point for station without MII at 100 Mb/s or less		Delay at MDI \leq (pause_quantum + 64) bits	MIIb: M	N/A [] M: Yes []
TIM4	Measurement point for station at 1000 Mb/s		Delay at MDI \leq (2 \times pause_quantum) bits	MIIc: M	N/A [] M: Yes []
TIM5	Measurement point for station at 10 Gb/s with PHY types other than 10GBASE-T		Delay at MDI \leq (60 \times pause_quantum) bits	MIId: M	N/A [] M: Yes []
TIM6	Measurement point for station at 10Gb/s with PHY type of 10GBASE-T.		Delay at MDI \leq (74 \times pause_quantum) bits	MIIE: M	N/A [] M: Yes []
<u>TIM7</u>	<u>Measurement point for station at 40 Gb/s</u>		<u>Delay at MDI \leq (118 \times pause_quantum) bits</u>	<u>MIIIf: M</u>	<u>N/A []</u> <u>M: Yes []</u>
<u>TIM8</u>	<u>Measurement point for station at 100 Gb/s</u>		<u>Delay at MDI \leq (394 \times pause_quantum) bits</u>	<u>MIIIg: M</u>	<u>N/A []</u> <u>M: Yes []</u>

Annex 69A

(normative)

Interference tolerance testing

69A.2.1 Pattern generator

Change second paragraph as follows:

For 10GBASE-KR ~~and 40GBASE-KR4~~, the peak-to-peak amplitude delivered by the pattern generator, as measured on a sequence of alternating ones and zeros, shall be no more than 800 mV, adjusted by a gain b_{TC} as defined in 69A.2.2, regardless of equalization setting.

69A.3 Test methodology

Change last paragraph as follows:

The interference ~~tolerance~~ tolerance test parameters are specified in Table 70-7 for 1000BASE-KX, in Table 71-7 for 10GBASE-KX4, and in Table 72-10 for 10GBASE-KR ~~and 40GBASE-KR4~~.

Annex 69B

(informative)

Interconnect characteristics

69B.4.1 Overview

Change Table 69B-1 as follows:

Table 69B-1—Insertion loss parameters

Parameter	1000BASE-KX	10GBASE-KX4	10GBASE-KR 40GBASE-KR4	Units
f_{\min}	0.05			GHz
f_{\max}	15.00			GHz
b_1	2.00×10^{-5}			
b_2	1.10×10^{-10}			
b_3	3.20×10^{-20}			
b_4	-1.20×10^{-30}			
f_1	0.125	0.312	1.000	GHz
f_2	1.250	3.125	6.000	GHz
f_a	0.100	0.100	0.100	GHz
f_b	1.250	3.125	5.15625	GHz

69B.4.3 Insertion loss

Replace Figure 69B-5 with the following:




Figure 69B-5—Insertion loss limit for 10GBASE-KR and 40GBASE-KR4

69B.4.4 Insertion loss deviation

Replace Figure 69B-6 with the following:

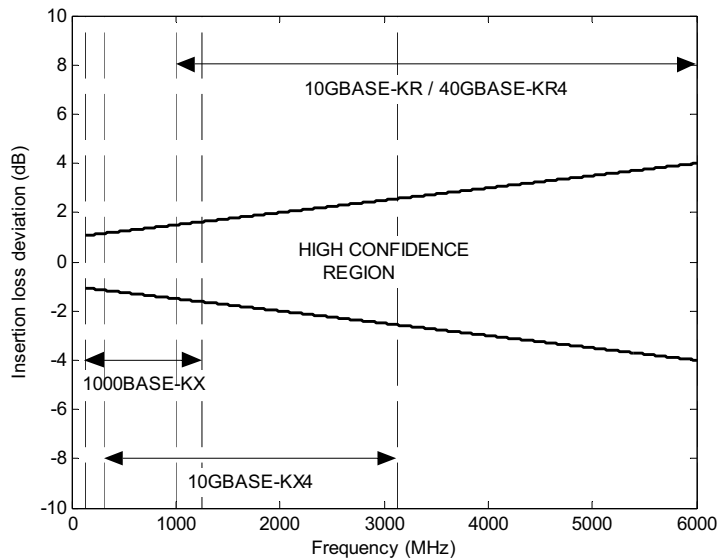


Figure 69B-6—Insertion loss deviation limits

69B.4.5 Return loss

Replace Figure 69B-7 with the following:




Figure 69B-7—Return loss limit

69B.4.6.4 Insertion loss to crosstalk ratio (ICR)

Replace Figure 69B-8 with the following:




Figure 69B-8—Insertion loss to crosstalk ratio limit

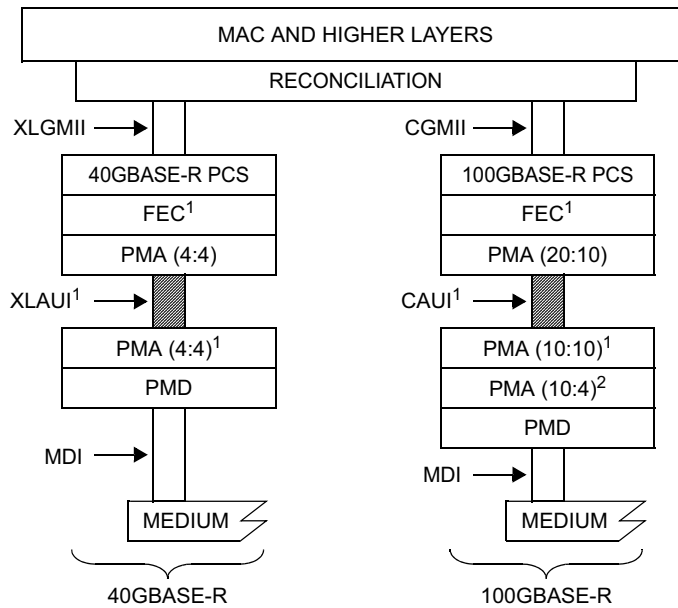
Annex 83A

(normative)

40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)

83A.1 Overview

This annex defines the functional and electrical characteristics for the optional 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI). Figure 83A–1 shows the relationships of the XLGMII, PMA, XLAUI, and PMD for 40 Gb/s and CGMII, PMA, CAUI, and PMD for 100 Gb/s.



CAUI = 100 Gb/s ATTACHMENT UNIT INTERFACE
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT
 XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
 NOTE 2—CONDITIONAL BASED ON PMD TYPE

Figure 83A–1—Example relationship of XLAUI and CAUI to IEEE 802.3 CSMA/CD LAN model

The purpose of the optional XLAUI or CAUI is to provide a flexible chip-to-chip and chip-module interconnect for 40 Gb/s or 100 Gb/s components. Annex 83A provides compliance requirements for XLAUI/CAUI transmitters and receivers while Annex 83B specifies the electrical requirements for the chip-module interconnection.

The XLAUI/CAUI allows interconnect distances of approximately 25 cm over printed circuit board including one connector, see 83A.4.

An example application of CAUI includes providing a physical connection between a ten-lane 100 Gb/s PMA and a 10:4 PMA mapping element. An example application of XLAUI is to provide lane extension for interfacing MAC and PHY components in a 40 Gb/s Ethernet system distributed across a circuit board.

The optional XLAUI/CAUI interface has the following characteristics:

- a) Independent transmit and receive data paths
- b) Differential AC coupled signaling with low voltage swing
- c) Self-timed interface
- d) Shared technology with other 40 Gb/s or 100 Gb/s interfaces
- e) Utilization of 64B/66B coding

83A.1.1 Summary of major concepts

The following is a list of the major concepts of XLAUI and CAUI:

- a) The optional XLAUI/CAUI interface can be inserted between PMA layers in the IEEE 802.3 CSMA/CD LAN model to transparently enable chip-to-chip communication
- b) The XLAUI is organized into four lanes, the CAUI is organized into ten lanes
- c) The XLAUI/CAUI interface is a parallel electrical interface with each lane running at a nominal rate of 10.3125 Gb/s

83A.1.2 Rate of operation

The XLAUI interface supports the 40 Gb/s data rate and the CAUI interface supports the 100 Gb/s data rate. For 40 Gb/s applications, the data stream shall be presented in four lanes as described in Clause 83. For 100 Gb/s applications, the data stream shall be presented in ten lanes as described in Clause 83. The data is 64B/66B coded. The nominal signaling rate for each lane in both 40 Gb/s and 100 Gb/s applications shall be 10.3125 Gb/s.

83A.2 XLAUI/CAUI link block diagram

XLAUI/CAUI link is illustrated in Figure 83A–2. XLAUI/CAUI channel is defined from the transmit pad to the receive pad including any AC coupling in the path.

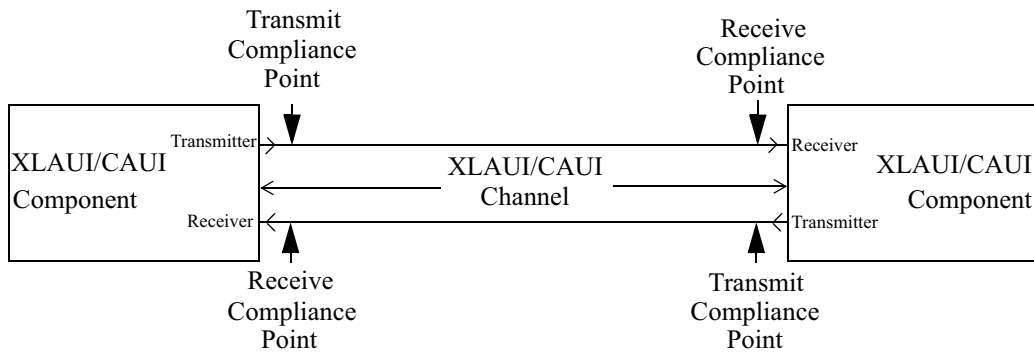


Figure 83A-2—Definition of transmit and receive compliance points

83A.2.1 Transmitter compliance points

The reference differential insertion loss, expressed in decibels, between the transmitter and the transmit compliance point is defined in Equation (83A-1) and illustrated in Figure 83A-3. The effects of differences between the actual insertion loss and the reference insertion loss are to be accounted in the measurements.

$$Insertion_loss(f) = -0.00086 + 0.2286\sqrt{f} + 0.08386f \text{ dB} \quad 0.01 \leq f \leq 11.1 \quad (83A-1)$$

where

$Insertion_loss(f)$ is the differential insertion loss at frequency f
 f is the frequency in GHz




Figure 83A-3—Insertion loss between transmitter and transmit compliance point

83A.2.2 Receiver compliance points

The reference differential insertion loss, expressed in decibels, between the receiver compliance point and the receiver is defined in Equation (83A–2) and illustrated in Figure 83A–4. The effects of differences between the actual insertion loss and the reference insertion loss are to be accounted in the measurements.

$$\text{Insertion_loss}(f) = -0.00086 + 0.2286\sqrt{f} + 0.08386f \quad \text{dB} \quad 0.01 \leq f \leq 11.1 \quad (83A-2)$$

where

$\text{Insertion_loss}(f)$ is the differential insertion loss at frequency f
 f is the frequency in GHz




Figure 83A–4—Insertion loss between receive compliance point and receiver

83A.3 XLAUI/CAUI electrical characteristics

The electrical characteristics of the XLAUI/CAUI interface are specified such that they can be applied within a variety of 40 Gb/s Ethernet or 100 Gb/s Ethernet equipment types. The electrical characteristics for XLAUI/CAUI shall meet the specifications defined in 83A.3.1, 83A.3.2, 83A.3.3, and 83A.3.4.

83A.3.1 Signal levels

The XLAUI/CAUI is a low-swing AC coupled differential interface. AC coupling allows for interoperability between components operating from different supply voltages. Differential signal swings are defined in the following subclauses, and depend on several factors such as transmitter de-emphasis and transmission line losses.

83A.3.2 Signal paths

The XLAUI/CAUI signal paths are point-to-point connections. Each path corresponds to a XLAUI/CAUI lane, and is comprised of two complementary signals making a balanced differential pair. For XLAUI, there are four differential paths in each direction for a total of eight pairs, or sixteen connections. For CAUI, there are ten differential paths in each direction for a total of twenty pairs, or 40 connections.

83A.3.3 Transmitter characteristics

The XLAUI/CAUI transmitter characteristics measured at the transmitter compliance point are specified in Table 83A–1. The XLAUI/CAUI signaling rate shall be the signaling rate defined in Table 83A–1.

Table 83A–1—Transmitter characteristics

Parameter	Subclause reference	Value	Units
Signaling rate per lane (range)	—	10.3125 ± 100 ppm	GBd
Single-ended output voltage maximum minimum	83A.3.3.1	4 –0.4	V V
Maximum differential output voltage, peak-to-peak	83A.3.3.1	760	mV
Minimum de-emphasis	83A.3.3.1	4.4	dB
Maximum de-emphasis	83A.3.3.1	7	dB
Minimum VMA	83A.3.3.1	See Equation (83A–4)	mV
Maximum termination mismatch at 1 MHz	86A.5.3.2	5	%
Maximum output AC common-mode voltage, RMS	86A.5.3.1	15	mV
Minimum output rise and fall time (20% to 80%)	83A.3.3.2	24	ps
Differential output return loss	83A.3.3.3	See Equation (83A–5)	dB
Common-mode output return loss	83A.3.3.4	See Equation (83A–6)	dB
Maximum Total Jitter	83A.3.3.5	0.32	UI
Maximum Deterministic Jitter	83A.3.3.5	0.17	UI
Transmitter eye mask definition X1	83A.3.3.5	0.16	UI
Transmitter eye mask definition X2	83A.3.3.5	0.38	UI
Transmitter eye mask definition Y1	83A.3.3.5	200	mV
Transmitter eye mask definition Y2	83A.3.3.5	380	mV

83A.3.3.1 Output amplitude

Driver differential output amplitude shall be less than the maximum differential output voltage defined in Table 83A–1 including any transmit de-emphasis. DC-referenced logic levels are not defined since the receiver is AC coupled. Single-ended output voltage shall be within the range specified in Table 83A–1 with respect to ground.

De-emphasis shall be greater than the minimum de-emphasis and less than the maximum de-emphasis defined in table Table 83A-1. De-emphasis is defined as the ratio between the amplitude following a transition and the amplitude during a non-transition bit as seen in Equation (83A-3). VMA is defined in 86A.5.3.5 using the square wave or PRBS9 (Pattern 4) defined in Table 86-11.

