

INTERNATIONAL STANDARD



**Information technology –Fibre channel –
Part 151: Fibre Channel BaseT (FC-BaseT)**

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017





THIS PUBLICATION IS COPYRIGHT PROTECTED
Copyright © 2017 ISO/IEC, Geneva, Switzerland

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either IEC or IEC's member National Committee in the country of the requester. If you have any questions about ISO/IEC copyright or have an enquiry about obtaining additional rights to this publication, please contact the address below or your local IEC member National Committee for further information.

IEC Central Office
3, rue de Varembe
CH-1211 Geneva 20
Switzerland

Tel.: +41 22 919 02 11
Fax: +41 22 919 03 00
info@iec.ch
www.iec.ch

About the IEC

The International Electrotechnical Commission (IEC) is the leading global organization that prepares and publishes International Standards for all electrical, electronic and related technologies.

About IEC publications

The technical content of IEC publications is kept under constant review by the IEC. Please make sure that you have the latest edition, a corrigenda or an amendment might have been published.

IEC Catalogue - webstore.iec.ch/catalogue

The stand-alone application for consulting the entire bibliographical information on IEC International Standards, Technical Specifications, Technical Reports and other documents. Available for PC, Mac OS, Android Tablets and iPad.

IEC publications search - www.iec.ch/searchpub

The advanced search enables to find IEC publications by a variety of criteria (reference number, text, technical committee,...). It also gives information on projects, replaced and withdrawn publications.

IEC Just Published - webstore.iec.ch/justpublished

Stay up to date on all new IEC publications. Just Published details all new publications released. Available online and also once a month by email.

Electropedia - www.electropedia.org

The world's leading online dictionary of electronic and electrical terms containing 20 000 terms and definitions in English and French, with equivalent terms in 16 additional languages. Also known as the International Electrotechnical Vocabulary (IEV) online.

IEC Glossary - std.iec.ch/glossary

65 000 electrotechnical terminology entries in English and French extracted from the Terms and Definitions clause of IEC publications issued since 2002. Some entries have been collected from earlier publications of IEC TC 37, 77, 86 and CISPR.

IEC Customer Service Centre - webstore.iec.ch/csc

If you wish to give us your feedback on this publication or need further assistance, please contact the Customer Service Centre: csc@iec.ch.

IECNORM.COM : Click to view the full PDF of ISO/IEC 2017-05-15:2017



ISO/IEC 14165-151

Edition 1.0 2017-11

INTERNATIONAL STANDARD



Information technology –Fibre channel –
Part 151: Fibre Channel BaseT (FC-BaseT)

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 35.200

ISBN 978-2-8322-5100-3

Warning! Make sure that you obtained this publication from an authorized distributor.



Information Technology — Fibre Channel — Part 151: BaseT

Technologies de l'information — Fibre Channel — Partie 151 (FC-BaseT)

Contents	Page
FOREWORD	7
INTRODUCTION	9
1 Scope	10
2 Normative references	10
3 Terms, definitions, abbreviations, symbols, and conventions	11
3.1 Terms and definitions	11
3.2 Editorial conventions	14
3.3 Abbreviations, acronyms, and symbols	14
3.4 Keywords	15
3.5 State Diagram notation	16
3.5.1 State Diagram conventions	16
3.5.2 State Diagram variables	17
3.5.3 State Diagram timers	17
3.5.4 State transitions	17
3.5.5 Operators	18
4 Structure and concepts	19
4.1 Overview	19
4.2 Relationship with other standards	20
4.3 FC-BaseT PHY logical model	21
4.4 FC-BaseT usage of XGMII	22
4.5 Operation of FC-BaseT	23
4.5.1 Overview	23
4.5.2 PCS overview	24
4.5.3 PMA overview	25
4.6 FC-BaseT service primitives and interfaces	25
4.6.1 Overview	25
4.6.2 PMA service interface	25
4.6.3 Management function interface	30
4.7 FC-BaseT Nomenclature	31
5 Physical Coding Sublayer (PCS)	32
5.1 Overview	32
5.2 PCS reset function	32
5.3 PCS transmit function	33
5.3.1 Overview	33
5.3.2 36/33 transcoding	34
5.3.3 Error detecting code	36
5.3.4 PCS scrambling	37
5.3.5 Schläfli Lattice coding	39
5.3.6 Trellis coding	40
5.3.7 Generation of PMA training sequences	41
5.4 PCS receive function	42
5.4.1 Overview	42
5.4.2 Decoding	43
5.4.3 PCS descrambling	44
5.4.4 33/36 transcoding	45
5.4.5 PCS synchronization	46
5.5 State Diagrams	46

6	PMA Sublayer and Medium Dependent Interface	48
6.1	PMA Overview	48
6.2	PMA Functions	49
6.2.1	PMA Reset function	49
6.2.2	PMA Transmit function	49
6.2.3	PMA Receive Function	49
6.2.4	PHY Control Function	50
6.2.5	Link Monitor Function	53
6.2.6	Clock Recovery Function	53
6.2.7	State Diagrams	53
6.3	PMA Electrical Specification	56
6.3.1	Isolation and EMC Requirements	56
6.3.2	Test Modes	56
6.3.3	Transmitter Electrical Specifications	59
6.3.4	Receiver Electrical Specifications	64
6.4	MDI Specification	66
6.4.1	Overview	66
6.4.2	MDI Mechanical Specification	66
6.4.3	Automatic MDI/MDI-X Configuration	67
6.4.4	MDI Electrical Specification	67
6.4.5	MDI Fault tolerance	67
6.5	Link Segment Characteristics	67
6.5.1	Overview	67
6.5.2	FC-BaseT Link Topology	68
6.5.3	FC-BaseT Cable Plant Requirements	68
7	Elasticity FIFO	69
7.1	Overview	69
7.2	Ordered sets processing	70
7.3	Clock skew compensation	71
8	PHY startup procedure	73
8.1	Overview	73
8.2	Host speeds determination	73
8.3	Host synchronization	73
8.4	FC-BaseT auto-negotiation	74
8.4.1	Overview	74
8.4.2	FC-BaseT support	75
8.4.3	Master-Slave relationship	76
8.4.4	Cable length estimation	77
8.4.5	Tentative Operating Speed Determination	78
8.4.6	Configuration resolution	79
8.5	Speed downshift function	80
8.6	State diagrams	80
9	Port management	86
9.1	Overview	86
9.2	FC-BaseT management registers	86
9.3	Control Register	88
9.4	Status Register	91
9.5	Speed Downshift Register	93
9.6	Test Register	94
9.7	SNR Margin Registers	94
9.8	Auto-Negotiation Pages Registers	95

Annex A (informative) Additional 36/33 Mappings	98
A.1 Overview	98
A.2 Notation	98
A.3 Additional mappings	98
Annex B (informative) Schläfli Lattice coding	100
Annex C (informative) Recommended 4GFC-BaseT power schedule	102
Bibliography	103

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

Figure	Page
Figure 1	State diagram notation example 16
Figure 2	FC-BaseT topology 19
Figure 3	FC-BaseT relationship to the OSI Reference Model, the Fibre Channel Levels, and the IEEE 802.3 LAN Model. 20
Figure 4	An FC-BaseT PHY logical model 21
Figure 5	PCS and PMA functional block diagram 24
Figure 6	PCS and PMA service interfaces 26
Figure 7	FC-BaseT nomenclature 31
Figure 8	PCS reference diagram 32
Figure 9	PCS transmit bit ordering 33
Figure 10	Side-stream scrambler for the Master PHY 37
Figure 11	Side-stream scrambler for the Slave PHY 38
Figure 12	4D PAM-8 Schläfli Lattice encoder. 39
Figure 13	4D PAM-8 Trellis encoder 40
Figure 14	Generation of PMA training PAM-2 sequences 41
Figure 15	PCS receive bit ordering 42
Figure 16	4D PAM-8 Schläfli Lattice decoder. 43
Figure 17	Side-stream descrambler for the Master PHY 44
Figure 18	Side-stream descrambler for the Slave PHY 44
Figure 19	PCS synchronization state diagram 47
Figure 20	PMA Reference Diagram 48
Figure 21	Example of Link Establishment 52
Figure 22	PHY Control State Diagram 55
Figure 23	Link Monitor State Diagram 56
Figure 24	Transmitter Test Fixture A 58
Figure 25	Transmitter Test Fixture B 58
Figure 26	Test Mode 1 Output (not to scale) 59
Figure 27	4GFC-BaseT Transmitter PSD at 5dBm 61
Figure 28	2GFC-BaseT Transmitter PSD at 5dBm 62
Figure 29	1GFC-BaseT Transmitter PSD at 5dBm 63
Figure 30	MDI Jack Connector. 66
Figure 31	Balanced Cabling Plug Connector 66
Figure 32	FC-BaseT Link Topology 68
Figure 33	FC-BaseT elasticity FIFO 69
Figure 34	Example of E-FIFO implementation 70
Figure 35	FC-BaseT clock skew compensation 71
Figure 36	Base page encoding. 75
Figure 37	Unformatted page encoding 76
Figure 38	Clock detect state diagram. 84
Figure 39	Clock loopback state diagram 85
Figure B.1	4D symbol error rate in function of SNR. 100

Table	Page
Table 1 ISO and American Conventions	14
Table 2 State Diagram Operators.	18
Table 3 FC-BaseT design goals for data rates and cable reaches for the cabling channels specified by ISO/IEC 11801:2002/AMD2:2010	20
Table 4 XGMII frequencies for FC-BaseT	22
Table 5 XGMII characters.	22
Table 6 FC-BaseT symbol and data rates	23
Table 7 Fibre Channel words	34
Table 8 XGMII representation of Fibre Channel words	35
Table 9 36/33 transcoding	35
Table 10 36/33 transcoding examples	35
Table 11 Bit to symbols mapping	40
Table 12 PMA training sequences	41
Table 13 PCS synchronization state variables.	46
Table 14 Idle2 and Idle3 Ordered Sets.	51
Table 15 PHY Control and Link Monitor State Variables	53
Table 16 PHY Control and Link Monitor Timers.	54
Table 17 Management Register Settings for Test Modes	57
Table 18 Disturbing Signal Frequency	59
Table 19 Droop Requirements	59
Table 20 Transmitter Power Schedule	64
Table 21 Transmitter Frequency Requirements.	64
Table 22 Receiver Frequency Requirements.	65
Table 23 Alien Noise Requirements	65
Table 24 Assignment of PMA Signals to MDI and MDI-X Contacts	67
Table 25 FC-BaseT MDI Electrical Requirements	67
Table 26 FC-BaseT Cable Plant Requirements	68
Table 27 Extended Fill Words	70
Table 28 Technology Ability Field.	75
Table 29 Unformatted code field of the first unformatted page	76
Table 30 Master-Slave seed assignment.	77
Table 31 Unformatted code field of the second unformatted page.	78
Table 32 Unformatted code field of the third unformatted page	79
Table 33 Cable length estimation state variables.	81
Table 34 Cable length estimation timers	83
Table 35 FC-BaseT registers	86
Table 36 FC-BaseT control register bits.	88
Table 37 FC-BaseT status register bits	91
Table 38 FC-BaseT speed downshift register bits	93
Table 39 FC-BaseT test register bits	94
Table 40 FC-BaseT SNR margin registers bits	94
Table 41 FC-BaseT auto-negotiation pages registers bits	96
Table A.1 Additional 36/33 mappings.	98
Table A.2 6-bit encoding of XGMII control characters	99
Table C.1 Recommended transmitter power schedule	102

INFORMATION TECHNOLOGY – FIBRE CHANNEL –

Part 151: Fibre Channel BaseT (FC-BaseT)

FOREWORD

- 1) ISO (International Organization for Standardization) and IEC (International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards. Their preparation is entrusted to technical committees; any ISO and IEC member body interested in the subject dealt with may participate in this preparatory work. International governmental and non-governmental organizations liaising with ISO and IEC also participate in this preparation.
- 2) In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.
- 3) The formal decisions or agreements of IEC and ISO on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC and ISO member bodies.
- 4) IEC, ISO and ISO/IEC publications have the form of recommendations for international use and are accepted by IEC and ISO member bodies in that sense. While all reasonable efforts are made to ensure that the technical content of IEC, ISO and ISO/IEC publications is accurate, IEC or ISO cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 5) In order to promote international uniformity, IEC and ISO member bodies undertake to apply IEC, ISO and ISO/IEC publications transparently to the maximum extent possible in their national and regional publications. Any divergence between any ISO/IEC publication and the corresponding national or regional publication should be clearly indicated in the latter.
- 6) ISO and IEC provide no marking procedure to indicate their approval and cannot be rendered responsible for any equipment declared to be in conformity with an ISO/IEC publication.
- 7) All users should ensure that they have the latest edition of this publication.
- 8) No liability shall attach to IEC or ISO or its directors, employees, servants or agents including individual experts and members of their technical committees and IEC or ISO member bodies for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication of, use of, or reliance upon, this ISO/IEC publication or any other IEC, ISO or ISO/IEC publications.
- 9) Attention is drawn to the normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.
- 10) Attention is drawn to the possibility that some of the elements of this International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

International Standard ISO/IEC 14165-151 was prepared by subcommittee 25: Interconnection of information technology equipment, of ISO/IEC joint technical committee 1: Information technology.

The list of all currently available parts of the ISO/IEC 14165 series, under the general title *Information technology – Fibre Channel*, can be found on the IEC web site.

This International Standard has been approved by vote of the member bodies and the voting results may

be obtained from the address given on the second title page.

This publication has been drafted in accordance with the ISO/IEC Directives, Part 2.

IMPORTANT - The 'colour inside' logo on the cover page of this publication indicates that it contains colours which are considered to be useful for the correct understanding of its contents. Users should therefore print this document using a colour printer.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

INTRODUCTION

This part of ISO/IEC 14165 describes extensions to the Fibre Channel signaling and physical layer requirements defined in ANSI INCITS 404-2005, Fibre Channel - Physical Interfaces 2, to transport Fibre Channel over the commonly available 4-pair balanced copper cabling specified in ISO/IEC 11801:2002 and TIA/EIA-568-B.2-2001. This document is one of the Fibre Channel family of standards.

[IECNORM.COM](https://www.iecnorm.com) : Click to view the full PDF of ISO/IEC 14165-151:2017

INFORMATION TECHNOLOGY — FIBRE CHANNEL —

Part 151: Fibre Channel BaseT (FC-BaseT)

1 Scope

This part of ISO/IEC 14165 describes extensions to the Fibre Channel signaling and physical layer requirements defined in ISO/IEC 14165-142, Fibre Channel - Physical Interfaces 2, to transport Fibre Channel over the commonly available 4-pair balanced copper cabling specified in ISO/IEC 11801:2002 and TIA/EIA-568-B.2-2001. This standard is one of the Fibre Channel family of standards.

2 Normative references

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

IEC 60603-7-2, *Connectors for electronic equipment – Part 7-2: Detail specification for 8-way, unshielded, free and fixed connectors, for data transmissions with frequencies up to 100 MHz*

IEC 60603-7-3, *Connectors for electronic equipment – Part 7-3: Detail specification for 8-way, shielded, free and fixed connectors, for data transmission with frequencies up to 100 MHz*

IEC 60603-7-4, *Connectors for electronic equipment – Part 7-4: Detail specification for 8-way, unshielded, free and fixed connectors, for data transmissions with frequencies up to 250 MHz*

IEC 60603-7-5, *Connectors for electronic equipment – Part 7-5: Detail specification for 8-way, shielded, free and fixed connectors, for data transmissions with frequencies up to 250 MHz*

ISO/IEC 14165-122, *Information technology - Fibre Channel - Part 122: Arbitrated Loop - 2 (FC-AL-2)* [ANSI INCITS 332-1999 including ANSI INCITS 332-1999/AM1-2003 and ANSI INCITS 332-1999/AM2-2006]

ISO/IEC 11801:2002, *Information technology - Generic cabling for customer premises*

ISO/IEC 11801:2002/AMD1:2008

ISO/IEC 11801:2002/AMD2:2010

ISO/IEC TR 24750, *Information technology - Assessment and mitigation of installed balanced cabling channels in order to support 10GBASE-T*

ANSI INCITS 404-2006, *Information technology - Fibre Channel - Part 142: Physical Interfaces - 2 (FC-PI-2)*

ANSI INCITS 424-2007, *Information technology - Fibre Channel - Framing and Signaling - 2 (FC-FS-2)*

IEEE Std 802.3-2005, *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*

IEEE Std 802.3an-2006, *Physical Layer and Management Parameters for 10 Gb/s Operation, Type 10GBASE-T*

3 Terms, definitions, abbreviations, symbols, and conventions

3.1 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

ISO and IEC maintain terminological databases for use in standardization at the following addresses:

- IEC Electropedia: available at <http://www.electropedia.org/>
- ISO Online browsing platform: available at <http://www.iso.org/obp>

3.1.1

10 Gigabit Media Independent Interface

XGMII

interface defined in IEEE 802.3-2005 for 10 Gbit/s operations but used in this standard for lower speeds operations (see 4.4)

3.1.2

Auto-Negotiation

algorithm that allows two devices at either end of a link segment to negotiate common data service functions

3.1.3

balanced cable

cable consisting of one or more metallic symmetrical cable elements (e.g., twisted pairs)

3.1.4

bit error rate

BER

ratio of the number of bits received in error to the total number of bits received

3.1.5

dimension

component of the PHY that generates locally and recovers from the link partner a logical sequence of symbols

Note 1 to entry: A transmit or receive dimension is a logical reference for a wire pair, independent from the specific wire pair where symbols are sent or received.

Note 2 to entry: Dimensions A, B, C, and D are respectively associated with sequences A_n , B_n , C_n , and D_n (see 4.3).

3.1.6

jitter

variations of signal transitions from their ideal positions in time. Jitter may be characterized by its spectral properties and its distribution in time

3.1.7

link

transmission path between any two interfaces of generic cabling

3.1.8**link segment**

point-to-point full-duplex medium connection between two and only two Medium Dependent Interfaces (MDIs)

3.1.9**Master PHY**

PHY that uses an external clock for generating its clock signals to determine the timing of transmitter and receiver operations

Note 1 to entry: It also uses the Master transmit scrambler generator polynomial for side-stream scrambling.

3.1.10**Medium Dependent Interface****MDI**

mechanical and electrical interface between the transmission medium and the Physical Layer device

3.1.11**Ordered Set**

word composed of a special character in its first (left-most) position and data characters in its remaining positions (see FC-FS-2)

3.1.12**Physical Coding Sublayer****PCS**

portion of the Physical Layer that couples the XGMII and the Physical Medium Attachment (PMA)

Note 1 to entry: The PCS contains the functions to encode data bits for transmission via the PMA and to decode the received conditioned signal from the PMA.

3.1.13**Physical Layer device****PHY**

portion of the Physical Layer between the Medium Dependent Interface (MDI) and the XGMII, consisting of the Physical Coding Sublayer (PCS) and the Physical Medium Attachment (PMA)

Note 1 to entry: The PHY contains the functions that transmit, receive, and manage the encoded signals that are impressed on and recovered from the physical medium.

3.1.14**Physical Medium Attachment sublayer****PMA sublayer**

portion of the Physical Layer that contains the functions for transmission, reception, clock recovery and skew alignment

3.1.15**Scrambler**

randomizing mechanism that is used to eliminate long strings of consecutive identical transmitted symbols and avoid the presence of spectral lines in the signal spectrum without changing the signaling rate

3.1.16**Side-stream scrambler**

scrambler in which the state of the scrambler is dependent only on the prior state of the scrambler and not on the transmitted data

3.1.17**Slave PHY**

PHY that recovers its clock from the received signal and uses it to determine the timing of transmitter operations

Note 1 to entry: It also uses the Slave transmit scrambler generator polynomial for side-stream scrambling.

3.1.18**Symbol**

smallest unit of data transmission on the medium

3.1.19**Symbol period**

time interval for transmission of one symbol

3.1.20**Symbol rate**

total number of symbols per second transferred to or from the Medium Dependent Interface (MDI) on a single wire pair

3.1.21**twisted-pair cable**

bundle of multiple twisted pairs within a single protective sheath

3.1.22**word**

<when used to indicate a size> 32 contiguous bits

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

3.2 Editorial conventions

In this standard, a number of conditions, mechanisms, sequences, parameters, events, states, or similar terms are printed with the first letter of each word in uppercase and the rest lowercase (e.g., Exchange, Class). Any lowercase uses of these words have the normal technical English meanings.

Lists sequenced by letters (e.g., a-red, b-blue, c-green) show no ordering relationship between the listed items. Numbered lists (e.g., 1-red, 2-blue, 3-green) show an ordering relationship between the listed items.

The ISO convention of numbering is used (i.e., the thousands and higher multiples are separated by a space and a comma is used as the decimal point.) A comparison of the American and ISO conventions are shown in table 1.

Table 1 – ISO and American Conventions

ISO	American
0,6	0.6
1 000	1,000
1 323 462,9	1,323,462.9

In case of any conflict between figure, table, and text, the text, then tables, and finally figures take precedence.

In all of the figures, tables, and text of this document, the most significant bit of a binary quantity is shown on the left side.

When the value of a bit or field is not relevant, x or xx appears in place of a specific value.

Unless stated otherwise, numbers that are not immediately followed by lower-case b or h are decimal values, numbers immediately followed by lower-case b (xxb) are binary values, and numbers or upper case letters immediately followed by lower-case h (xxh) are hexadecimal values.

A numeric range is indicated by listing the two extremes separated by “..” (e.g., “1 .. 6” indicates the range from 1 to 6, including 1 and 6).

3.3 Abbreviations, acronyms, and symbols

Abbreviations, acronyms, and symbols applicable to this standard are listed. Definitions of several of these items are included in 3.1.

⊕	The XOR operator
	concatenation symbol (e.g., A B represents the concatenation of A and B)
4D	4-dimensional
BER	bit error rate
dB	Decibel
E-FIFO	Elasticity FIFO
EDC	Error Detecting Code
EIA	Electronic Industries Association
EMC	Electromagnetic Compatibility
EMI	Electromagnetic Interference
FEXT	Far-end crosstalk
FLP	Fast Link Pulse
I2C	Inter-Integrated Circuit bus

IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
ISO	International Organization for Standardization
LAN	Local Area Network
MDI	Medium Dependent Interface
MDIO	Management Data Input/Output Interface
MII	Media Independent Interface
NEXT	Near-end crosstalk
NLP	Normal Link Pulse
OSI	Open Systems Interconnection
PAM	Pulse Amplitude Modulation
PAM-2	2-level Pulse Amplitude Modulation
PAM-8	8-level Pulse Amplitude Modulation
PCS	Physical Coding Sublayer
PHY	Physical Layer device
PMA	Physical Medium Attachment
ppm	Parts per million
PSD	Power Spectral Density
RF	Remote Fault
RX	Receiver
SNR	Signal to Noise Ratio
TIA	Telecommunication Industries Association
TX	Transmitter
XGMII	10 Gigabit Media Independent Interface

3.4 Keywords

3.4.1 expected: A keyword used to describe the behavior of the hardware or software in the design models assumed by this standard. Other hardware and software design models may also be implemented.

3.4.2 ignored: A keyword used to describe an unused bit, byte, word, field or code value. The contents or value of an ignored bit, byte, word, field or code value shall not be examined by the receiving device and may be set to any value by the transmitting device.

3.4.3 invalid: A keyword used to describe an illegal or unsupported bit, byte, word, field or code value. Receipt of an invalid bit, byte, word, field or code value shall be reported as an error.

3.4.4 mandatory: A keyword indicating an item that is required to be implemented as defined in this standard.

3.4.5 may: A keyword that indicates flexibility of choice with no implied preference (equivalent to “may or may not”).

3.4.6 may not: A keyword that indicates flexibility of choice with no implied preference (equivalent to “may or may not”).

3.4.7 obsolete: A keyword indicating that an item was defined in prior Fibre Channel standards but has been removed from this standard.

3.4.8 optional: A keyword that describes features that are not required to be implemented by this standard. However, if any optional feature defined by this standard is implemented, then it shall be implemented as defined in this standard.

3.4.9 reserved: A keyword referring to bits, bytes, words, fields and code values that are set aside for future standardization. A reserved bit, byte, word or field shall be set to zero, or in accordance with a future extension to this standard. Recipients are not required to check reserved bits, bytes, words or fields for zero values.

3.4.10 restricted: A keyword referring to bits, bytes, words, and fields that are set aside for use in other Fibre Channel standards. A restricted bit, byte, word, or field shall be treated as a reserved bit, byte, word or field for the purposes of the requirements defined in this standard.

3.4.11 shall: A keyword indicating a mandatory requirement. Designers are required to implement all such mandatory requirements to ensure interoperability with other products that conform to this standard.

3.4.12 should: A keyword indicating flexibility of choice with a strongly preferred alternative; equivalent to the phrase "it is strongly recommended".

3.4.13 x or xx: The value of the bit or field is not relevant.

3.5 State Diagram notation

3.5.1 State Diagram conventions

The operation of a protocol may be described by subdividing the protocol into a number of interrelated functions. The operation of the functions may be described by state diagrams. Each diagram represents the domain of a function and consists of a group of connected, mutually exclusive states. Only one state of a function is active at any given time (see figure 1).

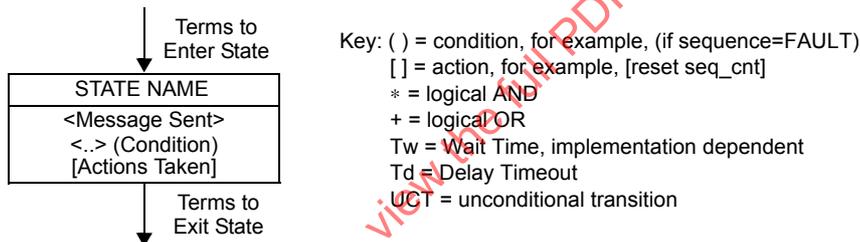


Figure 1 – State diagram notation example

Each state that the function is able to assume is represented by a rectangle. These are divided into two parts by a horizontal line. In the upper part the state is identified by a name in capital letters. The lower part contains the name of any ON signal that is generated by the function. Actions are described by short phrases and enclosed in brackets.

All permissible transitions between the states of a function are represented graphically by arrows between them. A transition that is global in nature (e.g., an exit condition from all states to the IDLE or RESET state) is indicated by an open arrow. Labels on transitions are qualifiers that shall be fulfilled before the transition is taken. The label UCT designates an unconditional transition. Qualifiers described by short phrases are enclosed in parentheses.

State transitions and sending and receiving of messages occur instantaneously. The actions inside a state block execute instantaneously. Actions inside state blocks are atomic (i.e., uninterruptible). When a state is entered and after performing all the actions listed in a state block one time, the state block then continuously evaluates its exit conditions until one is satisfied, at which point control passes through a transition arrow to

the next block. While the state awaits fulfillment of one of its exit conditions, the actions inside do not implicitly repeat.

Valid state actions may include .indication and .request messages. No actions are taken outside of any state block.

The following convention is used to describe a term-assignment statement that is associated with a transition:

- a) The character “:” (colon) is a delimiter used to denote that a term assignment statement follows; and
- b) The character “←” (left arrow) denotes assignment of the value following the arrow to the term preceding the arrow.

