

INTERNATIONAL
STANDARD

ISO/IEC/
IEEE
8802-15-3

First edition
2017-10

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

**Part 15-3:
High data rate wireless multi-media
networks**

*Technologies de l'information — Télécommunications et échange
d'information entre systèmes — Réseaux locaux et métropolitains —
Exigences spécifiques —*

Partie 15-3: Réseaux multimédias sans fil à haut débit de données



Reference number
ISO/IEC/IEEE 8802-15-3:2017(E)

© IEEE 2016

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017



COPYRIGHT PROTECTED DOCUMENT

© IEEE 2016

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO or IEEE at the address below or ISO's member body in the country of the requester.

ISO copyright office
Ch. de Blandonnet 8 • CP 401
CH-1214 Vernier, Geneva, Switzerland
Tel. +41 22 749 01 11
Fax +41 22 749 09 47
copyright@iso.org
www.iso.org

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

stds.ipr@ieee.org
www.ieee.org

Foreword

ISO (the International Organization for Standardization) and IEC (the 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 through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and nongovernmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

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

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of ISO/IEC JTC 1 is to prepare International Standards. 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.

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

ISO/IEC/IEEE 8802-15-3 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE Std. 802.15.3:2016). It was adopted by Joint Technical Committee ISO/IEC JTC 1, *Information technology, Subcommittee SC 6, Telecommunications and information exchange between systems*, in parallel with its approval by the ISO/IEC national bodies, under the “fast-track procedure” defined in the Partner Standards Development Organization cooperation agreement between ISO and IEEE. IEEE is responsible for the maintenance of this document with participation and input from ISO/IEC national bodies.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017

IEEE Std 802.15.3™-2016
(Revision of
IEEE Std 802.15.3-2003)

IEEE Standard for High Data Rate Wireless Multi-Media Networks

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 15 May 2016
IEEE-SA Standards Board

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017

Abstract: The protocol and compatible interconnection of data and multimedia communication equipment via 2.4 GHz and 60 GHz radio transmissions in a Wireless Personal Area Network (WPAN) using low power and multiple modulation formats to support scalable data rates is defined in this standard. The Medium Access Control (MAC) sublayer protocol supports both isochronous and asynchronous data types.

Keywords: ad hoc network, IEEE 802.15.3™, mobility, mm-wave, radio frequency, wireless

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

Copyright © 2016 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 25 July 2016. Printed in the United States of America.

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

PDF: ISBN 978-1-5044-2023-5 STD20940
Print: ISBN 978-1-5044-2024-2 STDPD20940

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

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

Important Notices and Disclaimers Concerning IEEE Standards Documents

IEEE documents are made available for use subject to important notices and legal disclaimers. These notices and disclaimers, or a reference to this page, appear in all standards and may be found under the heading “Important Notice” or “Important Notices and Disclaimers Concerning IEEE Standards Documents.”

Notice and Disclaimer of Liability Concerning the Use of IEEE Standards Documents

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

IEEE does not warrant or represent the accuracy or content of the material contained in its standards, and expressly disclaims all warranties (express, implied and statutory) not included in this or any other document relating to the standard, including, but not limited to, the warranties of: merchantability; fitness for a particular purpose; non-infringement; and quality, accuracy, effectiveness, currency, or completeness of material. In addition, IEEE disclaims any and all conditions relating to: results; and workmanlike effort. IEEE standards documents are supplied “AS IS” and “WITH ALL FAULTS.”

Use of an IEEE standard is wholly voluntary. The existence of an IEEE standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard.

In publishing and making its standards available, IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity nor is IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing any IEEE Standards document, should rely upon his or her own independent judgment in the exercise of reasonable care in any given circumstances or, as appropriate, seek the advice of a competent professional in determining the appropriateness of a given IEEE standard.

IN NO EVENT SHALL IEEE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO: PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE PUBLICATION, USE OF, OR RELIANCE UPON ANY STANDARD, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE AND REGARDLESS OF WHETHER SUCH DAMAGE WAS FORESEEABLE.

Translations

The IEEE consensus development process involves the review of documents in English only. In the event that an IEEE standard is translated, only the English version published by IEEE should be considered the approved IEEE standard.

Official statements

A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered or inferred to be the official position of IEEE or any of its committees and shall not be considered to be, or be relied upon as, a formal position of IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE.

Comments on standards

Comments for revision of IEEE Standards documents are welcome from any interested party, regardless of membership affiliation with IEEE. However, IEEE does not provide consulting information or advice pertaining to IEEE Standards documents. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Since IEEE standards represent a consensus of concerned interests, it is important that any responses to comments and questions also receive the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to comments or questions except in those cases where the matter has previously been addressed. For the same reason, IEEE does not respond to interpretation requests. Any person who would like to participate in revisions to an IEEE standard is welcome to join the relevant IEEE working group.

Comments on standards should be submitted to the following address:

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

Laws and regulations

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

Copyrights

IEEE draft and approved standards are copyrighted by IEEE under U.S. and international copyright laws. They are made available by IEEE and are adopted for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making these documents available for use and adoption by public authorities and private users, IEEE does not waive any rights in copyright to the documents.

Photocopies

Subject to payment of the appropriate fee, IEEE will grant users a limited, non-exclusive license to photocopy portions of any individual standard for company or organizational internal use or individual, non-commercial use only. To arrange for payment of licensing fees, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Updating of IEEE Standards documents

Users of IEEE Standards documents should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect.

Every IEEE standard is subjected to review at least every ten years. When a document is more than ten years old and has not undergone a revision process, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE standard.

In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE-SA Website at <http://ieeexplore.ieee.org/browse/standards/collection/ieee> or contact IEEE at the address listed previously. For more information about the IEEE SA or IEEE's standards development process, visit the IEEE-SA Website at <http://standards.ieee.org>.

Errata

Errata, if any, for all IEEE standards can be accessed on the IEEE-SA Website at the following URL: <http://standards.ieee.org/findstds/errata/index.html>. Users are encouraged to check this URL for errata periodically.

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken by the IEEE with respect to the existence or validity of any patent rights in connection therewith. If a patent holder or patent applicant has filed a statement of assurance via an Accepted Letter of Assurance, then the statement is listed on the IEEE-SA Website at <http://standards.ieee.org/about/sasb/patcom/patents.html>. Letters of Assurance may indicate whether the Submitter is willing or unwilling to grant licenses under patent rights without compensation or under reasonable rates, with reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to obtain such licenses.