See Figure 83A-5 for an illustration of absolute driver output voltage limits, definition of differential peak-to-peak amplitude, and definition of the parameters used to calculate de-emphasis. SLi<P> and SLi<N> are positive and negative sides of a differential signal pair for lane i (i = 0, 1, 2, 3 for XLAUI. For CAUI i = 0:9).

$$\text{De-emphasis (dB)} = 20\log_{10}\left(\frac{\text{Differential peak-to-peak amplitude}}{\text{VMA}}\right) \tag{83A-3}$$

$$\text{Minimum VMA (mV)} = (234.64 - 2.13x + 0.18x^2) \times 1.32(10^{-y/20}) \tag{83A-4}$$

where

- x is the rise or fall time (whichever is larger) in ps
- y is de-emphasis value in dB

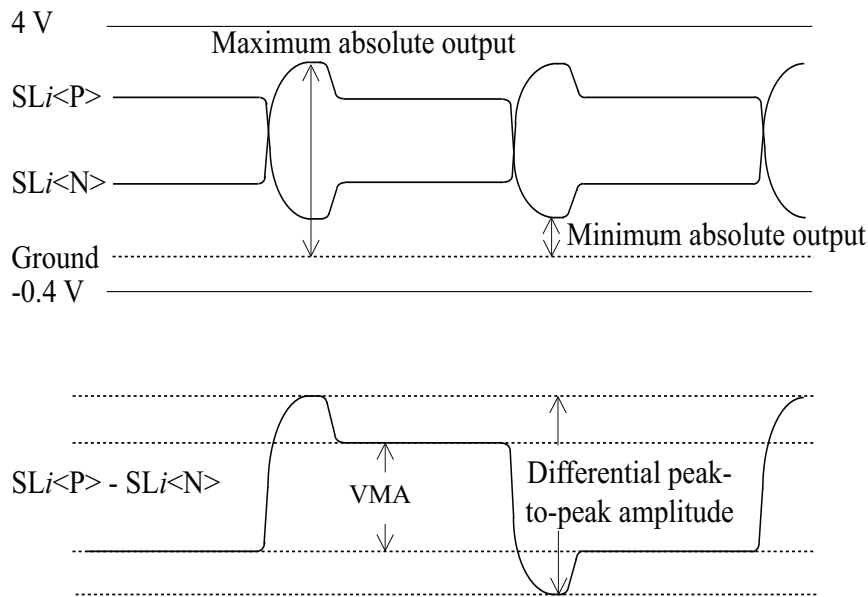


Figure 83A-5—Driver output voltage limits and definitions

83A.3.3.2 Rise/fall time

Differential rise/fall times shall be greater than the minimum output rise and fall time defined in Table 83A-1, as measured from the 20% to the 80% levels. Shorter transitions may result in excessive high-frequency components and increase EMI and crosstalk. The upper limit is defined by the transmit eye mask shown in Figure 83A-8. Rise/fall time is measured with de-emphasis off as defined in 83A.5.1. Rise/fall time measurements are taken using a square wave test pattern as defined in 83.5.10.

83A.3.3.3 Differential output return loss

Differential output return loss shall meet the requirements defined in Table 83A-1. Differential output return loss is given in Equation (83A-5) and is illustrated Figure 83A-6. Differential output return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the driver. This output impedance requirement applies to all valid output levels. The reference impedance for differential return loss measurements is 100 Ω. The maximum termination mismatch at 1 MHz shall be less than the requirement defined in Table 83A-1.

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 & 0.01 \leq f < 2.125 \\ 6.5 - 13.33 \log_{10} \left(\frac{f}{5.5} \right) & 2.125 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \tag{83A-5}$$

where

Return_loss(f) is the differential output return loss at frequency *f*
f is the frequency in GHz




Figure 83A-6—Differential output return loss

83A.3.3.4 Common-mode output return loss

Common-mode output return loss shall meet the requirements defined in Table 83A-1. Common-mode output return loss is given in Equation (83A-6) and is illustrated in Figure 83A-7. Common-mode output return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components

related to the driver. This output impedance requirement applies to all valid output levels. The reference impedance for common-mode return loss measurements is 25Ω .

$$\text{Return_loss}(f) \geq \left. \begin{array}{ll} 9 & 0.01 \leq f < 2.125 \\ 3.5 - 13.33 \log_{10} \left(\frac{f}{5.5} \right) & 2.125 \leq f < 7.1 \\ 2 & 7.1 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \quad (83A-6)$$

where

$\text{Return_loss}(f)$ is the common-mode output return loss at frequency f
 f is the frequency in GHz




Figure 83A-7—Common-mode output return loss

83A.3.3.5 Transmitter eye mask and transmitter jitter definition

The measured transmit signal at the transmit compliance point shall meet the eye template specified in Figure 83A-8 and Table 83A-1. The template measurement requirements are specified in 83A.5.1. The measured jitter at the transmit compliance point shall be less than the maximum Total Jitter as defined in Table 83A-1 and a maximum Deterministic Jitter as defined in Table 83A-1. Jitter and eye mask measurement requirements are described in 83A.5.1 and are conducted with de-emphasis off.




Figure 83A-8—Transmitter eye mask

83A.3.4 Receiver characteristics

Receiver characteristics at the receiver compliance point are specified in Table 83A-2 and detailed in 83A.3.4.1 through 83A.3.4.6.

Table 83A-2—Receiver characteristics

Parameter	Subclause reference	Value	Units
Signaling rate per lane (range)	—	10.3125 ± 100 ppm	GBd
Minimum input AC common-mode voltage tolerance, RMS	86A.5.3.1	20	mV
Minimum input rise and fall time tolerance (20% to 80%)	83A.3.3.2	24	ps
Differential input return loss	83A.3.4.3	See Equation (83A-7)	dB
Differential to common-mode input return loss	83A.3.4.4	See Equation (83A-8)	dB
Stressed receiver tolerance			
Minimum Total Input Jitter Tolerance	83A.3.4.2	0.62	UI
Minimum Deterministic Input Jitter Tolerance	83A.3.4.2	0.42	UI
Receiver eye mask definition X1	83A.3.4.2	0.31	UI
Receiver eye mask definition X2	83A.3.4.2	0.5	UI
Receiver eye mask definition Y1	83A.3.4.2	42.5	mV
Receiver eye mask definition Y2	83A.3.4.2	425	mV

83A.3.4.1 Bit error ratio

The receiver shall operate with a BER of better than 10^{-12} in the presence of a compliant input signal as defined in 83A.3.4.2.

NOTE—A transmitter capable of operating at $TJ = 0.3$ UI and $DJ = 0.16$ UI and receiver capable of operating at stress jitter tolerance of $TJ = 0.64$ UI and $DJ = 0.41$ UI would have sufficient margin for operation at approximately BER 10^{-15} .

83A.3.4.2 Input signal definition

A compliant input signal to a XLAUI/CAUI receiver has characteristics determined by a compliant XLAUI/CAUI driver and channel. The input signal definition satisfies the eye mask defined in Table 83A–2 and Figure 83A–9. Input signal jitter does not exceed the jitter tolerance requirements specified in 83A.3.4.6. Stressed receiver measurement requirements are specified in 83A.5.2.




Figure 83A–9—Receiver template

83A.3.4.3 Differential input return loss

Differential input return loss shall meet the requirements defined in Table 83A–2. Differential input return loss is given in Equation (83A–7) and is illustrated in Figure 83A–10. Differential input return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the receiver. The reference impedance for differential return loss measurements is 100Ω .

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 & 0.01 \leq f < 2.125 \\ 6.5 - 13.33 \log_{10} \left(\frac{f}{5.5} \right) & 2.125 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \quad (83A-7)$$

where

$Return_loss(f)$ is the differential input return loss at frequency f
 f is the frequency in GHz




Figure 83A-10—Differential input return loss

83A.3.4.4 Differential to common-mode input return loss

Differential to common-mode input return loss shall meet the requirements defined in Table 83A-2. Differential to common-mode input return loss is given in Equation (83A-8) and is illustrated in Figure 83A-11. Differential to common-mode input return loss includes contributions from on-chip circuitry, chip packaging, and any off-chip components related to the receiver. The reference impedance for differential return loss measurements is 100 Ω. The reference impedance for common-mode return loss measurements is 25 Ω.

$$Return_loss(f) \geq 15 \quad 0.01 \leq f < 11.1 \quad (\text{dB}) \quad (83A-8)$$

where

Return_loss(f) is the differential to common-mode input return loss at frequency *f*
f is the frequency in GHz




Figure 83A-11—Differential to common-mode input return loss

83A.3.4.5 AC coupling

The XLAUI/CAUI receiver shall be AC coupled to the XLAUI/CAUI transmitter to allow for maximum interoperability between various 10 Gb/s components. AC coupling is considered to be part of the receiver for the purposes of this specification except when interfacing with modules defined in Annex 83B or explicitly stated otherwise. It should be noted that there may be various methods for AC coupling in actual implementations.

83A.3.4.6 Jitter tolerance

The XLAUI/CAUI receiver shall have a peak-to-peak total jitter amplitude tolerance of at least the minimum total input jitter tolerance defined in Table 83A-2. This total jitter is composed of two components: deterministic jitter and random jitter. Deterministic jitter tolerance shall be at least the minimum deterministic input jitter tolerance defined in Table 83A-2. The XLAUI/CAUI receiver shall tolerate sinusoidal jitter with any frequency and amplitude defined by the mask of Figure 83A-12. This sub-component of deterministic jitter is intended to ensure margin for low-frequency jitter, wander, noise, crosstalk, and other variable system effects.

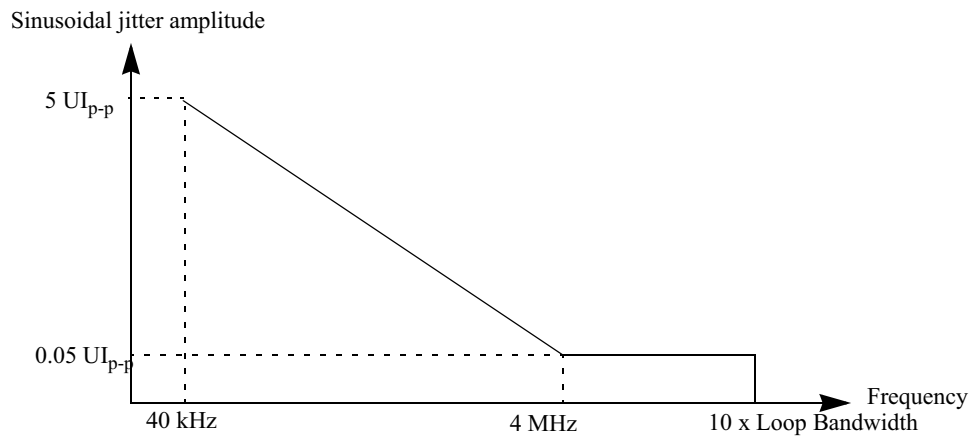


Figure 83A-12—Single-tone sinusoidal jitter mask

83A.4 Interconnect characteristics

This section describes recommended characteristics which are used to characterize a XLAUI/CAUI channel as shown in Figure 83A-2. The value for differential insertion loss is summarized in Equation (83A-9) and illustrated in Figure 83A-13. The value for minimum return loss is summarized in Equation (83A-10) and illustrated in Figure 83A-14. The channel is terminated with 100 Ω differential impedance. Other impairments such as crosstalk can have a material impact on the link performance and should be minimized.

$$Insertion_loss(f) < \left\{ \begin{array}{ll} 0.15 + 1.39\sqrt{f} + 1.4f & 0.01 \leq f < 7 \\ -15.86 + 4.2f & 7 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \tag{83A-9}$$

where

Insertion_loss(f) is the differential insertion loss at frequency *f*
f is the frequency in GHz

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12.5 & 0.01 \leq f < 5 \\ 12.5 - 27.5 \log_{10}\left(\frac{f}{5}\right) & 5 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \tag{83A-10}$$

where

Return_loss(f) is the differential input return loss at frequency *f*
f is the frequency in GHz




Figure 83A-13—Differential insertion loss




Figure 83A-14—Differential input return loss

The XLAUI/CAUI is primarily intended as a point-to-point interface of up to approximately 25 cm between integrated circuits using controlled impedance traces on low-cost printed circuit boards (PCBs). Longer reaches for the XLAUI/CAUI may be achieved by the use of better PCB materials, as the performance of an actual XLAUI/CAUI interconnect is highly dependent on the implementation.

83A.4.1 Characteristic impedance

The recommended differential characteristic impedance of circuit board trace pairs is $100 \Omega \pm 10\%$.

83A.5 Electrical parameter measurement methods

This subclause describes the measurement methodology that is to be used to verify XLAUI/CAUI compliance. Eye templates are measured with AC coupling and centered at 0 V differential. The signal waveform, eye, and jitter may be measured using a receiver with at least an equivalent 12 GHz low-pass filter response. Jitter values and eye masks are specified for BER 10^{-12} .

83A.5.1 Transmit jitter

Transmit jitter is defined with respect to a test procedure resulting in a BER bathtub curve such as that described in Annex 48B.3. For the purpose of jitter measurement, the effect of a single-pole high-pass filter with a 3 dB point at 4 MHz is applied to the jitter. The test pattern for jitter measurements shall be PRBS31 test pattern in 83.5.10 or scrambled idle in 82.2.10. Crossing times are defined with respect to the mid-point (0 V) of the AC coupled differential signal. De-emphasis shall be off during jitter testing. Transmit de-emphasis off state is defined by any setting that gives optimal performance for transmitter jitter and eye mask evaluation. All XLAUI/CAUI transmitter lanes shall be active and all XLAUI/CAUI receive lanes shall be receiving maximum amplitude and fastest rise time (as defined in Table 83A-1) during transmit jitter testing to ensure maximum lane-lane crosstalk is included in the jitter evaluation.

83A.5.2 Receiver tolerance

The XLAUI/CAUI jitter tolerance test setup in Figure 83A-15 or its functional equivalent shall meet the receiver eye mask defined in Table 83A-2. Applied jitter is measured using the methodology described in Annex 48B.3. Deterministic jitter is added to a clean test pattern by adding sinusoidal jitter as defined in 83A.3.4.6, along with low-pass filter stress, followed by a limiting function, and frequency-dependent attenuation. The low-pass filter stress is added until the 0.34 UI Deterministic Jitter is achieved. Frequency-dependent attenuation is then added using PCB trace or frequency-dependent attenuation which emulates PCB loss. Frequency-dependent attenuation is added until 0.42 UI Deterministic Jitter is achieved. Random jitter is added to the test signal using an interference generator which is a broadband noise source capable of producing white Gaussian noise with adjustable amplitude. The power spectral density shall be flat to ± 3 dB from 50 MHz to 6 GHz with a crest factor of no less than 5. The amplitude and output jitter of the filter stress plus limiter and random jitter injection shall meet the receiver eye mask defined in Table 83A-2. All XLAUI/CAUI lanes shall be active during jitter tolerance testing. The PRBS31 pattern defined in 83.5.10 or scrambled idle defined in 82.2.10 is used for evaluating XLAUI/CAUI jitter tolerance.