The state diagrams contain the authoritative statement of the functions they depict. When apparent conflicts between descriptive text and state diagrams arise, the state diagrams are to take precedence. However, such precedence does not override any explicit description in the text that has no parallel in the state diagrams.

The models presented by state diagrams are intended as the primary specifications of the functions to be provided. However, it is important to distinguish between a model and a real implementation. The models are optimized for simplicity and clarity of presentation, while any realistic implementation may place heavier emphasis on efficiency and suitability to a particular implementation technology. It is the functional behavior of any unit that shall match the standard, not its internal structure. The internal details of the model are useful only to the extent that they specify the external behavior clearly and precisely.

3.5.2 State Diagram variables

State variables are of two types, variables with no default value and variables with a default value. Variables with no default value, once set retain their value as long as succeeding blocks contain no references to them. Variables with default value evaluate to the variable default value in each state where the variable value is not explicitly set.

3.5.3 State Diagram timers

All timers operate in the same fashion. A timer is reset and starts counting upon entering a state where "start x_timer" is asserted. Time "x" after the timer has been started, "x_timer_done" is asserted and remains asserted until the timer is reset. At all other times, "x_timer_not_done" is asserted. When explicitly entering a state where "start x_timer" is asserted, the timer is reset and restarted even if the entered state is the same as the exited state. In addition, a timer is reset and stops counting upon entering a state where "stop timer" is asserted.

3.5.4 State transitions

The following terms are valid transition qualifiers:

- a) Boolean expressions;
- b) An event such as the expiration of a timer: timer_done;
- c) An event such as the reception of a message: PMA_UNITDATA.indication;
- d) An unconditional transition: UCT;

e) A branch taken when other exit conditions are not satisfied: ELSE.

Any open arrow (an arrow with no source block) represents a global transition. Global transitions are evaluated continuously whenever any state is evaluating its exit conditions. When a global transition becomes true, it supersedes all other transitions, including UCT, returning control to the block pointed to by the open arrow.

3.5.5 Operators

The state diagram operators are shown in table 2.

Table 2 – State Diagram Operators

Character	Meaning
*	Boolean AND
+	Boolean OR
^	Boolean XOR
!	Boolean NOT
<	Less than
≤	Less than or equal to
=	Equals (a test of equality)
≠	Not equals
≥	Greater than or equal to
>	Greater than
()	Indicates precedence
←	Assignment operator
∈	Indicates membership
∉	Indicates nonmembership
	Catenate
ELSE	No other state condition is satisfied

4 Structure and concepts

4.1 Overview

A vast infrastructure of installed twisted-pair copper cabling systems exists today. High volume has driven down the cost for copper PHYs and cabling systems to a fraction of their original cost. FC-BaseT is intended to leverage this economy of scale by permitting communication of Fibre Channel data streams at 1 Gbit/s, 2 Gbit/s, or 4 Gbit/s over balanced twisted-pair copper cabling systems.

1000BASE-T (see IEEE 802.3-2005) pioneered the use of coded multi-level baseband signaling to permit high data rates to be carried over the relatively narrow bandwidth available on balanced twisted-pair copper cabling. FC-BaseT uses a simplified version of the field proven and well understood line code of 1000BASE-T.

Twisted-pair cabling systems have limited usable bandwidth available for data transmission. FC data are de-multiplexed spatially across four physical wire pairs to reduce the data rate per wire by 75%. In addition, hybrids are employed to permit signals to be transmitted and received on the same wire pair enabling full-duplex point-to-point transmission. Figure 2 illustrates the full duplex, four wire configuration used for FC-BaseT.

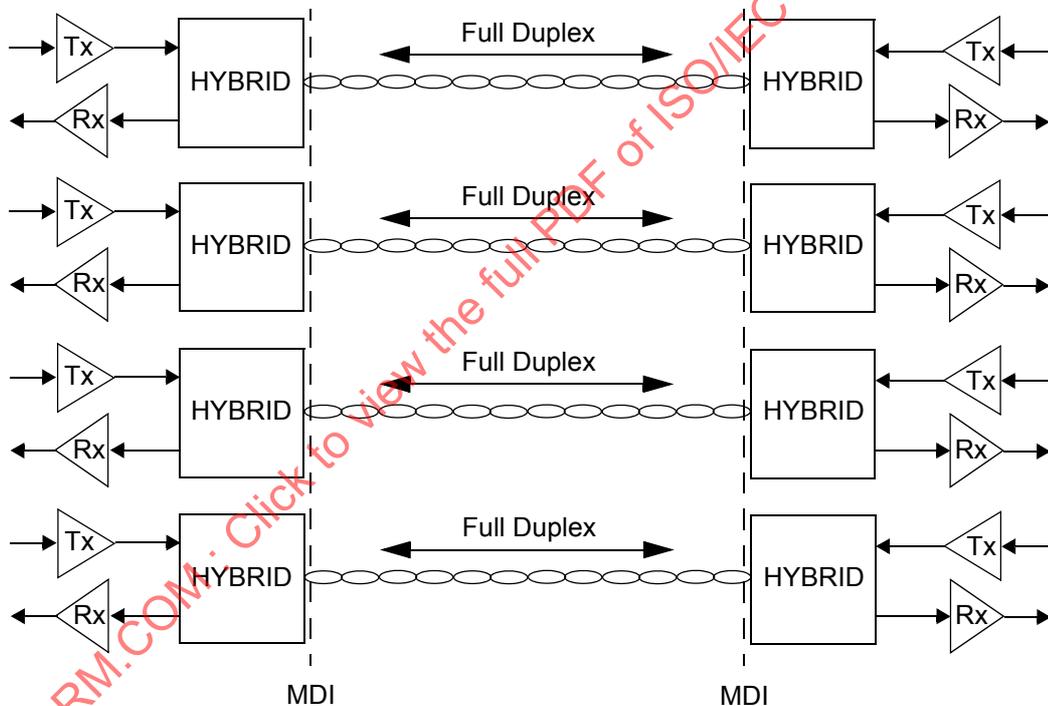


Figure 2 – FC-BaseT topology

Multi-level baseband signaling (i.e., Pulse Amplitude Modulation or PAM) is used on each of the wire pairs. In a multi-level transmission scheme, multiple bits are transmitted each clock interval. Data rate (i.e., bits per second) is transformed into symbol rate (i.e., symbols per second).

The following are the objectives of FC-BaseT:

- a) bit error rate (BER) better than 10^{-12} (e.g., $\leq 10^{-15}$);

- b) PHY delay lower than 1 μ s;
- c) Meet or exceed FCC Class A/CISPR EMC requirements;
- d) Support IEEE 802.3-2005 Auto-Negotiation;
- e) Support data rates, cable types and lengths as per table 3.

Table 3 – FC-BaseT design goals for data rates and cable reaches for the cabling channels specified by ISO/IEC 11801:2002/AMD2:2010

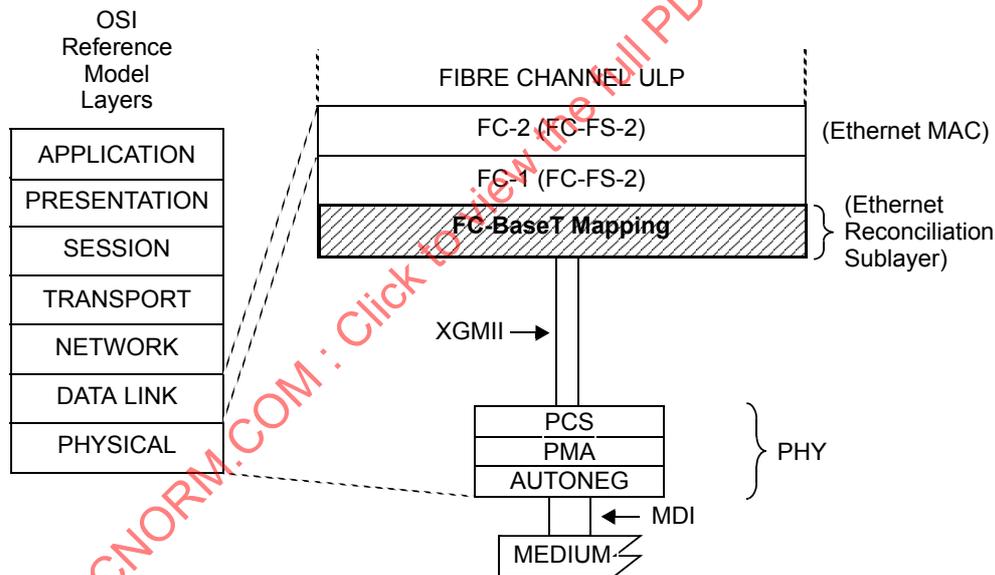
FC Data Rate	Class D Reach ^a	Class E Reach ^a	Class E _A to F _A Reach ^a
1GFC-BaseT	100 m	100 m	100 m
2GFC-BaseT	60 m	70 m	100 m
4GFC-BaseT	30 m ^b	40 m	100 m

^a Similar channels up to Class E_A are specified in TIA/EIA standards as Category 5e, Category 6 and Category 6a (see Annex D).

^b Expected performance, not guaranteed. 4GFC-BaseT exceeds the rated frequency performance of Class D channels (see ISO/IEC TR 24750).

4.2 Relationship with other standards

Figure 3 shows the model of the FC-BaseT FC-0 level, and its relationship with the ISO Reference model and the IEEE 802.3 LAN Model.



MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PHY = PHYSICAL LAYER DEVICE
 XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE (optional)

PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT
 AUTONEG = AUTONEGOTIATION

Figure 3 – FC-BaseT relationship to the OSI Reference Model, the Fibre Channel Levels, and the IEEE 802.3 LAN Model

An optional XGMII interface is used in this specification because it allows an easy representation of Fibre Channel data characters and control codes. Implementations are not required to expose nor use this interface. The XGMII interface is used in this standard as specified in 4.4.

4.3 FC-BaseT PHY logical model

The logical model of a possible FC-BaseT PHY implementation is shown in figure 4.

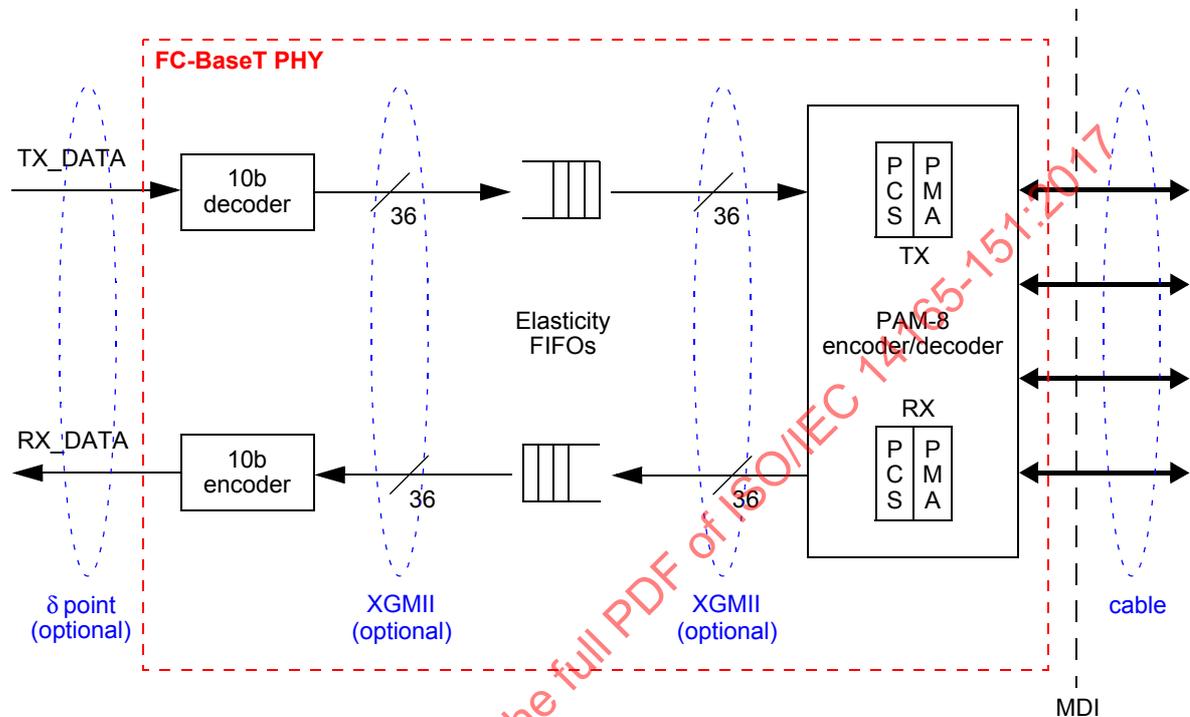


Figure 4 – An FC-BaseT PHY logical model

Fibre Channel is a full duplex continuously transmitting technology (i.e., when no data are to be sent, Idles and other Ordered Sets are transmitted by the upper layers). An FC-BaseT PHY transmits on the cable a continuous stream of data coming from the upper layers and receives a continuous stream of data from the cable. A Fibre Channel data stream is composed of frames and Ordered Sets. In some cases (e.g., when the particular Fibre Channel logical topology is an Arbitrated Loop, see FC-AL-2) Ordered Sets are used continuously in an information stream.

An FC-BaseT PHY may use the same electrical interface specified for the δ points (see FC-PI-2) for the various Fibre Channel speeds. In this case, on entering the FC-BaseT PHY the 10b encoding of the FC serial stream is removed and the FC data are converted in an internal stream of 36-bit XGMII words, as shown in figure 4.

The Elasticity FIFOs resolve potential clock skews between two connected PHYs and retime the data stream to be compliant with the FC-PI-2 jitter specification.

The PAM-8 encoder/decoder converts each 36-bit word on the XGMII interface in three 4-dimensional (4D) PAM-8 symbols. Each 4D symbol is a 4-tuple (A_n, B_n, C_n, D_n) of 1-dimensional signals, each taken from the set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$. Each 4D symbol is sent as an 8-level pulse amplitude modulated signal on each of the four wire pairs, called BI_DA, BI_DB, BI_DC, and BI_DD.

This standard defines the FC-BaseT γ point (see FC-PI-2) at the MDI, and does not specify requirements for the host to PHY interface. Implementation of δ points is optional (e.g., a PHY may implement a different host to PHY interface, or the PHY may be integrated in a protocol ASIC). When an FC-BaseT PHY is integrated in a protocol ASIC, it may be coupled directly with the FC-1 level. System operation from the perspective of signals at the MDI and management objects shall be independent from the host to PHY interface.

4.4 FC-BaseT usage of XGMII

The XGMII interface is defined in IEEE 802.3-2005 to support 10 Gbit/s operations, but is used in this standard to support the Fibre Channel speeds of 4 Gbit/s, 2 Gbit/s, and 1 Gbit/s with the clock frequencies shown in table 4.

Table 4 – XGMII frequencies for FC-BaseT

FC Speed	XGMII Frequency
4GFC-BaseT	53,125 MHz
2GFC-BaseT	26,562 5 MHz
1GFC-BaseT	13,281 25 MHz

XGMII is a 36-bit interface carrying four characters at a time. Each character is represented as eight data bits (called TXD in the transmit direction, RXD in the receive direction) and one control bit (called TXC in the transmit direction, RXC in the receive direction). The XGMII characters are summarized in table 5.

Table 5 – XGMII characters

XGMII TXC/RXC	XGMII TXD/RXD	10b Character	Description
0b	00h .. FFh	Dxx.y	Data characters
1b	07h	^a	XGMII Idle character (/I/)
1b	1Ch	K28.0	Reserved XGMII character
1b	3Ch	K28.1	Reserved XGMII character
1b	5Ch	K28.2	Primitive XGMII character (/P/)
1b	7Ch	K28.3	Reserved XGMII character
1b	9Ch	K28.4	Sequence XGMII character (/Q/)
1b	BCh	K28.5	Reserved XGMII character
1b	DCh	K28.6	Reserved XGMII character
1b	F7h	K23.7	Reserved XGMII character
1b	FBh	K27.7	Start XGMII character (/S/)
1b	FCh	K28.7	Reserved XGMII character
1b	FDh	K29.7	Terminate XGMII character (/T/)
1b	FEh	K30.7	Error XGMII character (/E/)
1b	all others	K30.7	Invalid XGMII character (/E/)

^a The XGMII Idle character, denoted as K07, does not have a 10b representation. IEEE 802.3-2005, clause 36.2.4.12, defines how to encode the K07 character for transmission in a 10b encoded serial stream.

4.5 Operation of FC-BaseT

4.5.1 Overview

FC-BaseT employs full duplex baseband transmission over four pairs of balanced cabling. The aggregate data rate is achieved by transmitting in each direction simultaneously on each wire pair, as shown in figure 2. Baseband 8-level PAM signaling is used on each of the wire pairs. Transmitted levels on each wire pair are selected from a 4-dimensional 8-level signal constellation. Each 4-dimensional symbol is a 4-tuple (A_n , B_n , C_n , D_n) of 1-dimensional signals, each taken from the set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$. This modulation scheme is called 4D PAM-8. The FC-BaseT data rate, symbol rate and symbol period per each of the supported FC-BaseT speeds is shown in table 6.

Table 6 – FC-BaseT symbol and data rates

FC Speed	FC0 Data Rate	FC-BaseT Data Rate	FC-BaseT Symbol Rate per Wire	Symbol Period
4GFC-BaseT	4,25 Gbit/s	3,4 Gbit/s	318,75 Ms/s	~3,137 ns ^a
2GFC-BaseT	2,125 Gbit/s	1,7 Gbit/s	159,375 Ms/s	~6,274 ns ^b
1GFC-BaseT	1,062 5 Gbit/s	0,85 Gbit/s	79,687 5 Ms/s	~12,549 ns ^c
<p>^a The exact symbol period is 160/51 ns. ^b The exact symbol period is 320/51 ns. ^c The exact symbol period is 640/51 ns.</p>				

Data and control characters are encoded at a rate of 2,75 information bits per PAM-8 symbol. Data and control characters are embedded in a framing scheme that runs continuously after startup of the link. PHY blocks consisting of three 4D PAM-8 symbols are continuously transmitted after link startup.

An FC-BaseT PHY may operate either as a Master PHY or as a Slave PHY. The Master-Slave relationship between two stations sharing a link segment is established during Auto-Negotiation (see 8.4). The Master PHY uses a local clock to determine the timing of transmitter operations. The Slave PHY recovers the clock from the received signal and uses it to determine the timing of transmitter operations (i.e., it performs loop timing).

An FC-BaseT PHY may be functionally decomposed in a Physical Coding Sublayer (PCS) and a Physical Medium Attachment (PMA) sublayer, summarized respectively in 4.5.2 and 4.5.3. Figure 5 shows the functional block diagram of the PCS and PMA sublayers.

All FC-BaseT PHY implementations shall be compatible at the MDI (see figure 4). Designers are free to implement circuitry within the PCS and PMA in an application-dependent manner provided that the MDI specification is met. Physical implementation of the XGMII is optional. System operation from the

perspective of signals at the MDI and management objects are identical whether the XGMII is implemented or not.

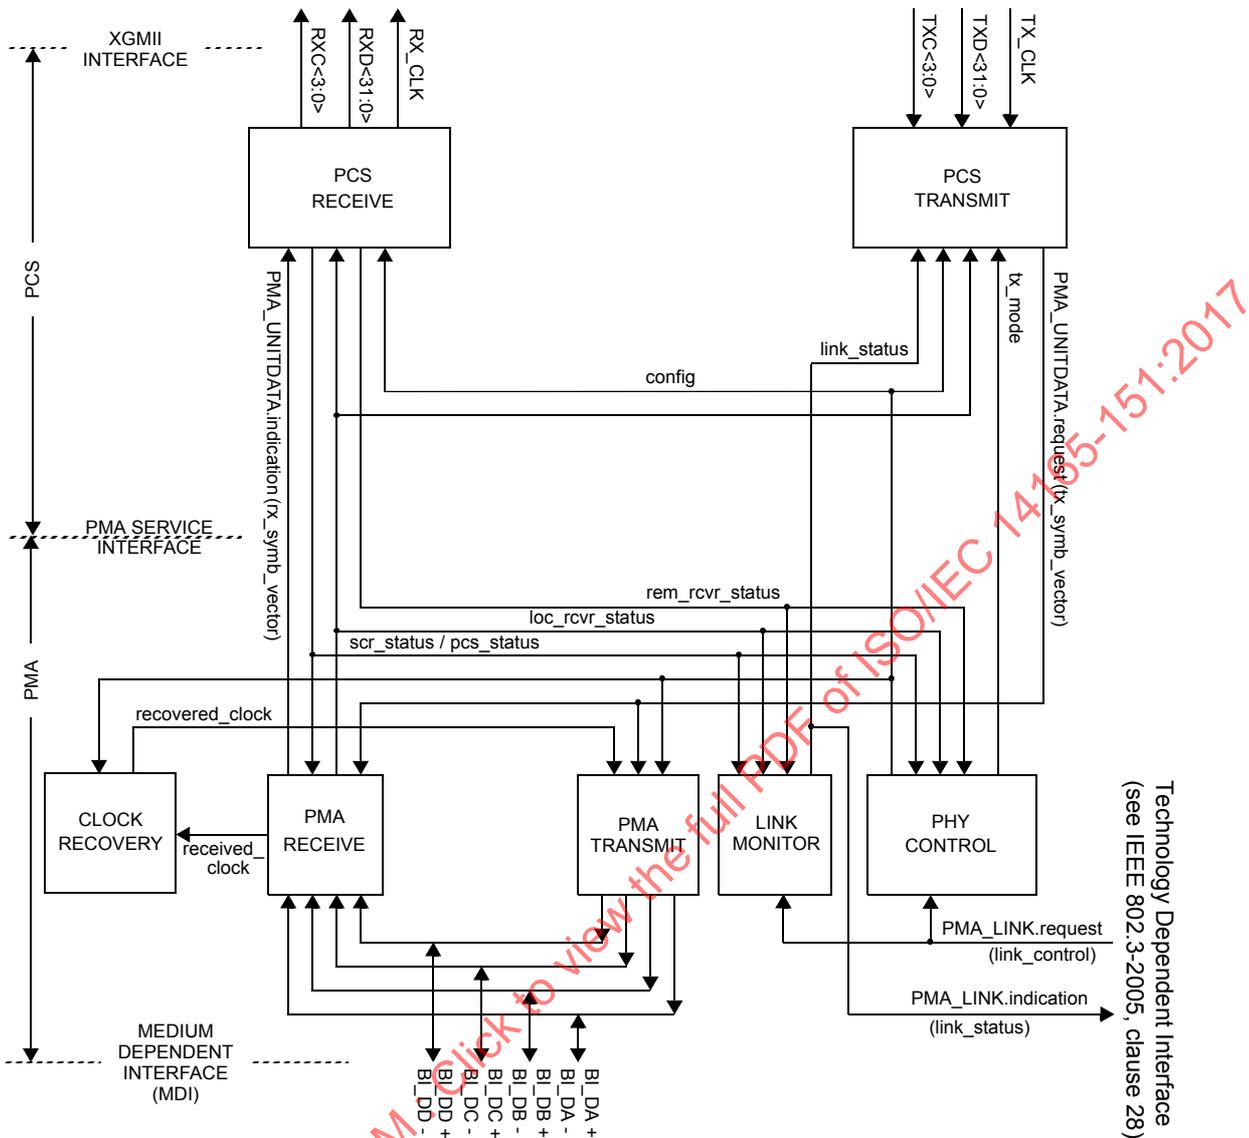


Figure 5 – PCS and PMA functional block diagram

4.5.2 PCS overview

The FC-BaseT Physical Coding Sublayer (PCS) couples the XGMII to the FC-BaseT Physical Medium Attachment (PMA) sublayer. The PCS supports a normal mode of operation and a training mode. Furthermore, the PCS contains a management interface.

In the transmit direction, in normal mode (see 5.3), the PCS transcodes each 36-bit XGMII word in a 33-bit block. Each 33-bit block is divided into three 11-bit transmission characters that are individually scrambled. Each 11-bit scrambled transmission character is then either Schläfli lattice or Trellis encoded in 12 bits (see 5.3.5 and 5.3.6). The defined Schläfli lattice coding defined in 5.3.5 is a proper subset of the Trellis coding defined in 5.3.6. The encoding to be used by the local PHY is dictated by the link partner PHY during

Auto-Negotiation. The 12 bits obtained after encoding are in turn divided in four groups of three bits, each identifying a symbol in the PAM-8 symbol set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$. The obtained 4D PAM-8 symbols are passed on to the PMA as a PMA_UNITDATA.request primitive.

In the receive direction, in normal mode (see 5.4), the PCS processes code-groups received from the remote PHY via the PMA and maps them to the XGMII service interface in the receive path. In this receive processing scheme, symbol clock synchronization is done by the PMA receive function.

4.5.3 PMA overview

The PMA sublayer couples messages from the PMA service interface onto the balanced cabling physical medium and provides the link management and PHY Control functions. The PMA provides full duplex communications over four pairs of balanced cabling.

The PMA Transmit function (see 6.2.2) comprises four independent transmitters to generate PAM-8 signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD.

The PMA Receive function (see 6.2.3) comprises four independent receivers for PAM-8 signals on each of the four pairs BI_DA, BI_DB, BI_DC and BI_DD. The receivers are responsible for acquiring symbol timing and, when operating in normal mode, for cancelling echo, near-end crosstalk, far-end crosstalk, and equalizing the signal. The 4D PAM-8 symbols are provided to the PCS receive function via the PMA_UNITDATA.indication primitive. The PMA also contains functions for Link Monitor.

The PHY Control function (see 6.2.4) generates signals that control the PCS and PMA sublayer operations. PHY Control begins following the completion of Auto-Negotiation and provides the start-up functions required for successful FC-BaseT operation. It determines whether the PHY operates in normal mode, enabling data transmission over the link segment, or whether the PHY sends special PAM-2 code-groups that are used in training mode.

4.6 FC-BaseT service primitives and interfaces

4.6.1 Overview

FC-BaseT transfers data and control information across the following four service interfaces:

- a) XGMII Service Interface (see 4.4);
- b) PMA Service Interface (see 4.6.2);
- c) Medium Dependent Interface (see 6.4); and
- d) Management Function Interface (see 4.6.3).

4.6.2 PMA service interface

4.6.2.1 Overview

FC-BaseT uses the following service primitives to exchange symbol vectors, status indications, and control signals across the PMA service interface:

- a) PMA_TXMODE.indication (tx_mode);
- b) PMA_CONFIG.indication (config);

- c) PMA_UNITDATA.request (tx_symb_vector);
- d) PMA_UNITDATA.indication (rx_symb_vector);
- e) PMA_SCRSTATUS.request (scr_status);
- f) PMA_PCSSTATUS.request (pcs_status);
- g) PMA_RXSTATUS.indication (loc_rcvr_status); and
- h) PMA_REMRXSTATUS.request (rem_rcvr_status).

These service primitives are summarized in figure 6.