Essential Patent Claims may exist for which a Letter of Assurance has not been received. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patent Claims, or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

At the time this standard was completed, the IEEE 802.15 working group had the following membership:

IEEE 802.15 Working group

Robert F. Heile, *Working Group Chair*

Rick Alfvén, *Working Group Vice-Chair*

Patrick Kinney, *Working Group Vice-Chair, Working Group Secretary*

James P. K. Gilb, *Working Group Technical Editor*

Benjamin A. Rolfe, *Working Group Treasurer*

James P. K. Gilb, *802.15.3m Chair, 802.15.3RevA Technical Editor*

Mohamed Abdallah	Yeong Min Jang	John Notor
Keiji Akiyama	Seong-Soon Joo	Hiroyo Ogawa
Hideki Aoyama	Volker Jungnickel	Mitsuaki Oshima
Arthur Astrin	Jeritt Kent	Chandrashekha P. S. Bhat
Tuncer Baykas	Jaehwan Kim	Glenn Parsons
Philip Beecher	Junhyeong Kim	Charles Perkins
Monique Brown	Young-Hoon Kim	Albert Petrick
Kiran Bynam	Youngsoo Kim	Clinton Powell
Edgar Callaway	Shoichi Kitazawa	Verotiana Rabarijaona
Chris Calvert	Tero Kivinen	Demir Rakanovic
Radhakrishna Canchi	Ryuji Kohno	Ivan Reede
Jaesang Cha	Fumihide Kojima	Richard Roberts
Soo-Young Chang	Keitarou Kondou	Noriyuki Sato
Clint Chaplin	Thomas Kuerner	Norihiko Sekine
Stephen Chasko	Amarjeet Kumar	Nikola Serafimovski
Sangsung Choi	Byung-Jae Kwak	Kunal Shah
Hee-Sang Chung	Namtuan Lee	Stephen Shellhammer
Hendricus De Ruijter	Hoosung Lee	Shusaku Shimada
Guido Dolmans	Jae Seung Lee	Gyung Chul Sihn
Igor Dotlic	Moon-Sik Lee	Gary Stuebing
Andrew Estrada	Myung Lee	Don Sturek
David Evans	Huan-Bang Li	Mineo Takai
Kiyoshi Fukui	Liang Li	Kou Togashi
Tim Godfrey	Michael Lynch	Kiyoshi Toshimitsu
Yoshiho Goto	Itaru Maekawa	Murat Uysal
Jussi Haapola	Hiroyuki Matsumura	Billy Verso
Shinsuke Hara	Kenichi Mori	Brian Weis
Marco Hernandez	Robert Moskowitz	Makoto Yaita
Ken Hiraga	Mohammad Nekoui	Peter Yee
Jay Holcomb	Chiu Ngo	Hidetoshi Yokota
Koji Horisaki	Trang Nguyen	Yu Zeng
Iwao Hosako	Paul Nikolic	Junwen Zhang
Bing Hu	Makoto Noda	Chunhui Zhu

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks

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

Tomoko Adachi	Noriyuki Ikeuchi	Satoshi Oyama
Thomas Alexander	Yasuhiko Inoue	Tomas Pajda
Richard Alfvén	Sergiu Iordanescu	Arumugam Paventhan
James D. Allen	Paul Isaacs	Dalibor Pokrajac
Nobumitsu Amachi	Akio Iso	Clinton Powell
Carol Ansley	Atsushi Ito	Venkatesha Prasad
Yusuke Asai	Adri Jovin	Demir Rakanovic
Alfred Asterjadhi	Shinkyo Kaku	R. K. Rannow
Gennaro Boggia	Piotr Karocki	James Reilly
Nancy Bravin	Jeritt Kent	Maximilian Riegel
William Byrd	Stuart Kerry	Robert Robinson
Radhakrishna Canchi	Yongbum Kim	Benjamin Rolfe
William Carney	Patrick Kinney	Naotaka Sato
Juan Carreon	Tero Kivinen	Frank Schewe
Yessenia Cevallos	Adrian Kliks	Reinhard Schrage
Minho Cheong	Yasushi Kudoh	Andy Scott
Peter Cibula	Thomas Kuerner	Michael Seaman
Charles Cook	Thomas Kurihara	Daniel Smolinski
Richard Doyle	Paul Lambert	Manikantan Srinivasan
Constantin Dumitrescu	Jeremy Landt	Dorothy Stanley
Sourav Dutta	Joseph Levy	Thomas Starai
Richard Edgar	Arthur H. Light	Eugene Stoudenmire
David Evans	William Lumpkins	Rene Struik
Liu Fangfang	Michael Lynch	Walter Struppler
Avraham Freedman	Elvis Maculuba	Mark Sturza
Yukihiro Fujimoto	Jeffery Maki	David Thompson
Yue Gao	Roger Marks	Pedro Tonhozi de Oliveira
Alexander Gelman	Jeffery Masters	Lorenzo Vangelista
James P. K. Gilb	Michael McInnis	Dmitri Varsanofiev
Eric W. Gray	Richard Mellitz	Prabodh Varshney
Randall Groves	Juichi Nakada	George Vlantis
Michael Gundlach	Nabil Nasser	Hung-Yu Wei
Robert F. Heile	Michael Newman	Andreas Wolf
Jerome Henry	Chiu Ngo	Chun Yu Charles Wong
Marco Hernandez	Nick S.A. Nikjoo	Oren Yuen
Werner Hoelzl	Makoto Noda	Zhen Zhou
Rita Horner	John Notor	

When the IEEE-SA Standards Board approved this standard on 15 May 2016, it had the following membership:

Jean-Philippe Faure, Chair
Ted Burse, Vice-Chair
John D. Kulick, Past Chair
Konstantinos Karachalios, Secretary

Chuck Adams
Masayuki Ariyoshi
Stephen Dukes
Jianbin Fan
J. Travis Griffith
Gary Hoffman

Ronald W. Hotchkiss
Michael Janezic
Joseph L. Koepfinger*
Hung Ling
Kevin Lu
Annette D. Reilly
Gary Robinson

Mehmet Ulema
Yingli Wen
Howard Wolfman
Don Wright
Yu Yuan
Daidi Zhong

*Member Emeritus

Introduction

This introduction is not part of IEEE Std 802.15.3-2016, IEEE Standard for High Data Rate Wireless Multi-Media Networks.

IEEE Std 802.15.3-2016 was originally developed to provide superior quality of service (QoS) over relatively medium range wireless links.

The idea of a high-rate addition to the IEEE 802.15 family of standards was first proposed in November 1999 at the IEEE Plenary meeting in Kaua'i, HI. The 802.15.3 task group began its official work at the March 2000 IEEE Plenary meeting in Albuquerque, NM, creating a criteria document and evaluation method. The down-selection of MAC and PHY proposals was completed at the November 2000 IEEE 802 Plenary meeting in Tampa, FL, and the writing of the draft began in December 2000. After working on the draft for one year, the document was ready for the task group ballot process in December 2001. The draft received final working group approval at the November 2002 IEEE Plenary meeting in Kaua'i, HI, and began the sponsor ballot process following the meeting. The draft went through one sponsor ballot and two recirculations before it was submitted to the IEEE Standards Association Standards Board (IEEE-SASB) for approval. The IEEE-SASB approved 802.15.3 as an IEEE standard in June 2003.