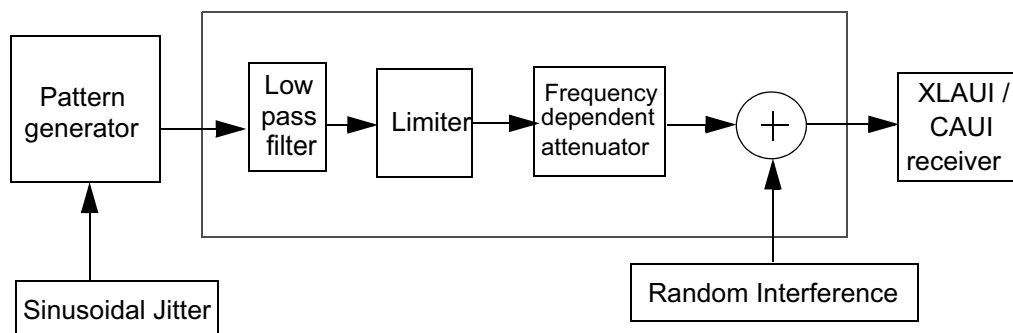


Figure 83A-15—Stressed-eye and jitter tolerance test setup

83A.6 Environmental specifications

83A.6.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

83A.6.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

83A.6.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

83A.6.4 Electromagnetic compatibility

A system integrating the XLAUI/CAUI shall comply with applicable local and national codes for the limitation of electromagnetic interference.

83A.6.5 Temperature and humidity

A system integrating the XLAUI/CAUI is 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 standard.

83A.7 Protocol implementation conformance statement (PICS) proforma for Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)²³

83A.7.1 Introduction

The supplier of a XLAUI/CAUI implementation that is claimed to conform to Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI), shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the same, can be found in Clause 21.

83A.7.2 Identification

83A.7.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations.	
NOTE 2—May be completed as appropriate in meeting the requirements for the identification.	
NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

83A.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Annex 83A, 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

²³Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

83A.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
NOL	XLAUI is organized into four lanes, CAUI is organized into ten lanes	83A.1.2	See Clause 83	M	Yes []
RATE	Each XLAUI/CAUI lane operates at 10.3125 Gb/s	83A.1.2	10.3125 Gb/s (nominal)	M	Yes []
IO	Meets XLAUI/CAUI Electrical Characteristics	83A.3	Supports transmit and receive compliance points	M	Yes []

83A.7.4 XLAUI/CAUI transmitter requirements

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Signaling rate	83A.3.3	10.3125 GBd \pm 100 ppm	M	Yes []
TC2	Single ended output voltage range	83A.3.3.1	-0.4 V to 4 V	M	Yes []
TC3	Minimum de-emphasis	83A.3.3.1	4.4 dB	M	Yes []
TC4	Maximum de-emphasis	83A.3.3.1	7 dB	M	Yes []
TC5	Maximum termination mismatch	83A.3.3.3	5%	M	Yes []
TC6	Maximum differential output voltage	83A.3.3.1	760 mV	M	Yes []
TC7	Minimum output rise and fall time (20% to 80%)	83A.3.3.2	24 ps	M	Yes []
TC8	Differential output return loss	83A.3.3.3	See Equation (83A-5)	M	Yes []
TC9	Common-mode output return loss	83A.3.3.4	See Equation (83A-6)	M	Yes []
TC10	Transmitter eye mask	83A.3.3.5	See Figure 83A-8	M	Yes []

83A.7.5 XLAUI/CAUI receiver requirements

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Receiver BER	83A.3.4.1	Better than 10^{-12}	M	Yes []
RC2	Differential input return loss	83A.3.4.3	See Equation (83A-7)	M	Yes []
RC3	Differential to common-mode input return loss	83A.3.4.4	See Equation (83A-8)	M	Yes []
RC4	Total Jitter Tolerance	83A.3.4.6	0.62 UI	M	Yes []
RC5	Deterministic Jitter Tolerance	83A.3.4.6	0.42 UI	M	Yes []
RC6	Sinusoidal Jitter Tolerance	83A.3.4.6	See Figure 83A-12	M	Yes []
RC7	Receiver AC coupling	83A.3.4.5	Present	M	Yes []

83A.7.6 Electrical measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
EM1	De-emphasis setting during jitter measurements	83A.5.1	Off	M	Yes []
EM2	Jitter Tolerance stressed input	83A.5.2	0.42 UI DJ, 0.2 UI RJ	M	Yes []
EM3	Random jitter interference generator	83A.5.2	± 3 dB from 50 MHz to 6 GHz with a crest factor of no less than 5	M	Yes []
EM4	Meet jitter tolerance requirement with all XLAUI/CAUI channels active	83A.5.2	Yes	M	Yes []
EM5	XLAUI/CAUI Jitter Tolerance Test Pattern	83A.5.2	PRBS31 or scrambled idle	M	Yes []

83A.7.7 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	Conformance to applicable sections of IEC 60950-1	83A.6.1	Yes	M	Yes []
ES2	Compliance with applicable local and national codes for the limitation of electromagnetic interference	83A.6.4	Yes	M	Yes []

Annex 83B

(normative)

Chip-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)

83B.1 Overview

This annex defines the functional and electrical characteristics for the optional chip-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI). The purpose of this annex is to provide electrical characteristics and associated compliance points for pluggable module applications that use the XLAUI/CAUI interface and shall use the same number of lanes and signaling rate defined in Annex 83A. Figure 83B–2 and Table 83B–1 summarize an example differential insertion loss budget associated with the chip-module application. The insertion loss of Equation (83A–9), excluding a 0.5 dB connector loss at 5.15625 GHz, is linearly scaled to 7.9 dB loss at 5.15625 GHz for the host XLAUI/CAUI component, and 2.1 dB loss at 5.15625 GHz for the module as per Table 83B–1 and Equation (83B–1) for the host and Equation (83B–2) for the module. Tradeoffs can be made between the host PCB insertion loss and the connector loss. Equation (83B–1) is illustrated in Figure 83B–1 and Equation (83B–2) is illustrated in Figure 83B–1.

$$Insertion_loss(f) \leq \left\{ \begin{array}{ll} 0.111 + 1.046\sqrt{f} + 1.05f & 0.01 \leq f < 7 \\ -11.82 + 3.15f & 7 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \quad (83B-1)$$

where

$Insertion_loss(f)$ is the differential insertion loss at frequency f
 f is the frequency in GHz

$$Insertion_loss(f) \leq \left\{ \begin{array}{ll} 0.03 + 0.278\sqrt{f} + 0.28f & 0.01 \leq f < 7 \\ -3.155 + 0.84f & 7 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \quad (83B-2)$$

where

$Insertion_loss(f)$ is the differential insertion loss at frequency f
 f is the frequency in GHz




Figure 83B-1—Module insertion loss




Figure 83B-2—Chip-module loss budget at 5.15625 GHz

Table 83B-1—Chip-module loss budget

Section	Differential insertion loss max. (at 5.15625 GHz)
Host XLAUI/CAUI component to connector	7.9 dB
Connector loss	0.5 dB
Connector to module XLAUI/CAUI component	2.1 dB

83B.2 Compliance point specifications for chip-module XLAUI/CAUI

The chip-module XLAUI/CAUI interface specifies compliance points around the module connector as depicted in Figure 83B-4 and Figure 83B-6. Chip-module devices shall meet the electrical characteristics defined in 83B.2, 83B.2.1, 83B.2.2, and 83B.2.3. Compliance points are defined to ensure interoperability between hosts and modules. A Module Compliance Board (MCB) is used to isolate the characteristics of the module, and a Host Compliance Board (HCB) is used to isolate the characteristics of the host. Figure 83B-4 and Figure 83B-6 include the loss associated with the HCB and MCB at 5.15625 GHz.

The reference differential insertion loss of the HCB PCB is given in Equation (83B-3) and illustrated in Figure 83B-3. The effects of differences between the insertion loss of an actual HCB and the reference insertion loss are to be accounted in the measurements.

$$\text{Insertion_loss}(f) = 0.017 + 0.5\sqrt{f} + 0.1836f \quad 0.01 \leq f \leq 11.1 \quad (83B-3)$$

where

$\text{Insertion_loss}(f)$ is the differential insertion loss at frequency f
 f is the frequency in GHz




Figure 83B-3—HCB insertion loss

The reference differential insertion loss of the MCB PCB is given in Equation (83B-4) and illustrated in Figure 83B-5. The effects of differences between the insertion loss of an actual MCB and the reference insertion loss are to be accounted in the measurements.




Figure 83B-4—Chip-module HCB insertion loss budget at 5.15625 GHz

$$Insertion_loss(f) = -0.00086 + 0.2286\sqrt{f} + 0.08386f \text{ dB} \quad 0.01 \leq f \leq 11.1 \quad (83B-4)$$

where

$Insertion_loss(f)$ is the differential insertion loss at frequency f
 f is the frequency in GHz




Figure 83B-5—MCB insertion loss

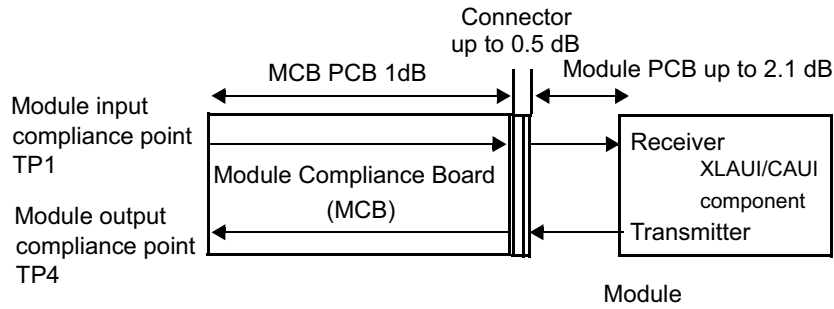


Figure 83B-6—Chip-module with MCB insertion loss budget at 5.15625 GHz

83B.2.1 Module specifications

A module that uses XLAUI/CAUI to interface with a host shall meet the characteristics outlined in Table 83B-2 and Table 83B-3 when measured using the MCB and HCB (where the HCB is used to calibrate inputs to the module). Table 83B-2 also lists the equivalent test points for the XLPPI/CPPI (see Figure 86-3).

Modules may support additional de-emphasis states, but the specification of additional states is outside the scope of this standard. De-emphasis shall be off during eye mask and jitter testing. Module electrical output de-emphasis off state is the optimal setting for module electrical output jitter and eye mask evaluation. AC coupling for both Tx and Rx paths shall be located in the module.

Table 83B-2—Specifications at module compliance points

Reference	Compliance point		Value	Unit
Minimum module differential input return loss	MCB input	TP1	See Equation (83B-5)	dB
Module input tolerance signal	HCB output	TP1a	See 83A.5.2	
Module output signal	MCB output	TP4	See Table 83B-3	
Minimum module differential output return loss	MCB output	TP4	See Equation (83B-5)	dB
Module output common-mode return loss	MCB output	TP4	See Equation (86A-2)	dB

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.01 \leq f < 2.19 \\ 5.56 - 8.7 \log_{10}\left(\frac{f}{5.5}\right) & 2.19 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \tag{83B-5}$$

where

Return_loss(f) is the differential input return loss at frequency *f*
f is the frequency in GHz

Minimum module differential input/output return loss is illustrated in Figure 83B–7.




Figure 83B–7—Module input/output return loss

Table 83B–3—Module electrical output

Parameter	Subclause reference	Value	Unit
Maximum differential output voltage, peak-to-peak	83A.3.3.1	760	mV
Minimum de-emphasis	83A.3.3.1	3.5	dB
Maximum de-emphasis	83A.3.3.1	6	dB
Minimum VMA	83A.3.3.1	See Equation (83B–6)	mV
Maximum termination mismatch at 1 MHz	86A.5.3.2	5	%
Maximum output AC common-mode voltage, RMS	86A.5.3.1	15	mV
Minimum output rise and fall time (20% to 80%)	83A.3.3.2	24	ps
Maximum Total Jitter	83A.3.3.5	0.4	UI
Maximum Deterministic Jitter	83A.3.3.5	0.25	UI
Module electrical output eye mask definition X1	83A.3.3.5	0.2	UI
Module electrical output eye mask definition X2	83A.3.3.5	0.5	UI
Module electrical output eye mask definition Y1	83A.3.3.5	136	mV
Module electrical output eye mask definition Y2	83A.3.3.5	380	mV

$$\text{Minimum VMA (mV)} = (-110 - 2.13x + 0.32x^2) \times (10^{-y/20}) \tag{83B-6}$$

where

- x is the rise or fall time (whichever is larger) in ps
- y is de-emphasis value in dB

83B.2.2 Host specifications

A host that uses XLAUI/CAUI to interface with a module shall meet the characteristics outlined in Table 83B-4 and Table 83B-5 when measured using the HCB and MCB (where MCB is used to calibrate inputs to the host). Table 83B-4 also lists the equivalent test points for the XLPPI/CPPI.

Table 83B-4—Specifications at host compliance points

Reference	Compliance point		Value	Unit
Host output signal specifications	HCB output	TP1a	See Table 83B-5	
Minimum host differential output return loss	HCB output	TP1a	See Equation (83B-7)	dB
Minimum host differential input return loss	HCB input	TP4a	See Equation (83B-7)	dB
Host input tolerance signal	MCB output	TP4	See 83B.2.3	

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.01 \leq f < 2.19 \\ 5.56 - 8.71 \log_{10}\left(\frac{f}{5.5}\right) & 2.19 \leq f \leq 11.1 \end{array} \right\} \text{ (dB)} \tag{83B-7}$$

where

- $Return_loss(f)$ is the differential input/output return loss at frequency f
- f is the frequency in GHz

Minimum host differential input/output return loss is illustrated in Figure 83B-8.