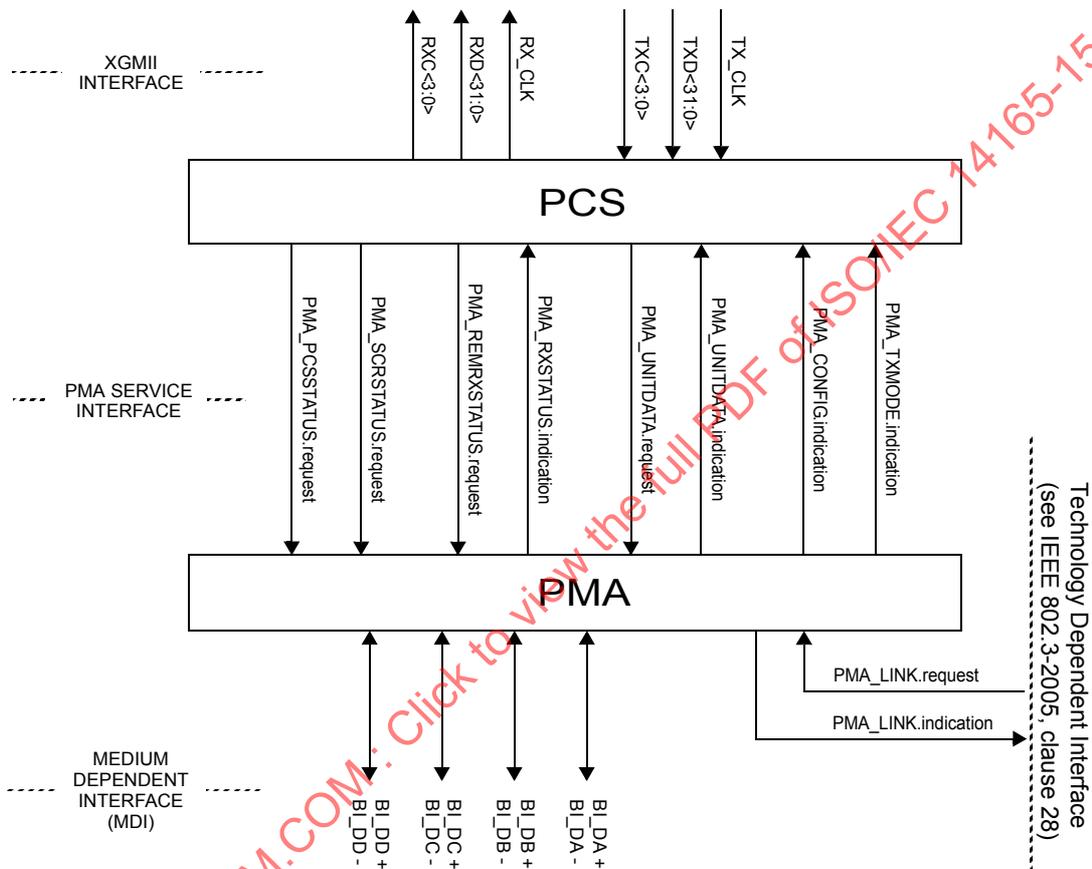


Figure 6 – PCS and PMA service interfaces

4.6.2.2 PMA_TXMODE.indication

Semantics of the primitive: PMA_TXMODE.indication (tx_mode)

This primitive specifies to PCS Transmit via the parameter tx_mode what sequence of code-groups the PCS should be transmitting.

The parameter tx_mode may assume one of the following values:

- a) SEND_Z: This value is continuously asserted when the transmitters are required to be silent;
- b) SEND_T1: This value is continuously asserted when transmission of sequences of 4D Type-1 PAM-2 symbols representing the training mode is to take place;
- c) SEND_T2: This value is continuously asserted when transmission of sequences of 4D Type-2 PAM-2 symbols representing the trained mode is to take place;
- d) SEND_I2: This value is continuously asserted when transmission of sequences of 4D PAM-8 symbols representing the scrambled Idle2 Ordered Set is to take place;
- e) SEND_I3: This value is continuously asserted when transmission of sequences of 4D PAM-8 symbols representing the scrambled Idle3 Ordered Set is to take place; or
- f) SEND_N: This value is continuously asserted when transmission of 4D PAM-8 symbols representing scrambled XGMII words or scrambled Idle3 Ordered Sets in normal mode is to take place.

When generated: The PHY Control function generates PMA_TXMODE.indication messages to indicate a change in tx_mode.

Effect of receipt: Upon receipt of this primitive, the PCS performs its Transmit function (see 5.3).

4.6.2.3 PMA_CONFIG.indication

Semantics of the primitive: PMA_CONFIG.indication (config)

This primitive specifies to the PCS and PMA Transmit functions via the parameter config whether the PHY operates as a Master PHY or as a Slave PHY. The Master-Slave configuration is determined during Auto-Negotiation.

The parameter config may assume one of the following values:

- a) MASTER: This value is continuously asserted when the PHY operates as a Master PHY; or
- b) SLAVE: This value is continuously asserted when the PHY operates as a Slave PHY.

When generated: The PMA generates PMA_CONFIG.indication messages to indicate a change in config.

Effect of receipt: PCS and PMA Clock Recovery perform their functions in Master or Slave configuration according to the value assumed by the parameter config.

4.6.2.4 PMA_UNITDATA.request

Semantics of the primitive: PMA_UNITDATA.request (tx_symb_vector)

This primitive defines the transfer of code-groups in the form of the tx_symb_vector parameter from the PCS to the PMA. The code-groups are obtained in the PCS Transmit function using the encoding rules defined in 5.3 to represent XGMII data and control streams or other sequences. The PMA_UNITDATA.request simultaneously conveys to the PMA via the parameter tx_symb_vector the value of the symbols to be sent over each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD.

The tx_symb_vector parameter takes on the form:

- a) SYMB_4D: A vector of four multi-level symbols, one for each of the four transmit pairs BI_DA, BI_DB, BI_DC, and BI_DD. Each symbol may assume one of the values in the set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$.

The symbols that are elements of tx_symb_vector are called, according to the pair on which each is transmitted, tx_symb_vector[BI_DA], tx_symb_vector[BI_DB], tx_symb_vector[BI_DC], and tx_symb_vector[BI_DD].

When generated: The PCS generates PMA_UNITDATA.request (SYMB_4D) synchronously with every transmit clock cycle.

Effect of receipt: Upon receipt of this primitive the PMA transmits on the MDI the signals corresponding to the indicated symbols. The parameter tx_symb_vector is also used by the PMA Receive function to process the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD for cancelling the echo and near-end crosstalk (NEXT).

4.6.2.5 PMA_UNITDATA.indication

Semantics of the primitive: PMA_UNITDATA.indication (rx_symb_vector)

This primitive defines the transfer of code-groups in the form of the rx_symb_vector parameter from the PMA to the PCS. The PMA_UNITDATA.indication simultaneously conveys to the PCS via the parameter rx_symb_vector the values of the symbols detected on each of the four receive pairs BI_DA, BI_DB, BI_DC, and BI_DD.

The rx_symb_vector parameter takes on the form:

- a) SYMB_4D: A vector of four multi-level symbols that is the receiver's best estimate of the symbols that were sent by the remote transmitter across the four pairs. Each symbol may assume one of the values in the set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$.

The symbols that are elements of rx_symb_vector are called, according to the pair upon which each symbol is received, rx_symb_vector[BI_DA], rx_symb_vector[BI_DB], rx_symb_vector[BI_DC], and rx_symb_vector[BI_DD].

When generated: The PMA generates PMA_UNITDATA.indication (SYMB_4D) messages synchronously with signals received at the MDI. The nominal rate of the PMA_UNITDATA.indication primitive is governed by the recovered clock.

Effect of receipt: The effect of receipt of this primitive is specified in 5.4.1.

4.6.2.6 PMA_SCRSTATUS.request

Semantics of the primitive: PMA_SCRSTATUS.request (scr_status)

This primitive is generated by PCS Receive to communicate the status of the descrambler for the local PHY. The parameter scr_status conveys to the PMA Receive and PHY Control functions the information that the descrambler has achieved synchronization.

The scr_status parameter may assume one of the following values:

- a) OK: The descrambler has achieved synchronization; or
 b) NOT_OK: The descrambler is not synchronized.

When generated: PCS Receive generates PMA_SCRSTATUS.request messages to indicate a change in scr_status.

Effect of receipt: The effect of receipt of this primitive is specified in 6.2.3 and 6.2.4.

4.6.2.7 PMA_PCSSTATUS.request

Semantics of the primitive: PMA_PCSSTATUS.request (pcs_status)

This primitive is generated by PCS Receive to indicate the fully operational state of the PCS for the local PHY. The parameter pcs_status conveys to the PMA Receive, PHY Control and Link Monitor functions the information that the PCS is operating reliably in data mode.

The pcs_status parameter may assume one of the following values:

- a) OK: The PCS is operating reliably in data mode; or
- b) NOT_OK: The PCS is not operating reliably in data mode.

When generated: PCS Receive generates PMA_PCSSTATUS.request messages to indicate a change in pcs_status.

Effect of receipt: The effect of receipt of this primitive is specified in 6.2.4.

4.6.2.8 PMA_RXSTATUS.indication

Semantics of the primitive: PMA_RXSTATUS.indication (loc_rcvr_status)

This primitive is generated by PMA Receive to indicate the status of the receive link at the local PHY. The parameter loc_rcvr_status conveys information on the status of the receive link to the PCS Transmit, PCS Receive, PHY Control, and Link Monitor functions. The criterion for setting the parameter loc_rcvr_status is left to the implementor. It may be based, for example, on observing the mean-square error at the decision point of the receiver and detecting errors during reception of symbol streams.

The loc_rcvr_status parameter may assume one of the following values:

- a) INIT: This value is asserted when operation of the receive link for the local PHY is unreliable;
- b) READY: This value is asserted when the convergence of the local PHY receiver's adaptive filter parameters is completed, but PCS block synchronization is not yet achieved or the link bit error rate has not yet been evaluated; or
- c) OK: This value is asserted when operation of the receive link for the local PHY is reliable.

When generated: PMA Receive generates PMA_RXSTATUS.indication messages to indicate a change in loc_rcvr_status.

Effect of receipt: The effect of receipt of this primitive is specified in 6.2.4.

4.6.2.9 PMA_REMRXSTATUS.request

Semantics of the primitive: PMA_REMRXSTATUS.request (rem_rcvr_status)

This primitive is generated by PCS Receive to indicate the status of the receive link at the remote PHY (see 6.2.4). The parameter rem_rcvr_status conveys to the PHY Control function the information on whether

reliable operation of the remote PHY is detected or not. Beyond the conditions specified in 6.2.4, the criterion for setting the parameter `rem_rcvr_status` is left to the implementor. It may be based, for example, on asserting `rem_rcvr_status` as INIT until `loc_rcvr_status` is OK and then asserting the detected value of `rem_rcvr_status` after proper PCS receive decoding is achieved.

The `rem_rcvr_status` parameter may assume one of the following values:

- a) INIT: This value is asserted when reliable operation of the receive link for the remote PHY is not detected;
- b) READY: This value is asserted when a sequence of code-groups representing a 4D Type-2 PAM-2 sequence is detected at the MDI; or
- c) OK: This value is asserted when reliable operation of the receive link for the remote PHY is detected.

When generated: The PCS generates `PMA_REMRXSTATUS.request` messages to indicate a change in `rem_rcvr_status` on the basis of signals received at the MDI.

Effect of receipt: The effect of receipt of this primitive is specified in 6.2.4.

4.6.3 Management function interface

4.6.3.1 Overview

FC-BaseT uses the following service primitives to exchange status indications and control signals across the Technology Dependent Interface as specified in IEEE 802.3-2005, clause 28:

- a) `PMA_LINK.request` (`link_control`); and
- b) `PMA_LINK.indication` (`link_status`).

4.6.3.2 `PMA_LINK.request`

See IEEE 802.3-2005, clause 28.2.6.2.

4.6.3.3 `PMA_LINK.indication`

See IEEE 802.3-2005, clause 28.2.6.1.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

4.7 FC-BaseT Nomenclature

The nomenclature for the FC-BaseT physical variants is specified in figure 7.

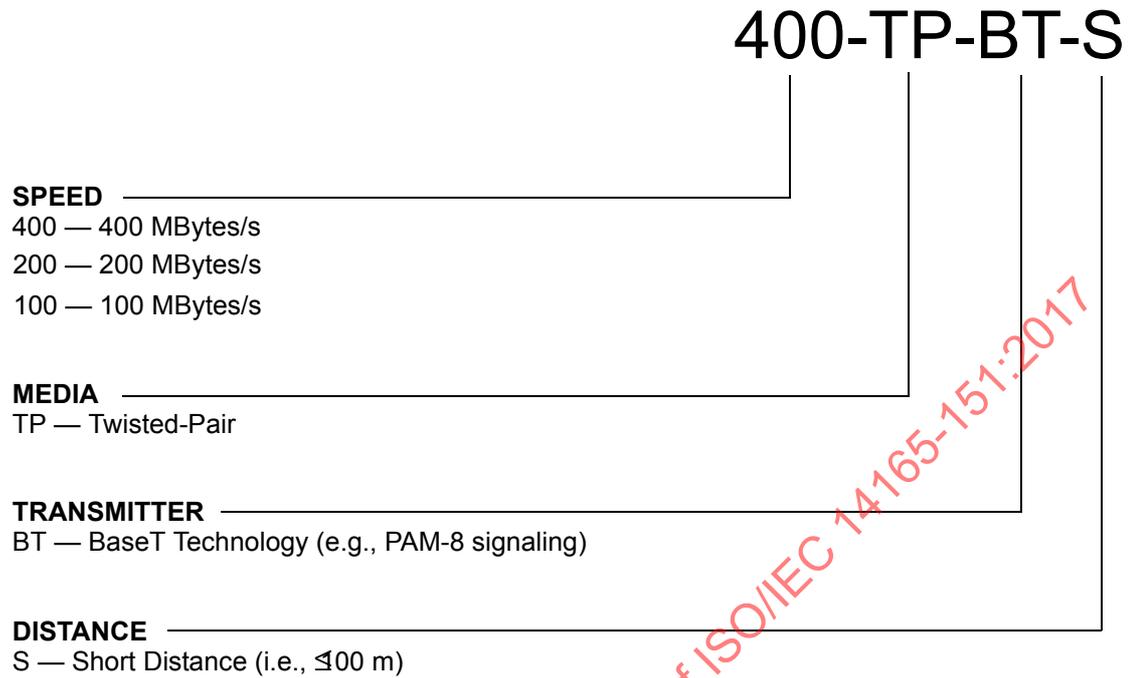


Figure 7 – FC-BaseT nomenclature

5 Physical Coding Sublayer (PCS)

5.1 Overview

The FC-BaseT Physical Coding Sublayer (PCS) couples the XGMII to the FC-BaseT Physical Medium Attachment (PMA) sublayer. The PCS supports a normal mode of operation and a training mode.

The PCS comprises one PCS Reset function and two simultaneous and asynchronous operating functions. The PCS operating functions are PCS Transmit and PCS Receive. All operating functions start immediately after the successful completion of the PCS Reset function.

The PCS reference diagram, shown in figure 8, illustrates how the two operating functions relate to the messages of the PCS-PMA interface. Connections from the management interface to other layers are pervasive and are not shown in figure 8.

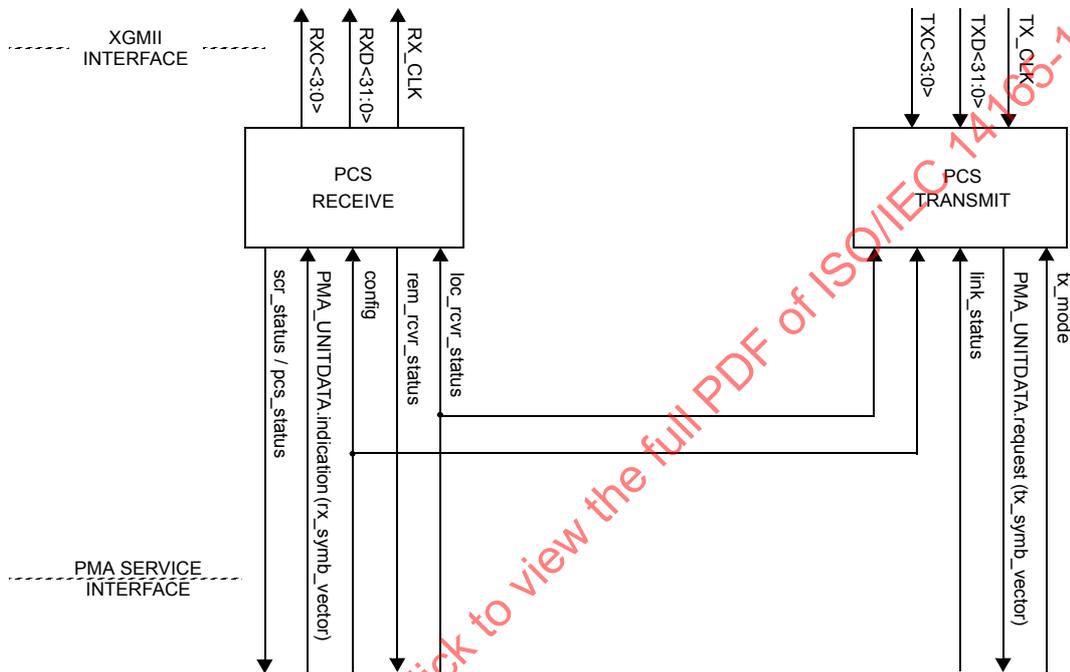


Figure 8 – PCS reference diagram

5.2 PCS reset function

PCS Reset initializes all PCS functions. The PCS Reset function shall be executed whenever one of the following conditions occur:

- Power on; and
- The receipt of a request for reset from the management entity.

PCS Reset sets pcs_reset=ON while any of the above reset conditions hold true. The reference diagrams do not explicitly show the PCS Reset function.

5.3 PCS transmit function

5.3.1 Overview

The PCS Transmit function maps XGMII words, composed by data and control characters, from the XGMII signals TXD<31:0> and TXC<3:0> to the PMA service interface format. The resulting code-groups are passed to the PMA Transmit function via the parameter tx_symb_vector.

The PCS Transmit bit ordering shall be as shown in figure 9, that shows the mapping from XGMII to a 33B block for an XGMII word composed by four data characters.

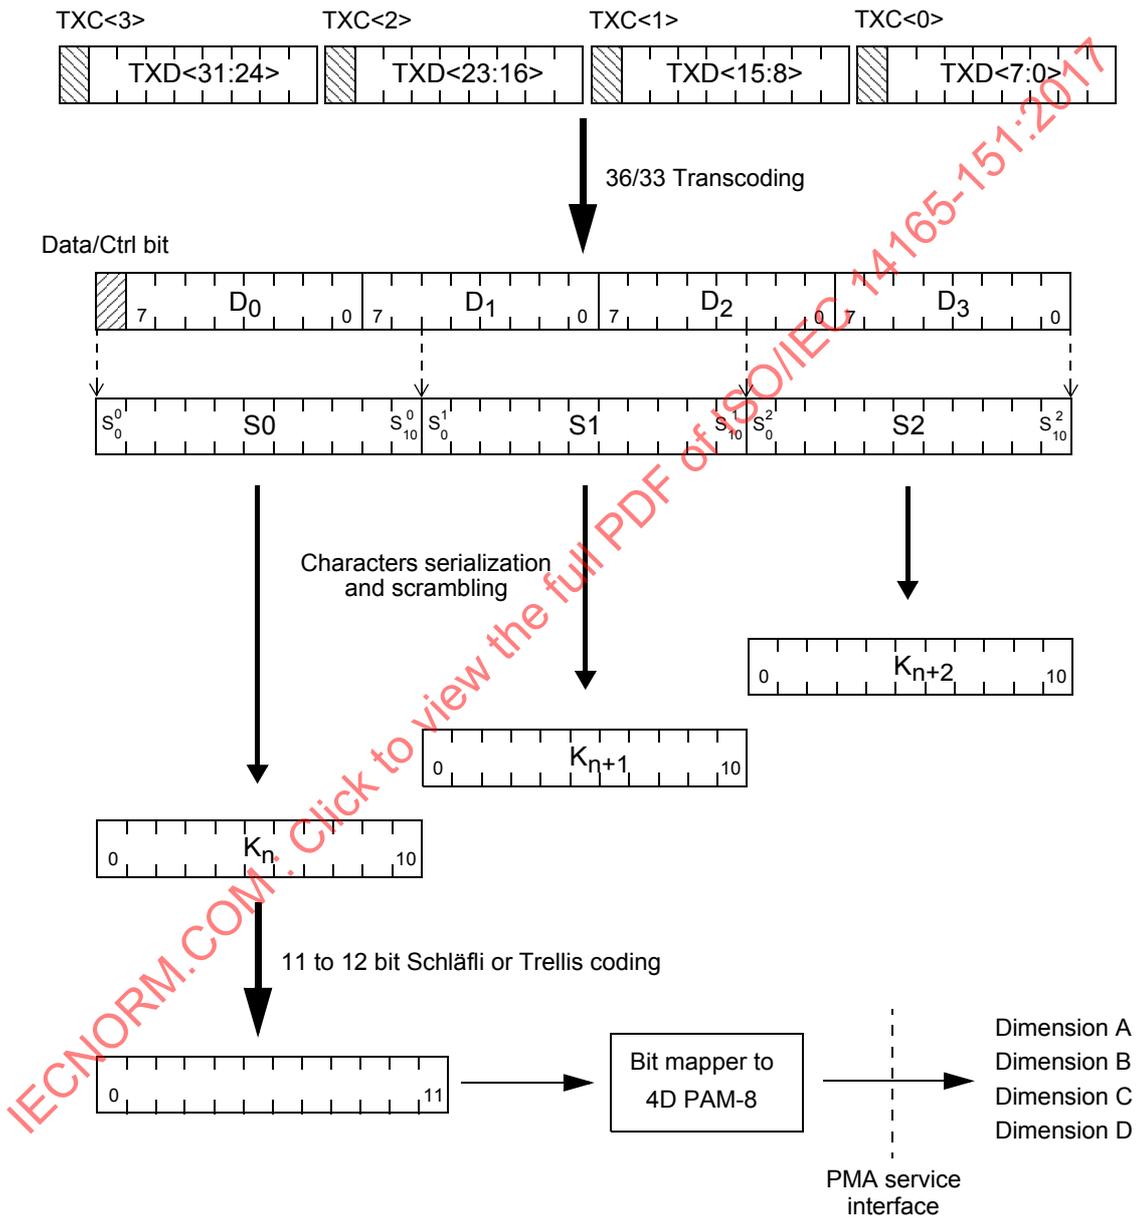


Figure 9 – PCS transmit bit ordering

When operating in normal mode, the PCS Transmit process continuously generates 33-bit (33B) blocks based upon the TXD <31:0> and TXC <3:0> signals on the XGMII. Each 33B block is divided into three 11-bit transmission characters that are individually scrambled. Each 11-bit scrambled transmission character is converted in a 4D PAM-8 symbol, transferred to the PMA.

In each symbol period, when communicating with the PMA, PCS Transmit generates a code-group (A_n , B_n , C_n , D_n) that is transferred to the PMA via the PMA_UNITDATA.request primitive. The PMA transmits symbols A_n , B_n , C_n , D_n over dimensions A, B, C, and D respectively. The integer, n , is a time index that is introduced to establish a temporal relationship between different symbol periods. A symbol period, T , is as specified in table 6.

If a PMA_TXMODE.indication message has the value SEND_Z, PCS Transmit keeps the transmitters silent.

If a PMA_TXMODE.indication message has the value SEND_T1 or SEND_T2, PCS Transmit generates the sequences of code-groups (TA_n , TB_n , TC_n , TD_n) defined in 5.3.7 and passes them to the PMA via the PMA_UNITDATA.request primitive. These code-groups are used for training mode and only transmit the values $\{-5, +5\}$ to keep the transmit power in the training mode roughly the same as the transmit power in normal mode.

If a PMA_TXMODE.indication message has the value SEND_I2 or SEND_I3, PCS Transmit generates sequences of 4D PAM-8 symbols representing respectively the scrambled Idle2 Ordered Set or the scrambled Idle3 Ordered Set (see 6.2.4), symbols generated as in normal mode of operation.

In the normal mode of operation, the PMA_TXMODE.indication message has the value SEND_N, and the PCS Transmit function uses a 36/33 transcoding technique to generate 33B blocks that represent data or control. Each 33B block is divided into three 11-bit transmission characters, S0, S1, and S2. S0 is transmitted first, followed by S1 and after by S2. Each 11-bit transmission character is scrambled and then either Schläfli lattice or Trellis encoded, with 12 bits. The 12 bits are in turn divided in four groups of three bits, each identifying a symbol in the PAM-8 symbol set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$. The four identified symbols form a 4D PAM-8 symbol, transferred to the PMA.

Either 4D PAM-8 Trellis encoding (see 5.3.6) or 4D PAM-8 Schläfli lattice encoding (see 5.3.5) is employed. Schläfli lattice is the densest 4-dimensional lattice, using half points (i.e., 2048) of a 4096 points 4D PAM-8 constellation. This encoding ensures that the minimum squared Euclidean distance between points of the Schläfli lattice is 2 times the minimum squared Euclidean distance between points of the complete 4D PAM-8 constellation. For both Trellis and Schläfli lattice encoding the resulting modulation rate is of 2,75 bits per symbol.

5.3.2 36/33 transcoding

Fibre Channel uses two types of words, data words and Ordered Sets, summarized in table 7 using a 40-bit representation.