IEEE Std 802.15.3-2003 defined an ad-hoc MAC that enabled simple and fast network formation, excellent QoS, strong security with 128-bit AES encryption in CCM mode and methods to coexist with other wireless networks in the band. The PHY operated in the license exempt 2.4 GHz band with data rates of 11 Mb/s to 55 Mb/s at distances of greater than 70 m.

Interest in working on an amendment to fix and enhance IEEE Std 802.15.3 began in the September 2003 meeting in Singapore. A study group was formed at the November 2003 meeting in Albuquerque, which completed a project authorization request that was approved in March 2004. Work progressed quickly, and the first task group letter ballot was successfully completed in March 2005. After one recirculation, the draft began sponsor ballot in August 2005. The ballot was successful, and after a recirculation ballot to validate some minor changes, IEEE Std 802.15.3b was approved by the IEEE Standards Board on 6 December 2005, just over two years after the study group started.

IEEE Std 802.15.3b-2005 fixed mistakes and added enhancements that improve the efficiency of the base standard. Some of the key changes/additions in this amendment are as follows:

- An improved definition of the medium access control (MAC) layer management entity (MLME) service access point (SAP).
- A new acknowledgment (ACK) policy, implied-ACK, which allows polling and a more efficient use of channel time.
- A method for relinquishing time in a channel time allocation (CTA) to allow another device (DEV) time to transmit data.
- The ability to assign device identifiers (DEVIDs) to group addresses to allow multicast connections.
- Faster recovery of network operations when the piconet coordinator (PNC) abruptly disconnects\ with the conditional handover and the next PNC processes.
- Multiple contention periods during a superframe.

In the July 2003 meeting in San Francisco, an interest group was formed to consider a mmWave PHY for the existing standard. A study group was formally created in the March 2004 IEEE 802 plenary meeting in Orlando and developed a project authorization request that was approved in March 2005. The first meeting as a task group was in May 2005 in Cairns, Australia and the group worked steadily developing channel models and evaluation documents. The PHY modes were selected in November 2007 at the Atlanta meeting

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks

and draft progressed rapidly, entering working group letter ballot in June 2008. After three working group recirculation ballots, sponsor ballot started in March 2009. A total of three sponsor recirculation ballots were held, leading to approval of IEEE Std 802.15.3c-2009 by the IEEE-SA Standards Board on 11 September 2009.

Some of the key features of IEEE Std 802.15.3c-2009 and additions are as follows:

- Operation in the 60 GHz band.
- New data rates, with the highest greater than 5 Gb/s.
- Beamforming negotiation for the transmitter to increase the communication range.
- MAC packet aggregation and the acknowledgment of individual subpackets to reduce retransmission overhead.

The PHY specifies three modes and one common mode. The three PHY modes are as follows:

- Single carrier (SC) mode optimized for low power and low complexity.
- High-speed interface (HSI) mode optimized for low-latency bidirectional data transfer.
- Audio/video (AV) mode optimized for the delivery of uncompressed, high-definition video and audio.

In 2015, there were two new projects that wanted to add additional optional PHYs. To support the development of these new amendments, the base standard needed to be revised to create a single document that incorporated the prior two amendments. A PAR was generated in July 2015 at the IEEE Plenary meeting in Waikoloa Village, HI. The PAR was approved by the IEEE SASB in September 2015 and the initial draft was completed soon thereafter (by the same Technical Editor who worked on the previous three documents). After two working group ballots, the sponsor ballot began in December 2015 and completed in February 2016 after two Sponsor ballot recirculations. IEEE Std 802.15.3-2016 was approved by the IEEE-SA Standards Board in May 2016.

Contents

1.	Overview.....	26
1.1	Scope.....	26
1.2	Purpose.....	26
2.	Normative references.....	27
3.	Definitions, acronyms, and abbreviations.....	28
3.1	Definitions.....	28
3.2	Acronyms and abbreviations.....	29
4.	General description.....	33
4.1	What is a piconet?.....	33
4.2	Components of an IEEE 802.15.3 piconet.....	33
4.3	Overview of medium access control (MAC) functionality.....	33
4.3.1	Coordination.....	34
4.3.1.1	Starting a piconet.....	34
4.3.1.2	Handing over control of the piconet.....	34
4.3.1.3	Creating a child piconet.....	34
4.3.1.4	Creating a neighbor piconet.....	35
4.3.2	Ending a piconet.....	35
4.3.3	Ending a piconet with a dependent piconet involved.....	35
4.3.4	Association and disassociation.....	36
4.3.5	Security overview.....	36
4.3.6	The IEEE 802.15.3 superframe.....	36
4.3.7	Channel time management.....	37
4.3.8	Data communications between DEVS.....	38
4.3.9	Information discovery in the piconet.....	38
4.3.10	Dynamic channel selection.....	39
4.3.11	Power management.....	39
4.3.12	Controlling transmit power in the piconet.....	40
4.3.13	Superframe structure using quasi-omni mode.....	40
4.3.14	Frame aggregation.....	41
4.3.15	Beam forming.....	41
4.3.16	Channel probing.....	41
4.3.17	Unequal error protection (UEP).....	41
4.4	Characteristics of the 2.4 GHz PHY.....	41
4.4.1	General characteristics.....	41
4.4.2	Coexistence and interoperability.....	42
4.5	Characteristics of the mmWave PHY.....	42
4.5.1	mmWave PHY characteristics.....	42
4.5.2	Piconets using mmWave PHY modes.....	43
4.6	Conventions.....	44
4.7	Concept of primitives.....	44
5.	Layer management.....	46
5.1	Overview of management model.....	46
5.2	Generic management primitives.....	47
5.2.1	MLME-GET.request.....	48