Figure 83B-8—Host input/output return loss

Table 83B-5—Host electrical output

Parameter	Subclause reference	Value	Units
Maximum output AC common-mode voltage, RMS	86A.5.3.1	20	mV
Minimum output rise and fall time (20% to 80%)	83A.3.3.2	24	ps
Maximum Total Jitter	83A.3.3.5	0.62	UI
Maximum Deterministic Jitter	83A.3.3.5	0.42	UI
Host electrical output eye mask definition X1	83A.3.3.5	0.31	UI
Host electrical output eye mask definition X2	83A.3.3.5	0.5	UI
Host electrical output eye mask definition Y1	83A.3.3.5	42.5	mV
Host electrical output eye mask definition Y2	83A.3.3.5	425	mV

83B.2.3 Host input signal tolerance

Host XLAUI/CAUI jitter tolerance evaluation shall be defined by a stressed input signal that comprises 0.25 UI Deterministic Jitter, and 0.15 UI random jitter for BER 10^{-12} . Deterministic Jitter is added to a clean test pattern as sinusoidal jitter defined in 83A.3.4.6. The limited low-pass filter stress is added until the 0.22 UI Deterministic Jitter is achieved. Frequency-dependent attenuation is then added using PCB trace or frequency-dependent attenuation that emulates PCB loss. Frequency-dependent attenuation is added until 0.25 UI Deterministic Jitter is achieved. Random jitter is added to the test signal using an interference generator, which is a broadband noise source capable of producing white Gaussian noise with adjustable

amplitude. The power spectral density shall be flat to ± 3 dB from 50 MHz to 6 GHz with a crest factor of no less than 5. Figure 83B–9 depicts a XLAUI/CAUI jitter tolerance test setup. The amplitude and output jitter of the filter stress plus limiter and random jitter injection shall meet the receiver eye mask illustrated in Figure 83A–9 with the values for X1, X2, Y1, Y2 given in Table 83B–3. All XLAUI/CAUI lanes shall be active during jitter tolerance testing. The PRBS31 pattern defined in 83.5.10 or scrambled idle defined in 82.2.10 is used for evaluating XLAUI/CAUI jitter tolerance.

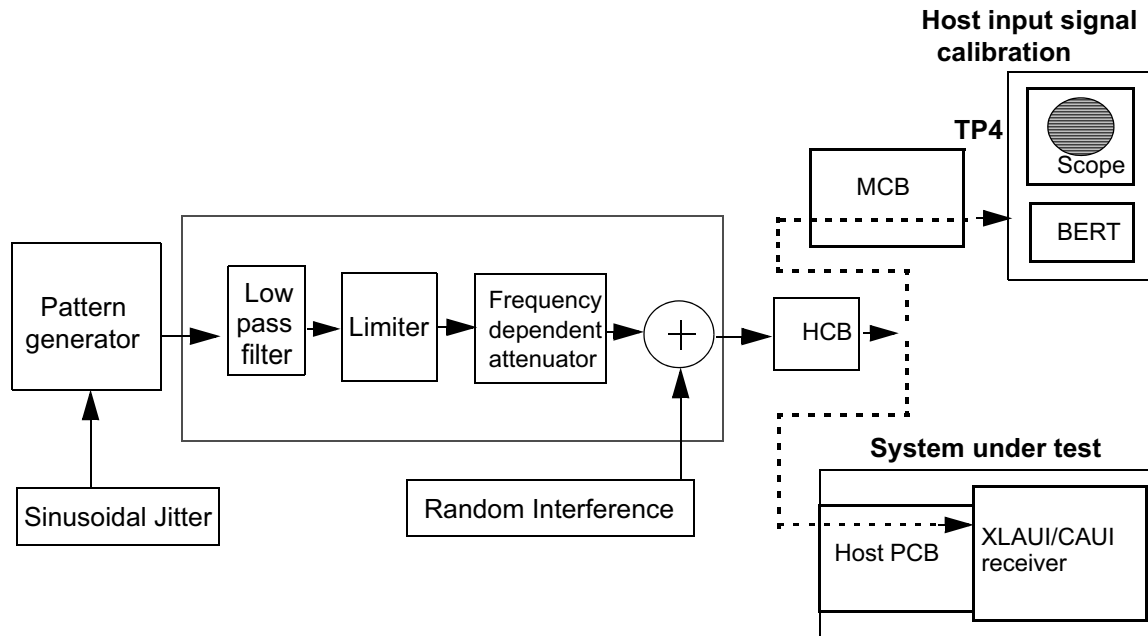


Figure 83B–9—Stressed-eye and jitter tolerance test setup

83B.3 Environmental specifications

83B.3.1 General safety

All equipment subject to this clause shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

83B.3.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

83B.3.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

83B.3.4 Electromagnetic compatibility

A system integrating the XLAUI/CAUI shall comply with applicable local and national codes for the limitation of electromagnetic interference.

83B.3.5 Temperature and humidity

A system integrating the XLAUI/CAUI is 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 standard.

83B.4 Protocol implementation conformance statement (PICS) proforma for Annex 83B, Chip-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)²⁴

83B.4.1 Introduction

The supplier of a chip to module XLAUI/CAUI implementation that is claimed to conform to Annex 83B, Chip-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI), shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the same, can be found in Clause 21.

83B.4.2 Identification

83B.4.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
<p>NOTE 1—Required for all implementations.</p> <p>NOTE 2—May be completed as appropriate in meeting the requirements for the identification.</p> <p>NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).</p>	

83B.4.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010, Annex 83B, Chip-module 40 Gb/s Attachment Unit Interface (XLAUI) and 100 Gb/s Attachment Unit Interface (CAUI)
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
<p>Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)</p>	

Date of Statement	
-------------------	--

²⁴Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

83B.4.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
NOL	XLAUI is organized into four lanes, CAUI is organized into ten lanes	83A.1.2	See Clause 83	M	Yes []
RATE	Each XLAUI/CAUI lane operates at 10.3125 Gb/s	83A.1.2	10.3125 Gb/s (nominal)	M	Yes []
IO	Meets chip-module XLAUI/CAUI Electrical Characteristics	83B.2	Supports host/module compliance points	M	Yes []

83B.4.4 Module requirements

Item	Feature	Subclause	Value/Comment	Status	Support
MC1	XLAUI/CAUI compliant module	83B.2.1	Meets requirements defined in Table 83B-2 and Table 83B-3	M	Yes []
MC2	De-emphasis setting during module jitter evaluation	83B.2.1	Off	M	Yes []
MC3	AC coupling for Tx and Rx	83B.2.1	Present in module	M	Yes []

83B.4.5 Host requirements

Item	Feature	Subclause	Value/Comment	Status	Support
HC1	XLAUI/CAUI compliant host	83B.2.2	Meets requirements defined in Table 83B-4 and Table 83B-5	M	Yes []
HC2	All XLAUI/CAUI lanes active during jitter evaluation	83B.2.3	Yes	M	Yes []
HC3	Applied jitter for host jitter tolerance evaluation	83B.2.3	0.25UI Deterministic Jitter, 0.15UI random jitter	M	Yes []

83B.4.6 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	Conformance to applicable sections of IEC 60950-1	83B.3.1	Yes	M	Yes []
ES2	Compliance with applicable local and national codes for the limitation of electromagnetic interference	83B.3.4	Yes	M	Yes []

Annex 83C

(informative)

PMA sublayer partitioning examples

The following subclauses provide various partitioning examples. Partitioning guidelines and MMD numbering conventions are described in 83.1.4.

83C.1 Partitioning examples with FEC

The example of FEC implemented in a separate device from either the PCS or the PMD is illustrated in Figure 83–2.

83C.1.1 FEC implemented with PCS

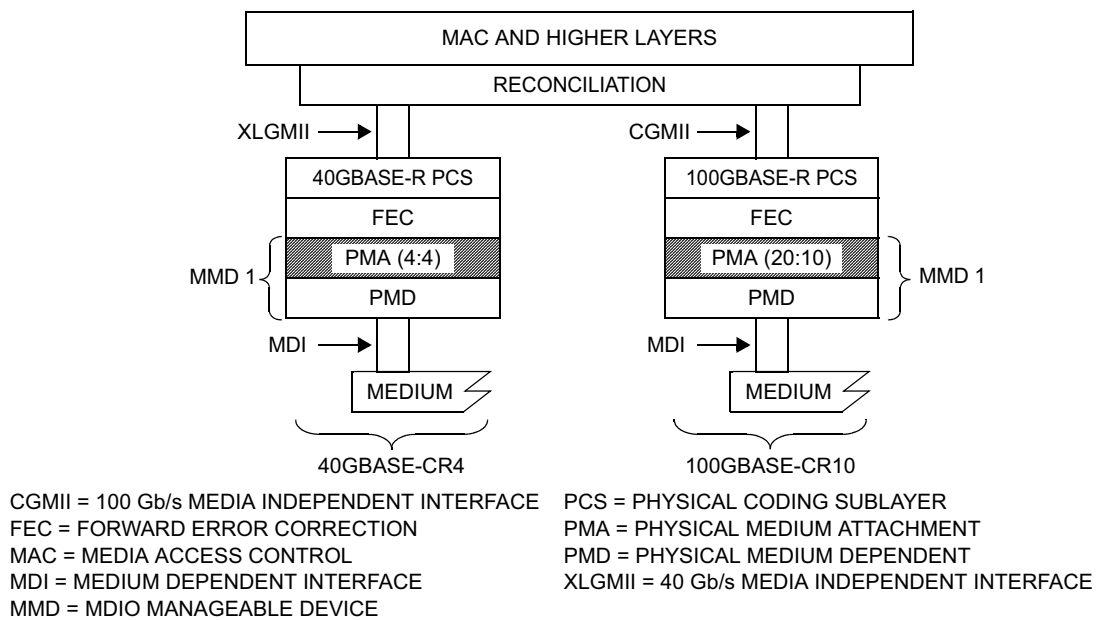
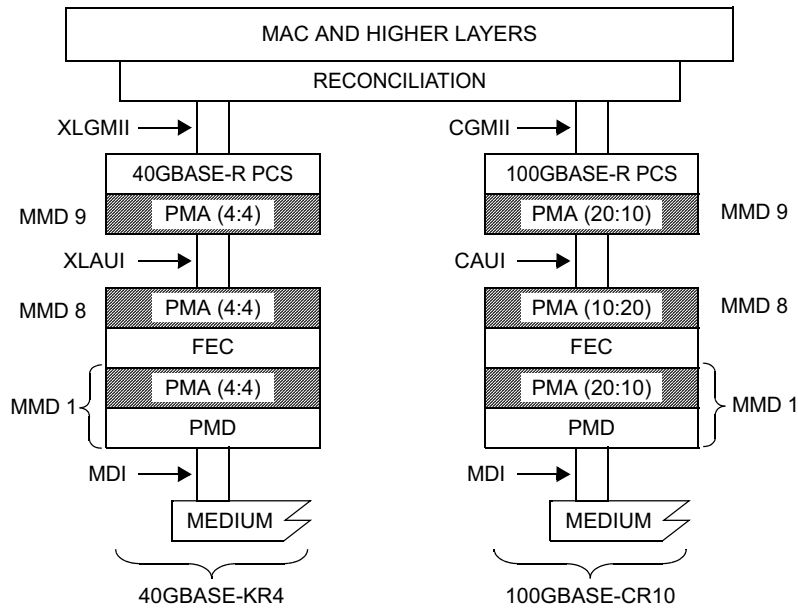


Figure 83C–1—Example FEC implemented with PCS

83C.1.2 FEC implemented with PMD



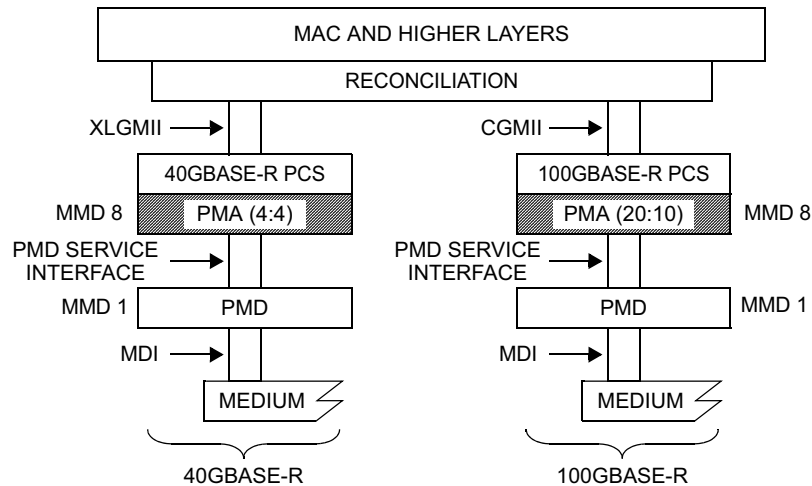
CAUI = 100 Gb/s ATTACHMENT UNIT INTERFACE
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLAUI = 40 Gb/s ATTACHMENT UNIT INTERFACE
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C-2—Example FEC implemented with PMD

83C.2 Partitioning examples without FEC

83C.2.1 Single PMA sublayer without FEC



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 MMD = MDIO MANAGEABLE DEVICE

PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE

Figure 83C-3—Example single PMA sublayer without FEC

83C.2.2 Single XLAUI/CAUI without FEC

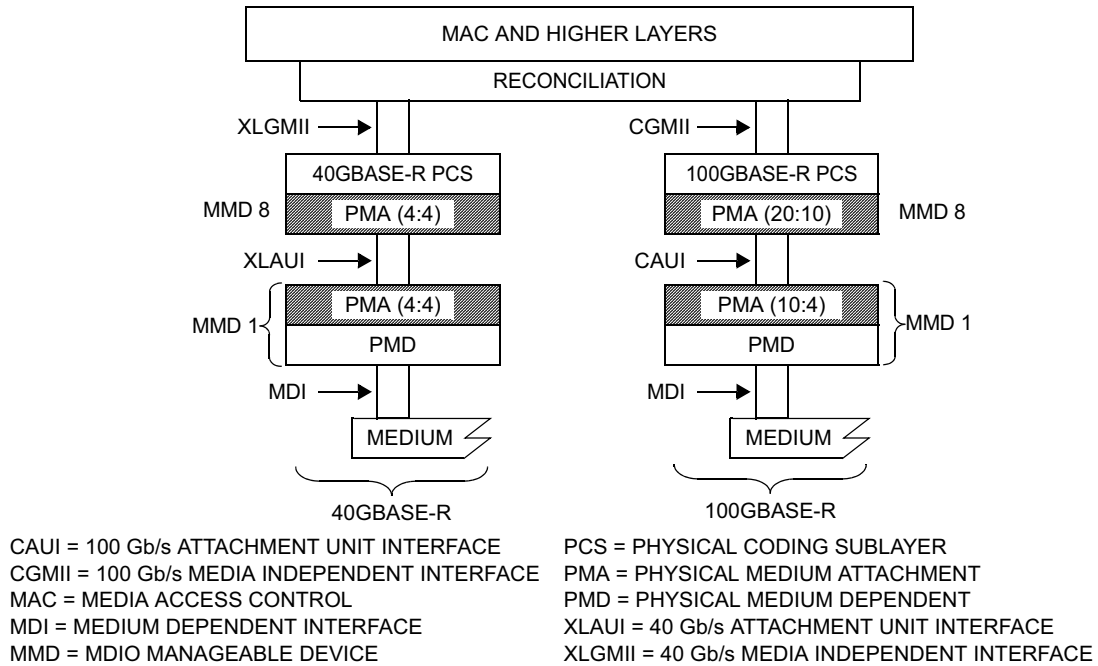


Figure 83C-4—Example single XLAUI/CAUI without FEC

83C.2.3 Separate SERDES for optical module interface

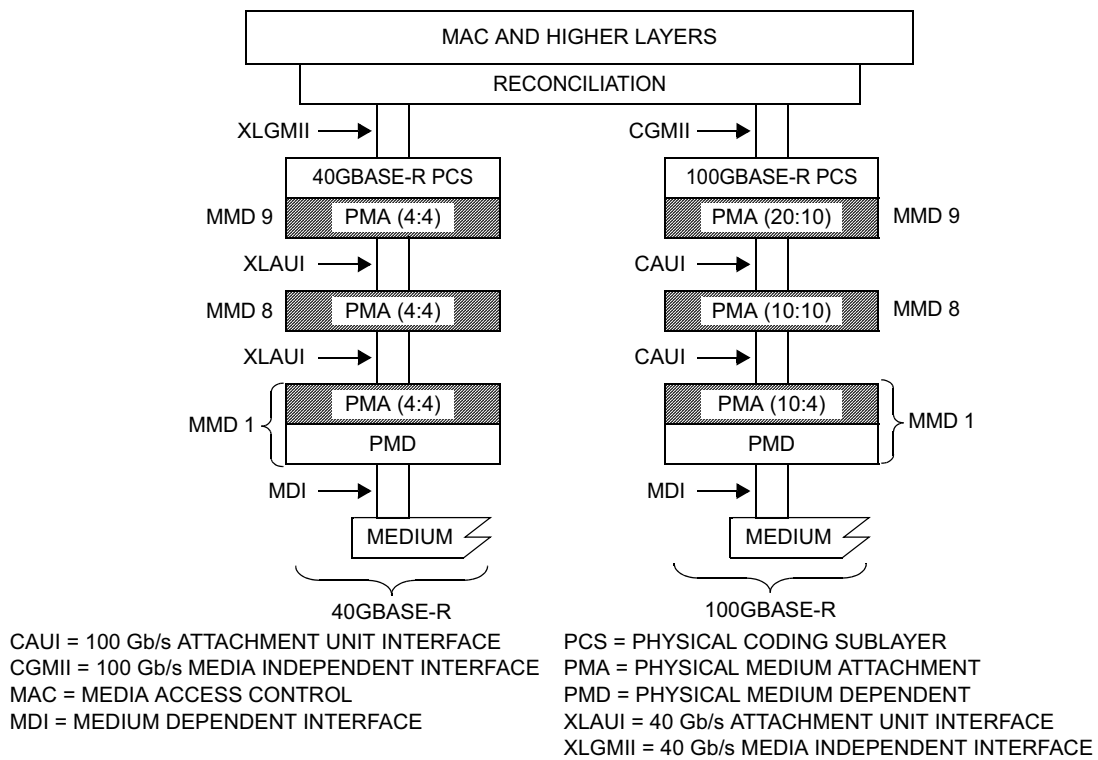


Figure 83C-5—Separate SERDES for optical module interface

Annex 85A

(informative)

40GBASE-CR4 and 100GBASE-CR10 TP0 and TP5 test point parameters

85A.1 Overview

Annex 85A provides information on parameters associated with test points TP0 and TP5 that may not be testable in an implemented system. TP0 and TP5 test points are illustrated in the 40GBASE-CR4 and 100GBASE-CR10 link block diagram of Figure 85–2.

85A.2 Transmitter characteristics at TP0

The specifications at TP0 are summarized in Table 85A–1

Table 85A–1—Transmitter characteristics at TP0 summary

Parameter	Subclause reference	Value	Units
Signaling rate, per lane	85.8.3.8	10.3125 ± 100 ppm	GBd
Unit interval nominal	85.8.3.8	96.969697	ps
Differential peak-to-peak output voltage (max) with Tx disabled	85.8.3.3	30	mV
Common-mode voltage limits	72.7.1.4	0 to 1.9	V
Differential output return loss (min)	72.7.1.5	See Equation (72–4) and Equation (72–5)	dB
Common-mode output return loss (min)	72.7.1.6	See Equation (72–6) and Equation (72–7)	dB
Common-mode AC output voltage (max, RMS)		30	mV
Transition time (20%–80%)	72.7.1.7	24 to 47	ps
Max output jitter (peak-to-peak) ^a	72.7.1.8		
Random jitter ^b		0.15	UI
Deterministic jitter		0.15	UI
Duty Cycle Distortion ^c		0.035	UI
Total jitter		0.28	UI

^aJitter is measured with emphasis off.

^bJitter is specified at BER 10^{-12} .

^cDuty Cycle Distortion is considered part of the deterministic jitter distribution.

85A.3 Receiver characteristics at TP5

The receiver characteristics at TP5 are summarized in Table 85A–2.

Table 85A–2— Receiver characteristics at TP5 summary

Parameter	Subclause reference	Value	Units
Bit error ratio	85.8.4.3	10^{-12}	
Signaling rate, per lane	85.8.4.4	10.3125 ± 100 ppm	GBd
Unit interval (UI) nominal	85.8.4.4	96.969697	ps
Receiver coupling	85.8.4.5	AC	
Differential input peak-to-peak amplitude (max)	72.7.2.4	1200 ^a	mV
Differential input return loss (min) ^b	72.7.2.5	See Equation (72–4) and Equation (72–5)	dB
Differential to common-mode input return loss	85.8.4	10 min from 10 MHz to 10 GHz	dB

^aThe receiver shall tolerate amplitudes up to 1600 mV without permanent damage.

^bRelative to 100 Ω differential.

85A.4 Transmitter and receiver differential printed circuit board trace loss

With the insertion loss TP0 to TP2 or TP3 to TP5 given in 85.8.3.4 and an assumed mated connector loss of 1.74 dB, the maximum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane (i.e., the maximum value of the sum of the insertion losses from TP0 to the MDI host receptacle and from TP5 to the MDI host receptacle) are determined using Equation (85A–1). The maximum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards is 7 dB at 5.15625 GHz. The maximum insertion loss for the transmitter or the receiver differential controlled impedance printed circuit board is one half of the maximum insertion loss $IL_{PCBmax}(f)$.

$$IL_{PCB}(f) \leq IL_{PCBmax}(f) = (0.3)[20\log_{10}(e)(b_1\sqrt{f} + b_2f + b_3f^2 + b_4f^3)] \quad (\text{dB}) \quad (85A-1)$$

for $10 \text{ MHz} \leq f \leq 7500 \text{ MHz}$.

where

f	is the frequency in Hz
$IL_{PCB}(f)$	is the insertion loss for the transmitter and receiver PCB
$IL_{PCBmax}(f)$	is the maximum insertion loss for the transmitter and receiver PCB
b_1	2×10^{-5}
b_2	1.1×10^{-10}
b_3	3.2×10^{-20}
b_4	-1.2×10^{-30}
e	≈ 2.71828

The minimum insertion loss allocation for the transmitter and receiver differential controlled impedance printed circuit boards for each differential lane (i.e., the minimum value of the sum of the insertion losses from TP0 to MDI receptacle and TP5 to MDI receptacle) are determined using Equation (85A–2) and the coefficients b_1 through b_4 are given in Equation (85A–1). The minimum insertion loss for the transmitter or the receiver differential controlled impedance printed circuit board is one half of the minimum insertion loss $IL_{PCBmin}(f)$.

$$IL_{PCB}(f) \geq IL_{PCBmin}(f) = (0.0574)[20\log_{10}(e)(b_1\sqrt{f} + b_2f + b_3f^2 + b_4f^3)] \text{ (dB)} \quad (85A-2)$$

for $10 \text{ MHz} \leq f \leq 7500 \text{ MHz}$.

where

f	is the frequency in Hz
$IL_{PCB}(f)$	is the insertion loss for the transmitter and receiver PCB
$IL_{PCBmin}(f)$	is the minimum insertion loss for the transmitter and receiver PCB

85A.5 Channel insertion loss

This subclause provides information on channel insertion losses for intended topologies ranging from 0.5 m to 7 m in length. The maximum channel insertion loss associated with the 7 m topology is determined using Equation (85A–3). The channel insertion loss associated with the 0.5 m topology and a maximum host channel is determined by Equation (85A–4). The channel insertion loss budget at 5.15625 GHz is illustrated in Figure 85A–1.

The maximum channel insertion loss is determined using Equation (85A–3). The maximum channel insertion loss is 24.44 dB at 5.15625 GHz.

$$IL_{Chmax}(f) = IL_{Camax}(f) + 2IL_{Host}(f) - 2IL_{MatedTF}(f) \text{ (dB)} \quad (85A-3)$$

for $50 \text{ MHz} \leq f \leq 7500 \text{ MHz}$.

where

f	is the frequency in MHz
$IL_{Chmax}(f)$	is the maximum channel insertion loss between TP0 and TP5
$IL_{Camax}(f)$	is the maximum cable assembly insertion loss using Equation (85–19)
$IL_{Host}(f)$	is the maximum insertion loss from TP0 to TP2 or TP3 to TP5 using Equation (85–14)
$IL_{MatedTF}(f)$	is the maximum insertion loss of the mated test fixture using Equation (85–36)

The channel insertion loss between TP0 and TP5 representative of a 0.5 m cable assembly and a maximum host channel is determined using Equation (85A–4).

$$(IL_{Ch0.5m}(f) = 0.275IL_{Camax}(f) + 2IL_{Host}(f) - 2IL_{MatedTF}(f) \text{ (dB)} \quad (85A-4)$$

for $50 \text{ MHz} \leq f \leq 7500 \text{ MHz}$.

where

f	is the frequency in MHz
-----	-------------------------

- $IL_{Ch0.5m}(f)$ is the channel insertion loss between TP0 and TP5 representative of a 0.5 m cable assembly and a maximum host channel
- $IL_{C_{max}}(f)$ is the maximum cable assembly insertion loss from TP1 to TP4 using Equation (85-19)
- $IL_{Host}(f)$ is the maximum insertion loss from TP0 to TP2 or TP3 to TP5 using Equation (85-14)
- $IL_{MatedTF}(f)$ is the maximum insertion loss of the mated test fixture using Equation (85-36)




Figure 85A-1—Illustration of channel insertion loss budget at 5.15625 GHz

85A.6 Channel return loss

The return loss of each lane of the 40GBASE-CR4 or 100GBASE-CR10 channel is recommended to meet the values determined using Equation (85-25).

85A.7 Channel insertion loss deviation (ILD)

The channel insertion loss deviation is the difference between the channel insertion loss and the fitted channel insertion loss is determined using Equation (85A-5).

$$ILD(f) = IL_{Ch}(f) - IL_{fitted}(f) \text{ (dB)} \tag{85A-5}$$

where

f	is the frequency in MHz
$ILD(f)$	is the channel insertion loss deviation at frequency f
$IL_{\text{fitted}}(f)$	is defined in Equation (85–19)

Given the channel insertion loss is at N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 7500 MHz with a maximum frequency spacing of 10 MHz, the coefficients of the fitted channel insertion loss are determined using Equation (85–20) and Equation (85–21).

The channel insertion loss deviation is recommended to be within the region defined by Equation (85–23) and Equation (85–24) for all frequencies from 50 MHz to 7500 MHz.

85A.8 Channel integrated crosstalk noise (ICN)

Since four lanes or ten lanes are used to transfer data between PMDs, the near-end crosstalk (NEXT) that is coupled into a victim receiver will be from the four or ten adjacent transmitters. The channel multiple disturber NEXT loss, $MDNEXT_loss_1(f)$, is specified using the individual NEXT losses as shown in Equation (85–26).

In addition, the far-end Crosstalk (FEXT) that is coupled into a receiver will be from the three or nine other transmitters adjacent to the victim transmitter. The channel multiple disturber FEXT loss, $MDFEXT_loss(f)$, is specified using the individual FEXT losses as shown in Equation (85–27).

Given the channel $MDNEXT_loss(f)$ and $MDFEXT_loss(f)$ measured over N uniformly-spaced frequencies f_n spanning the frequency range 50 MHz to 10000 MHz with a maximum frequency spacing of 10 MHz, the RMS value of the integrated crosstalk noise is determined using Equation (85–28) through Equation (85–32) and the parameters shown in Table 85–11.

The total integrated crosstalk RMS noise voltage of the channel is recommended to meet the values determined using Equation (85A–6) illustrated in Figure 85A–2.

$$\sigma_{x, ch} \leq \left\{ \begin{array}{ll} 10 & 3 \leq IL \leq 7.5 \\ 13.4 - 0.45IL & 7.5 < IL \leq 24.44 \end{array} \right\} \quad (\text{mV}) \quad (85A-6)$$

where IL is the value of the channel insertion loss in dB at 5.15625 GHz.




Figure 85A-2—Channel integrated crosstalk noise

Annex 86A

(normative)

Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI)

86A.1 Overview

The Parallel Physical Interface (nPPI) is an optional instantiation of the PMD service interface for the PMDs in Clause 86, which is described in 86.2 and 80.3, or for the 40GBASE-LR4 PMD service interface described in 87.2 and 80.3. It allows the construction of compact optical transceiver modules for 40GBASE-SR4, 40GBASE-LR4, or 100GBASE-SR10 with no clock and data recovery circuits inside. The 40 Gb/s Parallel Physical Interface (XLPPI, four lanes) is used with 40GBASE-SR4 or 40GBASE-LR4, and the 100 Gb/s Parallel Physical Interface (CPPI, ten lanes) with 100GBASE-SR10. The term “nPPI” denotes either or both.

The PMD and PMA attached to the nPPI are required to comply with the delay, Skew, and Skew Variation requirements in 86.3, 83.5.3, and 83.5.4 as appropriate. The PMD MDIO function mapping given in 86.4 may apply. The PMD functional specifications are as given in 86.5.

This annex is arranged as follows: following the overview, a reminder of the block diagram, a brief introduction to the test points, and lane assignments, 86A.4 contains the electrical specifications for nPPI from host to module (Tx side) and then module to host (Rx side). Test points, compliance boards, and electrical parameters are defined in 86A.5. The host PCB (“channel”) response recommendation is provided in 86A.6. Safety, installation, environment, and labeling is addressed in 86A.7 and PICS is provided in 86A.8. This annex is very closely related to Clause 86.

86A.2 Block diagram and test points

The PMD block diagram is shown in Figure 86–2. Figure 86–3 shows the test points.