Table 7 – Fibre Channel words

Word Type	10b Representation
Data Word	$D_{xx.y_0} D_{xx.y_1} D_{xx.y_2} D_{xx.y_3}$
FC Ordered Set	$K28.5 D_{xx.y_1} D_{xx.y_2} D_{xx.y_3}$

5.3.3 Error detecting code

Referring to table 9, the bits of the 11-bit transmission characters S0, S1 and S2 are respectively denoted as $S_0^0 \dots S_{10}^0$, $S_0^1 \dots S_{10}^1$, $S_0^2 \dots S_{10}^2$. With this notation the bits specifying the error detecting code are S_1^0 , S_0^1 , S_1^1 , S_0^2 , and S_1^2 . These bits shall be computed as a function of the bits of the 33B block as follows:

$$S_1^0 = S_3^0 \oplus S_4^0 \oplus S_5^0 \oplus S_6^0 \oplus S_{10}^0 \oplus S_2^1 \oplus S_3^1 \oplus S_7^1 \oplus S_9^1 \oplus S_{10}^1 \oplus S_4^2 \oplus S_6^2 \oplus S_7^2 \oplus S_8^2 \oplus S_9^2$$

$$S_0^1 = S_2^0 \oplus S_3^0 \oplus S_6^0 \oplus S_7^0 \oplus S_9^0 \oplus S_{10}^0 \oplus S_3^1 \oplus S_4^1 \oplus S_6^1 \oplus S_7^1 \oplus S_8^1 \oplus S_9^1 \oplus S_3^2 \oplus S_4^2 \oplus S_5^2 \oplus S_6^2 \oplus S_9^2 \oplus S_{10}^2$$

$$S_1^1 = S_2^0 \oplus S_4^0 \oplus S_5^0 \oplus S_7^0 \oplus S_8^0 \oplus S_{10}^0 \oplus S_2^1 \oplus S_4^1 \oplus S_5^1 \oplus S_7^1 \oplus S_8^1 \oplus S_{10}^1 \oplus S_2^2 \oplus S_4^2 \oplus S_5^2 \oplus S_7^2 \oplus S_8^2 \oplus S_{10}^2$$

$$S_0^2 = S_5^0 \oplus S_6^0 \oplus S_7^0 \oplus S_8^0 \oplus S_9^0 \oplus S_{10}^0 \oplus S_2^1 \oplus S_3^1 \oplus S_4^1 \oplus S_5^1 \oplus S_6^1 \oplus S_7^1 \oplus S_2^2 \oplus S_3^2 \oplus S_4^2 \oplus S_8^2 \oplus S_9^2 \oplus S_{10}^2$$

$$S_1^2 = S_5^0 \oplus S_6^0 \oplus S_7^0 \oplus S_2^1 \oplus S_3^1 \oplus S_4^1 \oplus S_8^2 \oplus S_9^2 \oplus S_{10}^2$$

The EDC bits are computed in this way to satisfy the following equation, based on a parity check matrix meant to detect as many as possible small errors:

$$\begin{bmatrix} 0 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \oplus \begin{bmatrix} S_0^0 \\ \dots \\ S_{10}^0 \\ S_0^1 \\ \dots \\ S_{10}^1 \\ S_0^2 \\ \dots \\ S_{10}^2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

The FC-BaseT error detecting code detects the following errors over a single transmission character:

- a) All 1-bit errors;
- b) All 2-bit errors.
 - A) two 1-bit errors over 2 dimensions; and
 - B) one 2-bit error over 1 dimension;
- c) The most probable 3-bit errors:
 - A) three 1-bit errors over 3 dimensions; and

- B) one 2-bit error over 1 dimension + one 1-bit error over 1 dimension;
- d) The most probable 4-bit errors:
 - A) four 1-bit errors over 4 dimensions; and
 - B) one 2-bit error over 1 dimension + two 1-bit errors over 2 dimensions.

The FC-BaseT error detecting code detects the following errors over multiple transmission characters in the same 33-bit block:

- a) All 1-bit errors;
- b) All 2-bit errors:
 - A) two 1-bit errors over two characters;
- c) Most 3-bit errors:
 - A) three 1-bit errors over a single dimension of the three characters; and
 - B) most two 1-bit errors over two dimensions of a character + one 1-bit error on an adjacent character;
- d) Most trapezoid error patterns:
 - A) two 1-bit errors in the same dimension of two adjacent characters + two 1-bit errors in each of two other dimensions of the same characters.

5.3.4 PCS scrambling

The FC-BaseT PCS Transmit function employs side-stream scrambling. The scrambler for the Master PHY shall produce the same result as the implementation shown in figure 10, that implements the scrambler polynomial $G_M(x) = 1 + x^{13} + x^{33}$ using linear-feedback shift registers.

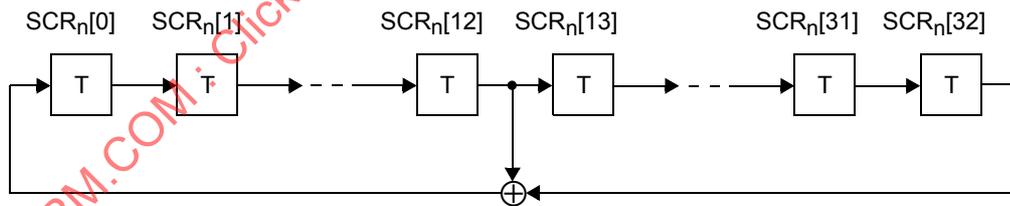


Figure 10 – Side-stream scrambler for the Master PHY

The scrambler for the Slave PHY shall produce the same result as the implementation shown in figure 11, that implements the scrambler polynomial $G_M(x) = 1 + x^{20} + x^{33}$ using linear-feedback shift registers.

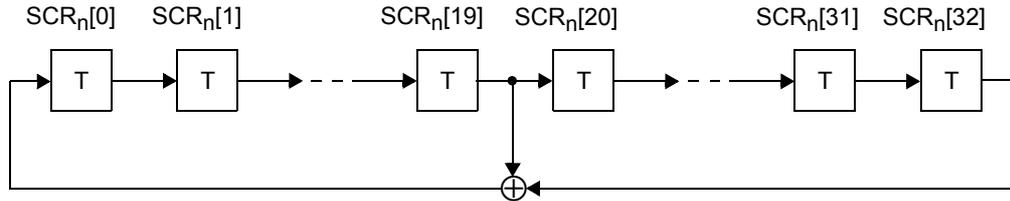


Figure 11 – Side-stream scrambler for the Slave PHY

NOTE 2 - These scrambling polynomials are the same used in 1000BASE-T (see 802.3-2005, clause 40).

The bits stored in the shift register delay line at time n are denoted by $Scr_n[0 .. 32]$. At each symbol period, the shift register is advanced by one bit, and one new bit $Scr_n[0]$ is generated.

The transmitter side-stream scrambler is reset upon execution of the PCS Reset function. If PCS Reset is executed, all bits of the 33-bit vector representing the side-stream scrambler state are arbitrarily set. The initialization of the scrambler state is left to the implementor. In no case the scrambler state shall be initialized to all zeros.

The bits stored in the shift register delay line at time n are used to construct an 11-bit scrambling vector $Q_n[0 .. 10]$. By construction, the eleven bits of the scrambling vector $Q_n[0 .. 10]$ are derived from elements of the same maximum-length shift register sequence of length $2^{33}-1$ as $Scr_n[0]$, but shifted in time by varying delays. The associated delays are all large and different so that there is no apparent correlation among the $Q_n[0 .. 10]$ bits. For both Master and Slave PHYs, the scrambling vector $Q_n[0 .. 10]$ is obtained by the following linear combinations of bits stored in the transmit scrambler shift register delay line:

$$Q_n[0] = Scr_n[0]$$

$$Q_n[1] = Scr_n[3] \oplus Scr_n[8]$$

$$Q_n[2] = Scr_n[6] \oplus Scr_n[16]$$

$$Q_n[3] = Scr_n[9] \oplus Scr_n[14] \oplus Scr_n[19] \oplus Scr_n[24]$$

$$Q_n[4] = Scr_n[1] \oplus Scr_n[5]$$

$$Q_n[5] = Scr_n[4] \oplus Scr_n[8] \oplus Scr_n[9] \oplus Scr_n[13]$$

$$Q_n[6] = Scr_n[7] \oplus Scr_n[11] \oplus Scr_n[17] \oplus Scr_n[21]$$

$$Q_n[7] = Scr_n[4] \oplus Scr_n[6]$$

$$Q_n[8] = Scr_n[7] \oplus Scr_n[9] \oplus Scr_n[12] \oplus Scr_n[14]$$

$$Q_n[9] = Scr_n[10] \oplus Scr_n[12] \oplus Scr_n[20] \oplus Scr_n[22]$$

$$Q_n[10] = Scr_n[13] \oplus Scr_n[15] \oplus Scr_n[18] \oplus Scr_n[20] \oplus Scr_n[23] \oplus Scr_n[25] \oplus Scr_n[28] \oplus Scr_n[30]$$

NOTE 3 - The bits of the scrambling vector $Q_n[0 \dots 10]$ are computed in the same manner as bits $Sy_n[3:0]$, $Sg_n[2:0]$, and $Sx_n[3:0]$ in IEEE 802.3-2005, clause 40.3.1.3.2.

The scrambling vector of time n is used to scramble the 11-bit transmission character to be transmitted at time n (e.g., if S_0 has to be transmitted at time n , Q_n is used to scramble S_0 , Q_{n+1} is used to scramble S_1 , and Q_{n+2} is used to scramble S_2). The result of the scrambling is an 11-bit scrambled transmission character K_n . Each bit of K_n is obtained by performing a 'XOR' operation between each bit of the unscrambled transmission character S_x and the corresponding bit of the scrambling vector at time n (e.g., $K_n[0] = S_0^0 \oplus Q_n[0]$; $K_n[1] = S_1^0 \oplus Q_n[1]$; ... ; $K_n[10] = S_{10}^0 \oplus Q_n[10]$).

5.3.5 Schläfli Lattice coding

If selected (see 8.4), each 11-bit scrambled transmission character K_n is Schläfli lattice encoded with 12 bits, using the encoder shown in figure 12.

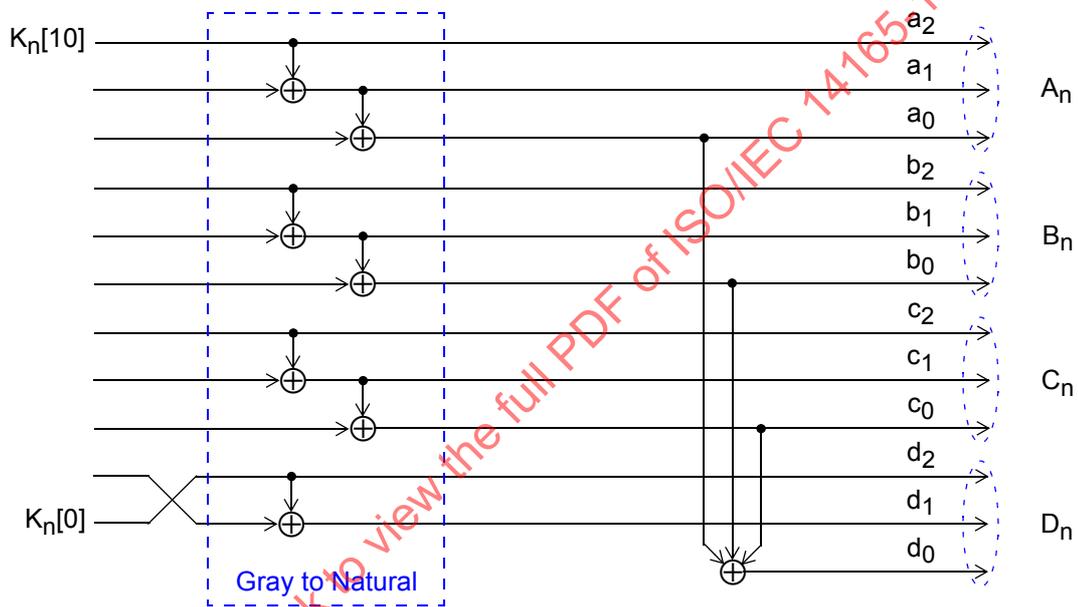


Figure 12 – 4D PAM-8 Schläfli Lattice encoder

The 11 bits of each scrambled transmission character K_n are processed in the order shown in figure 12, from bit 0 to bit 10. $K_n[0]$ is mapped to bit d_2 in figure 12, to ensure that bit 0 of transmission character S_0 (i.e., the Data/Ctrl bit) is mapped into the MSB of the PAM-8 symbol transmitted on dimension D. Bit a_2 , b_2 , and c_2 are respectively generated equal to bits $K_n[10]$, $K_n[7]$, and $K_n[4]$. Bit a_1 , b_1 , and c_1 are respectively generated by XORing bits a_2 with $K_n[9]$, bits b_2 with $K_n[6]$, and bit c_2 with $K_n[3]$. Bit a_0 , b_0 , and c_0 are respectively generated by XORing bits a_1 with $K_n[8]$, bits b_1 with $K_n[5]$, and bit c_1 with $K_n[2]$. Bit d_1 is generated by XORing bit d_2 with $K_n[1]$. Bit d_0 is generated by XORing bits a_0 , b_0 , and c_0 .

Each group of three bits $x_2 x_1 x_0$ (where x may be a, b, c, or d) identifies a symbol in the PAM-8 symbol set $\{-7, -5, -3, -1, +1, +3, +5, +7\}$. The bits to PAM-8 symbol mapping is shown in table 11.

Table 11 – Bit to symbols mapping

Bits ($x_2 x_1 x_0$)	Symbol
111b	+7
110b	+5
101b	+3
100b	+1
011b	-1
010b	-3
001b	-5
000b	-7

As an example, the 11 bits 10010100111b are Schläfli lattice coded to the 4-dimensional PAM-8 symbol (+7, +5, -5, +1).

5.3.6 Trellis coding

If selected (see 8.4), each 11-bit scrambled transmission character K_n is Trellis encoded with 12 bits, using the 8-states Trellis encoder shown in figure 13.

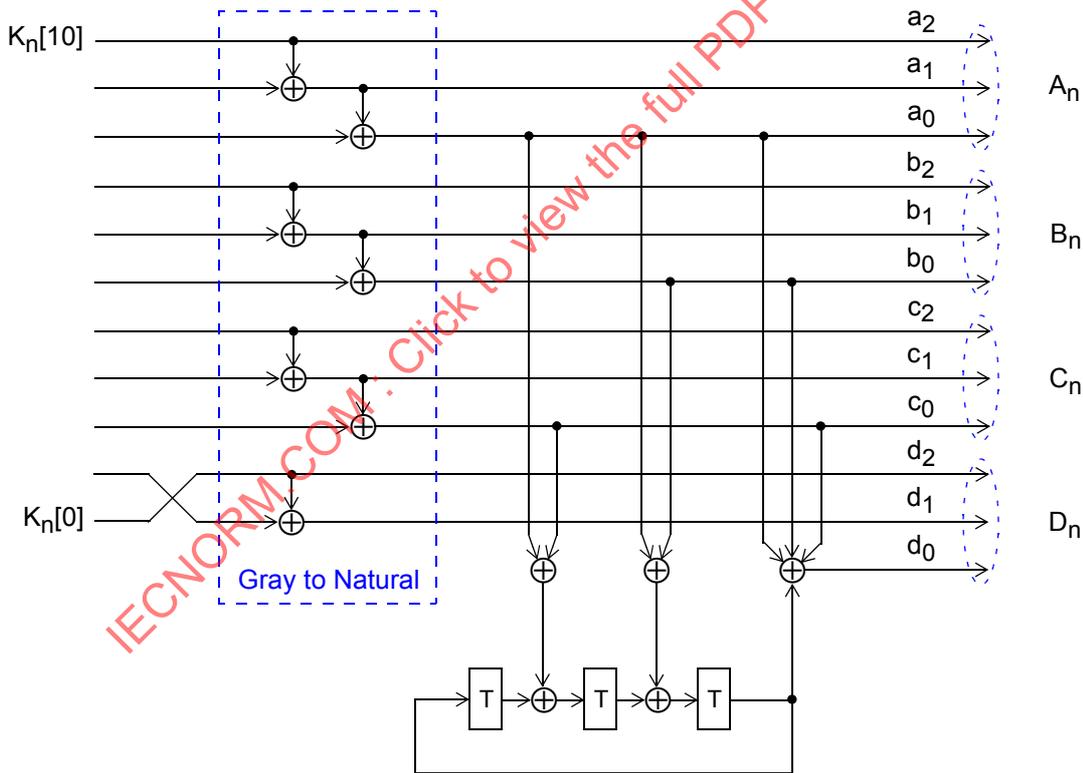


Figure 13 – 4D PAM-8 Trellis encoder

5.3.7 Generation of PMA training sequences

The FC-BaseT PMA training provides for pair swap, polarity, and skew detection and allows the convergence of decision feedback equalizers. PMA training is based on sequences of PAM-2 uncorrelated symbols having identical power spectral density as data signals, within 0,75 dB. The PCS Transmit function generates PAM-2 PMA training sequences as shown in figure 14.

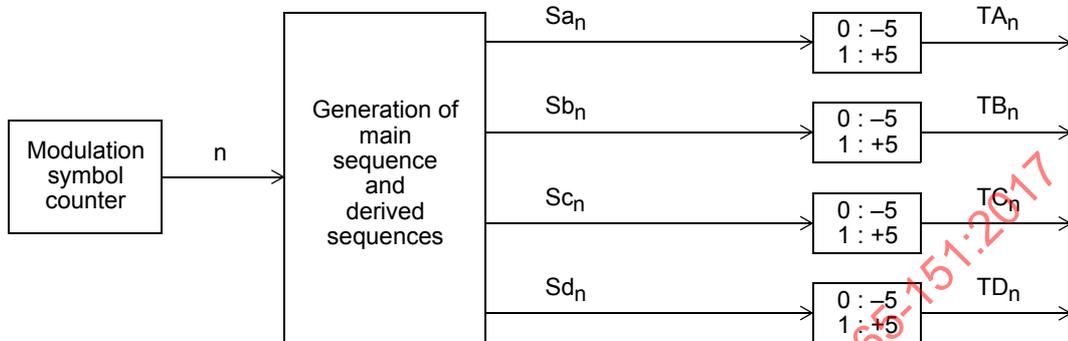


Figure 14 – Generation of PMA training PAM-2 sequences

PMA training signals are based on the generation, at time n , of the four bits Sa_n , Sb_n , Sc_n , Sd_n . The four bits Sa_n , Sb_n , Sc_n , Sd_n are mapped to a 4D symbol (TA_n , TB_n , TC_n , TD_n) by selecting the value -5 if the bit is zero, or the value $+5$ if the bit is one, as shown in figure 14.

Two PMA training sequences are defined, denoted Type-1 PAM-2 and Type-2 PAM-2. The four bits Sa_n , Sb_n , Sc_n , Sd_n are generated from the scrambling vector Q_n (see 5.3.4) for both Type-1 and Type-2 PAM-2 training sequences as shown in table 12.

Table 12 – PMA training sequences

Type-1 PAM-2	Type-2 PAM-2
$Sa_n = Q_n[0]$	$Sa_n = Q_n[0]$
$Sb_n = Q_n[1]$	$Sb_n = Q_n[4]$
$Sc_n = Q_n[2]$	$Sc_n = Q_n[2]$
$Sd_n = Q_n[3]$	$Sd_n = Q_n[3]$

For the Type-1 PAM-2 training sequence, two PMA training modes are defined, normal and periodic. The PMA training mode is selected during Auto-Negotiation.

In periodic mode the scrambler state is reinitialized after every 16 384 symbol periods, to generate a periodically repeating pattern with period 16 384. The initial (i.e., when $n = 0$) 33 bit values of the scrambler state shall be generated by combining the 22 low order bits from the value 39A422h for the 22 MSBs and the random value SB10 .. SB0 from table 41 generated by the local device for the 11 LSBs. The scrambler state shall be reset to this same 33 bit values whenever $n \bmod 16\ 384 = 0$.

The scrambler synchronization obtained by the PMA training function is used in the normal mode of operation.

5.4 PCS receive function

5.4.1 Overview

The PCS Receive function accepts received code-groups provided by the PMA Receive function via the parameter rx_symb_vector. The PCS Receive function uses knowledge of the encoding rules to correctly align the 33B blocks (see 5.4.4). The received 33B blocks are converted to the signals RXD<31:0> and RXC<3:0> for transmission to the XGMII.

The PCS Receive bit ordering shall be as shown in figure 15, that shows the mapping from a 33B block to XGMII for a block containing four data characters.

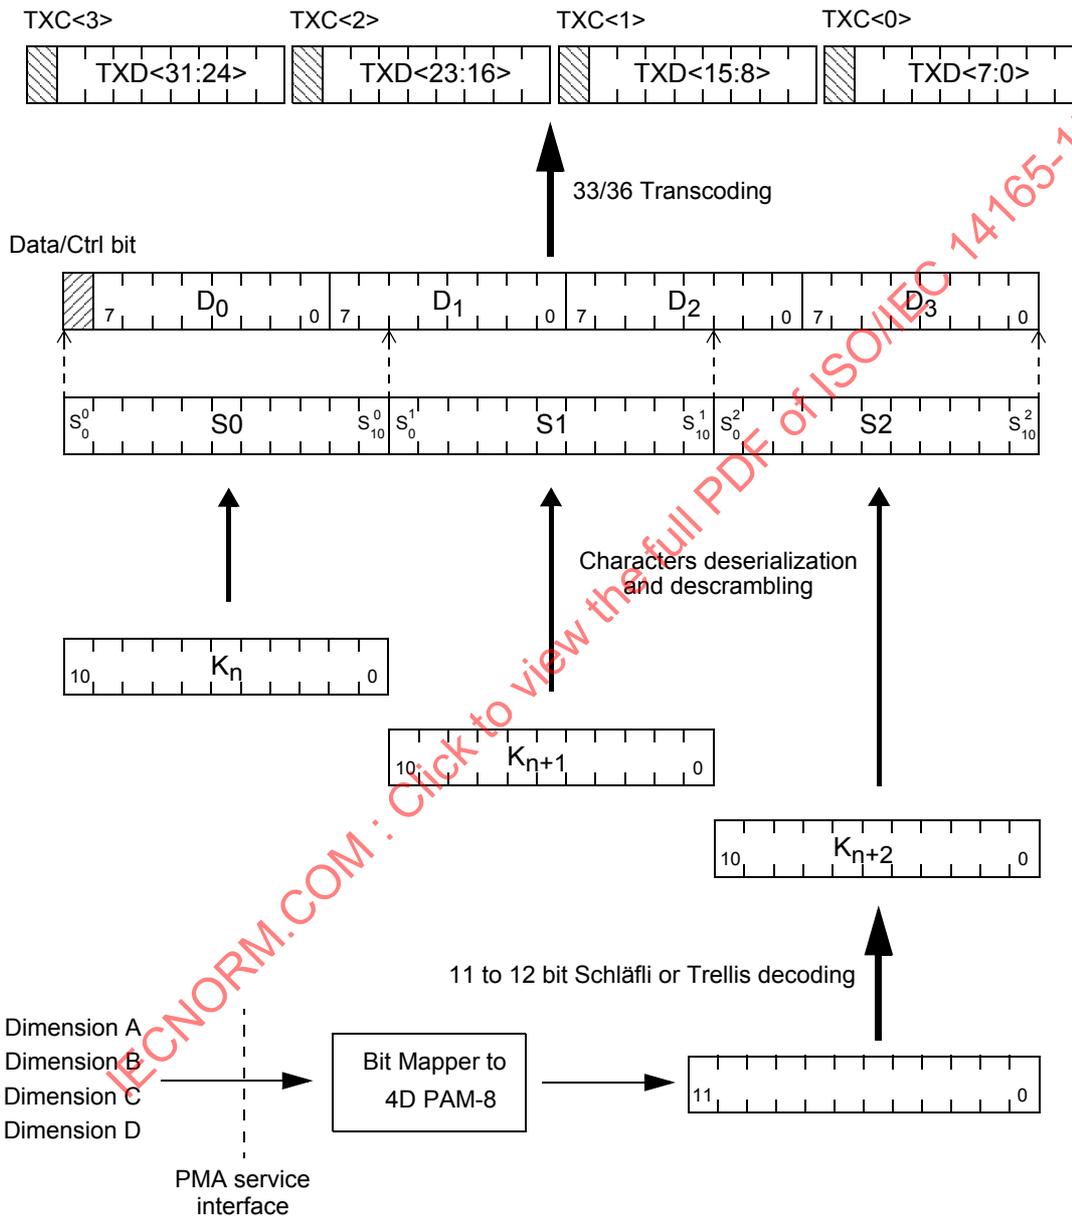


Figure 15 – PCS receive bit ordering

During PMA training mode, the PCS Receive function synchronizes to the remote scrambler and signals the reliable acquisition of the descrambler state by setting the parameter `scr_status` to OK.

When the PCS Synchronization process has obtained synchronization (see 5.4.5), the PCS Receive function decodes in a 11-bit scrambled transmission character the 4D PAM-8 symbol transferred to the PCS by the `PMA_UNITDATA.indication` primitive via the parameter `rx_symb_vector`. Each 11-bit scrambled character is descrambled and three consecutive descrambled characters are assembled in a 33B block. The 33B block is evaluated and transcoded, as specified in 5.4.4, in a 36-bit XGMII word, to be sent over the XGMII receive interface.

5.4.2 Decoding

The PCS Receive function accepts received code-groups provided by the PMA Receive function via the parameter `rx_symb_vector`. If the link partner is using Schläfli lattice encoding, each received 4D PAM-8 symbol is Schläfli lattice decoded in a 11-bit scrambled character as shown in figure 16.

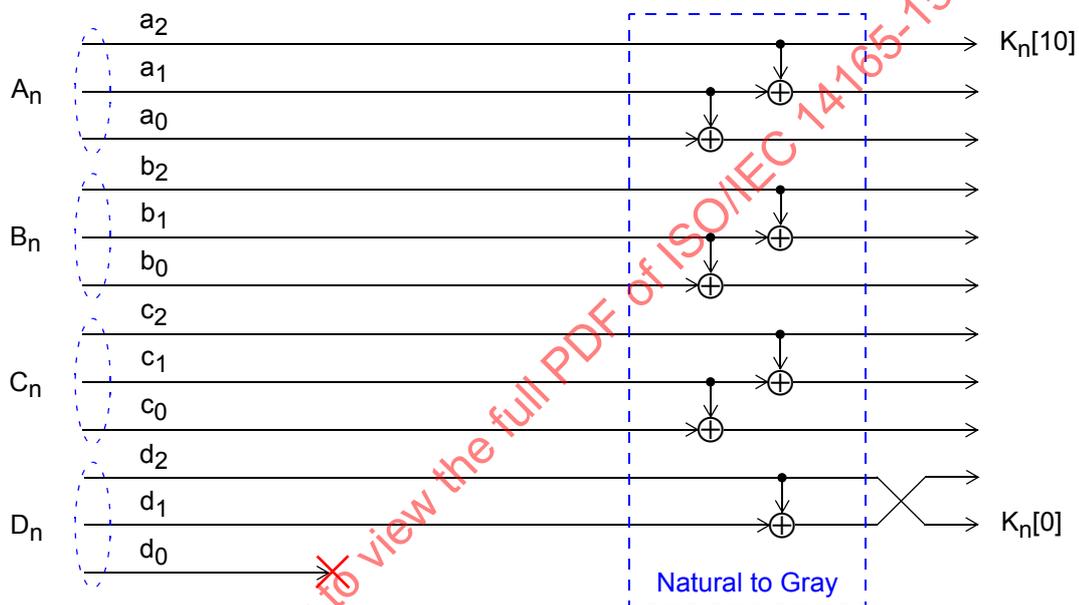


Figure 16 – 4D PAM-8 Schläfli Lattice decoder

If the link partner is using Trellis encoding, each 4D PAM-8 symbol is decoded in a 11-bit scrambled character with a 8-state Viterbi decoder. The Viterbi decoder is self-synchronizing, therefore there is no need to specify an initial state for the Trellis encoder. This standard does not define the 8-state Viterbi decoder.

NOTE 4 - Information on Viterbi decoding may be found in the book "Error-Correction Coding for Digital Communications", of George C. Clark, Jr. and J. Bibb Cain, 1981. ISBN: 0306406152.

After performing the decoding, each received 11-bit scrambled character K_n is descrambled.

5.4.3 PCS descrambling

The FC-BaseT PCS Receive function employs side-stream descrambling. The descrambler for the Master PHY shall produce the same result as the implementation shown in figure 17, that implements the descrambler polynomial $G'_M(x) = 1 + x^{20} + x^{33}$ using linear-feedback shift registers.