5.2.2	MLME-GET.confirm	48
5.2.3	MLME-SET.request	48
5.2.4	MLME-SET.confirm	48
5.3	MLME SAP interface	49
5.3.1	Resetting the MAC/MLME	50
5.3.1.1	MLME-RESET.request	50
5.3.1.2	MLME-RESET.confirm	50
5.3.2	Scanning for piconets	51
5.3.2.1	MLME-SCAN.request	53
5.3.2.2	MLME-SCAN.confirm	53
5.3.2.3	MLME-SCAN.indication	54
5.3.3	Starting a piconet	54
5.3.3.1	MLME-START.request	55
5.3.3.2	MLME-START.confirm	55
5.3.4	Stopping a piconet	55
5.3.4.1	MLME-STOP.request	56
5.3.4.2	MLME-STOP.confirm	56
5.3.5	Associating with a piconet	57
5.3.5.1	MLME-ASSOCIATE.request	58
5.3.5.2	MLME-ASSOCIATE.confirm	58
5.3.5.3	MLME-ASSOCIATE.indication	58
5.3.6	Disassociation from a piconet	58
5.3.6.1	MLME-DISASSOCIATE.request	59
5.3.6.2	MLME-DISASSOCIATE.confirm	59
5.3.6.3	MLME-DISASSOCIATE.indication	60
5.3.7	Security management	60
5.3.7.1	MLME-MEMBERSHIP-UPDATE.request	62
5.3.7.2	MLME-MEMBERSHIP-UPDATE.confirm	62
5.3.7.3	MLME-SECURITY-ERROR.indication	62
5.3.7.4	MLME-SECURITY-MESSAGE.request	62
5.3.7.5	MLME-SECURITY-MESSAGE.confirm	63
5.3.7.6	MLME-SECURITY-MESSAGE.indication	63
5.3.8	PNC handover	63
5.3.8.1	MLME-PNC-HANDOVER.indication	63
5.3.8.2	MLME-PNC-HANDOVER.response	64
5.3.8.3	MLME-NEW-PNC.indication	64
5.3.9	Requesting DEV information from the PNC	64
5.3.9.1	MLME-DEV-INFO.request	65
5.3.9.2	MLME-DEV-INFO.confirm	65
5.3.9.3	MLME-DEV-INFO.indication	65
5.3.10	Security information retrieval	65
5.3.10.1	MLME-SECURITY-INFO.request	65
5.3.10.2	MLME-SECURITY-INFO.confirm	66
5.3.10.3	MLME-SECURITY-INFO.indication	67
5.3.10.4	MLME-SECURITY-INFO.response	67
5.3.11	Application specific data management	67
5.3.11.1	MLME-APPLICATION-SPECIFIC.request	68
5.3.11.2	MLME-APPLICATION-SPECIFIC.confirm	69
5.3.11.3	MLME-APPLICATION-SPECIFIC.indication	69
5.3.12	Piconet services management	69
5.3.12.1	MLME-ANNOUNCE-SERVICE.request	70
5.3.12.2	MLME-ANNOUNCE-SERVICE.confirm	70
5.3.12.3	MLME-PICONET-SERVICES.request	71
5.3.12.4	MLME-PICONET-SERVICES.confirm	71

5.3.12.5	MLME-PICONET-SERVICES.indication	71
5.3.13	Stream management	71
5.3.13.1	MLME-CREATE-STREAM.request	73
5.3.13.2	MLME-CREATE-STREAM.confirm	74
5.3.13.3	MLME-CREATE-STREAM.indication	74
5.3.13.4	MLME-MODIFY-STREAM.request	74
5.3.13.5	MLME-MODIFY-STREAM.confirm	75
5.3.13.6	MLME-TERMINATE-STREAM.request	75
5.3.13.7	MLME-TERMINATE-STREAM.confirm	75
5.3.13.8	MLME-TERMINATE-STREAM.indication	75
5.3.14	Piconet parameter management	76
5.3.14.1	MLME-BSID-CHANGE.request	76
5.3.14.2	MLME-BSID-CHANGE.confirm	76
5.3.14.3	MLME-PICONET-PARM-CHANGE.indication	77
5.3.15	Power management	77
5.3.15.1	MLME-PS-SET-INFO-RATION.request	79
5.3.15.2	MLME-PS-SET-INFO-RATION.confirm	79
5.3.15.3	MLME-SPS-CONFIGURE.request	79
5.3.15.4	MLME-SPS-CONFIGURE.confirm	80
5.3.15.5	MLME-PM-MODE-CHANGE.request	80
5.3.15.6	MLME-PM-MODE-CHANGE.confirm	80
5.3.15.7	MLME-PM-MODE-CHANGE.indication	80
5.3.15.8	MLME-MONITOR-PM-MODE.request	81
5.3.15.9	MLME-MONITOR-PM-MODE.confirm	81
5.3.15.10	MLME-MONITOR-PM-MODE.indication	81
5.3.16	Multicast operations	82
5.3.16.1	MLME-MULTICAST-CONFIGURATION.request	83
5.3.16.2	MLME-MULTICAST-CONFIGURATION.confirm	83
5.3.16.3	MLME-MULTICAST-RX-SETUP.request	83
5.3.16.4	MLME-MULTICAST-RX-SETUP.confirm	84
5.3.17	Timing synchronization	84
5.3.17.1	MLME-BEACON-EVENT.request	84
5.3.17.2	MLME-BEACON-EVENT.confirm	84
5.3.17.3	MLME-BEACON-EVENT.indication	85
5.3.18	Transmit switched diversity (TSD)	85
5.3.18.1	MLME-TXDIV.request	85
5.3.18.2	MLME-TXDIV.confirm	86
5.4	MAC management	86
5.4.1	MAC PIB PNC group	86
5.4.2	MAC PIB characteristic group	87
5.5	MAC SAP	88
5.5.1	MAC-ASYNC-DATA.request	91
5.5.2	MAC-ASYNC-DATA.confirm	91
5.5.3	MAC-ASYNC-DATA.indication	91
5.5.4	MAC-ISOCH-DATA.request	92
5.5.5	MAC-ISOCH-DATA.confirm	92
5.5.6	MAC-ISOCH-DATA.indication	93
6.	MAC frame formats	94
6.1	Frame format conventions	94
6.2	General frame format	95
6.2.1	Frame Control field	96
6.2.1.1	Protocol Version field	96

6.2.1.2	Frame Type field.....	97
6.2.1.3	SEC field.....	97
6.2.1.4	ACK Policy field, Imp-ACK Request field, and Blk-ACK field	97
6.2.1.5	Retry field	97
6.2.1.6	More Data field	98
6.2.1.7	Imp-ACK Negative Acknowledgment (NAK) field.....	98
6.2.1.8	CTA Relinquish field.....	98
6.2.2	Piconet ID (PNID) field	99
6.2.3	SrcID and DestID fields	99
6.2.4	Fragmentation Control field	99
6.2.4.1	MSDU Number field	99
6.2.4.2	Fragment Number field.....	100
6.2.4.3	Last Fragment Number field.....	100
6.2.5	Stream Index field	100
6.2.6	MAC header validation	100
6.2.7	MAC Frame Body field	101
6.2.7.1	Frame Payload field.....	101
6.2.7.2	Secure session ID (SECID) field	101
6.2.7.3	Secure Frame Counter (SFC) field	101
6.2.7.4	Secure Payload field	101
6.2.7.5	Integrity Code field.....	101
6.2.7.6	Frame Check Sequence (FCS) field.....	101
6.2.8	SC and HSI aggregated frame format	102
6.2.8.1	Standard aggregation format.....	102
6.2.8.1.1	Non-secure standard aggregation format.....	102
6.2.8.1.2	Secure standard aggregation format	105
6.2.8.2	Low-latency aggregation formats	105
6.2.8.2.1	Non-secure low-latency aggregation format.....	105
6.2.8.2.2	Secure low-latency aggregation format	108
6.2.9	AV aggregated frame format	109
6.2.9.1	Extended MAC header format.....	109
6.2.9.1.1	Extended Control Header field	110
6.2.9.1.2	MAC Extension Header field	111
6.2.9.1.3	Security Header field	112
6.2.9.1.4	Video Header field.....	112
6.2.9.2	Subframe format	113
6.2.9.3	AV aggregated	115
6.2.9.4	Regular.....	115
6.2.9.5	Directional ACK.....	116
6.3	Format of individual frame types.....	116
6.3.1	Beacon frame	116
6.3.1.1	Non-secure Beacon frame.....	116
6.3.1.2	Secure Beacon frame	118
6.3.2	Acknowledgment frames	119
6.3.2.1	Imm-ACK frame.....	119
6.3.2.2	Delayed ACK (Dly-ACK) frame.....	120
6.3.3	Command frame	121
6.3.3.1	Non-secure Command frame.....	121
6.3.3.2	Secure Command frame.....	122
6.3.4	Data frame	123
6.3.4.1	Non-secure Data frame	123
6.3.4.2	Secure Data frame.....	123
6.3.5	Multi-protocol Data frame	124
6.3.5.1	Non-secure Multi-protocol Data frame.....	124