The nPPI is standardized at the test points described in 86.8.1. The transmit side electrical signal (host electrical output and module electrical input) is defined at the output of the Host Compliance Board (TP1a), and other specifications of the module electrical input port are defined at the input of the Module Compliance Board (TP1). The receive side electrical signal (module electrical output and host electrical input) is defined at the output of the Module Compliance Board (TP4), and other specifications of the host electrical input port are defined at the input of the Host Compliance Board (TP4a). Test points and compliance boards are defined more thoroughly in 86A.5.1.

86A.3 Lane assignments

There are no lane assignments (within a group of transmit or receive lanes) for XLPPI or CPPI. While it is expected that a PMD will map electrical lane i to optical lane i and vice versa, there is no need to define the physical ordering of the lanes, as the PCS is capable of receiving the lanes in any arrangement.

86A.4 Electrical specifications for nPPI

The signaling rate for a lane of an XLPPi or CPPI interface shall be as defined in Table 86–2. 86A.4.1 and 86A.4.2 specify the host to module (Tx side) and module to host (Rx side) respectively of the nPPI. Parameters are defined in 86A.5 and 86.8. A recommended PCB (“channel”) response for the host (PMA) is provided in 86A.6. Test points are defined in 86A.5.1.

86A.4.1 nPPI host to module electrical specifications

Each output lane and signal of the nPPI host (PMA), if measured at TP1a (see 86A.5.1) with the specified crosstalk signals applied on all input lanes, shall meet the specifications of Table 86A–1 per the definitions in 86A.5. Each lane of the nPPI module (PMD) electrical input, if measured at TP1 and TP1a with all Rx lanes (module output) operating, shall meet the specifications of Table 86A–2 per the definitions in 86A.5. The module electrical input shall be AC coupled, i.e., it shall present a high DC common-mode impedance at TP1. There may be various methods for AC coupling in actual implementations.

Table 86A–1—nPPI host electrical output specifications at TP1a

Parameter description	Min	Max	Units	Conditions
Single ended output voltage	–0.3	4	V	Referred to signal common
AC common-mode output voltage	—	15	mV	RMS
Termination mismatch at 1 MHz	—	5	%	
Differential output return loss	See 86A.4.1.1	—	dB	
Common-mode output return loss	See 86A.4.1.2	—	dB	
Output transition time, 20% to 80%	28	—	ps	
J2 Jitter output	—	0.17	UI	
J9 Jitter output	—	0.29	UI	
Data Dependent Pulse Width Shrinkage (DDPWS)	—	0.07	UI	
Q_{sq}	45	—	V/V	
	Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	0.11, 0.31 95, 350		UI mV	Hit ratio = 5×10^{-5}
Crosstalk source VMA, each input lane	700		mV	At TP4
Crosstalk source transition times, 20% to 80%	34		ps	At TP4

Table 86A–2—nPPI module electrical input specifications at TP1 and TP1a

Parameter description	Test point	Min	Max	Units	Conditions
Single ended input voltage tolerance ^a	TP1a	–0.3	4	V	Referred to TP1 signal common
AC common-mode input voltage tolerance	TP1a	15	—	mV	RMS
Differential input return loss	TP1	See 86A.4.1.1	—	dB	10 MHz to 11.1 GHz
Differential to common-mode input return loss	TP1	10	—	dB	10 MHz to 11.1 GHz
J2 Jitter tolerance	TP1a	0.17	—	UI	
J9 Jitter tolerance	TP1a	0.29	—	UI	
Data Dependent Pulse Width Shrinkage (DDPWS) tolerance	TP1a	0.07	—	UI	
		Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	TP1a	0.11, 0.31 95, 350		UI mV	Hit ratio = 5×10^{-5}
Crosstalk calibration signal VMA	TP4	850		mV	While calibrating compliance signal ^b
Crosstalk calibration signal transition times, 20% to 80%	TP4	34		ps	

^a The single ended input voltage tolerance is the allowable range of the instantaneous input signals.

^b The crosstalk calibration signals are applied to the mated HCB-MCB at TP4a and measured at TP4 following the same principles as the host electrical input calibration (see 86A.5.3.8.5). They are removed before testing.

86A.4.1.1 Differential return losses at TP1 and TP1a

From 10 MHz to 11.1 GHz, the magnitude in decibels of the module differential input return loss at TP1 and the host differential output return loss at TP1a (see 86A.5.1) shall not exceed the limit given in Equation (86A–1) and illustrated in Figure 86A–1.

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.01 \leq f < 4.11 \\ 6.3 - 13 \log_{10}\left(\frac{f}{5.5}\right) & 4.11 \leq f \leq 11.1 \end{array} \right\} \text{ dB} \quad (86A-1)$$

where

$Return_loss(f)$ is the return loss at frequency f
 f is the frequency in GHz

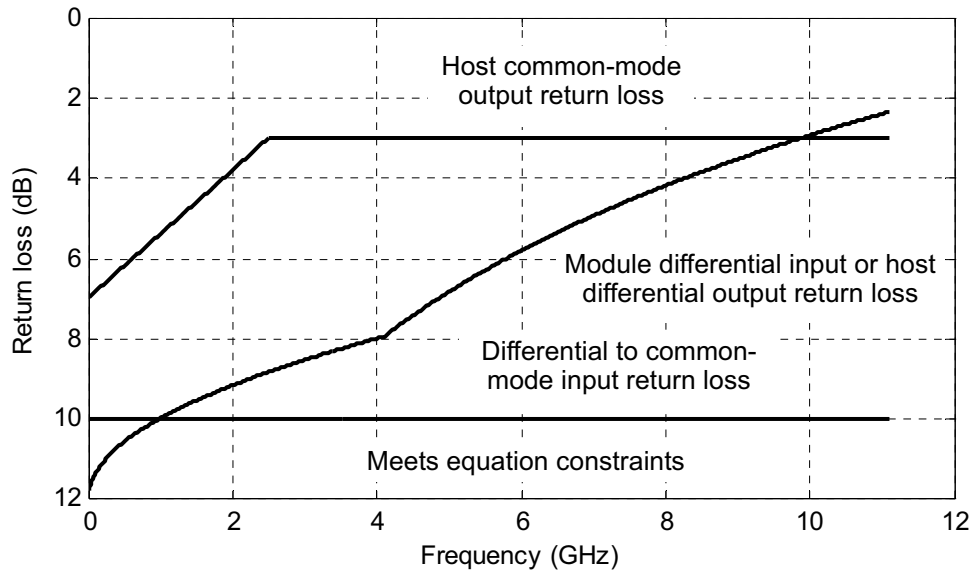


Figure 86A-1—Return loss specifications

86A.4.1.2 Common-mode output return loss at TP1a

From 10 MHz to 11.1 GHz, the magnitude in decibels of the host common-mode output return loss at TP1a shall not exceed the limit given in Equation (86A-2) and illustrated in Figure 86A-1.

$$Return_loss(f) \geq \begin{cases} 7 - 1.6f & 0.01 \leq f < 2.5 \\ 3 & 2.5 \leq f \leq 11.1 \end{cases} \text{ dB} \quad (86A-2)$$

where

$Return_loss(f)$ is the return loss at frequency f
 f is the frequency in GHz

86A.4.2 nPPI module to host electrical specifications

Each electrical output lane and signal of the nPPI module (PMD), if measured at TP4 shall meet the specifications of Table 86A-3 per the definitions in 86A.5 while the specified crosstalk sources are applied to the module's electrical input. Each lane of the nPPI host (PMA) input shall meet the specifications of Table 86A-4 at TP4 and/or TP4a per the definitions in 86A.5. The module electrical output shall be AC coupled, i.e., it shall present a high DC common-mode impedance at TP4. There may be various methods for AC coupling in actual implementations.

86A.4.2.1 Differential return losses at TP4 and TP4a

From 10 MHz to 11.1 GHz, the magnitude in decibels of the module differential output return loss at TP4 and the host differential input return loss at TP4a shall not exceed the limit given in Equation (86A-3) and illustrated in Figure 86A-1.

Table 86A–3—nPPI module electrical output specifications at TP4

Parameter description	Min	Max	Units	Conditions
Single ended output voltage tolerance	–0.3	4	V	Referred to signal common
AC common-mode output voltage (RMS)	—	7.5	mV	
Termination mismatch at 1 MHz	—	5	%	
Differential output return loss	See 86A.4.2.1	—	dB	10 MHz to 11.1 GHz
Common-mode output return loss	See 86A.4.2.2	—	dB	10 MHz to 11.1 GHz
Output transition time, 20% to 80%	28	—	ps	
J2 Jitter output	—	0.42	UI	
J9 Jitter output	—	0.65	UI	
	Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	0.29, 0.5 150, 425		UI mV	Hit ratio = 5×10^{-5}
Crosstalk source VMA, each lane	700		mV	At TP1a
Crosstalk source transition times, 20% to 80%	37		ps	At TP1a

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 12 - 2\sqrt{f} & 0.01 \leq f < 4.11 \\ 6.3 - 13 \log_{10}\left(\frac{f}{5.5}\right) & 4.11 \leq f \leq 11.1 \end{array} \right\} \text{ dB} \quad (86A-3)$$

where

$Return_loss(f)$ is the return loss at frequency f
 f is the frequency in GHz

86A.4.2.2 Common-mode output return loss at TP4

From 10 MHz to 11.1 GHz, the magnitude in decibels of the module common-mode output return loss at TP4 shall not exceed the limit given in Equation (86A–2) and illustrated in Figure 86A–1.

86A.5 Definitions of electrical parameters and measurement methods

Test points are defined in 86A.5.1, compliance boards in 86A.5.1.1, test patterns in 86A.5.2 and parameters in 86A.5.3 and 86.8. Multi-lane testing considerations are given in 86.8.2.1.

86A.5.1 Test points and compliance boards

Figure 86–3 shows the six test points for 40GBASE-SR4 and 100GBASE-SR10. These are TP1, TP1a, TP2, TP3, TP4, and TP4a; four of these are Skew points SP2, SP3, SP4, and SP5 as shown. For 40GBASE-LR4,

Table 86A-4—nPPI host electrical input specifications at TP4 and TP4a

Parameter description	Test point	Min	Max	Units	Conditions
Single ended input voltage ^a	TP4	−0.3	4	V	Referred to signal common
AC common-mode input voltage tolerance	TP4	7.5	—	mV	RMS
Differential input return loss	TP4a	See 86A.4.2.1	—	dB	
Differential to common-mode input return loss	TP4a	10	—	dB	10 MHz to 11.1 GHz
Host input signal tolerance, interface BER limit		—	$10^{-1/2}$	—	
Conditions of host electrical receiver signal tolerance test: ^b					
		Specification values			
Eye mask coordinates: X1, X2 Y1, Y2	TP4	0.29, 0.5 150, 425		UI mV	Hit ratio = 5×10^{-5}
Transition time, 20% to 80%		34		ps	
J2 Jitter	TP4	0.42		UI	
J9 Jitter	TP4	0.65		UI	
Data Dependent Pulse Width Shrinkage (DDPWS)		0.34		UI	
VMA of aggressor lanes	TP4	850		mV	
Crosstalk calibration signal VMA	TP1a	700		mV	
Crosstalk calibration signal transition times, 20% to 80%	TP1a	37		ps	

^a The host is required to tolerate (work correctly with) input signals with instantaneous voltages anywhere in the specified range.

^b The specification values are test conditions for measuring signal tolerance and are not characteristics of the host (see 86A.5.3.8).

points TP1, TP1a, TP4, TP4a, SP2, and SP5 are in equivalent positions. Figure 86-3 also shows the substitution of compliance boards for module (PMD) or host (PMA). These compliance boards are defined to connect generic test equipment to the module and host for test purposes. The module can be plugged into the Module Compliance Board (MCB), which has specified electrical parameters. The Host Compliance Board (HCB), which also has specified electrical parameters, can plug into the host. The MCB and the HCB can be plugged together for calibration of compliance signals and to check the electrical parameters of the boards. Table 86A-5 shows the parameters or signals measured at each point. Also, TP0 and TP5 define the host ends of the electrical channel, at the PMA IC.

All electrical measurements and parameters are defined as through HCB and/or MCB as appropriate, with corrections if necessary.

CAUTION

A PMD with an nPPI interface is AC coupled; however, an HCB is not. The user should take care that the test equipment does not improperly load or damage a host under test.

Table 86A–5—Parameters defined at each test point

Test point	Direction	Parameter
TP1	Looking downstream into module transmitter input	Module transmitter input return loss
TP1a	Looking upstream into host transmitter output	Host transmitter output signal and output return loss, module transmitter compliance signal calibration, host receiver compliance crosstalk signal calibration
TP4	Looking upstream into module receiver output	Module receiver output signal and output return loss, host receiver compliance signal calibration
TP4a	Looking downstream into host receiver input	Host receiver input return loss

86A.5.1.1 Compliance board parameters

The electrical characteristics of the HCB and MCB are given in 86A.5.1.1.1 and 86A.5.1.1.2. If boards are used that do not match the specifications given, the measurement results for nPPI shall be corrected for the differences between the actual HCB or MCB's properties and the reference differential insertion losses given in 86A.5.1.1.1. As it may be impractical to correct eye measurements for a board with differential insertion loss outside the limits given in 86A.5.1.1.2, such boards shall not be used. Boards that do not meet the specifications for mated HCB-MCB in 86A.5.1.1.2 shall not be used.

86A.5.1.1.1 Reference insertion losses of HCB and MCB

The reference differential insertion loss in decibels of the HCB and of the MCB, excluding the module connector, are given in Equation (86A–4) and Equation (86A–5) respectively, and illustrated in Figure 86A–2.

For the HCB,

$$Insertion_loss(f) = 0.01 + 0.3\sqrt{f} + 0.11f \quad \text{dB} \quad 0.01 \leq f \leq 11.1 \quad (86A-4)$$

where

$Insertion_loss(f)$ is the insertion loss at frequency f
 f is the frequency in GHz

For the MCB,

$$Insertion_loss(f) = 0.0006 + 0.16\sqrt{f} + 0.0587f \quad \text{dB} \quad 0.01 \leq f \leq 11.1 \quad (86A-5)$$

where

$Insertion_loss(f)$ is the insertion loss at frequency f
 f is the frequency in GHz




Figure 86A-2—Reference differential insertion losses of HCB, MCB excluding connector

86A.5.1.1.2 Electrical specifications of mated HCB and MCB

The limits on the differential insertion loss in decibels of the mated HCB and MCB (in either direction) are given in Equation (86A-6) (which defines the minimum insertion loss) and Equation (86A-7) (which defines the maximum insertion loss). These limits are illustrated in Figure 86A-3.