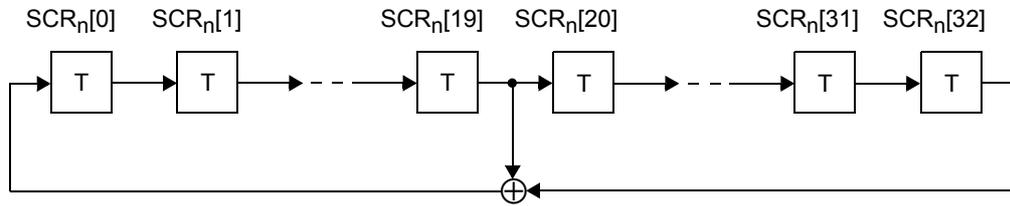


Figure 17 – Side-stream descrambler for the Master PHY

The descrambler for the Slave PHY shall produce the same result as the implementation shown in figure 18, that implements the descrambler polynomial $G'_S(x) = 1 + x^{13} + x^{33}$ using linear-feedback shift registers.

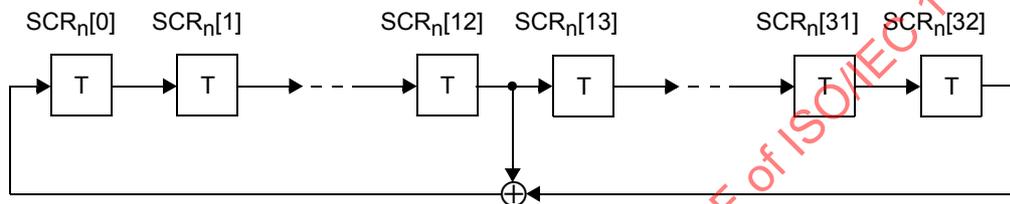


Figure 18 – Side-stream descrambler for the Slave PHY

NOTE 5 - The descrambler for the Master PHY is the same as the scrambler for the Slave PHY and the descrambler for the Slave PHY is the same as the scrambler for the Master PHY.

The bits stored in the shift register delay line at time n are denoted by $Scr_n[0 .. 32]$. At each symbol period, the shift register is advanced by one bit, and one new bit $Scr_n[0]$ is generated. The bits stored in the shift register delay line at time n are used to construct an 11-bit descrambling vector $Q'_n[0 .. 10]$ in the same manner as the scrambling vector defined in 5.3.4.

The descrambling vector of time n is used to descramble the 11-bit transmission character K_n received at time n . The result of the descrambling is an 11-bit unscrambled transmission character S_x (i.e., S_0, S_1 or S_2). Each bit of S_x is obtained by performing a 'XOR' operation between each bit of the received transmission character K_n and the corresponding bit of the descrambling vector at time n (e.g.,

$$S_0^0 = K_n[0] \oplus Q'_n[0]; S_1^0 = K_n[1] \oplus Q'_n[1]; \dots; K_n[10] = S_{10}^0 = K_n[10] \oplus Q'_n[10]).$$

5.4.4 33/36 transcoding

The PCS Receive function shall evaluate each received 33B block as follows:

- a) DATA: if the 33B block is a data block (i.e., if the Data/Ctrl bit is set to zero, see table 9);
- b) VALID: if the 33B block is a control block (i.e., if the Data/Ctrl bit is set to one) and the EDC check is successful; or
- c) INVALID: if the 33B block is a control block (i.e., if the Data/Ctrl bit is set to one) and the EDC check is unsuccessful.

The EDC shall be checked by computing the EDC bits as specified in 5.3.3 and verifying that they are the same as the ones received in the 33B control block.

When operating in normal mode (i.e., when PCS Receive achieved block synchronization, see 5.4.5):

- a) each DATA or VALID 33B block shall be transcoded to its 36-bit XGMII representation by reversing the mapping specified in 5.3.2; and
- b) each INVALID 33B block shall be transcoded to the current Fill Word (see 7.2) when the Invalid Block Masking bit of the management Control register (see 9.3) is set to one, or to the XGMII error word (i.e., K30.7 || K30.7 || K30.7 || K30.7) when the Invalid Block Masking bit of the management Control register is set to zero.

An implementation may attempt to classify additional 33B blocks as INVALID using the knowledge of the structure of a Fibre Channel stream of traffic. The error detecting code specified in 5.3.3 is defined only for 33B control blocks and therefore does not protect from an error on the Data/Ctrl bit. In particular, if the Data/Ctrl bit of a 33B control block is erroneously received as zero, the 33B block is interpreted as a data block and no EDC check is performed.

To detect these errors an FC-BaseT receiver may use the knowledge of the structure of a Fibre Channel stream of traffic, and in particular of the following rules (see FC-FS-2):

- a) data words occur only inside a Fibre Channel frame;
- b) the minimum size of a Fibre Channel frame is seven words;
- c) Fibre Channel frames are delimited by the SOFx and EOFx Ordered Sets; and
- d) the minimum inter-frame gap at a receiver is of two Fill Words.

These rules mean that a stream of Fibre Channel traffic is valid only if:

- a) there are at least seven consecutive data words following a transition from Ordered Sets to data words; and
- b) there are at least four Ordered Sets following a transition from data words to Ordered Sets.

An FC-BaseT receiver may then detect an error when:

- a) there are less than seven consecutive 33B data blocks following a transition from 33B control blocks to 33B data blocks; and

- b) there are less than four 33B control blocks following a transition from 33B data blocks to 33B control blocks.

In particular, a 33B data block preceded and followed by 33B control blocks may be evaluated as an INVALID 33B block.

5.4.5 PCS synchronization

The PCS Receive function shall operate according to the state diagram shown in figure 19 in order to achieve 33B block synchronization.

PCS Receive begins its operations in the NO_PCS_SYNC state and unconditionally transitions in the VERIFY_BLOCK_1 state, where each received 33B block is evaluated as specified in 5.4.4. The valid_block_count counter is incremented on receiving a VALID block and set to zero on receiving an INVALID block. PCS Synchronization is achieved, and PCS Receive transitions in the PCS_SYNC state, when valid_block_count exceeds the parameter U.

When PCS Synchronization is achieved, PCS Receive continues to evaluate each received 33B block in the VERIFY_BLOCK_2 state. The invalid_block_count counter is incremented on receiving an INVALID block and decremented on receiving a VALID block. PCS Synchronization is lost, and PCS Receive transitions in the NO_PCS_SYNC state, when invalid_block_count exceeds the parameter U.

The purpose of the PCS Synchronization process is the timely detection of persistent error conditions at the physical layer (e.g., a loss of descrambler synchronization). The value of the parameter U is vendor specific and should be chosen in a way that avoids detecting transient errors as persistent ones. The value chosen for the parameter U should be in the range 16 .. 64.

5.5 State Diagrams

A state diagram is provided for the PCS Synchronization process. The variables used in the state diagram are shown in table 13.

Table 13 – PCS synchronization state variables

Variable	Description
pcs_status	See 4.6.2.7.
block_status	Specifies the status of a received 33B block. Each received 33B block is evaluated as specified in 5.4.4. This variable may assume the following values: UNKNOWN: status not yet evaluated; VALID: the 33B block is valid; INVALID: the 33B block is invalid; or DATA: the 33B block is a data block.
valid_block_count	Used to count the number of valid 33B blocks
invalid_block_count	Used to count the number of invalid 33B blocks

The PCS Synchronization state diagram is shown in figure 19.

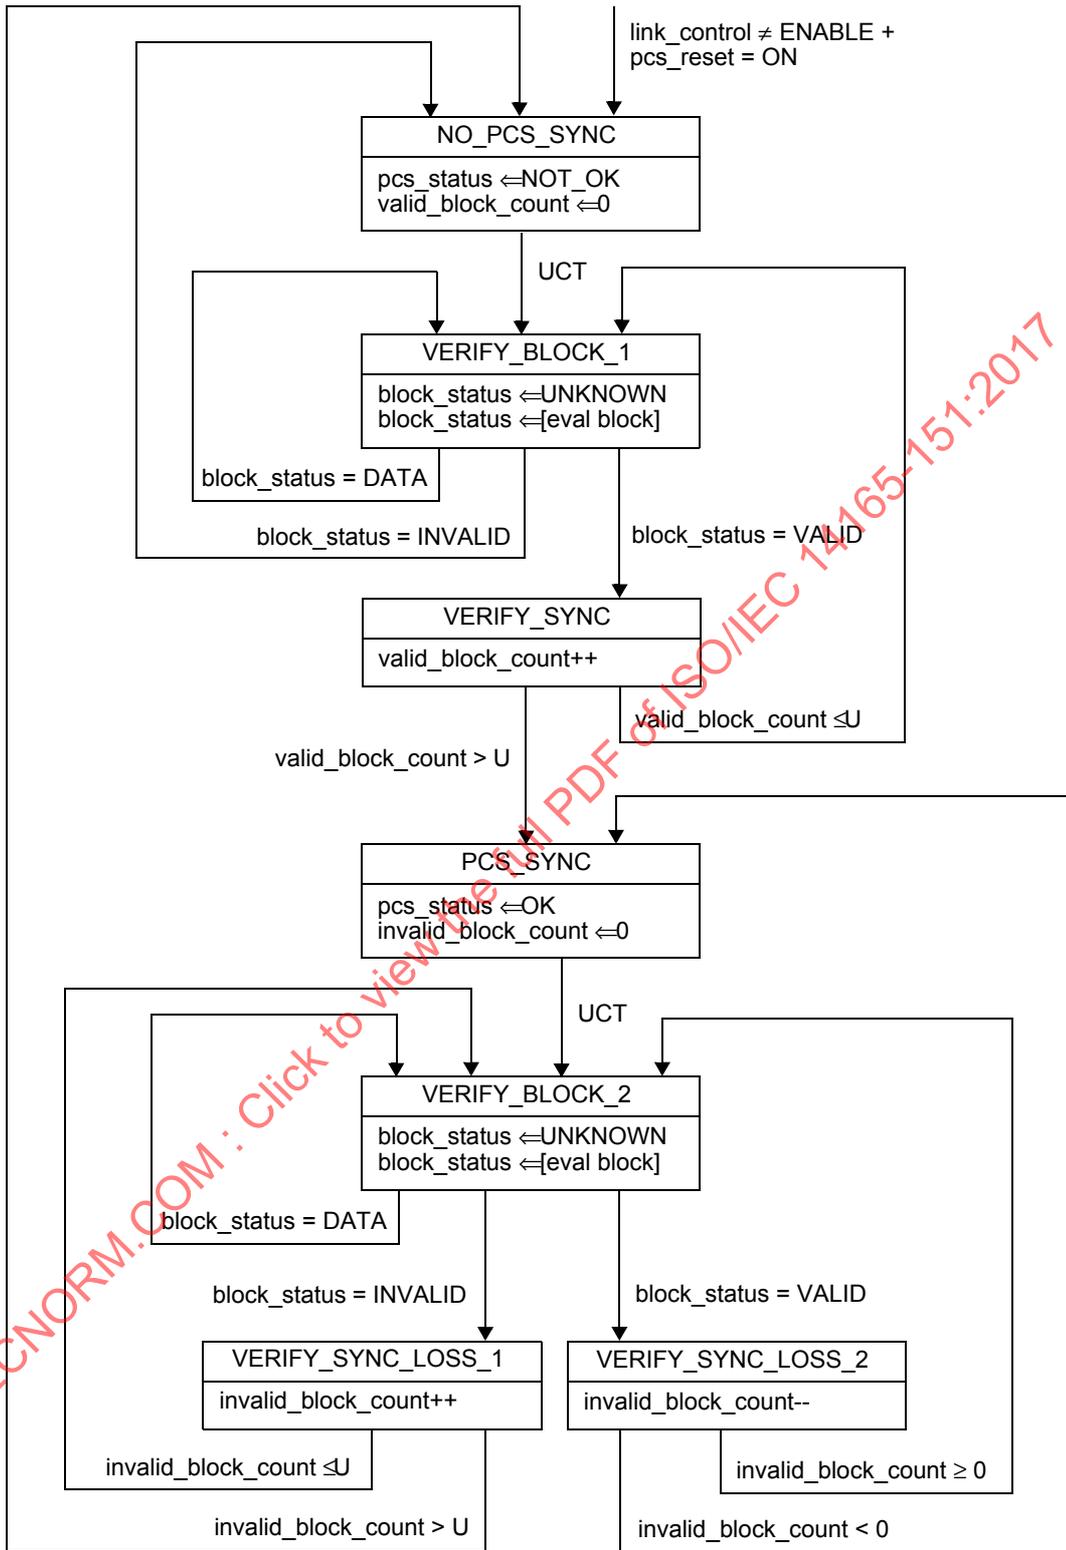


Figure 19 – PCS synchronization state diagram

6 PMA Sublayer and Medium Dependent Interface

6.1 PMA Overview

The PMA sublayer couples messages from the PMA service interface specified in 4.6.2 to the baseband medium via the Medium Dependent Interface (MDI).

The PMA sublayer comprises one PMA Reset function and five simultaneous and asynchronous operating functions. The PMA operating functions are PHY Control, PMA Transmit, PMA Receive, Link Monitor, and Clock Recovery. All operating functions are started immediately after the successful completion of the PMA Reset function.

The PMA reference diagram, shown in figure 20, illustrates how the five operating functions relate to the messages of the PMA Service interface and the signals of the MDI. Connections from the management interface to other layers are pervasive and are not shown in figure 20.

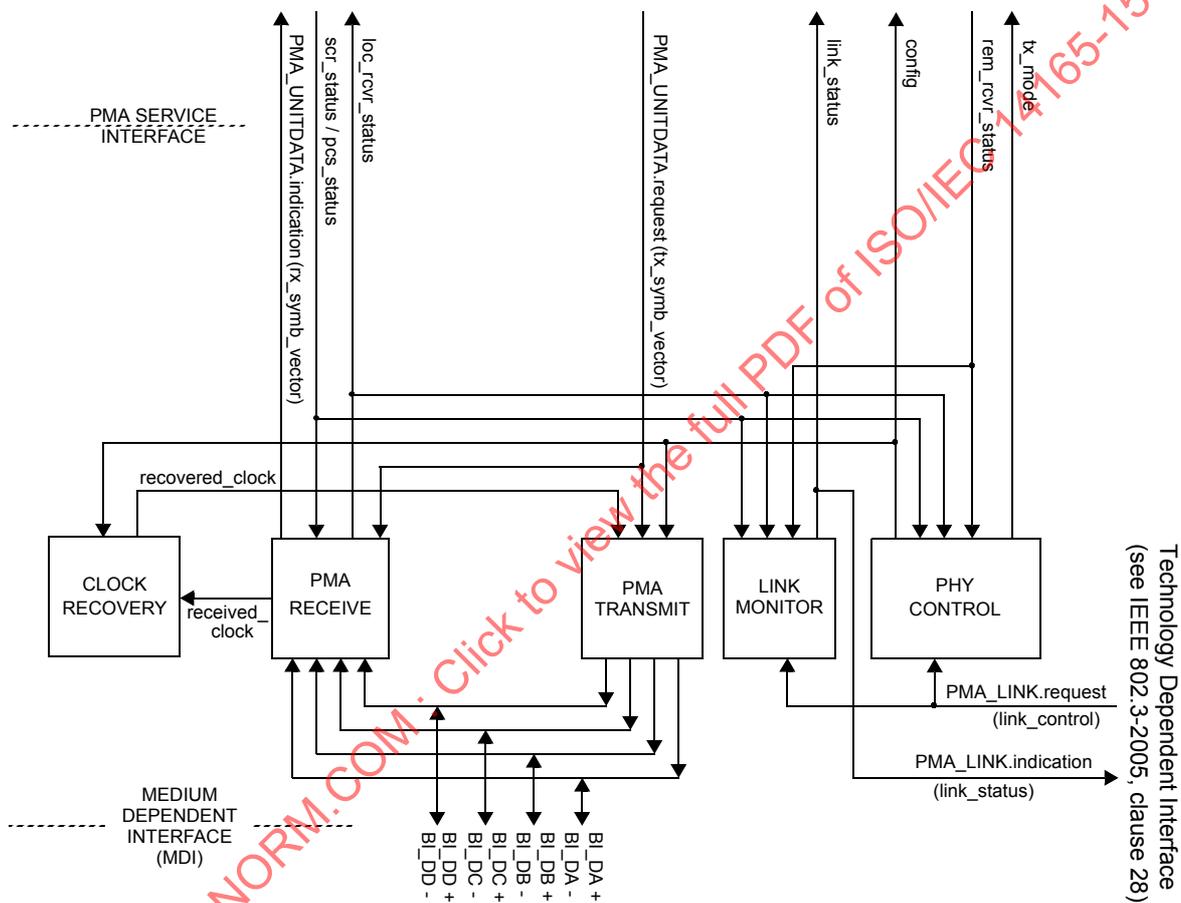


Figure 20 – PMA Reference Diagram

6.2 PMA Functions

6.2.1 PMA Reset function

The PMA Reset function shall be executed whenever one of the two following conditions occur:

- a) Power on; and
- b) The receipt of a request for reset from the management entity.

PMA Reset sets `pma_reset=ON` while any of the above reset conditions hold true. The reference diagrams do not explicitly show the PMA Reset function.

6.2.2 PMA Transmit function

The PMA Transmit function comprises four synchronous transmitters to generate four PAM-8 signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. The PMA Transmit function shall continuously transmit onto the MDI pulses modulated by the PAM-8 symbols given by `tx_symb_vector[BI_DA]`, `tx_symb_vector[BI_DB]`, `tx_symb_vector[BI_DC]`, and `tx_symb_vector[BI_DD]`, respectively. The four transmitters shall be driven by the same transmit clock, TX_TCLK. The signals generated by PMA Transmit shall comply with the electrical specifications given in 6.3.

When the `PMA_CONFIG.indication` parameter `config` is MASTER, the PMA Transmit function shall source TX_TCLK from a local clock source. When the `PMA_CONFIG.indication` parameter `config` is SLAVE, the PMA Transmit function shall source TX_TCLK from the recovered clock of 6.2.6.

6.2.3 PMA Receive Function

The PMA Receive function comprises four independent receivers for PAM-8 signals on each of the four pairs BI_DA, BI_DB, BI_DC, and BI_DD. PMA Receive contains the circuits necessary to both detect PAM-8 symbol sequences from the signals received at the MDI over receive pairs BI_DA, BI_DB, BI_DC, and BI_DD and to present these sequences to the PCS Receive function. The PMA shall translate the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD over a channel meeting the requirements of 6.5 into the `PMA_UNITDATA.indication` parameter `rx_symb_vector`, with a symbol error rate compatible with the BER specified by the BER requirement bits of the Management Control Register (see 9.3). When Schläfli lattice coding is used, this process involves finding the point in the 4D PAM-8 Schläfli lattice having minimum Euclidean distance from the signals received on pairs BI_DA, BI_DB, BI_DC, and BI_DD.

To achieve the indicated performance, it is highly recommended that PMA Receive include the functions of signal equalization, echo and crosstalk cancellation, and, when Trellis coding is used, sequence estimation. The sequence of code-groups assigned to `tx_symb_vector` is needed to perform echo and self near-end crosstalk cancellation.

The PMA Receive function uses the `scr_status` parameter and the state of the equalization, cancellation, and estimation functions to determine the quality of the receiver performance, and generates the `loc_rcvr_status` variable accordingly. The precise algorithm for generation of `loc_rcvr_status` is implementation dependent.

The PMA Receive function shall use the sequence of symbols received during PMA training to detect and correct for pair swaps and crossovers. The receiver pairs BI_DA, BI_DB, BI_DC, and BI_DD may be connected in any manner described in 6.4.2 to the corresponding transmit pairs. The PMA Receive function shall also detect and correct for polarity mismatches on any pairs and corrects for differential delay variations across the wire-pairs.

6.2.4 PHY Control Function

PHY Control generates the control actions that are needed to bring the PHY into the normal mode of operation. The PHY Control function shall operate according to the state diagram shown in figure 22. An example of the resulting link establishment process is shown in figure 21.

During Auto-Negotiation, PHY Control is in the DISABLE_TRANSMITTER state and the transmitters are disabled. When the Auto-Negotiation process asserts link_control=ENABLE, PHY Control enters the SILENT state.

Upon entering the SILENT state, maxwait_timer, minwait_timer, and loc_rcvr_timer are started, while transmitters are kept silent by setting tx_mode=SEND_Z. The Slave PHY performs most of its receiver convergence functions in the SILENT state, converging its decision feedback equalizer (DFE) and achieving descrambler synchronization when the Master PHY is in the TRAINING state.

In Master mode, PHY Control transitions to the TRAINING state upon expiration of the minwait_timer. In Slave mode, PHY Control transitions to the TRAINING state after the Slave PHY has converged its DFE, acquired timing, achieved descrambler synchronization (i.e., set scr_status=OK), the minwait_timer has expired, and the loc_rcvr_timer has not expired. If the loc_rcvr_timer expires, the Slave PHY remains in the SILENT state until a subsequent Auto-Negotiation is performed.

Upon entering the TRAINING state, the minwait_timer is started and transmission of a Type-1 PAM-2 training sequence is forced by setting tx_mode=SEND_T1. The Master PHY performs its receiver convergence functions in the TRAINING state, converging echo and NEXT cancellers while the Slave PHY is still in the SILENT state, and converging its DFE and achieving descrambler synchronization (i.e., set scr_status=OK) when the Slave PHY is in the TRAINING state. The Slave PHY completes the convergence of the adaptive filter parameters in the TRAINING state, with the convergence of the echo and NEXT cancellers, at which time it sets loc_rcvr_status=READY.

In Master mode, PHY Control transitions to the TRAINED state when the convergence of the adaptive filter parameters is completed (i.e., loc_rcvr_status=READY), the minwait_timer has expired, and the loc_rcvr_timer has not expired. If loc_rcvr_timer expires, the Master PHY remains in the TRAINING state until a subsequent Auto-Negotiation is performed. In Slave mode, PHY Control transitions to the TRAINED state when the convergence of the adaptive filter parameters is completed (i.e., loc_rcvr_status=READY) and the minwait_timer has expired.

Upon entering the TRAINED state, the minwait_timer is started and transmission of a Type-2 PAM-2 training sequence is forced by setting tx_mode=SEND_T2. The TRAINED state is used to explicitly signal to the link partner the completion of the convergence of its adaptive filter parameters (i.e., to signal that loc_rcvr_status=READY).

PHY Control transitions to the LINK_QUALIFICATION state when the PCS detects a Type-2 PAM-2 training sequence, at which time it sets rem_rcvr_status=READY, the minwait_timer has expired, and loc_rcvr_status=READY.

Upon entering the LINK_QUALIFICATION state, the linkeval_timer is started and transmission of Idle2 Ordered Sets is forced by setting tx_mode=SEND_I2. When in the LINK_QUALIFICATION state, both Master and Slave PHYs achieve PCS block synchronization (i.e., set pcs_status=OK) and performs SNR estimations to verify that the bit error rate requirements specified in the BER requirement bits of the Management Control register may be satisfied.

PHY Control transitions to the LINK_QUALIFIED state when PCS block synchronization is achieved (i.e., pcs_status=OK), the receiver has determined that it satisfies the bit error rate requirements specified in the BER requirement bits of the Management Control register, and the linkeval_timer has expired. PHY Control

sets `loc_rcvr_status=OK` when PCS block synchronization is achieved and the bit error rate is satisfactory. If the bit error rate is determined to be not satisfactory, the PHY remains in the `LINK_QUALIFICATION` state until a subsequent Auto-Negotiation is performed.

Upon entering the `LINK_QUALIFIED` state, `minwait_timer` is started and transmission of Idle3 Ordered Sets is forced by setting `tx_mode=SEND_I3`. The `LINK_QUALIFIED` state is used to explicitly signal to the link partner that PCS block synchronization is achieved and the bit error rate is satisfactory (i.e., to signal that `loc_rcvr_status=OK`).

PHY Control transitions to the `SEND_DATA` state from the `LINK_QUALIFIED` state when the PCS detects Idle3 Ordered Sets from the link partner (at which time it sets `rem_rcvr_status=OK`), `loc_rcvr_status` is set to `OK`, `pcs_status` is set to `OK`, and the `minwait_timer` has expired.

Upon entering the `SEND_DATA` state, the `maxwait_timer` is stopped and transmission of data or Idle3 Ordered Sets is forced by setting `tx_mode=SEND_N`. The `SEND_DATA` state corresponds to the normal mode of operation of an FC-BaseT PHY. In the `SEND_DATA` state the data (i.e., XGMII words) received from the host are sent to the link partner, otherwise Idle3 Ordered Sets are sent when no data is received from the host. Situations in which no data are received from the hosts include the case in which the host is not yet synchronized with the FC-BaseT PHY at the selected link operating speed.

When in the `SEND_DATA` state, if the `minwait_timer` has expired and `pcs_status` becomes different than `OK` or `loc_rcvr_status` becomes different than `OK` or `rem_rcvr_status` becomes different than `OK`, PHY Control transitions to the `SILENT` state. PHY Control remains in the `SILENT` state until the expiration of the `minwait_timer`. If the link partner was in the `SEND_DATA` state, this silent period causes the link partner's `rem_rcvr_status` to the value `INIT`, forcing the link partner's to enter the `SILENT` state. PHY Control then attempts to re-acquire communication with the same Auto-Negotiation parameters until the `maxwait_timer` expires. The Auto-Negotiation process begins again upon expiration of the `maxwait_timer`.

PHY Control may force the transmit scrambler state to be initialized to an arbitrary value by requesting the execution of the PCS Reset function defined in 5.2.

The Idle2 and Idle3 Ordered Sets are shown in table 14.

Table 14 – Idle2 and Idle3 Ordered Sets

Name	Ordered Set	Transcoded Representation (see 5.3.2)		
		S0	S1	S2
Idle2	K28.5 D7.0 D9.1 D9.1	1000000001b	01110010100b	01100101001b
Idle3	K28.5 D7.0 D9.5 D9.5	1000000001b	00111010100b	00110101001b

Master Actions		Slave Actions	
RX	TX	TX	RX
DISABLE_TRANSMITTER state		DISABLE_TRANSMITTER state	
loc_rcvr_status=INIT rem_rcvr_status=INIT			loc_rcvr_status=INIT rem_rcvr_status=INIT
SILENT state		SILENT state	
	SEND_Z		
TRAINING state		SEND_Z	On detecting PAM-2 signals: Converge equalizers Acquire descrambler synch (i.e., set scr_status=OK)
Converge echo/NEXT cancellers	SEND_T1		TRAINING state
On detecting PAM-2 signals: Converge equalizers Acquire descrambler synch (i.e., set scr_status=OK) loc_rcvr_status=READY		SEND_T1	Converge echo/NEXT cancellers loc_rcvr_status=READY
On PCS detecting T2 sequence: rem_rcvr_status=READY			TRAINED state
TRAINED state		SEND_T2	On PCS detecting T2 sequence: rem_rcvr_status=READY
	SEND_T2		
LINK_QUALIFICATION state		SEND_I2	On detecting PAM-8 signals: Achieve PCS synch (i.e., set pcs_status=OK) Ensure BER is satisfactory loc_rcvr_status=OK
On detecting PAM-8 signals: Achieve PCS synch (i.e., set pcs_status=OK) Ensure BER is satisfactory loc_rcvr_status=OK	SEND_I2		On detecting PAM-8 signals: Achieve PCS synch (i.e., set pcs_status=OK) Ensure BER is satisfactory loc_rcvr_status=OK
On PCS detecting Idle3: rem_rcvr_status=OK			LINK_QUALIFIED state
LINK_QUALIFIED state		SEND_I3	On PCS detecting Idle3: rem_rcvr_status=OK
	SEND_I3		
SEND_DATA state		SEND_DATA state	
	SEND_N	SEND_N	

time

Figure 21 – Example of Link Establishment

6.2.5 Link Monitor Function

Link Monitor determines the status of the underlying receive channel and communicates it via the variable `link_status`. Failure of the underlying receive channel typically causes the PMA's clients to suspend normal operation. The Link Monitor function shall operate according to the state diagram shown in figure 23.