6.3.5.2	Secure Multi-protocol Data frame	125
6.3.6	Sync frame	125
6.4	Information elements (IEs)	126
6.4.1	CTA IE	128
6.4.2	BSID IE	129
6.4.3	Parent Piconet IE	129
6.4.4	DEV Association IE	129
6.4.5	PNC shutdown IE	130
6.4.6	Piconet Parameter Change IE	131
6.4.7	AS IE	132
6.4.8	Pending channel time map (PCTM) IE	132
6.4.9	PNC Handover IE	133
6.4.10	CTA Status IE	133
6.4.11	Capability IE	134
6.4.12	Transmit Power Parameters IE	140
6.4.13	PS Status IE	141
6.4.14	CWB IE	141
6.4.15	Overlapping PNID IE	142
6.4.16	Piconet Services IE	142
6.4.17	Vendor Defined IE	143
6.4.18	Group ID IE	143
6.4.19	Stream Renew IE	144
6.4.20	Next PNC IE	144
6.4.21	Piconet Channel Status IE	144
6.4.22	Synchronization IE	145
6.4.23	TSD IE	147
6.4.24	UEP Specific IE	148
6.4.25	IFS IE	148
6.4.26	CTA Relinquish Duration IE	149
6.4.27	Feedback IE	149
6.4.28	Mapping IE	151
6.4.29	BST Clustering IE	152
6.4.30	PET Clustering IE	152
6.4.31	Beam PET IE	153
6.4.32	HRS Beam PET IE	154
6.4.33	PET Amplitude IE	155
6.4.34	PET Phase IE	155
6.4.35	Sync Frame Frequency IE	156
6.4.36	Directional Peer IE	156
6.5	MAC commands	157
6.5.1	Association and disassociation commands	159
6.5.1.1	Association Request command	159
6.5.1.2	Association Response command	160
6.5.1.3	Disassociation Request command	161
6.5.2	Security commands	161
6.5.2.1	Request Key command	161
6.5.2.2	Request Key Response command	161
6.5.2.3	Distribute Key Request command	162
6.5.2.4	Distribute Key Response command	162
6.5.3	PNC handover commands	162
6.5.3.1	PNC Handover Request command	163
6.5.3.2	PNC Handover Response command	163
6.5.3.3	PNC Handover Information command	164
6.5.4	Information request commands	164

6.5.4.1	PNC Information Request command	164
6.5.4.2	PNC Information command	164
6.5.4.3	Security Information Request command	166
6.5.4.4	Security Information command	166
6.5.4.5	Probe Request command	167
6.5.4.6	Probe Response command	169
6.5.5	Information announcement commands	171
6.5.5.1	Piconet Services command	171
6.5.5.2	Announce command	171
6.5.5.3	Announce Response command	173
6.5.6	Channel time allocation request, modification, and termination commands	173
6.5.6.1	Channel Time Request command	173
6.5.6.2	Channel Time Response command	176
6.5.7	Channel status commands	177
6.5.7.1	Channel Status Request command	177
6.5.7.2	Channel status response	177
6.5.7.3	Remote Scan Request command	178
6.5.7.4	Remote Scan Response command	178
6.5.7.5	Transmit Power Change command	180
6.5.8	Power management commands	180
6.5.8.1	PS Set Information Request command	180
6.5.8.2	PS Set Information Response command	180
6.5.8.3	SPS Configuration Request command	181
6.5.8.4	SPS Configuration Response command	182
6.5.8.5	PM Mode Change command	183
6.5.8.6	PM Mode Change Response command	183
6.5.9	Special commands	184
6.5.9.1	Security Message command	184
6.5.9.2	Vendor Defined command	184
6.5.9.3	AS IE Request command	185
6.5.9.4	AS IE Response command	185
6.5.10	Multicast configuration commands	186
6.5.10.1	Multicast Configuration Request command	186
6.5.10.2	Multicast Configuration Response command	186
7.	MAC functional description	187
7.1	Introduction	187
7.2	Starting, maintaining, and stopping piconets	188
7.2.1	Scanning through channels	188
7.2.2	Starting a piconet	189
7.2.3	PNC handover	190
7.2.4	Preliminary handover	193
7.2.5	Next PNC	194
7.2.6	Dependent PNC handover	195
7.2.7	Child piconet	198
7.2.8	Neighbor piconet	200
7.2.9	Stopping piconet operations	202
7.2.9.1	Stopping an independent or parent piconet	203
7.2.9.2	Parent PNC stopping a dependent piconet	203
7.2.9.3	Dependent PNC termination of a dependent piconet	204
7.2.10	Non-PNC-capable DEVs	204
7.3	Association and disassociation with a piconet	204
7.3.1	Association	204