Figure 86A-3—Differential insertion loss limits of mated HCB-MCB

$$Insertion_loss(f) \geq -0.11 + 0.46\sqrt{f} + 0.16f \text{ dB} \quad 0.01 \leq f \leq 11.1 \quad (86A-6)$$

$$Insertion_loss(f) \leq \left\{ \begin{array}{ll} 0.029 + 0.861\sqrt{f} + 0.158f & 0.01 \leq f < 5.5 \\ 0.2 + 0.65f & 5.5 \leq f \leq 11.1 \end{array} \right\} \text{ dB} \quad (86A-7)$$

where

$Insertion_loss(f)$ is the insertion loss at frequency f
 f is the frequency in GHz

The limit on the differential return loss in decibels of the mated HCB and MCB, looking into the HCB or looking into the MCB, is given in Equation (86A-8), and illustrated in Figure 86A-4.




Figure 86A-4—Return loss limits of mated HCB-MCB

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 20 - 2f & 0.01 \leq f < 2.5 \\ 15 & 2.5 \leq f < 5 \\ 13.8 - 28.85 \log_{10}\left(\frac{f}{5.5}\right) & 5 \leq f \leq 11.1 \end{array} \right\} \text{ dB} \quad (86A-8)$$

where

$Return_loss(f)$ is the return loss at frequency f
 f is the frequency in GHz

The limit on the common-mode return loss in decibels of the mated HCB and MCB is given in Equation (86A-9) and illustrated in Figure 86A-4.

$$Return_loss(f) \geq \begin{cases} 12 - 2.8f & 0.01 \leq f < 2.5 \\ 5.2 - 0.08f & 2.5 \leq f \leq 15 \end{cases} \text{ dB} \quad (86A-9)$$

where

$Return_loss(f)$ is the return loss at frequency f
 f is the frequency in GHz

The limit on the differential to common-mode conversion loss in decibels of the mated HCB and MCB, for input to HCB and output from MCB, or input to MCB and output from HCB, is given in Equation (86A-10) and illustrated in Figure 86A-5.




Figure 86A-5—Differential to common-mode conversion loss limit of mated HCB-MCB

$$Mode_conversion_loss(f) \geq \begin{cases} 30 - 2.91f & 0.01 \leq f < 5.5 \\ 14 & 5.5 \leq f \leq 15 \end{cases} \text{ dB} \quad (86A-10)$$

where

$Mode_conversion_loss(f)$ is the mode conversion loss at frequency f
 f is the frequency in GHz

The limits on integrated crosstalk noise of the mated HCB and MCB are as specified in 85.10.9.5 with the exception that the frequency range is 0.01 GHz to 12 GHz.

86A.5.2 Test patterns and related subclauses

While compliance is to be achieved in normal operation, specific test patterns are defined for measurement consistency and to enable measurement of some parameters. Table 86-11 lists the defined test patterns, and Table 86A-6 gives the test patterns to be used in each measurement, unless otherwise specified, and also lists references to the subclauses in which each parameter is defined. Multi-lane testing considerations are given in 86.8.2.1. As Pattern 3 is more demanding than Pattern 5 (which itself is the same or more

Table 86A–6—Test patterns and related subclauses

Parameter	Pattern	Related subclause
J2 Jitter	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.3.1
J9 Jitter	3 or 5	86.8.3.3.2
Data Dependent Pulse Width Shrinkage (DDPWS)	4	86A.5.3.4
AC common-mode voltage	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86A.5.3.1
Transition time	Square wave or 4	86A.5.3.3
Electrical waveform (eye mask)	3, 5, or valid 40GBASE-SR4 or 100GBASE-SR10 signal	86.8.3.2, 86A.5.3.6
Q_{sq}	Square wave or 4	86A.5.3.5
Host electrical receiver signal tolerance	3 or 5	86A.5.3.8

demanding than other 40GBASE-R or 100GBASE-R bit streams) an item which is compliant using Pattern 5 is considered compliant even if it does not meet the required limit using Pattern 3.

86A.5.3 Parameter definitions

In addition to the parameter definitions in the following subclause, some definitions with dual use (both optical and electrical) are given in 86.8.3.

86A.5.3.1 AC common-mode voltage

The common-mode voltage of a differential signal at any time is the average of signal+ and signal– at that time. RMS AC common-mode voltage may be calculated by applying the histogram function over 1 UI to the common-mode signal. As AC common-mode generation is very sensitive to the cable or oscilloscope delay mismatch, it is recommended to delay match the oscilloscope inputs for any measurements.

86A.5.3.2 Termination mismatch

Termination mismatch is the percentage difference between the two low-frequency impedances to common of a differential electrical port. Termination mismatch is defined as shown in Equation (86A–11).

$$\Delta Z_M = 2 \times \frac{|Z_p - Z_n|}{Z_p + Z_n} \times 100 \% \quad (86A-11)$$

Termination mismatch can be measured by applying a low-frequency test tone to the differential inputs as shown in Figure 86A–6. The test frequency must be high enough to overcome the high pass effects of any AC coupling capacitors. The measured differential output or input impedance is designated by Z_{diff} .

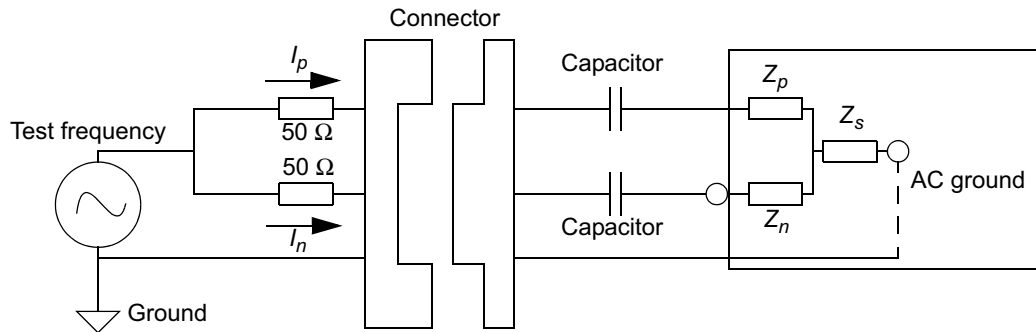


Figure 86A-6—Measurement of AC termination mismatch

For a 100 Ω port, low-frequency termination mismatch is then given by Equation (86A-12).

$$\Delta Z_M = 2 \times \frac{|I_p - I_n|}{I_p + I_n} \times \frac{Z_{diff} + 100}{Z_{diff}} \times 100 \% \quad (86A-12)$$

where I_p and I_n are the currents flowing into the port as shown in Figure 86A-6.

Z_s is the effective series impedance between the terminations Z_p and Z_n and the AC ground.

86A.5.3.3 Transition time

In this annex, transition times (rise and fall times) are defined as the time between the 20% and 80% times, or 80% and 20% times, respectively, of isolated edges.

If the test pattern is the square wave with eight ones and eight zeros, the 0% level and the 100% level are as defined by the OMA measurement procedure (see 68.6.2).

If the test pattern is PRBS9, the transitions within sequences of five zeros and four ones, and nine ones and five zeros, respectively, are measured. These are bits 10 to 18 and 1 to 14, respectively, where bits 1 to 9 are the run of nine zeros. In this case, the 0% level and the 100% level may be estimated as ZeroLevel and ZeroLevel + MeasuredOMA in the TWDP code (see 68.6.6.2), or by the average signal within windows from $-3 UI$ to $-2 UI$ and from $2 UI$ to $3 UI$ relative to the edge.

For electrical signals, the waveform is observed through a 12 GHz low-pass filter response (such as a Bessel-Thomson response).

NOTE—This definition is not the same as the rise and fall times typically reported by an oscilloscope from an eye diagram, which take all the edges into account.

86A.5.3.4 Data Dependent Pulse Width Shrinkage (DDPWS)

An oscilloscope, time interval analyzer, or other instrument with equivalent capability may be used to measure DDPWS. A clock recovery unit (CRU) as defined in 86.8.3.2 is used to trigger the oscilloscope. A repeating PRBS9 pseudo-random test pattern, 511 bits long, is used. For electrical jitter measurements, the measurement bandwidth is 12 GHz (such as a Bessel-Thomson response). If the measurement bandwidth affects the result, it can be corrected for by post-processing. However, a bandwidth above 12 GHz is expected to have little effect on the results.

The crossing level is the average value of the entire waveform being measured. The instrument is synchronized to the pattern repetition frequency and the waveforms or the crossing times are averaged

sufficiently to remove the effects of random jitter and noise in the system. The PRBS9 pattern has 128 positive-going transitions and 128 negative-going transitions. The crossing times t_1 to t_{256} of each transition of the averaged waveform (when the averaged waveform crosses the crossing level) are found as shown in Figure 86A-7.




Figure 86A-7—Data Dependent Pulse Width Shrinkage test method

The DDPWS is the difference between one symbol period and the minimum of all the differences between pairs of adjacent transitions as follows in Equation (86A-13).

$$DDPWS = T - \min(t_2 - t_1, t_3 - t_2, \dots, t_{256+1} - t_{256}) \quad (86A-13)$$

where T is one symbol period.

Note that the difference from the next edge in the repeating sequence, t_{256+1} , is also considered.

86A.5.3.5 Signal to noise ratio Q_{sq}

Q_{sq} is a measure of signal to noise ratio. For an electrical signal, it is analogous to Q_{sq} defined for an optical signal in 68.6.7, and it relates the low-frequency signal amplitude to the noise in an electrical -3 dB bandwidth of 12 GHz. It is defined with all co-propagating and counter-propagating crosstalk sources active, using one of patterns 3, 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal. The input lanes of the item under test are receiving signals that are asynchronous to those being output.

Q_{sq} may be measured using an oscilloscope as follows:

- The Voltage Modulation Amplitude (VMA) of the output lane is measured, using a square wave (8 ones, 8 zeros) or PRBS9 (Pattern 4). VMA is the difference between the 0% level and the 100% level defined in 86A.5.3.3; it is defined by analogy to OMA (see 68.6.2) but with a 12 GHz observation bandwidth.
- Using the same pattern, the RMS noise over flat regions of the logic one and logic zero portions of the signal (see Figure 68-4 for the analogous optical measurement), is measured, compensating for noise in the measurement system. If possible, means should be used to prevent noise of frequency less than 1 MHz from affecting the result.

- c) Q_{sq} is given by Equation (86A-14).

$$Q_{sq} = VMA / (n1 + n0) \quad (86A-14)$$

where

- $n1$ is the RMS noise of logic one
 $n0$ is the RMS noise of logic zero

86A.5.3.6 Eye mask for TP1a and TP4

The eye mask is defined by parameters X1, X2, Y1, and Y2. Unlike the optical eye mask, the vertical dimensions are fixed rather than scaled to the signal. Figure 83A-8 (an example of a hexagonal eye mask such as at TP1a) and Figure 83A-9 (a diamond mask such as at TP4) show the meaning of the parameters X1, X2, Y1 and Y2. The eye is defined as measured using a receiver with an electrical -3 dB bandwidth of 12 GHz (such as a Bessel-Thomson response). Further requirements are given in 86.8.3.2.

86A.5.3.7 Reference impedances for electrical measurements

The reference impedance for differential electrical measurements is 100 Ω and the reference impedance for common-mode electrical measurements is 25 Ω .

86A.5.3.8 Host input signal tolerance

To be compliant the host input signal tolerance shall satisfy the requirements of 86A.5.3.8.1 to 86A.5.3.8.6.

86A.5.3.8.1 Introduction

This subclause provides guidance for jitter tolerance testing at the host input (PMA) compliance point TP4/TP4a. Compliance is required with input jitter, vertical eye closure (Y1), and vertical peak level (Y2) as specified in Table 86A-4. Compliance is defined at an interface BER (the average of the BERs of all the lanes when stressed) of 10^{-12} . There are two test conditions: one each for the sensitivity and overload vertical eye parameters conditions. The reference test procedure is described in detail for a single stressed lane. Each Rx lane is tested in turn while all lanes are operated. Aggressor lanes are operated with the VMA specified in Table 86A-4.

86A.5.3.8.2 Test equipment and setup

A jitter tolerance test configuration is shown in Figure 86A-8. A test source is used to continuously generate the test signal. The test signal is conditioned within the guidelines outlined in 86A.5.3.8.3 to exhibit the appropriate jitter stress.

An RF attenuator or other output amplitude control of the test source may be required to set the vertical eye opening of the stressed eye.

The test equipment measured at TP4a looking towards the test signal source has better than 20 dB output return loss up to 12 GHz. The output return loss of the test system when measured at TP4 with the Module Compliance Board (looking towards the test signal source) is 2 dB better than the specifications of Table 86A-3 up to 8 GHz and 1 dB better up to 11 GHz.

86A.5.3.8.3 Stressed eye jitter characteristics

This subclause describes required test signal characteristics along with considerations and suggested approaches for test signal generation. The test signal is generated by the functions shown in Figure 86A-8 or

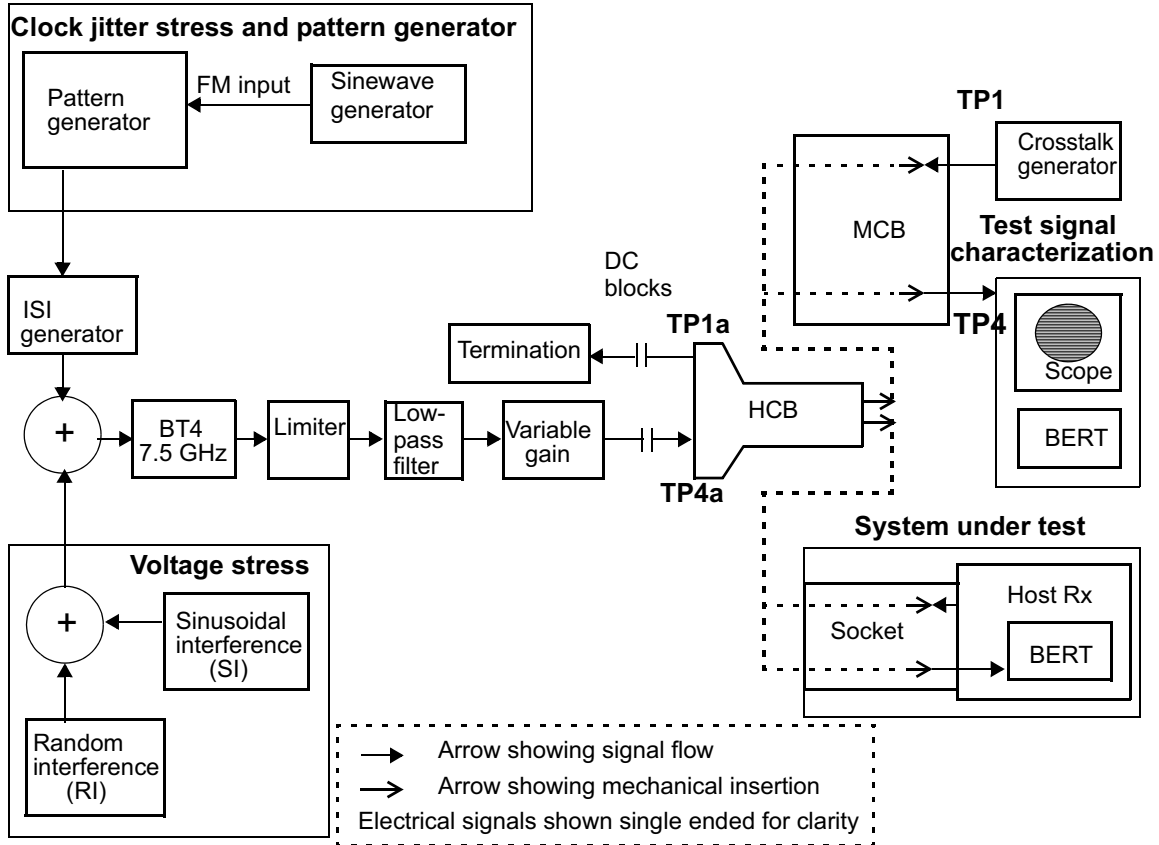


Figure 86A-8—Example jitter tolerance test configuration

by equivalent means. Figure 86A-9 illustrates how the jitter parameters in Table 86A-4 map to the jitter components in the stressed-eye test signal.