Upon power on, reset, or release from power down, the Auto-Negotiation algorithm sends fast link pulses to signal its presence to a remote station. If the presence of a remote station is sensed through reception of fast link pulses, the Auto-Negotiation algorithm exchanges Auto-Negotiation information with the remote station. During this period, `link_status=FAIL` is asserted. If a remote FC-BaseT station is detected, the Auto-Negotiation algorithm permits full operation by setting `link_control=ENABLE`. As soon as reliable transmission is achieved, the variable `link_status=OK` is asserted, upon which further PHY operations may take place.

Link Monitor begins its operations in the `LINK_DOWN` state with the same conditions that cause PHY Control to enter the `DISABLE_TRANSMITTER` state. Link Monitor remains in the `LINK_DOWN` state until PHY Control enters the `SEND_DATA` state. Link Monitor transitions to the `LINK_UP` state with the same conditions that cause PHY Control to enter the `SEND_DATA` state. Link Monitor transitions to state `LINK_DOWN` if the `maxwait_timer` has expired and `pcs_status≠OK` or `loc_rcvr_status≠OK` or `rem_rcvr_status≠OK`.

6.2.6 Clock Recovery Function

The Clock Recovery function shall provide clocks suitable for signal sampling on each line so that the symbol-error rate indicated in 6.2.3 is achieved. The received clock signal shall be stable and ready for use when training has been completed (`loc_rcvr_status=OK`). The received clock signal is supplied to the PMA Transmit function by the variable `received_clock`.

The Clock Recovery function may provide independent clock phases for sampling the signals on each of the four pairs.

6.2.7 State Diagrams

State diagrams are provided for the PHY Control and the Link Monitor functions. The variables used in the state diagrams are shown in table 15.

Table 15 – PHY Control and Link Monitor State Variables

Variable	Description
<code>tx_mode</code>	See 4.6.2.2.
<code>config</code>	See 4.6.2.3.
<code>scr_status</code>	See 4.6.2.6.
<code>pcs_status</code>	See 4.6.2.7.
<code>loc_rcvr_status</code>	See 4.6.2.8.
<code>rem_rcvr_status</code>	See 4.6.2.9.
<code>link_control</code>	See IEEE 802.3-2005, clause 28.2.6.2.
<code>link_status</code>	See IEEE 802.3-2005, clause 28.2.6.1.

The timers used in the state diagrams are shown in table 16.

Table 16 – PHY Control and Link Monitor Timers

Timer	Description
maxwait_timer	Used to limit the amount of time during which an attempt is made to establish link connectivity via the PHY Control function described in 6.2.4. This timer shall expire $(2\ 000 \pm 2)$ ms after being started. This timer is used in both the PHY Control and Link Monitor state diagrams. The maxwait_timer is tested by the Link Monitor function to force link_status to be set to FAIL if the timer expires and pcs_status \neq OK or loc_rcvr_status \neq OK or rem_rcvr_status \neq OK.
minwait_timer	Used to determine the minimum amount of time the PHY Control function dwells in most of its states. This timer shall expire $(1 \pm 0,1)$ ms after being started.
loc_rcvr_timer	Used to ensure that the startup time budget is divided appropriately, so that both the Master and Slave PHYs have a defined minimum useful training period to converge their respective receivers. For a Master PHY this timer shall expire $(1\ 884 \pm 2)$ ms after being started. For a Slave PHY this timer shall expire (693 ± 2) ms after being started.
linkeval_timer	Used to ensure a minimum period for which Idle2 Ordered Sets are transmitted, in order to achieve PCS block synchronization and to perform BER estimations. This timer shall expire (100 ± 2) ms after being started.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

The PHY Control state diagram is shown in figure 22.

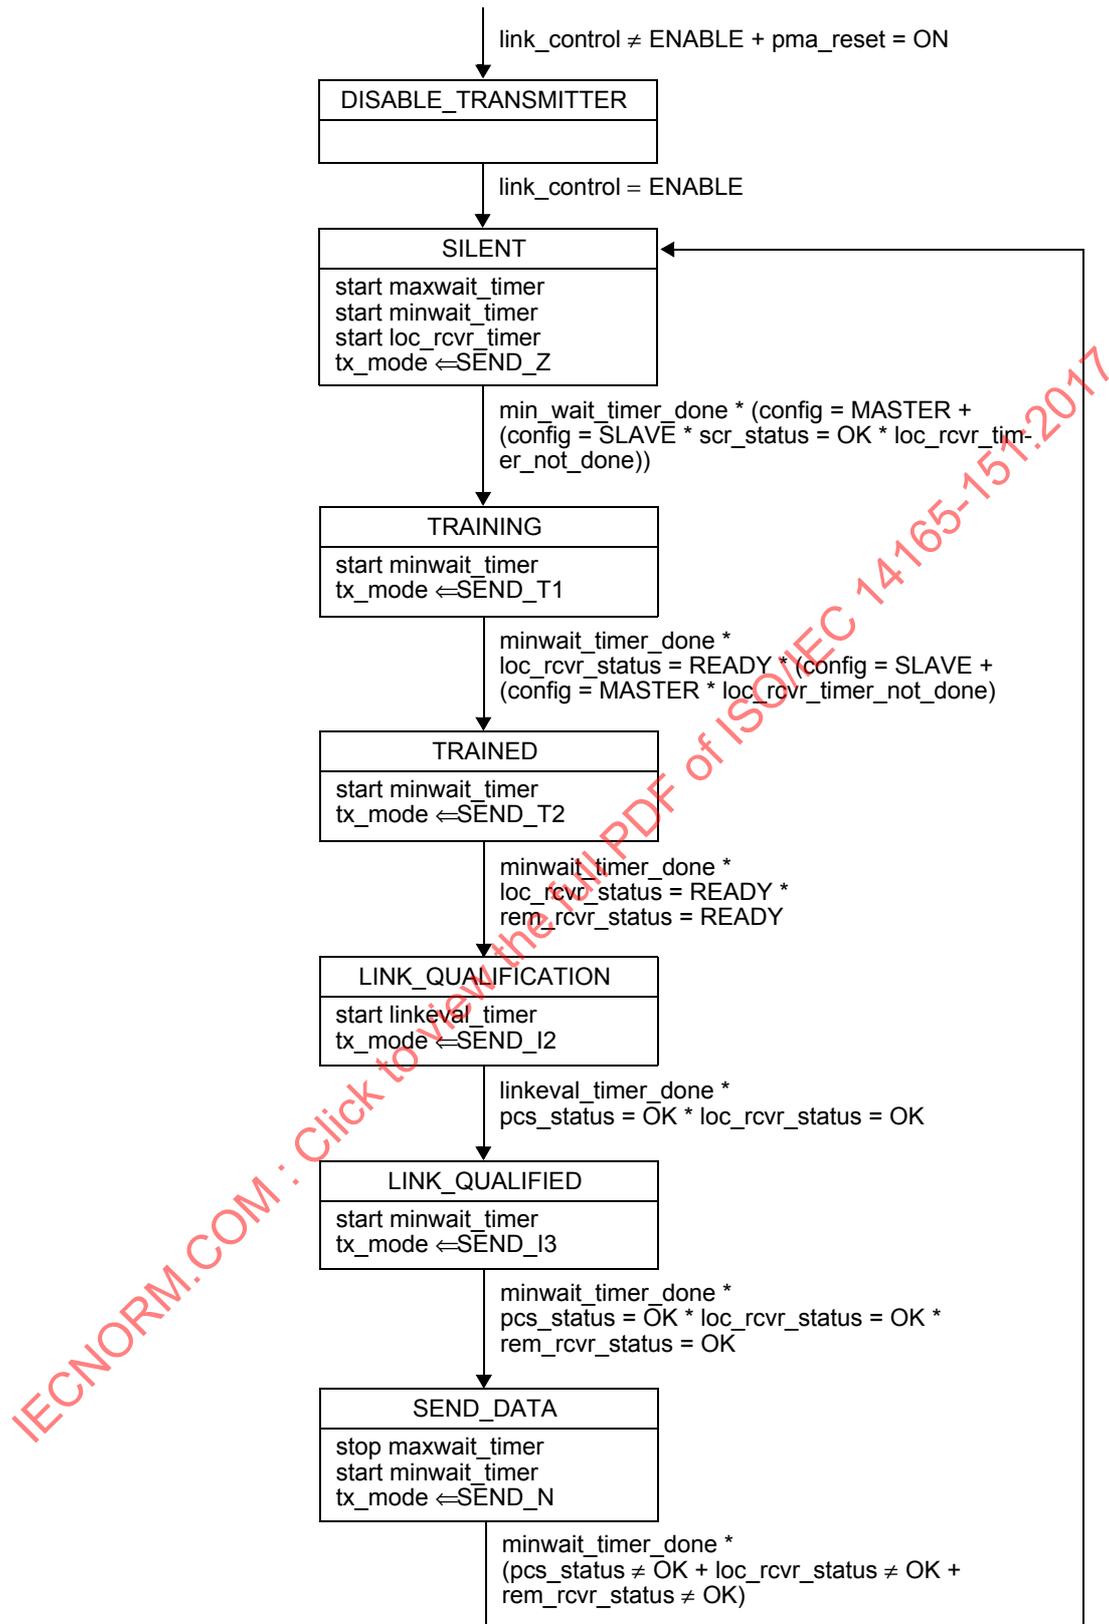


Figure 22 – PHY Control State Diagram

The Link Monitor state diagram is shown in figure 23.

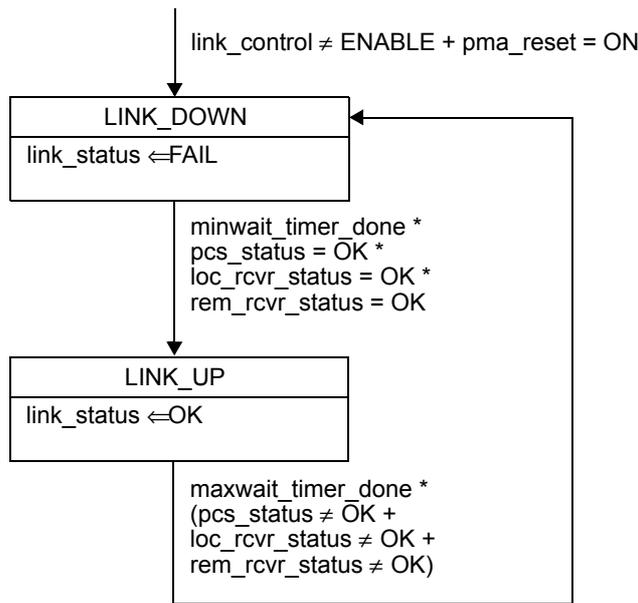


Figure 23 – Link Monitor State Diagram

6.3 PMA Electrical Specification

6.3.1 Isolation and EMC Requirements

An FC-BaseT PHY shall provide electrical isolation as defined in IEEE 802.3-2005, clause 40.6.1.1.

Systems containing FC-BaseT PHYs should be able to meet the EMC susceptibility test defined in EN61000-4-6 (i.e., conducted immunity up to 80 MHz) and EN61000-4-3 (i.e., radiated immunity above 80 MHz) to levels specified in CISPR 24 and/or EN55024. Systems containing FC-BaseT PHYs should be able to meet CISPR/FCC Class A EMC emissions requirements.

6.3.2 Test Modes

6.3.2.1 Overview

The test modes described in this subclause shall be provided to allow for testing of the transmitter waveform, transmitter distortion, transmitted jitter, and transmitter droop.

These test modes shall only change the data symbols provided to the transmitter circuitry and not alter the electrical and jitter characteristics of the transmitter and receiver from those of normal operation. These modes shall be enabled by setting bits 2 .. 0 of the Test register (i.e., management register 16, see table

35), as shown in table 17.

Table 17 – Management Register Settings for Test Modes

Value of bits 2 .. 0 of Register 16	Mode
000b	Normal operation
001b	Test mode 1 - Transmit droop test mode
010b	Test mode 2 - Transmit jitter test in MASTER mode
011b	Test mode 3 - Transmit jitter test in SLAVE mode
100b	Test mode 4 - Transmitter distortion test
101b	Test mode 5 - Normal operation with no power backoff, used for the PSD mask and power level test
110b	Reserved, operations not defined
111b	Reserved, operations not defined

Test mode 1 is for testing transmitter droop. When test mode 1 is enabled, the PHY shall transmit sixty-four +7 symbols followed by sixty-four -7 symbols. This sequence is repeated continuously.

When test mode 2 is enabled, the PHY shall transmit the data symbol sequence {+7, -7} repeatedly on all channels. The transmitter shall time the transmitted symbols from a symbol rate clock in the Master timing mode.

When test mode 3 is enabled, the PHY shall transmit the data sequence {+7, -7} repeatedly on all channels. The transmitter shall time the transmitted symbols from a symbol rate clock in the Slave timing mode.

When test mode 4 is enabled, the PHY shall transmit at its maximum supported power the sequence of symbols generated by the scrambler generator polynomial $g_{s1} = 1 + x^9 + x^{11}$. The maximum-length shift register used to generate the sequences defined by this polynomial shall be updated once per symbol interval. The bits stored in the shift register delay line at a particular time n are denoted by $Scr_n[0 .. 10]$. At each symbol period the shift register is advanced by one bit and one new bit represented by $Scr_n[0]$ is generated. Bits $Scr_n[8]$ and $Scr_n[10]$ are XORed together to generate the next $Scr_n[0]$ bit. The PAM-8 symbols to be transmitted shall be obtained as specified in table 11 from the bit sequences x_{0n} , x_{1n} , and x_{2n} , generated from the following combinations of the scrambler bits:

$$X_{0n} = Scr_n[0]$$

$$X_{1n} = Scr_n[1] \oplus Scr_n[4]$$

$$X_{2n} = Scr_n[2] \oplus Scr_n[4]$$

The PAM-8 symbol sequence shall be presented simultaneously to all transmitters. The transmitter shall time the transmitted symbols from a symbol clock in the Master timing mode.

Test mode 5 is for checking whether the transmitter is compliant with the transmit PSD mask and the transmit power level. When test mode 5 is enabled, the PHY shall transmit as in normal operation at its maximum supported power. The sequence of 4D PAM-8 symbols to be transmitted in this mode shall be constructed from the Master side-stream scrambler, shown in figure 10. The initial state of the scrambler shall be any arbitrary non-zero value. The 11-bit scrambling vector $Q_n[0 .. 10]$ (see 5.3.4) shall be Schläfli lattice encoded as shown in figure 12 to form four 3-bit values A_n , B_n , C_n , D_n . These values shall be mapped to PAM-8 symbols as specified in table 11 and transmitted while in test mode 5.

6.3.2.2 Test Fixtures

The fixture shown in figure 24, or an equivalent one, shall be used for measuring the transmitter specifications defined in 6.3.3 for droop, jitter and PSD mask.

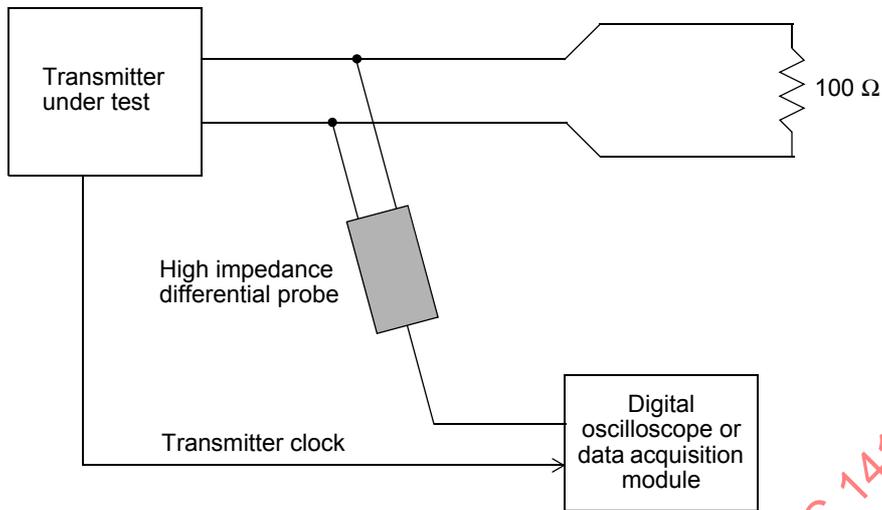


Figure 24 – Transmitter Test Fixture A

The fixture shown in figure 25, or an equivalent one, shall be used for measuring the transmitter specifications defined in 6.3.3 for distortion.

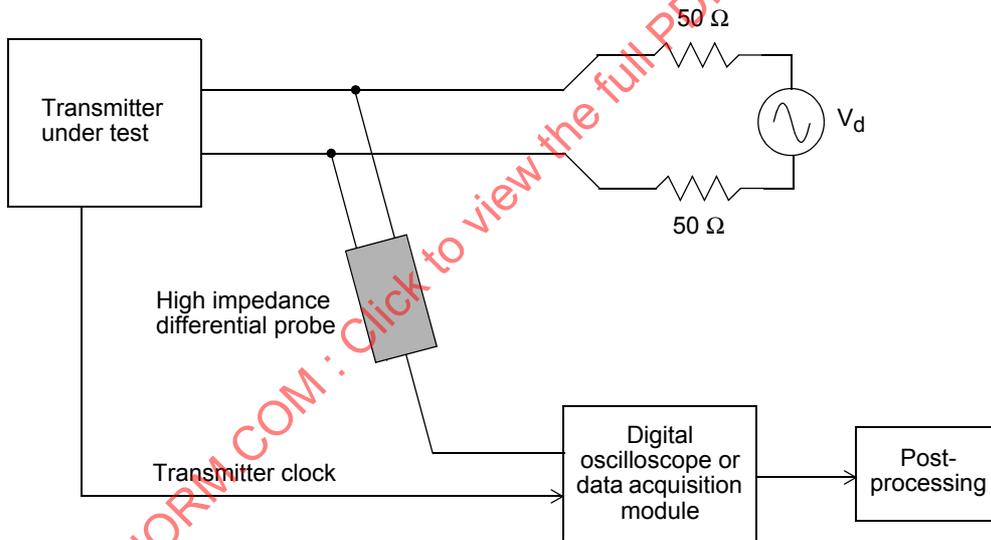


Figure 25 – Transmitter Test Fixture B

To allow for measurement of transmitted jitter in Master and Slave modes, the PHY shall provide access to the transmit symbol clock that times the transmitted symbols. The PHY shall provide a means to enable this clock output if it is not normally enabled.

The disturbing signal V_d shall have a differential amplitude of 4 Vpp and a frequency as shown in table 18.

Table 18 – Disturbing Signal Frequency

Mode	Frequency
4GFC-BaseT	54 MHz
2GFC-BaseT	27 MHz
1GFC-BaseT	13,5 MHz

The generator of the disturbing signal shall have sufficient linearity and range in order to not introduce any appreciable distortion when connected to the transmitter output.

6.3.3 Transmitter Electrical Specifications

6.3.3.1 Overview

The PMA shall operate with AC coupling to the MDI. Where a load is not specified, the transmitter shall meet the requirements of this subclause with a 100 ohm resistive differential load connected to each transmitter output.

6.3.3.2 Transmitter Output Droop

When the transmitter is operating in test mode 1, using the transmitter test fixture A, the magnitude of both the positive and negative droop measured with respect to an initial peak value after the zero crossing and final value before the next zero crossing shall be as shown in table 19.

Table 19 – Droop Requirements

Mode	Droop
4GFC-BaseT	<12%
2GFC-BaseT	<23%
1GFC-BaseT	<39%

Figure 26 shows an example of the output of the test mode 1. Symbol periods are shown in table 6.

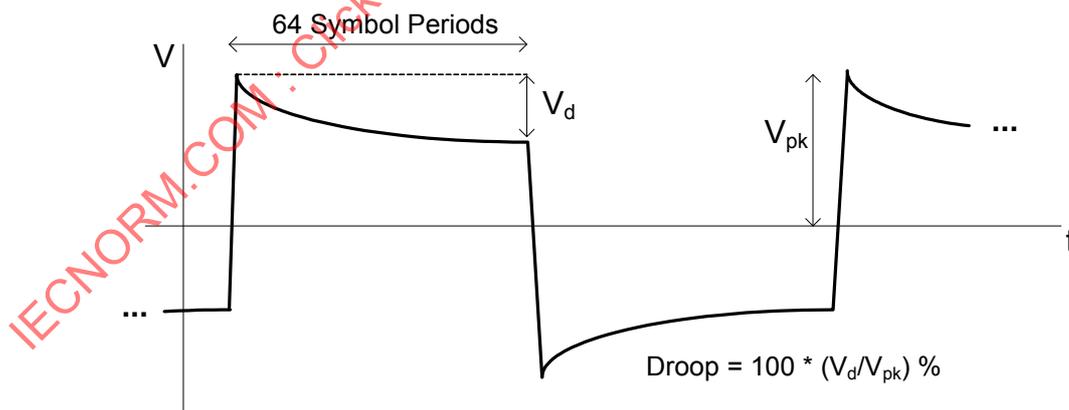


Figure 26 – Test Mode 1 Output (not to scale)

6.3.3.3 Transmitter Distortion

The transmitter peak distortion is determined by sampling the differential signal output with the symbol rate clock at an arbitrary phase and processing a block of any 2 047 consecutive samples with the MATLAB code given in IEEE 802.3-2005, clause 40.6.1.2.4. Optionally, the data may first be processed by a digital filter meant to model the receive filter (including high pass). The MATLAB code removes the disturbing signal from the measurement and measures the peak distortion.

When in test mode 4 and observing the differential signal output at the MDI using transmitter test fixture B, for each pair, with no intervening cable, the peak distortion as defined in this subclause shall be less than 10 mV in all supported modes.

NOTE 6 - 10 mV are sufficient for interoperability with a far end receiver but may not be sufficient for local echo cancellation. Local echo cancellation is a receiver issue and is left to the implementor.

6.3.3.4 Transmitter Timing Jitter

When in test modes 2 and 3, the PHY transmits the levels {+7, -7} continuously at 5 dBm. Jitter shall be measured over an integration time interval of not less than 1 ms. RMS jitter is defined as the root mean square period difference from the average period ($T - T_{avg}$), accumulated over a sample size of at least 1 ms in duration, as defined by the following equation:

$$\tau_j = \sqrt{\frac{\sum (T - T_{avg})^2}{\text{Sample Size}}}$$

The RMS period jitter at either the Master or Slave MDI output shall be less than 20 ps in all supported modes.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

6.3.3.5 Transmitter Power Spectral Density (PSD) and Power Level

In test mode 5, the average transmit power shall be $(5 \pm 0,25)$ dBm.

For 4GFC-BaseT the power spectral density of the transmitter, measured into a 100Ω load, shall be between the following upper and lower masks, where f is in MHz:

$$\text{Upper PSD}(f) \leq \begin{cases} -76 \text{ dBm/Hz} & 0 < f \leq 100 \text{ MHz} \\ -76 - (f - 100) \times \frac{37}{750} \text{ dBm/Hz} & 100 \text{ MHz} < f \leq 850 \text{ MHz} \\ -113 \text{ dBm/Hz} & 850 \text{ MHz} < f \end{cases}$$

$$\text{Lower PSD}(f) \geq \begin{cases} -79 \text{ dBm/Hz} & 300 \text{ kHz} < f \leq 40 \text{ MHz} \\ -79 - \frac{(f - 40)}{30} \text{ dBm/Hz} & 40 \text{ MHz} < f \leq 160 \text{ MHz} \end{cases}$$

The 4GFC-BaseT upper and lower masks are shown in figure 27, with an example of transmitter PSD at 5 dBm.

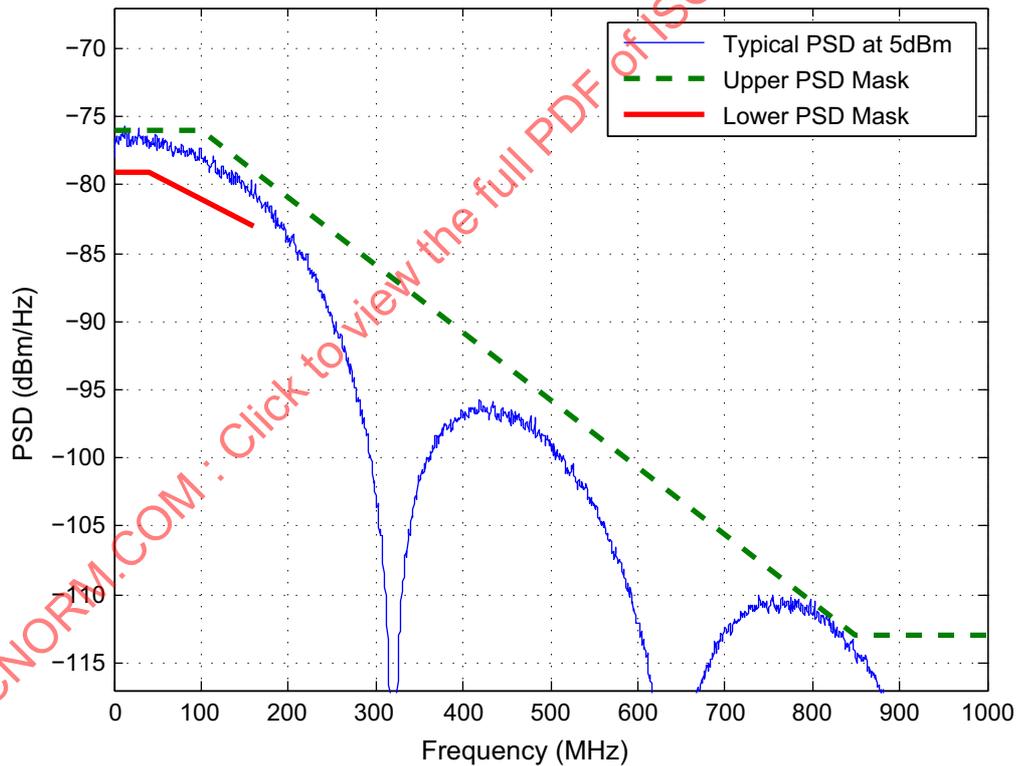


Figure 27 – 4GFC-BaseT Transmitter PSD at 5dBm

For 2GFC-BaseT the power spectral density of the transmitter, measured into a 100 Ω load, shall be between the following upper and lower masks, where f is in MHz:

$$\text{Upper PSD}(f) \leq \begin{cases} -73 \text{ dBm/Hz} & 0 < f \leq 50 \text{ MHz} \\ -73 - \frac{(f-50)}{10} \text{ dBm/Hz} & 50 \text{ MHz} < f \leq 450 \text{ MHz} \\ -113 \text{ dBm/Hz} & 450 \text{ MHz} < f \end{cases}$$

$$\text{Lower PSD}(f) \geq \begin{cases} -76 \text{ dBm/Hz} & 300 \text{ kHz} < f \leq 20 \text{ MHz} \\ -76 - \frac{(f-20)}{15} \text{ dBm/Hz} & 20 \text{ MHz} < f \leq 80 \text{ MHz} \end{cases}$$

The 2GFC-BaseT upper and lower masks are shown in figure 28, with an example of transmitter PSD at 5 dBm.