7.3.2	Piconet services	206
7.3.3	Broadcasting piconet information	208
7.3.4	Disassociation	209
7.4	Channel access	210
7.4.1	Interframe space (IFS)	211
7.4.2	Carrier sense multiple access with collision avoidance (CSMA/CA)	211
7.4.3	Channel time allocation period channel access	213
7.4.3.1	Channel time allocations (CTAs)	213
7.4.3.2	Channel time allocation (CTA) and channel time usage	215
7.4.3.3	Management CTAs	216
7.4.3.4	Slotted aloha access	216
7.4.3.5	Allocation of MCTAs	217
7.4.3.6	Guard time	217
7.4.3.7	Calculating channel time requests	219
7.4.3.8	Relinquishing CTA time to another DEV	220
7.5	Channel time management	222
7.5.1	Isochronous stream management	222
7.5.1.1	Isochronous stream creation	223
7.5.1.2	Isochronous stream modification	226
7.5.1.3	Isochronous stream termination	229
7.5.2	Asynchronous channel time reservation and termination	231
7.5.2.1	Asynchronous channel time reservation	231
7.5.2.2	Asynchronous channel time termination	233
7.5.3	Multicast group configuration	234
7.6	Synchronization	235
7.6.1	Time accuracy	236
7.6.2	Beacon generation	236
7.6.3	Beacon reception	237
7.6.4	Beacon information announcement	237
7.6.5	Superframe support for directional PHYs	238
7.6.5.1	Quasi-omni beacon	239
7.6.5.2	Directional CAP	240
7.6.5.3	Directional association	241
7.6.5.4	DEV quasi-omni transmit direction training and tracking	242
7.7	Fragmentation and defragmentation	243
7.8	Aggregation	244
7.8.1	Standard aggregation	244
7.8.2	Low-latency aggregation	245
7.9	Acknowledgment and retransmission	246
7.9.1	No acknowledgment (no-ACK)	246
7.9.2	Immediate acknowledgment (Imm-ACK)	246
7.9.3	Delayed acknowledgment (Dly-ACK)	247
7.9.4	Implied acknowledgment (Imp-ACK)	248
7.9.5	Block acknowledgment (Blk-ACK)	250
7.9.5.1	Blk-ACK for standard aggregation	250
7.9.5.2	Blk-ACK for low-latency aggregation	251
7.9.6	Retransmissions	252
7.9.7	Duplicate detection	252
7.10	Peer discovery	252
7.10.1	PNC information request	252
7.10.2	Probe request and response	253
7.10.3	Announce	254
7.10.4	Channel status request	254
7.10.5	Remote scan	254

7.10.6	PNC channel scanning	255
7.10.7	Channel probing	255
7.11	Changing piconet parameters.....	256
7.11.1	Moving beacon	256
7.11.2	Changing superframe duration	256
7.11.3	Setting the PNID or BSID	257
7.11.4	Maintaining synchronization in dependent piconets	258
7.12	Interference mitigation.....	258
7.12.1	Dynamic channel selection	259
7.12.2	Transmit power control	260
7.12.2.1	Maximum transmitter power for the CAP, beacon, and directed MCTAs	260
7.12.2.2	Adjustable transmitter power in a CTA	261
7.13	Multi-rate support	261
7.14	Power management.....	262
7.14.1	Piconet synchronized power save (PSPS) mode	263
7.14.2	Device synchronized power save (DSPS) mode	264
7.14.2.1	Creation, use, and management of DSPS sets	264
7.14.2.2	Changing DSPS mode and operation.....	265
7.14.2.3	CTA timing in DSPS mode	266
7.14.3	Asynchronous power save (APS) mode	268
7.14.4	MSCs for power save modes	269
7.15	AS IE operation	272
7.16	MAC sublayer parameters	275
7.17	Unequal error protection (UEP).....	278
7.18	Sync frame transmission and virtually dependent piconet	278
8.	Security	280
8.1	Security mechanisms	280
8.1.1	Security membership and key establishment	280
8.1.2	Key transport	280
8.1.3	Data encryption	280
8.1.4	Data integrity	280
8.1.5	Beacon integrity protection	280
8.1.6	Command integrity protection	281
8.1.7	Freshness protection	281
8.2	Security modes.....	281
8.2.1	Security mode 0	281
8.2.2	Security mode 1	281
8.3	Security support.....	281
8.3.1	PNC handover	281
8.3.2	Changes in the piconet group data key	282
8.3.3	Joining a secure piconet	282
8.3.4	Membership update	282
8.3.5	Secure frame generation	283
8.3.6	Updating CurrentTimeToken	284
8.3.7	Secure frame reception	284
8.3.8	Selecting the SECID for a new key	285
8.3.9	Key selection	285
8.4	Protocol details	289
8.4.1	Security information request and distribution	289
8.4.2	Key distribution protocol	291
8.4.3	Key request protocol	291

9.	Security specifications	292
9.1	Modes for security	292
9.2	Symmetric cryptography building blocks	292
9.2.1	Notational conventions	292
9.2.2	CCM combined encryption and data authentication	292
9.2.3	CCM parameters	292
9.2.4	Nonce value	292
9.2.5	AES encryption	293
9.3	Symmetric cryptography implementation	293
9.3.1	Symmetric cryptography data formats	293
9.3.2	Symmetric cryptographic operations	293
9.4	CCM mode	295
9.4.1	Inputs	295
9.4.2	Data authentication	296
9.4.3	Encryption	298
9.4.4	Output	298
9.4.5	Decryption	299
9.4.6	Restrictions	299
9.4.7	List of symbols	299
10.	PHY specification for the 2.4 GHz band	301
10.1	Overview of the 2.4 GHz PHY	301
10.2	General requirements	302
10.2.1	Operating frequency range	302
10.2.2	Radio frequency (RF) power measurements	302
10.2.3	Channel assignments	302
10.2.4	Scanning channels	302
10.2.5	Unwanted emissions	303
10.2.6	Operating temperature range	303
10.2.7	PHY layer timing	303
10.2.7.1	Interframe space	303
10.2.7.2	Receive-to-transmit turnaround time	303
10.2.7.3	Transmit-to-receive turnaround time	304
10.2.7.4	Time between successive transmissions	304
10.2.7.5	Channel switch time	304
10.2.8	Data size restrictions	304
10.2.8.1	Maximum frame length	304
10.2.8.2	Maximum transfer unit size	304
10.2.8.3	Minimum fragment size	304
10.2.9	Header check sequence	304
10.2.10	Channel access methods	306
10.3	Modulation and coding	306
10.3.1	Base data rate	307
10.3.2	Signal constellations	307
10.3.3	DQPSK modulation	308
10.3.4	QPSK and 16/32/64-QAM with trellis coding	309
10.4	PHY frame format	316
10.4.1	Frame format	316
10.4.2	PHY preamble	317
10.4.3	Header modulation	319
10.4.4	Scrambling	319
10.4.5	PHY header	321