Figure 86A-9—Stressed eye jitter components

The 0.05 UI Sinusoidal Jitter (SJ) component of J2 Jitter is defined for frequencies much higher than the CDR bandwidth (e.g., ~20 MHz). At lower frequencies, the CDR must track additional applied SJ as detailed in the relevant specifications (see Figure 86A-10 and 52.8.1).

The balance of the J2 Jitter is created by the following mechanisms: an ISI generator, sinusoidal interference (SI), and random interference (RI), all passed through a limiting function.

The test signal at TP4 has DDPWS as defined by Table 86A–4. Any duty cycle distortion (DCD) in the test signal does not exceed 0.02 UI. DCD is the difference between the mean position of all falling edges and the mean position of all rising edges of a waveform. It is measured at the average level of the waveform in a 12 GHz bandwidth, using Pattern 3 (PRBS31), Pattern 4 (PRBS9), Pattern 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal.

The ISI generator may be implemented by a low-pass filter, length of FR4 trace, length of coax cable, or other equivalent method. It is required that the resulting signal be passed through a limiter function. A suitable limiter function may be implemented using a discrete limiting amplifier followed by a low-pass filter and an attenuator. The low-pass filter emulates the bandwidth and/or slew rate of a practical limiter. The attenuator is used to set the output amplitude to minimum and maximum values allowed by the eye mask coordinates of Table 86A–4.

A voltage stress is to be applied before the limiter function. This stress is composed of a single tone sinusoidal interferer (SI) in the frequency range 100 MHz to 2 GHz and a broadband noise source (RI) with a minimum power spectrum –3 dB point of 6 GHz and minimum crest factor of 7. It is the intent that this combination of voltage stress and limiting function introduce pulse-shrinkage jitter behavior. However, no more than 20% of the J2 Jitter is created by the sinusoidal interferer.

Jitter generation mechanisms for the pattern generator are typically based on phase modulation of the clock source, edge modulation of a variable delay line or a combination thereof.

Any approach that modulates or creates the appropriate levels and frequencies of the jitter components is acceptable.

86A.5.3.8.4 Calibration

Calibration of the test signal is performed using the guidelines for test setup in 86A.5.3.8.2 and illustrated in Figure 86A–8. The aim of the calibration is to achieve a test signal exhibiting jitter stress in accordance with Table 86A–4.

The test signal is calibrated differentially into standard instrumentation loads. If complementary single-ended signals are used, they are carefully matched in both amplitude and phase.

For accurate calibration, it is imperative that all elements in the signal path (cables, DC blocks, etc.) have wide and smooth frequency response as well as linear phase response throughout the spectrum of interest. Baseline wander and overshoot/undershoot are minimized.

An AC coupling –3 dB corner frequency of 20 kHz is expected to be adequate to eliminate baseline wander effects; however, high-frequency performance is critical and must not be sacrificed by the AC coupling.

Given random jitter and the nature of the long test patterns, low probability jitter events will likely be present. It is recommended for jitter calibration that a technique that can accurately measure low probability events be used to avoid overly stressful test conditions.

It is recommended that the actual compliance test pattern be used during calibration. For jitter stress calibration it is permissible, however, to use any appropriate test pattern that still results in the creation of a compliance test pattern with the appropriate jitter stress.

86A.5.3.8.5 Calibration procedure

The vertical eye opening and peak level are set approximately to the levels specified in Table 86A–4.

With an applied calibration test pattern and no additional jitter stress applied, the intrinsic J2 Jitter and J9 Jitter of the test source due to intrinsic noise and finite bandwidth effects are measured. At this stage, J2 Jitter is less than 0.15 UI and J9 Jitter less than 0.25 UI.

SJ is added until the J2 Jitter increases by 0.05 UI above this intrinsic J2 Jitter level. The SJ jitter frequency is well above the CDR bandwidth and asynchronous to the characteristic frequencies of the signal.

Next, additional high probability jitter as specified in 86A.5.3.8.3 is added by the ISI generator until at least 80% of the J2 Jitter has been created. The sinusoidal interferer amplitude is then turned on and adjusted until the required level of J2 Jitter is achieved. The frequency of any sinusoidal interferer is asynchronous to the characteristic frequencies of the signal.

A compliant test signal exhibits Data Dependent Pulse Width Shrinkage (defined in 86A.5.3.4) as specified in Table 86A-4. This is measured with noise and clock-jitter sources turned off.

Once the required level of J2 Jitter has been achieved, turn on the crosstalk source that is set such that at the output of the Host Compliance Board, the amplitude and the transition time are as given in Table 86A-4. The crosstalk pattern is Pattern 3 (PRBS31), Pattern 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal, and is asynchronous with the test signal. Then the RI (random interference) voltage stress is added until the specified value of J9 Jitter is achieved.

If necessary, the sine interferer is readjusted to obtain the required level of J2 Jitter and if the sinusoidal interferer is changed then the random interferer is readjusted to obtain the required level of J9 Jitter. Iterative adjustments of the sinusoidal interferer and random interferer are made until the required values of both J2 Jitter and J9 Jitter are achieved.

If necessary, the vertical eye opening is readjusted to required levels.

The vertical eye opening and peak level specifications are verified.

Care must be taken when characterizing the signal used to make receiver tolerance measurements. The intrinsic noise and jitter introduced by the calibration measurement equipment (e.g., filters, oscilloscope and BERT) must be accounted for and controlled. If equipment imperfections affect the results materially, corrections such as root-sum-of-squares deconvolution of Gaussian noise and jitter are used.

86A.5.3.8.6 Test procedure

Testing is performed differentially through a Host Compliance Board (see 86A.5.1).

Using a test signal arranged according to 86A.5.3.8.2 and calibrated according to 86A.5.3.8.5, operate the system with the test pattern specified in Table 86A-6. Each lane is tested in turn while all are operated. Aggressor lanes are operated with the VMA specified in Table 86A-4. The BERs of all the lanes when stressed are averaged to form the interface BER. See 86.8.2.1.

All signals and reference clocks that operate during normal operation are active during the test including all the other host lanes in both directions. The test signal and the host's transmitted signals are asynchronous. The host transmits Pattern 3 (PRBS31), Pattern 5, or a valid 40GBASE-SR4 or 100GBASE-SR10 signal. The sinusoidal jitter is stepped across the frequency and amplitude range according to Table 86A-7 and illustrated in Figure 86A-10, while monitoring the BER of the lane(s). The interface BER of a compliant host receiver remains below 10^{-12} .




Figure 86A-10—Mask of the sinusoidal component of jitter tolerance

Table 86A-7—Applied sinusoidal jitter

Frequency range	Sinusoidal jitter, peak-to-peak, (UI)
$f < 40 \text{ kHz}$	Not specified
$40 \text{ kHz} \leq f < 4 \text{ MHz}$	$2 \times 10^5 / f$
$4 \text{ MHz} \leq f \leq 10 \text{ LB}^a$	0.05

^aLB = loop bandwidth; upper frequency bound for added sine jitter should be at least 10 times the loop bandwidth of the receiver being tested.

86A.6 Recommended electrical channel

The recommended limits for the differential insertion loss in decibels of the host PCB and connector mated to the HCB, between the PMA IC (TP0 or TP5) and TP1a or TP4a, are given in Equation (86A-15) and Equation (86A-16), and illustrated in Figure 86A-11. It is recommended that

$$\text{Insertion_loss}(f) \leq \left\{ \begin{array}{ll} 0.5 & 0.01 \leq f < 0.11 \\ 0.114 + 0.8914\sqrt{f} + 0.846f & 0.11 \leq f < 7 \\ -35.91 + 6.3291f & 7 \leq f < 8 \\ 14.72 & 8 \leq f \leq 11.1 \end{array} \right\} \text{ dB} \tag{86A-15}$$

and

$$\text{Insertion_loss}(f) \geq \left\{ \begin{array}{ll} -0.22 + 0.46f & 0.01 \leq f < 7 \\ 3 & 7 \leq f \leq 11.1 \end{array} \right\} \text{ dB} \tag{86A-16}$$

where

Insertion_loss(*f*) is the insertion loss at frequency *f*
f is the frequency in GHz




Figure 86A-11—Recommended insertion loss limits of host PCB, connector and HCB

The recommended maximum loss of the host PCB only (without connector or HCB) at 5.15625 GHz is 4.4 dB.

86A.7 Safety, installation, environment, and labeling

86A.7.1 General safety

All equipment subject to this annex shall conform to IEC 60950-1.

86A.7.2 Installation

It is recommended that proper installation practices, as defined by applicable local codes and regulation, be followed in every instance in which such practices are applicable.

86A.7.3 Environment

The 40GBASE-SR4 and 100GBASE-SR10 operating environment specifications are as defined in 52.11, as defined in 52.11.1 for electromagnetic emission, and as defined in 52.11.2 for temperature, humidity, and handling.

86A.7.4 PMD labeling

The 40GBASE-SR4 and 100GBASE-SR10 labeling recommendations and requirements are as defined in 52.12.

86A.8 Protocol implementation conformance statement (PICS) proforma for Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPi) and 100GBASE-SR10 (CPPI)²⁵

86A.8.1 Introduction

The supplier of an XLPPi/CPPI implementation that is claimed to conform to Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPi) and 100GBASE-SR10 (CPPI), shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the same, can be found in Clause 21.

86A.8.2 Identification

86A.8.2.1 Implementation identification

Supplier ¹	
Contact point for enquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations.	
NOTE 2—May be completed as appropriate in meeting the requirements for the identification.	
NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).	

86A.8.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ba-2010 Annex 86A, Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPi) and 100GBASE-SR10 (CPPI)
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21; the answer Yes means that the implementation does not conform to IEEE Std 802.3ba-2010.)	

Date of Statement	
-------------------	--

²⁵Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

86A.8.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*R4	40GBASE-R4	86A.4	Can operate as part of 40GBASE-SR4 or 40GBASE-LR4 PMA or PMD	O.1	Yes [] No []
*R10	100GBASE-SR10	86A.4	Can operate as part of 100GBASE-SR10 PMA or PMD	O.1	Yes [] No []
*MO	Module		Items marked with MO are applicable to a module (PMD)	O.2	Yes [] No []
*HO	Host		Items marked with HO are applicable to a host (PMA)	O.2	Yes [] No []
*MD	MDIO capability	86A.1	Registers and interface supported	O	Yes [] No []

86A.8.4 PICS proforma tables for Parallel Physical Interface (nPPI) for 40GBASE-SR4 and 40GBASE-LR4 (XLPPI) and 100GBASE-SR10 (CPPI)

86A.8.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SF1	Compatible with 40GBASE-R or 100GBASE-R PCS and PMA and 40GBASE-SR4, 40GBASE-LR4 or 100GBASE-SR10 PMD	86A.1		M	Yes []

86A.8.4.2 Electrical specifications for nPPI

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Signaling rate per Table 86–2	86A.4	10.3125 GBd \pm 100 ppm	M	Yes []
S2	Host output signal per Table 86A–1	86A.4.1	Per definitions in 86A.5	HO:M	Yes [] N/A []
S3	Module electrical input per Table 86A–2	86A.4.1	Per definitions in 86A.5	MO:M	Yes [] N/A []
S4	Module electrical input AC coupled	86A.4.1		MO:M	Yes [] N/A []
S5	Module electrical differential input return loss at TP1	86A.4.1.1	Per Equation (86A–1)	MO:M	Yes [] N/A []
S6	Host differential output return loss at TP1a	86A.4.1.1	Per Equation (86A–1)	HO:M	Yes [] N/A []
S7	Host common-mode output return loss at TP1a	86A.4.1.2	Per Equation (86A–2)	HO:M	Yes [] N/A []
S8	Module electrical output per Table 86A–3	86A.4.2	Per definitions in 86A.5	MO:M	Yes [] N/A []
S9	Host input per Table 86A–4	86A.4.2	Per definitions in 86A.5	HO:M	Yes [] N/A []
S10	Module electrical output AC coupled	86A.4.2		MO:M	Yes [] N/A []
S11	Module electrical output return loss at TP4	86A.4.2.1	Per Equation (86A–3)	MO:M	Yes [] N/A []
S12	Host differential input return loss at TP4a	86A.4.2.1	Per Equation (86A–3)	HO:M	Yes [] N/A []
S13	Module electrical common-mode output return loss at TP4	86A.4.2.2	Per Equation (86A–2)	MO:M	Yes [] N/A []

86A.8.4.3 Definitions of parameters and measurement methods

Item	Feature	Subclause	Value/Comment	Status	Support
SEM1	Compliance boards	86A.5.1	Use compliance boards, correct as necessary	M	Yes []
SEM2	Compliance boards	86A.5.1.1	Results corrected	M	Yes []
SEM3	Compliance boards	86A.5.1.1	Individual insertion losses per 86A.5.1.1.2	M	Yes []
SEM4	Compliance boards	86A.5.1.1	Mated HCB-MCB per 86A.5.1.1.2	M	Yes []
SEM5	Host input signal tolerance	86A.5.3.8	As specified	HO:M	Yes [] N/A []

86A.8.4.4 Environmental and safety specifications

Item	Feature	Subclause	Value/Comment	Status	Support
SES1	General safety	86A.7.1	Conforms to IEC 60950-1	M	Yes []
SES3	Electromagnetic interference	86A.7.3	Complies with applicable local and national codes for the limitation of electromagnetic interference	M	Yes []