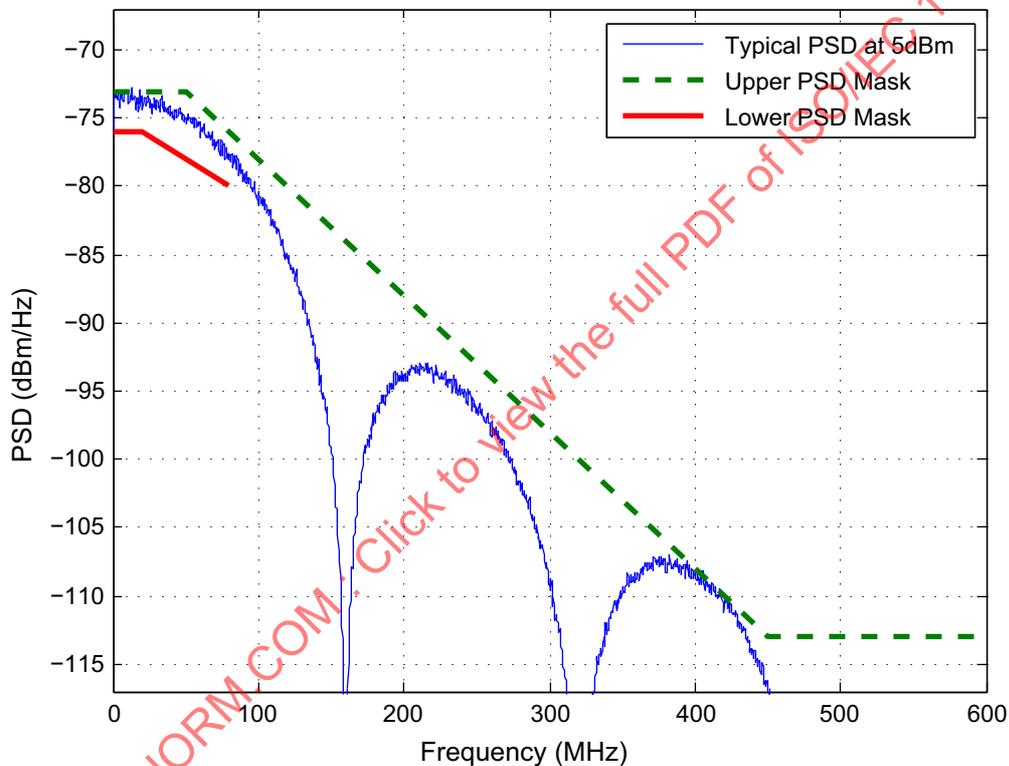


Figure 28 – 2GFC-BaseT Transmitter PSD at 5dBm

For 1GFC-BaseT the power spectral density of the transmitter, measured into a 100 Ω load, shall be between the following upper and lower masks, where f is in MHz:

$$\text{Upper PSD}(f) \leq \begin{cases} -70 \text{ dBm/Hz} & 0 < f \leq 25 \text{ MHz} \\ -70 - (f - 25) \times \frac{43}{225} \text{ dBm/Hz} & 25 \text{ MHz} < f \leq 250 \text{ MHz} \\ -113 \text{ dBm/Hz} & 250 \text{ MHz} < f \end{cases}$$

$$\text{Lower PSD}(f) \geq \begin{cases} -73 \text{ dBm/Hz} & 300 \text{ kHz} < f \leq 10 \text{ MHz} \\ -73 - (f - 10) \times \frac{2}{15} \text{ dBm/Hz} & 10 \text{ MHz} < f \leq 40 \text{ MHz} \end{cases}$$

The 1GFC-BaseT upper and lower masks are shown in figure 29, with an example of transmitter PSD at 5 dBm.

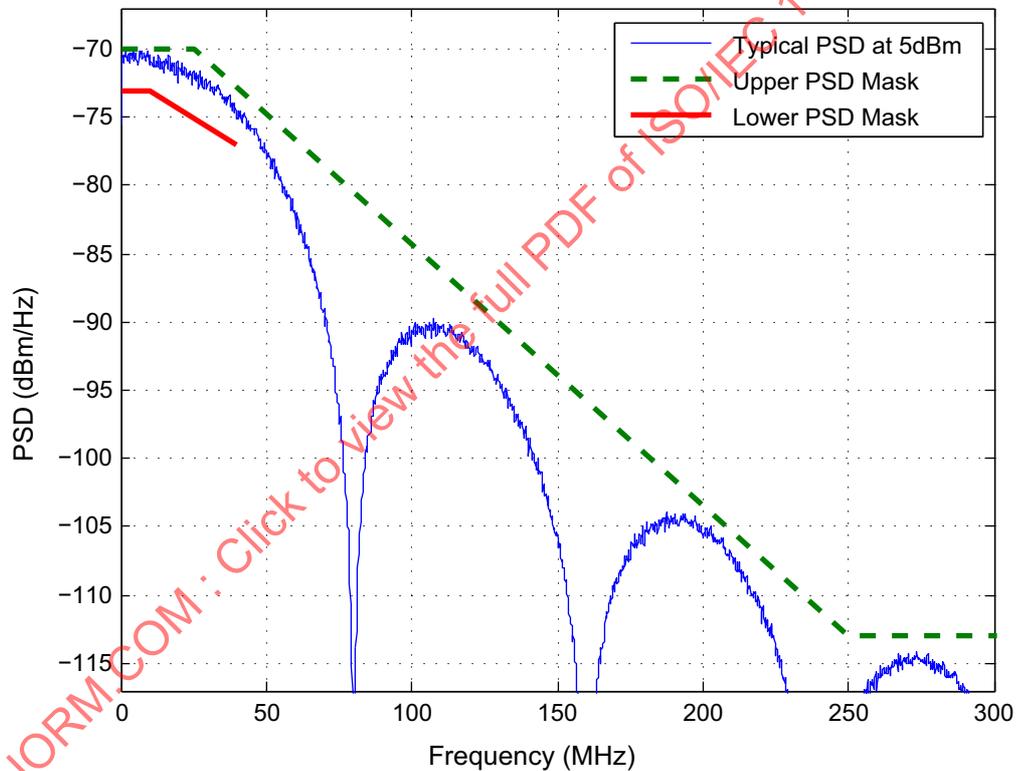


Figure 29 – 1GFC-BaseT Transmitter PSD at 5dBm

6.3.3.6 Transmitter Power Schedule

The nominal power of the PMA transmitted signal shall be varied with the cable length according to the transmitter power schedule shown in table 20, that shows the maximum transmitter power.

Table 20 – Transmitter Power Schedule

Mode	Maximum Transmitter Power	Cable Length
4GFC-BaseT ^a	2,50 dBm	0m ≤ L < 10m
	3,00 dBm	10m ≤ L < 20m
	4,00 dBm	20m ≤ L < 40m
	5,25 dBm	40m ≤ L
2GFC-BaseT	2,50 dBm	0m ≤ L < 30m
	4,00 dBm	30m ≤ L < 50m
	5,25 dBm	50m ≤ L
1GFC-BaseT	2,50 dBm	0m ≤ L < 60m
	5,25 dBm	60m ≤ L

^a Transmitting at the maximum transmitter power in 4GFC-BaseT mode may create some interference issues in certain situations. See Annex C for a recommended power schedule.

6.3.3.7 Transmit Clock Frequency

The symbol transmission rate on each pair of the Master PHY shall be within the range shown in table 21.

Table 21 – Transmitter Frequency Requirements

Mode	Frequency
4GFC-BaseT	318,75 MHz ± 50ppm
2GFC-BaseT	159,375 MHz ± 50ppm
1GFC-BaseT	79,6 875 MHz ± 50ppm

6.3.4 Receiver Electrical Specifications

6.3.4.1 Overview

The PMA shall provide the PMA Receive function specified in 6.2.3 in accordance with the electrical specifications of 6.3.4 and using the interconnecting cabling and hardware that is within the limits specified in 6.5.

6.3.4.2 Receiver Differential Input Signals

Differential signals received at the MDI that were transmitted from a remote transmitter within the specifications of 6.3.3 and have passed through a link having the characteristics specified in 6.5 shall be received with a BER less than the one specified by the BER requirement bits of the Management Control Register (see 9.3) and sent to the PCS. This specification shall be satisfied by a frame error ratio for the FC-BaseT speeds as follows:

- 4GFC-BaseT: frame error ratio less than 4 096 times the specified BER, for 512 byte frames;
- 2GFC-BaseT: frame error ratio less than 2 048 times the specified BER, for 256 byte frames; and

- c) 1GFC-BaseT: frame error ratio less than 1 024 times the specified BER, for 128 byte frames.

6.3.4.3 Receiver Frequency Tolerance

The receiver shall properly receive incoming data with a symbol rate within the range shown in table 22.

Table 22 – Receiver Frequency Requirements

Mode	Frequency
4GFC-BaseT	318,75 MHz \pm 50ppm
2GFC-BaseT	159,375 MHz \pm 50ppm
1GFC-BaseT	79,6 875 MHz \pm 50ppm

6.3.4.4 Alien Crosstalk Noise Rejection

While receiving data from a transmitter compliant with the specifications in 6.3.3 through a link segment compliant with the specifications in 6.5, a receiver shall operate with a BER lower than the one specified by the BER requirement bits of the Management Control Register (see 9.3) with four noise sources at the specified levels representing alien crosstalk, one connected to each of the four pairs.

Independent noise sources should be injected into each MDI input using couplers that do not significantly alter the link segment characteristics. Each noise source shall have a flat noise spectrum, with 3 dB bandwidth of at least 2 MHz to 250 MHz, and a power spectral density such that at the MDI port of the devices under test the power spectral density of the injected noise shall be as shown in table 23.

Table 23 – Alien Noise Requirements

Mode	Alien Noise PSD
4GFC-BaseT	-129 dBm/Hz
2GFC-BaseT	-123 dBm/Hz
1GFC-BaseT	-120 dBm/Hz

A flat noise source is chosen to model the sum of all alien noise sources (see IEEE 802.3an-2006, clause 55.5.4.4).

6.4 MDI Specification

6.4.1 Overview

This subclause defines the FC-BaseT MDI. Each FC-BaseT device shall conform with the mated connector requirements to allow interoperability within an FC-BaseT environment.

6.4.2 MDI Mechanical Specification

The mechanical interface to the balanced cabling used by FC-BaseT shall be an eight-pin connector meeting the requirements of:

- a) IEC 60603-7-2 or IEC 60603-7-3 for 1GFC-BaseT and 2GFC-BaseT; and
- b) IEC 60603-7-4 or IEC 60603-7-5 for 4GFC-BaseT.

The jack connector shall be used on the FC-BaseT PHY and the plug connector shall be used on the balanced cabling. These connectors are informatively shown in figure 30 and figure 31.

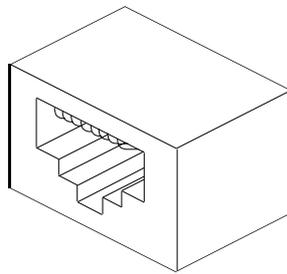


Figure 30 – MDI Jack Connector

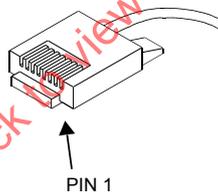


Figure 31 – Balanced Cabling Plug Connector

The assignment of PMA signals to MDI and MDI-X (i.e., crossover MDI, see 6.4.3) contacts for the FC-BaseT PHY is shown in table 24.

Table 24 – Assignment of PMA Signals to MDI and MDI-X Contacts

Contact	PMA Signal in MDI	PMA Signal in MDI-X
1	BI_DA+	BI_DB+
2	BI_DA–	BI_DB–
3	BI_DB+	BI_DA+
4	BI_DC+	BI_DD+
5	BI_DC–	BI_DD–
6	BI_DB–	BI_DA–
7	BI_DD+	BI_DC+
8	BI_DD–	BI_DC–

6.4.3 Automatic MDI/MDI-X Configuration

Automatic MDI/MDI-X configuration is intended to eliminate the need for crossover cables between similar devices. Automatic MDI/MDI-X configuration is required for FC-BaseT devices and shall comply with IEEE 802.3-2005, clauses 40.4.4.1 and 40.4.4.2.

6.4.4 MDI Electrical Specification

The FC-BaseT MDI electrical requirements when the MDI jack connector is mated with a balanced cabling plug connector are defined by the references shown in table 25.

Table 25 – FC-BaseT MDI Electrical Requirements

FC Speed	FEXT	Return Loss	Impedance Balance
4GFC-BaseT	IEEE 802.3an-2006, Clause 55.8.2 ^a	IEEE 802.3an-2006, Clause 55.8.2.1 ^a	IEEE 802.3an-2006, Clause 55.8.2.2 ^a
2GFC-BaseT	IEEE 802.3-2005, Clause 40.8.3	IEEE 802.3-2005, Clause 40.8.3.1	IEEE 802.3-2005, Clause 40.8.3.2
1GFC-BaseT	IEEE 802.3-2005, Clause 40.8.3	IEEE 802.3-2005, Clause 40.8.3.1	IEEE 802.3-2005, Clause 40.8.3.2

^a The requirements for 4GFC-BaseT are limited to a maximum frequency of 250 MHz.

6.4.5 MDI Fault tolerance

The FC-BaseT MDI shall meet the MDI fault tolerance requirements defined in IEEE 802.3-2005, clause 40.8.3.4.

6.5 Link Segment Characteristics

6.5.1 Overview

The link distances specified in table 3 are based on insuring interoperability between FC-BaseT PHYs. It is the implementor's responsibility to ensure that the cable assemblies are within the operating limits defined in this subclause. Greater link distances may be obtained by specifically engineering a link based on

knowledge of the technology characteristics and conditions under which the link is installed and operated. However, these link distance extensions are outside the scope of this standard.

6.5.2 FC-BaseT Link Topology

An FC-BaseT link consists of a horizontal cable component, two connection end points, and two patch cords connected to the FC-BaseT PHYs. The FC-BaseT link is between the connection points shown in figure 32 and include the permanent link as defined in ISO/IEC 11801:2002 and TIA/EIA-568-B.2-2001.

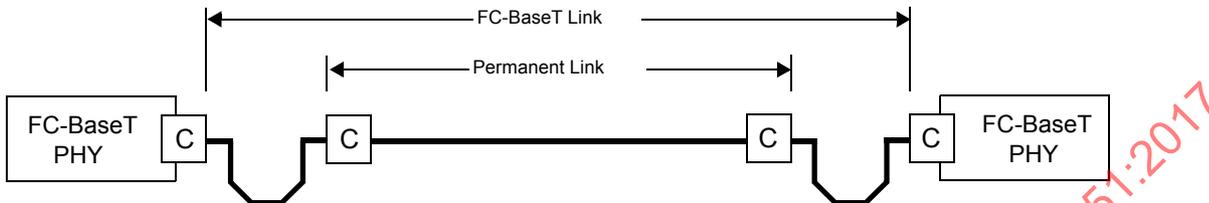


Figure 32 – FC-BaseT Link Topology

6.5.3 FC-BaseT Cable Plant Requirements

The FC-BaseT cable plant requirements for the different cable types are defined by the references shown in table 26.

Table 26 – FC-BaseT Cable Plant Requirements

Cable Type	Internal Impairments (Insertion Loss, NEXT, ELFEXT, Return Loss, Delay, Delay Skew)	Alien Crosstalk (Alien NEXT and Alien ELFEXT ^c)
Class D / Category 5e	ISO/IEC 11801:2002, Annex A ^a	IEEE 802.3-2005, Clause 40.7.6 ^d
Class E / Category 6	ISO/IEC 11801:2002, Annex A ^a	ISO/IEC TR 24750 or TIA/EIA TSB-155 ^e
Class E _A / Category 6a	ISO/IEC 11801:2002+A1:2008 or TIA/EIA-568-B.2-10 ^b	ISO/IEC 11801:2002+A1:2008 or TIA/EIA-568-B.2-10

^a For determining insertion loss the equations in Annex A are used. The length of the two patch cords times 1,2 shall be included in the length calculation in Annex A. The 20% increase is to account for the increased insertion loss of patch cords. The FC-BaseT link should include two connectors in the insertion loss calculation.

^b The length used in these references assume a 100 meter maximum length.

^c The Alien crosstalk is specified as a power sum of all disturbers. In ISO/IEC the term PS AACR-F is used for Power Sum Alien ELFEXT.

^d The noise specified includes both Alien NEXT and Alien FEXT.

^e The Alien crosstalk may require mitigating procedures to meet the requirements defined in these references.

7 Elasticity FIFO

7.1 Overview

FC-BaseT uses Elasticity FIFOs (E-FIFO) to solve potential clock skews between two connected PHYs and to retime the data stream in order to be compliant with the FC-PI-2 jitter specification. The E-FIFO performs speed matching between the incoming word stream and a local reference clock that may differ by ± 100 ppm. The operations of the Elasticity FIFO shall be compliant with FC-FS-2 and FC-AL-2. Figure 33 shows the Elasticity FIFOs in an FC-BaseT PHY.

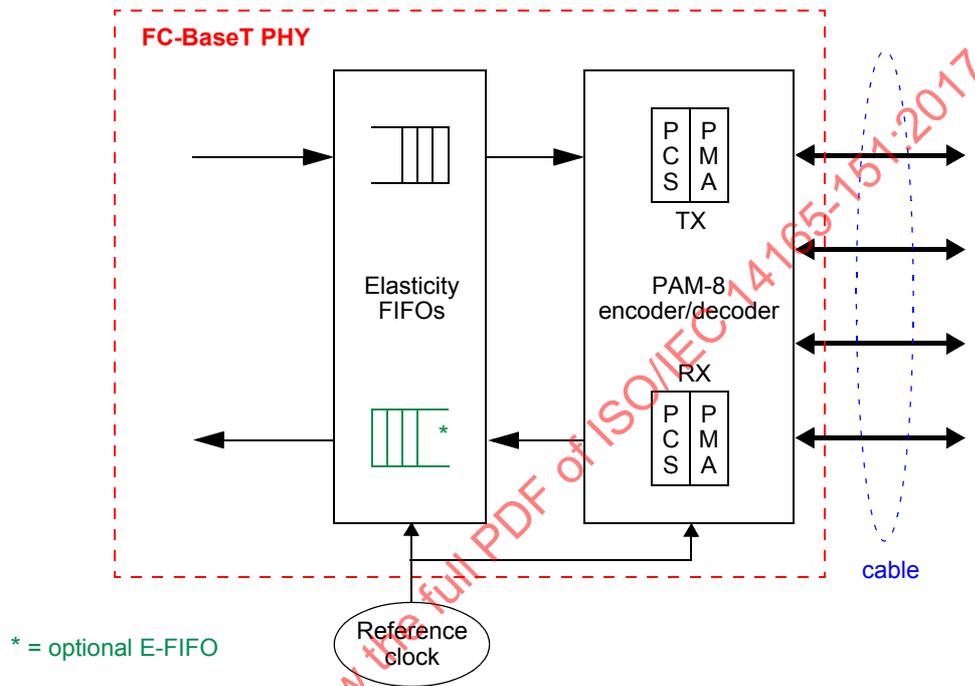


Figure 33 – FC-BaseT elasticity FIFO

The E-FIFOs are driven by a local reference clock in order to perform speed matching and reset the jitter. Speed matching is achieved by insertion or deletion of FC Idles or of other FC Ordered Sets (see 7.2). The E-FIFO on the FC-BaseT PHY receiving path, denoted with an asterisk in figure 33, is not required for speed matching but may be required for jitter compliance.

7.2 Ordered sets processing

An Elasticity FIFO recognizes and processes the Ordered Sets shown in table 27 (see FC-FS-2) as Extended Fill Words. Speed matching between the incoming word stream and a local reference clock is performed by inserting or removing Extended Fill Words as needed.

Table 27 – Extended Fill Words

Abbreviation	Description
Idle	Idle
ARBff	Arbitrate
ARByx	Arbitrate
ARB(val)	Arbitrate
NOS	Not_Operational
OLS	Offline
LR	Link_Reset
LRR	Link_Reset_Response
LIP(F7,F7)	Loop Initialization - F7,F7
LIP(F8,F7)	Loop Initialization - F8,F7
LIP(F7,x)	Loop Initialization - F7,x
LIP(F8,x)	Loop Initialization - F8,x
LIPyx	Loop Initialization - reset
LIPfx	Loop Initialization - reset all
LIPba	Loop Initialization - reserved LIPba
LPByx	Loop Port Bypass
LPBfx	Loop Port Bypass all
LPEyx	Loop Port Enable
LPEfx	Loop Port Enable all

Figure 34 shows an example of Elasticity FIFO implementation.

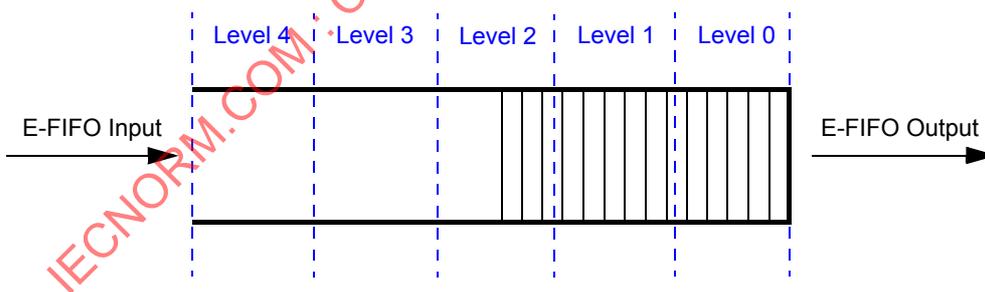


Figure 34 – Example of E-FIFO implementation

This implementation has the following 5 states, depending on how full the buffer is:

- a) Level 0 - High Priority Insertion Pending;

- b) Level 1 - Low Priority Insertion Pending;
- c) Level 2 - Quiescent;
- d) Level 3 - Low Priority Deletion Pending; and
- e) Level 4 - High Priority Deletion Pending.

When the E-FIFO operates in Level 2, no action is needed. When the E-FIFO operates in Level 1 or Level 3, it respectively inserts or removes FC Idles. When the E-FIFO operates in Level 0 or Level 4, it respectively inserts or removes Extended Fill Words (see table 27). The insertion or removal of Extended Fill Words shall not corrupt the FC data traffic.

Different E-FIFO implementations are allowed by this standard if they meet the requirements of FC-FS-2 and FC-AL-2.

7.3 Clock skew compensation

Symbols on an FC-BaseT link are transmitted and received over the same wire pairs in a synchronous manner. During Auto-Negotiation (see clause 8) one of the two connected FC-BaseT PHY becomes the Master PHY and the other one the Slave PHY. The Master PHY sets the clock signaling on the cable, the Slave PHY recovers the clock information from the signals on the cable. However, the frequency with which an FC-BaseT PHY transmit 33B blocks over the cable may differ from the frequency with which it receives the XGMII words to be transmitted by ± 100 ppm. Figure 35 shows how the Elasticity FIFO solves this issue.

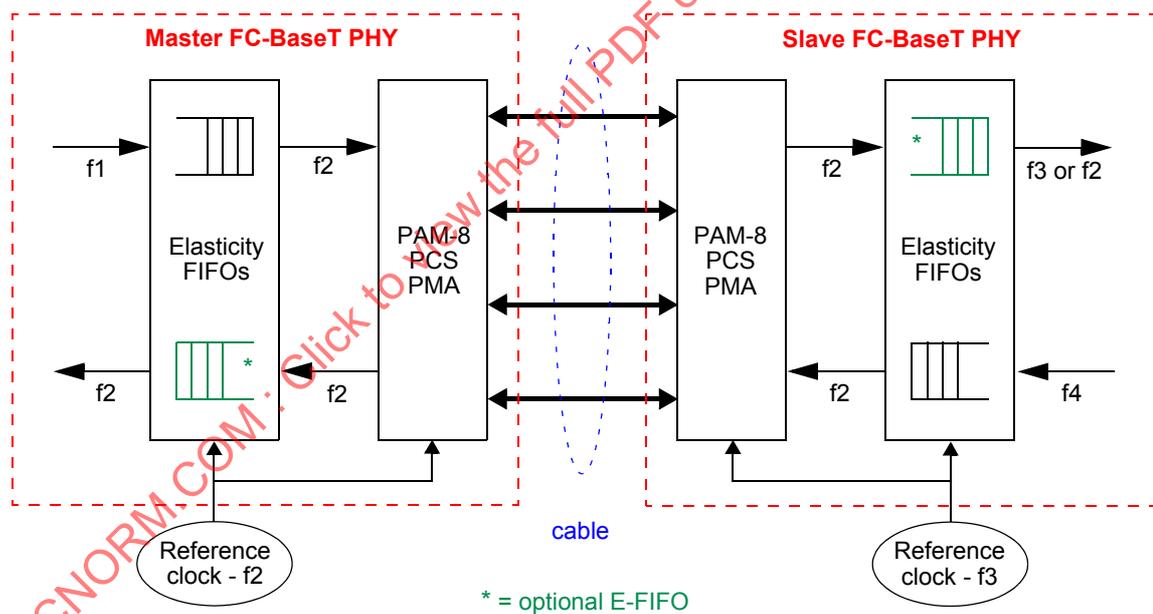


Figure 35 – FC-BaseT clock skew compensation

The Master FC-BaseT PHY transmits 33B blocks to and receives 33B blocks from the cable with the frequency f_2 of its reference clock. On the transmission path of the Master PHY the E-FIFO absorbs possible differences between the reference clock frequency f_2 and the frequency f_1 with which the XGMII words are sent to the E-FIFO after the 10b decoding. On the receiving path of the Master PHY the XGMII

words arrive with the reference clock frequency f_2 . If an E-FIFO is present on the receiving path, the received XGMII words are re-timed to the frequency f_2 of the local reference clock.

The Slave FC-BaseT PHY transmit 33B blocks to and receives 33B blocks from the cable with the frequency f_2 recovered from the signals on the cable. On the transmission path of the Slave PHY the E-FIFO absorbs possible differences between the frequency f_2 and the frequency f_4 with which the XGMII words are sent to the E-FIFO after the 10b decoding. On the receiving path of the Slave PHY the XGMII words arrive with the reference clock frequency f_2 . If an E-FIFO is present on the receiving path, the received XGMII words are re-timed to the frequency f_3 of the local reference clock, otherwise the received XGMII words have the frequency f_2 .