10.4.6	Stuff bits	322
10.4.7	Tail symbols	322
10.5	Transmitter specifications	325
10.5.1	Error vector magnitude definition	325
10.5.2	EVM calculated values	326
10.5.3	Transmit power spectral density (PSD) mask	326
10.5.4	Transmit center frequency tolerance	326
10.5.5	Symbol rate	326
10.5.6	Clock synchronization	327
10.5.7	Transmit power-on and power-down ramp	327
10.5.8	RF carrier suppression	327
10.5.9	Transmit power	327
10.6	Receiver specifications	328
10.6.1	Error rate criterion	328
10.6.2	Receiver sensitivity	328
10.6.3	Receiver maximum input level	329
10.6.4	Receiver jamming resistance	329
10.6.5	Receiver CCA performance	330
10.6.6	Receiver RSSI	330
10.6.7	Link quality indication	330
10.7	PHY management	330
11.	PHY specification for millimeter wave	333
11.1	General requirements	333
11.1.1	Regulatory information	333
11.1.2	RF power measurements	334
11.1.3	Unwanted emissions	334
11.1.4	Operating temperature range	334
11.1.5	RF channelization	334
11.1.6	Transmit PSD mask	334
11.1.7	Error vector magnitude calculation	334
11.1.7.1	SC PHY	334
11.1.7.2	AV and HSI PHYs	335
11.1.8	Common PHY management for mmWave PHY modes	336
11.1.8.1	Supported MCSs	336
11.1.8.2	Preferred fragment size	336
11.1.8.3	Receive Status field	337
11.1.9	Requirements for mmWave PNCs	338
11.1.10	CP operation	338
11.1.11	mmWave PHY mode usage in CTA	339
11.1.12	Common mode signaling (CMS)	339
11.1.12.1	Forward error correction for CMS	340
11.1.12.2	Code spreading for CMS	340
11.1.12.3	Scrambling for CMS	340
11.1.12.4	PHY preamble for CMS	341
11.1.12.5	Frame header for CMS	341
11.1.12.6	PHY header for CMS	342
11.1.12.7	PHY Payload field for CMS	343
11.1.12.8	Receiver clear channel assessment performance for CMS	343
11.1.13	mmWave PHY PIB	343
11.2	Single Carrier Mode of mmWave PHY	344
11.2.1	PHY operating specifications of SC PHY	345
11.2.1.1	Channelization	345

11.2.1.2	Scanning channels.....	345
11.2.2	Modulation, forward error correction, and spreading	345
11.2.2.1	MCS-dependent parameters.....	345
11.2.2.2	Header rate-dependent parameters.....	347
11.2.2.3	Timing-related parameters.....	348
11.2.2.4	Frame-related parameters	348
11.2.2.5	Modulation.....	349
11.2.2.5.1	p/2 BPSK	349
11.2.2.5.2	p/2 QPSK	349
11.2.2.5.3	p/2 8-PSK.....	350
11.2.2.5.4	p/2 16-QAM.....	351
11.2.2.6	FEC	351
11.2.2.6.1	Reed-Solomon block codes in GF(28).....	351
11.2.2.6.2	Irregular LDPC codes	352
11.2.2.6.3	Rate 14/15 LDPC code	353
11.2.2.7	Stuff bits.....	353
11.2.2.8	Code spreading	356
11.2.2.8.1	Golay sequences	356
11.2.2.8.2	PRBS generation with LFSR	356
11.2.2.9	Unequal error protection.....	357
11.2.2.10	Scrambling	358
11.2.3	SC PHY frame format	359
11.2.3.1	PHY preamble.....	360
11.2.3.1.1	Sync field	360
11.2.3.1.2	SFD field.....	360
11.2.3.1.3	Channel estimation sequence (CES).....	360
11.2.3.2	Frame header.....	360
11.2.3.2.1	SC PHY header.....	362
11.2.3.2.2	Base header HCS	363
11.2.3.2.3	Base header FEC.....	363
11.2.3.2.4	MAC subheader HCS	363
11.2.3.2.5	MAC subheader FEC.....	364
11.2.3.3	SC PHY Payload field	364
11.2.3.3.1	SC PHY Payload scrambling.....	364
11.2.3.3.2	Modulation.....	364
11.2.3.3.3	FEC	364
11.2.3.3.4	Code spreading	364
11.2.3.4	Pilot word and PCES	365
11.2.3.4.1	Subblocks and pilot word	365
11.2.3.4.2	Pilot channel estimation sequence (PCES).....	365
11.2.4	Transmitter specifications	366
11.2.4.1	Error vector magnitude	366
11.2.4.2	Transmit center frequency tolerance.....	366
11.2.4.3	Symbol rate.....	366
11.2.4.4	Transmit power-on and power-down ramp	366
11.2.5	Receiver specifications	366
11.2.5.1	Error rate criterion	366
11.2.5.2	Receiver sensitivity.....	367
11.2.5.3	Receiver maximum input level.....	367
11.2.5.4	Receiver clear channel assessment performance	367
11.2.6	PHY layer timing	368
11.2.6.1	Interframe space.....	368
11.2.6.2	Receive-to-transmit turnaround time	368
11.2.6.3	Transmit-to-receive turnaround-time.....	368

11.2.6.4	Time between successive transmissions	368
11.2.6.5	Channel switch.....	369
11.2.7	PHY management for SC PHY	369
11.2.7.1	Maximum frame size	369
11.2.7.2	Maximum transfer unit size	369
11.2.7.3	Minimum fragment size.....	369
11.2.8	Optional OOK/DAMI modes	369
11.2.8.1	OOK.....	369
11.2.8.2	DAMI.....	370
11.2.8.3	FEC.....	370
11.2.8.4	OOK/DAMI spreading	370
11.3	High Speed Interface mode of mmWave PHY	370
11.3.1	General operating specifications	370
11.3.1.1	Operating frequency bands	370
11.3.2	HSI PHY modulation, forward error correction, and spreading	370
11.3.2.1	MCS-dependent parameters.....	370
11.3.2.2	HSI PHY timing-related parameters.....	372
11.3.2.3	HSI PHY frame-related parameters.....	373
11.3.2.4	HSI PHY FEC.....	373
11.3.2.4.1	LDPC block code.....	374
11.3.2.4.2	EEP data multiplexer	374
11.3.2.4.3	UEP data multiplexer.....	374
11.3.2.4.4	Bit interleaver	375
11.3.2.5	Stuff bits.....	376
11.3.2.6	Constellation mapping	376
11.3.2.7	HSI spreader	378
11.3.2.7.1	Spreader for spreading factor of 1	379
11.3.2.7.2	Spreader for spreading factor of 48	379
11.3.2.8	Tone interleaver	379
11.3.2.9	HSI PHY OFDM modulator	380
11.3.2.9.1	Pilot subcarriers	381
11.3.2.9.2	Guard subcarriers.....	382
11.3.2.10	PCES insertion.....	382
11.3.2.11	HSI PHY scrambling	382
11.3.3	HSI PHY frame format	383
11.3.3.1	PHY preamble.....	383
11.3.3.2	Frame header.....	384
11.3.3.3	HSI PHY header	386
11.3.3.4	Header check sequences	387
11.3.3.4.1	Main Header HCS.....	387
11.3.3.4.2	MAC Subheader HCS.....	387
11.3.3.5	PHY Payload field	388
11.3.4	Transmitter specifications	388
11.3.4.1	EVM requirement	388
11.3.4.2	Chip rate and clock alignment	389
11.3.5	Receiver specifications	389
11.3.5.1	Receiver sensitivity.....	389
11.3.5.2	Receiver CCA performance.....	389
11.3.5.3	Receiver maximum input level.....	389
11.3.5.4	PHY layer timing	389
11.3.5.5	Interframe spacing	389
11.3.5.6	Receive-to-transmit turnaround time	390
11.3.5.7	Transmit-to-receive turnaround time	390
11.3.5.8	Time between successive transmissions.....	390