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

8 PHY startup procedure

8.1 Overview

The FC-BaseT startup procedure is based on IEEE 802.3-2005 Auto-Negotiation and considers the properties of the cable in the speed selection. Each PHY maintains a list of speed capabilities that it is able to support (e.g., 4GFC-BaseT, 2GFC-BaseT, and 1GFC-BaseT), each up to a certain reach. The procedure operates as follows:

- 1) Each PHY determines the speeds supported by the host and keeps the intersection with its speed capabilities (see 8.2);
- 2) Each PHY performs Auto-Negotiation (see 8.4). As part of the Auto-Negotiation process, the involved PHYs, after determining that they both support FC-BaseT, perform a cable length estimation. Each PHY advertises to the other one the speeds that it is able to support at the estimated cable length. The higher speed in the intersection between the advertised speeds is selected as tentative link operating speed;
- 3) When the tentative link operating speed is selected, each PHY attempts to establish the link according to the process described in 6.2.4. The link establishment process includes training of the PMA and qualification of the link. If link qualification fails (i.e., a PHY determines that the available SNR margin is not enough to achieve the specified bit error rate), the PHY retries to establish the link or goes back to step 2, to perform again Auto-Negotiation, as described in 8.5. In this subsequent Auto-Negotiation process the PHY advertises speeds lower than the tentative link operating speed that failed the link qualification. If link qualification succeeds, the successful tentative link operating speed is selected as the link operating speed; and
- 4) When the link operating speed is selected, the PHY begins to operate in normal mode. The host synchronizes its speed with the selected link operating speed (see 8.3) and the startup procedure is completed.

This scheme allows the building of completely interoperable implementations differing only in their reach at different speeds, assuming the same reach for the lowest speed (i.e., 1GFC-BaseT). A lower reach at higher speeds allows power efficient implementations. Support for 100m reach at the 1GFC-BaseT speed is mandatory, in order to guarantee interoperability among all FC-BaseT implementations.

8.2 Host speeds determination

An FC-BaseT PHY determines the speeds supported by the host to which it is connected from the Host Speeds bits (i.e., bits 2 .. 0) of the Management Control register. A host supporting FC-BaseT PHYs is required to report its supported speeds to the PHY by writing in the Management Control register (see 9.3).

8.3 Host synchronization

When an FC-BaseT PHY operates in normal mode, the selected link operating speed is encoded in the Speed Resolution bits (i.e., bits 8 .. 7) of the Status register (see 9.4). A host may use this information to synchronize with the PHY.

Alternatively, if an FC-BaseT PHY implements δ points, a host may synchronize with the PHY using the clock information extracted from the words transmitted by the PHY to the host at the selected link operating speed when the Host Synch Mode bits of the Control register are not set to zero (see 9.3).

It is strongly recommended that an FC-BaseT PHY implementing δ points supplies the content of the Rx_LOS bit of the Status register (see 9.4) also as a signal, to provide a feature analogous to the Rx_LOS signal defined in FC-PI-2.

8.4 FC-BaseT auto-negotiation

8.4.1 Overview

The FC-BaseT Auto-Negotiation shall be performed on the cable according to the rules specified in IEEE 802.3-2005, clause 28, as modified by this subclause and with the following qualifications:

- a) FC-BaseT Auto-Negotiation is mandatory;
- b) FC-BaseT Auto-Negotiation supports the NLP Receive Link Integrity Test function but does not implement a 10BASE-T PMA;
- c) FC-BaseT Auto-Negotiation does not support parallel detection of either 100BASE-TX or 100BASE-T4 PMAs;
- d) FC-BaseT Auto-Negotiation does not support the IEEE 802.3-2005 clause 22 MII;
- e) Registers comparable to MII registers 0 .. 8 (see IEEE 802.3-2005, clause 22), that are used by IEEE 802.3-2005 clause 28 Auto-Negotiation, are implemented by the FC-Base-T Management Registers 0 .. 8;
- f) The link_fail_inhibit_timer (see IEEE 802.3-2005, clause 28.3.2) shall have a default value of 2 192 ms; and
- g) All link test pulses in the FLP Burst sequence shall meet the template requirements of IEEE 802.3-2005, figure 14-12, when measured across a 100 ohm \pm 1% resistive test load; both with the load connected directly to the MDI and with the load connected through all of the cable types and distances supported by the advertised capabilities.

The FC-BaseT Auto-Negotiation is based on the exchange of four 16-bit Pages, one Base Page and three Unformatted Pages, according to the following steps:

- 1) FC-BaseT support. The Base Page exchange verifies that both PHYs support FC-BaseT (see 8.4.2). If one of the two PHYs does not support FC-BaseT, the Auto-Negotiation process ends (see 8.4.2). Further exchanges are possible, but outside the scope of this standard. If both PHYs support FC-BaseT, Auto-Negotiation continues;
- 2) Master-Slave relationship. The first Unformatted Page exchange is used to pass the information required to determine the Master-Slave relationship between the two PHYs (see 8.4.3);
- 3) Cable length estimation. The second Unformatted Page exchange is used to pass the cable length estimation information (see 8.4.4); and
- 4) Tentative operating speed determination. The third Unformatted Page exchange allows the two PHYs to exchange the speeds compatible with the estimated cable length and other parameters (see 8.4.5). The best common speed exchanged is selected as tentative link operating speed.

The Acknowledge 2 bit (see figure 37) is not utilized and has no meaning when used for the FC-BaseT message page exchange.

8.4.2 FC-BaseT support

Support for FC-BaseT is determined via a Base Page exchange. Figure 36 shows the Base Page encoding.

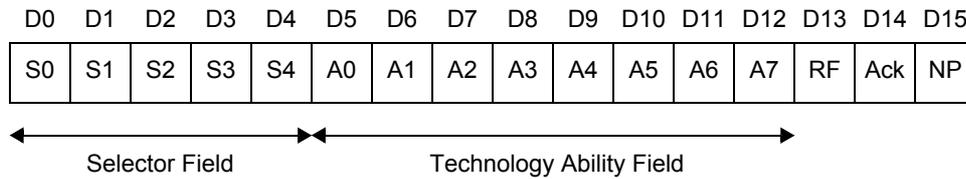


Figure 36 – Base page encoding

The fields composing the Base Page shall be used as follows:

Selector field: shall be set to the value 00101b.

NOTE 7 - This value indicates INCITS, as specified in <http://www.ieee802.org/3/selectors/selectors.html>.

Technology Ability field: the bits A0 .. A7 shall have the semantic shown in table 28.

Table 28 – Technology Ability Field

Bit	Description
A0 .. A5	Reserved to INCITS
A6	1b: FC-BaseT supported 0b: FC-BaseT not supported
A7	1b: Extended Next Pages (XNP) supported ^a 0b: Extended Next Pages (XNP) not supported
^a See IEEE 802.3an-2006, clause 28.2.1.2.3, for the definition of Extended Next Pages.	

Remote Fault (RF): see IEEE 802.3-2005, clause 28.2.1.2.3.

Acknowledge (Ack): see IEEE 802.3-2005, clause 28.2.1.2.4.

Next Page (NP): see IEEE 802.3-2005, clause 28.2.1.2.5.

Extended Next Pages may be supported by a PHY, but they are not used by the FC-BaseT Auto-Negotiation procedure.

A PHY indicates support for FC-BaseT by setting the Selector field to the value 00101b and bit A6 of the Technology Ability field to one.

If the Selector Field of the Link Partner does not match the one of the local PHY and the Next Page bit of the Link Partner Base Page is set to one, the local PHY shall transmit Null Message Pages after its Base Page and shall not transmit the three FC-BaseT Unformatted Next Pages. An FC-BaseT PHY may restart Auto-Negotiation at any time.

8.4.3 Master-Slave relationship

All FC-BaseT PHYs are capable of operating as Master or Slave. All FC-BaseT PHYs are single-port devices. The Master-Slave relationship is determined during Auto-Negotiation by exchanging two seed values. The seed values are exchanged via Unformatted Pages. Figure 37 shows the Unformatted Page encoding.

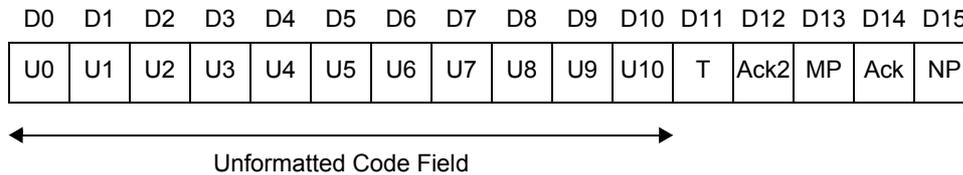


Figure 37 – Unformatted page encoding

The fields composing the Unformatted Page shall be used as follows:

Toggle (T): see IEEE 802.3-2005, clause 28.2.3.4.6.

Acknowledge 2 (Ack2): see IEEE 802.3-2005, clause 28.2.3.4.5. This bit is not used by FC-BaseT.

Message Page (MP): see IEEE 802.3-2005, clause 28.2.3.4.4. For use with FC-BaseT this bit shall be set to zero, to indicate that the page is an Unformatted Page.

Acknowledge (Ack): see IEEE 802.3-2005, clause 28.2.3.4.3.

Next Page (NP): see IEEE 802.3-2005, clause 28.2.3.4.2.

The bits of the Unformatted Code field of the first Unformatted Page have the semantic shown in table 29.

Table 29 – Unformatted code field of the first unformatted page

Bit	Description
U0	FC-BaseT MASTER-SLAVE Seed Bit 0 (SB0, LSB)
U1	FC-BaseT MASTER-SLAVE Seed Bit 1 (SB1)
U2	FC-BaseT MASTER-SLAVE Seed Bit 2 (SB2)
U3	FC-BaseT MASTER-SLAVE Seed Bit 3 (SB3)
U4	FC-BaseT MASTER-SLAVE Seed Bit 4 (SB4)
U5	FC-BaseT MASTER-SLAVE Seed Bit 5 (SB5)
U6	FC-BaseT MASTER-SLAVE Seed Bit 6 (SB6)
U7	FC-BaseT MASTER-SLAVE Seed Bit 7 (SB7)
U8	FC-BaseT MASTER-SLAVE Seed Bit 8 (SB8)
U9	FC-BaseT MASTER-SLAVE Seed Bit 9 (SB9)
U10	FC-BaseT MASTER-SLAVE Seed Bit 10 (SB10, MSB)

When receiving the first Unformatted Page after the Auto-Negotiation signal ability_match is set to true (see IEEE 802.3-2005, figure 28-16), the PHY determines a tentative Master-Slave resolution. This intermediate resolution is not to be confused with the actual Master-Slave resolution that is performed in the FLP LINK

GOOD CHECK state of the Auto-Negotiation Arbitration state diagram (see IEEE 802.3-2005, figure 28-16), even though the resolution mechanisms are the same. The tentative resolution assignment:

- a) is based on decoding the FLP bursts of the first Unformatted Page that caused ability_match to be set to true; and
- b) is ignored by the Auto-Negotiation process. It is generated solely for the purpose of synchronizing the cable length estimation algorithm with Auto-Negotiation and it only has meaning in the context of cable length estimation (see 8.4.4).

The Master-Slave Seed is assigned as follows. A random or pseudo random number, rand_val, belonging to a sequence of independent, identically distributed integer numbers with a uniform distribution in the range of 0 to $2^{11}-1$, is generated. The algorithm used to generate the number should be designed to minimize the correlation between the number generated by any two devices at any given time. The Master-Slave Seed is assigned based on the value of rand_val and the values of the Master-Slave Test Configuration Enable bit (i.e., bit 4) and Master-Slave Test Configuration Value bit (i.e., bit 3) of the management Test register (see 9.6), as shown in table 30.

Table 30 – Master-Slave seed assignment

Test Register bits 4 .. 3 value	Master-Slave Seed	Description
11b	rand_val OR 10000000000b	Preference for PHY to be configured as Master
10b	rand_val AND 01111111111b	Preference for PHY to be configured as Slave
0xb	rand_val	No preference for PHY to be configured as Master or Slave

8.4.4 Cable length estimation

Cable length estimation is performed in parallel and synchronously with the FC-BaseT Auto-Negotiation after the tentative Master-Slave resolution as follows:

- 1) Tentative Slave PHY enables 'Clock Loopback' mode on transmit pair BI_DC: it detects clock Normal Link Pulses (NLPs) within the Auto-Negotiation Fast Link Pulse (FLP) bursts (see IEEE 802.3-2005, clause 28) and retransmits them after a specified delay on transmit pair BI_DC (see figure 39);
- 2) Tentative Master PHY enables 'Clock Detect' mode on receive pairs BI_DC and BI_DD: it measures the time duration for the return of Auto-Negotiation clock NLPs (see figure 38);
- 3) Tentative Master PHY determines cable length estimate before end of Auto-negotiation state 'COMPLETE ACKNOWLEDGE' (see IEEE 802.3-2005, clause 28); and
- 4) Tentative Slave PHY disables 'Clock Loopback' mode at end of Auto-Negotiation state 'COMPLETE ACKNOWLEDGE' (see IEEE 802.3-2005, clause 28).

NOTE 8 - Transmit pair BI_DA of one PHY has to be connected to the receive pair BI_DA of its link partner, and vice versa, in order to perform Auto-Negotiation. However this does not guarantee that transmit pair BI_DC of one PHY is connected to the receive pair BI_DC of its partner, because the cable pairs associated with BI_DC and BI_DD may not be crossed over in the same way as the cable pairs associated with BI_DA and BI_DB. Therefore a PHY looking for NLPs transmitted by its link partner PHY on BI_DC should monitor the receive pairs BI_DC and BI_DD to detect the transmitted NLPs.

With this algorithm the tentative Master PHY measures the cable round trip delay. For the purposes of this standard, the cable length in meters shall be computed from this delay estimate using the following formula:

$$\text{cable length} = \text{roundtrip_cable_delay} \times \frac{21}{64}$$

NOTE 9 - This formula assumes a nominal velocity of propagation of a signal over a cable of 0,6973c. Measurements have shown that the value of the velocity of propagation for a specific cable may differ from this value by up to $\pm 7,5\%$.

When the measure of the round trip delay is complete, the tentative Master sends the measured delay to the tentative Slave in the second Unformatted Page. The Unformatted Page encoding is shown in figure 37. The bits of the Unformatted Code field of the second Unformatted Page have the semantic shown in Table 31.

Table 31 – Unformatted code field of the second unformatted page

Bit	Description
U0	FC-BaseT Round Trip Cable Delay Bit 0 (RTCDB0, LSB)
U1	FC-BaseT Round Trip Cable Delay Bit 1 (RTCDB1)
U2	FC-BaseT Round Trip Cable Delay Bit 2 (RTCDB2)
U3	FC-BaseT Round Trip Cable Delay Bit 3 (RTCDB3)
U4	FC-BaseT Round Trip Cable Delay Bit 4 (RTCDB4)
U5	FC-BaseT Round Trip Cable Delay Bit 5 (RTCDB5)
U6	FC-BaseT Round Trip Cable Delay Bit 6 (RTCDB6)
U7	FC-BaseT Round Trip Cable Delay Bit 7 (RTCDB7)
U8	FC-BaseT Round Trip Cable Delay Bit 8 (RTCDB8)
U9	FC-BaseT Round Trip Cable Delay Bit 9 (RTCDB9, MSB)
U10	Reserved to T11

For the tentative Master, the value reported in the Unformatted Code field is the value of the 10-bit roundtrip_cable_delay variable (see table 33). For the tentative Slave, the value reported in the Unformatted Code field is 514.

8.4.5 Tentative Operating Speed Determination

Each PHY shall advertise in the third Unformatted Page the speeds it supports at the estimated cable length. These speeds shall be a subset of the speeds indicated by the Fallback Speed bits of the Speed Downshift register (see 9.5). The subset is selected to eliminate the speeds that the PHY does not support based on the estimated cable length.

The Unformatted Page encoding is shown in figure 37. The bits of the Unformatted Code field of the third Unformatted Page have the semantic shown in table 32.

Table 32 – Unformatted code field of the third unformatted page

Bit	Description
U0	1b: 1GFC-BaseT supported 0b: 1GFC-BaseT not supported
U1	1b: 2GFC-BaseT supported 0b: 2GFC-BaseT supported
U2	1b: 4GFC-BaseT supported 0b: 4GFC-BaseT not supported
U7 .. U3	Reserved to T11
U8	Requested link partner PCS Transmit encoding: 1b: Trellis coding requested 0b: Schläfli Lattice coding requested
U9	Requested link partner PMA training mode: 1b: Periodic PMA training requested 0b: Normal PMA training requested
U10	1b: Other Data carried 0b: Fibre Channel Data carried

8.4.6 Configuration resolution

Since a local FC-BaseT PHY and a link partner may support multiple configurations, a mechanism to resolve which one to use is necessary. This mechanism is called Configuration Resolution and is made on entrance to the FLP LINK GOOD CHECK state of the Arbitration state diagram (see IEEE 802.3-2005, figure 28-16). Configuration Resolution resolves the Master-Slave configuration, the transmitter power schedule, the Other Data configuration and the tentative link operating speed as follows:

- a) Master-Slave configuration resolution: If the local and link partner Master-Slave Seed values are not equal, the device with the higher seed value is configured as Master, the other one as Slave. If the two seeds are equal the Master-Slave configuration resolution fails;
- b) Transmitter power schedule resolution: The transmitter power schedule shall be selected from table 20 using the cable length estimate described in 8.4.4 for the cable length. If no valid cable length estimation has been determined, the transmitter power schedule resolution fails;
- c) Other Data configuration resolution: If the value of the link partner Other Data bit is equal to the local Other Data bit, the Other Data resolves to the common value. If the value of link partner Other Data bit is different to the local Other Data bit then the Other Data resolution fails; and
- d) Tentative link operating speed resolution: The highest advertised common speed shall be selected as the tentative link operating speed. Speeds are ordered from highest to lowest as follows:
 - I) 4GFC-BaseT;
 - II) 2GFC-BaseT; and
 - III) 1GFC-BaseT.

If the resolution of any of the Master-Slave configuration, the transmitter power schedule, the Other Data configuration or the tentative link operating speed fails, Auto-Negotiation fails.

8.5 Speed downshift function

An FC-BaseT PHY uses an Auto-Negotiation Speed Downshift function to progressively fall back to advertise lower speeds than its maximum capability if it fails to establish a link (see 6.2.4) after a successful Auto-Negotiation. For Speed Downshift, the PHY first tries Auto-Negotiation by advertising its speed capabilities in accordance with the Fallback Speeds bits in the Management Speed Downshift register and the cable length estimate acquired during Auto-Negotiation, as described in 8.4.5. After three failed attempts to establish the link, the PHY sets the highest supported speed of the Fallback Speeds bits to zero and restarts the Auto-Negotiation process. This process continues through all supported speeds of the Fallback Speeds bits to the minimum supported speed. When there are no lower speeds to try the process starts again with the Fallback Speeds bits set to the values of the Host Speeds bits of the Management Control register.

NOTE 10 - A failed attempt to establish a link in the context of the Speed Downshift function means that Auto-Negotiation successfully completed but the PHY was unable to establish the link (see 6.2.4) at the tentative link operating speed.

8.6 State diagrams

State diagrams are provided for the Clock Detect and Clock Loopback functions.

$T_{318.75 \text{ MHz}}$ is the period of a 318,75 MHz clock with a maximum frequency offset of ± 50 ppm.

$N_{\text{delay linkpulse}}$ is the delay to the nearest integer between the arrival of the positive edge of an Auto-Negotiation clock link pulse and the assertion of the linkpulse signal to true. The unit of measurement is $T_{318.75 \text{ MHz}}$.

IECNORM.COM : Click to view the full PDF of ISO/IEC 14165-151:2017

The variables used in the state diagrams are shown in table 33.

Table 33 – Cable length estimation state variables (part 1 of 3)

Variable	Description
ability_match	See IEEE 802.3-2005, clause 28.3.
ack_finished	See IEEE 802.3-2005, clause 28.3.
linkpulse	See IEEE 802.3-2005, clause 28.3.
TD_AUTONEG	See IEEE 802.3-2005, clause 28.3.
check_complete ^a	Controls the exit from the CHECK CONSISTENCY state of the Clock Detect state diagram. It is set to false on entry to the CHECK CONSISTENCY state. It is subsequently set to true in this state once checking has been completed. This variable may assume the following values: false: consistency check is not started or is not complete (default); or true: indicates that the last received link pulse has been checked for consistency and consistency_ok has been updated accordingly.
clen_clock_detect_enable ^{a, b}	Controls the enabling of the cable length estimation clock detection circuit on receive pairs BI_DC and BI_DD. Its value is true only when the PHY is engaged in FC-BaseT Auto-Negotiation, the PHY has resolved that it is the tentative Master, the signal ability_match has been set to true for the first Unformatted Page and the signal ack_finished has not yet been set to true for the first Unformatted Page. This variable may assume the following values: false: disable cable length estimation clock detect circuit; or true: enable cable length estimation clock detect circuit.
clen_clock_loopback_enable ^a	Controls the enabling of the cable length estimation clock loopback circuit. Its value is true only when the PHY is engaged in FC-BaseT Auto-Negotiation, the PHY has resolved that it is the tentative Slave, the signal ability_match has been set to true for the first Unformatted Page and the signal ack_finished has not yet been set to true for the first Unformatted Page. This variable may assume the following values: false: disable cable length estimation clock loopback circuit; or true: enable cable length estimation clock loopback circuit.
linkpulse_on_cd ^a	Indicates that a valid link pulse has been detected on receive pair BI_DC or BI_DD. This variable may assume the following values: false: linkpulse_on_cd is set to false after any Clock Detect State Diagram state transition (default); or true: linkpulse_on_cd is set to true when a link pulse has been received.
^a Variable set by its definition; not explicitly set in the state diagrams. ^b The clock loopback and clock detect circuits are enabled using signals derived from Auto-Negotiation. The enable signals have been selected so that they synchronize the measurements between the two PHYs and guarantee a minimum overlap of 5 FLP bursts with clock NLPs that may be looped back and detected.	

Table 33 – Cable length estimation state variables (part 2 of 3)

Variable	Description
consistency_ok ^a	<p>This variable is evaluated in the CHECK CONSISTENCY state and is only valid if the signal check_complete = true. It is used to indicate the consistency of the last five detected link pulses. Link pulses are consistent if the absolute value of the difference between the mean delay of the pulses and the maximum and minimum delay of the pulses is less than or equal to $3 \times T_{318.75}$ MHz. The delay of a link pulse is calculated by measuring the time interval (in units of $T_{318.75}$ MHz) between the positive edge of a transmitted Auto-Negotiation clock NLP and positive edge of the looped back clock NLP. This interval is then adjusted by $640 \times T_{318.75}$ MHz to remove the fixed delay added by the link partner PHY.</p> <p>This variable may assume the following values: false: less than five link pulses have been detected or five or more link pulses have been detected and the last five are not consistent (default); or true: the last five link pulses are consistent.</p>
nlp_is_clock ^a	<p>This variable qualifies the Auto-negotiation linkpulse variable indicating whether the current NLP is a clock pulse or not. This variable is only valid when linkpulse is true.</p> <p>This variable may assume the following values: false: indicates that the link pulse signaled by linkpulse is not a clock pulse; or true: indicates that the link pulse signaled by linkpulse is a clock pulse.</p>
roundtrip_cable_delay ^a	<p>If a valid estimate of the cable length has been made, this variable is a measure of the round trip cable delay in units of $T_{318.75}$ MHz. In this case the value of the variable is the rounded and saturated mean value of the delay of the five link pulses that caused consistency_ok to be set to true. Rounding is to the nearest positive integer. Mean values greater than 511 shall be saturated to 511.</p> <p>If no valid estimate of the cable length has been made this variable is used to indicate the reason why no estimate is available.</p> <p>This variable is conveyed to the link partner in the second Unformatted Page. This variable may assume the following values: 0 .. 511: valid estimate; 512: no_resolution. This is the value reported by a PHY if it was unable to resolve the tentative Master-Slave relationship; 513: incomplete_estimate. This is the value reported by the tentative Master to indicate unsuccessful cable length estimation; 514: no_estimate. This is the value reported by the Slave PHY; or 515 .. 1023: reserved.</p>
<p>^a Variable set by its definition; not explicitly set in the state diagrams.</p> <p>^b The clock loopback and clock detect circuits are enabled using signals derived from Auto-Negotiation. The enable signals have been selected so that they synchronize the measurements between the two PHYs and guarantee a minimum overlap of 5 FLP bursts with clock NLPs that may be looped back and detected.</p>	

Table 33 – Cable length estimation state variables (part 3 of 3)

Variable	Description
TD_CLEN_EST	Controls the signal sent by cable length estimation loopback circuit to the link pulse generation circuit on transmit pair BI_DC. This variable may assume the following values: idle: cable length estimation prevents transmission of all link pulses on the MDI for BI_DC; or link_test_pulse: cable length estimation causes a single link pulse to be transmitted on the MDI for transmit pair BI_DC.
transmit_state_is_tx_clock_bit ^a	Indicates if the Auto-Negotiation Transmit state machine (see IEEE 802.3-2005, figure 28-14) is in the state TRANSMIT CLOCK BIT. This variable may assume the following values: false: the Auto-negotiation Transmit state machine is not in the state TRANSMIT CLOCK BIT; or true: the Auto-negotiation Transmit state machine is in the state TRANSMIT CLOCK BIT.
^a Variable set by its definition; not explicitly set in the state diagrams. ^b The clock loopback and clock detect circuits are enabled using signals derived from Auto-Negotiation. The enable signals have been selected so that they synchronize the measurements between the two PHYs and guarantee a minimum overlap of 5 FLP bursts with clock NLPs that may be looped back and detected.	

The timers used in the state diagrams are shown in table 34.

Table 34 – Cable length estimation timers

Timer	Description
clock_pulse_detect_max_timer	Timer for the maximum time between transmitting a clock link pulse on transmit pair BI_DA and detecting the return link pulse on receive pair BI_DC or BI_DD. This timer shall expire $(1\ 372 - N_{\text{delay linkpulse}}) \times T_{318.75 \text{ MHz}}$ after being started.
clock_pulse_detect_min_timer	Timer for the minimum time between transmitting a clock link pulse on transmit pair BI_DA and detecting the return link pulse on receive pair BI_DC or BI_DD. This timer shall expire $(630 - N_{\text{delay linkpulse}}) \times T_{318.75 \text{ MHz}}$ after being started.
clock_pulse_stopwatch_timer	Timer for measuring the time duration between transmitting a clock link pulse on transmit pair BI_DA and detecting the return link pulse on receive pair BI_DC or BI_DD. The timer shall expire $(1\ 372 - N_{\text{delay linkpulse}}) \times T_{318.75 \text{ MHz}}$ after being started.
delay_pulse_timer	Timer to ensure a fixed delay between the rising positive edge of a received NLP clock pulse and the rising positive edge of a retransmitted NLP pulse. The timer shall take into account the latency of the link pulse detection and the transmitter so that there is a fixed latency of $(640 \pm 2) \times T_{318.75 \text{ MHz}}$ between the positive edge of the link pulse appearing at the MDI and the positive edge of the link pulse being retransmitted at the MDI.