11.3.5.9	Channel switch time.....	390
11.3.6	HSI PHY management	390
11.3.6.1	PHY supported data rate encoding	390
11.3.6.2	HSI PHY fragment size encoding.....	390
11.3.6.3	Maximum frame length	391
11.3.6.4	Maximum transfer unit size	391
11.3.6.5	Minimum fragment size.....	391
11.4	Audio/Visual mode of mmWave PHY	391
11.4.1	General requirements	392
11.4.1.1	AV PHY channelization	392
11.4.1.2	PHY layer timing	393
11.4.1.2.1	Interframe space.....	393
11.4.1.2.2	Receive-to-transmit turnaround time.....	394
11.4.1.2.3	Transmit-to-receive turnaround time.....	394
11.4.1.2.4	Time between successive transmissions.....	394
11.4.1.2.5	Channel switch time.....	394
11.4.1.2.6	CCA detect time.....	394
11.4.1.3	Data size restrictions	394
11.4.1.3.1	Maximum frame length	394
11.4.1.3.2	Maximum transfer unit size	395
11.4.1.3.3	Minimum fragment size.....	395
11.4.1.4	Header check sequence	395
11.4.2	AV PHY modulation and forward error correction	395
11.4.2.1	AV PHY base rate.....	396
11.4.2.2	Repetition coding and spatial diversity.....	397
11.4.2.3	Stuff bits.....	397
11.4.2.4	HRP splitter and scrambler	398
11.4.2.5	LRP scrambler	398
11.4.2.6	HRP outer code	399
11.4.2.7	HRP outer interleaver	399
11.4.2.8	Convolutional encoder	400
11.4.2.9	Puncturing	400
11.4.2.10	HRP data multiplexer and bit interleaver	400
11.4.2.10.1	EEP data multiplexer	401
11.4.2.10.2	UEP coding data multiplexer	402
11.4.2.10.3	UEP mapping data multiplexer.....	404
11.4.2.10.4	HRP header data multiplexer	405
11.4.2.11	Bit reversal tone interleaver	406
11.4.2.12	Signal constellations	406
11.4.2.13	AV PHY non-data subcarriers	407
11.4.2.14	AV PHY OFDM modulation.....	408
11.4.3	AV PHY frame formats	409
11.4.3.1	HRP preamble.....	410
11.4.3.2	HRP header	411
11.4.3.3	LRP preamble sequences	411
11.4.3.4	Long omni LRP preamble.....	412
11.4.3.5	Short omni LRP preamble	413
11.4.3.6	Omni LRP header	414
11.4.3.7	Directional LRP preamble	414
11.4.3.8	Directional LRP header and payload	415
11.4.4	AV PHY transmitter requirements	416
11.4.4.1	TX mask.....	416
11.4.4.2	EVM requirement	417
11.4.4.3	Symbol timing.....	417

11.4.4.4	TX frequency accuracy	418
11.4.4.5	TX power ramp on and off	418
11.4.5	AV PHY Receiver characteristics	418
11.4.5.1	Error rate criterion	418
11.4.5.2	Sensitivity	418
11.4.5.3	Maximum input level	418
11.4.6	Preambles and training symbols	418
12.	Beam forming	423
12.1	Introduction	423
12.2	Beam forming terminology	423
12.2.1	Quasi-omni patterns	424
12.2.2	Sectors	424
12.2.3	Beams	425
12.2.4	Clusters	426
12.3	Beam forming codebooks	427
12.3.1	Beam forming codebooks	428
12.4	Beam forming reference model	429
12.5	Beam forming protocol	431
12.5.1	Two-level training mechanism	431
12.5.1.1	Sector-level training	431
12.5.1.1.1	AAS sector-level training	431
12.5.1.1.2	SAS sector-level training	434
12.5.1.1.3	Sector-level training failure remedy	437
12.5.1.2	Beam-level training	437
12.5.1.2.1	AAS beam-level training	437
12.5.1.2.2	SAS beam-level training	440
12.5.2	Beam tracking	443
12.6	On-demand beam forming	445
12.7	Proactive beam forming	445
12.8	TSD	445
Annex A	(informative) Bibliography	447
Annex B	(normative) Frame convergence sublayer	448
B.1	Generic convergence sublayer	448
B.2	EPD FCSL PDU classification	449
B.3	EPD FCSL	449
B.3.1	EPD FCSL QoS support	450
B.3.2	Data entity inter-relationships	450
B.4	EPD FCSL SAP	451
B.4.1	MA-UNITDATA.request	452
B.4.2	MA-UNITDATA.indication	452
B.4.3	MA-UNITDATA-STATUS.indication	453
B.5	Stream SAP	454
B.5.1	Stream creation, modification, and deletion	454
B.5.1.1	STREAM_INITIATE.request	456
B.5.1.2	STREAM_INITIATE.confirm	456
B.5.1.3	STREAM_INITIATE.indication	456
B.5.1.4	STREAM_MODIFY.request	457
B.5.1.5	STREAM_MODIFY.confirm	457
B.5.1.6	STREAM_END.request	457

B.5.1.7	STREAM_END.confirm.....	458
B.5.1.8	STREAM_END.indication	458
B.5.2	Stream data interface	458
B.5.2.1	STREAM_DATA.request.....	459
B.5.2.2	STREAM_DATA.confirm.....	460
B.5.2.3	STREAM_DATA.indication	460
Annex C	(informative) Security considerations	461
C.1	Background assumptions	461
C.1.1	Physical assumptions	461
C.1.2	Network assumptions.....	461
C.1.3	Attack model assumptions.....	462
C.1.4	Security key lifecycle issues	462
C.1.4.1	Key lifecycle	462
C.1.4.2	Membership lifecycle.....	462
C.1.4.3	Group membership change rekey.....	462
C.2	Claimed security services	462
C.2.1	Beacon protection protocol	462
C.2.2	Distribute key protocol	463
C.2.3	Key request protocol	463
C.2.4	Data protection protocol	464
C.3	Properties of the IEEE 802.15.3 security suite	464
C.3.1	Key usage.....	464
C.3.2	Replay prevention	464
Annex D	(informative) Coexistence, interoperability, and interference.....	466
D.1	Interoperability.....	466
D.1.1	Interoperability with IEEE 802.11 DSSS and IEEE 802.11b.....	466
D.1.2	IEEE 802.11 FHSS and IEEE 802.15.1.....	467
D.2	Coexistence.....	467
D.2.1	Coexistence with IEEE Std 802.11b.....	467
D.2.1.1	Passive scanning.....	468
D.2.1.2	Dynamic channel selection	468
D.2.1.3	The ability to request channel quality information	468
D.2.1.4	Link quality and RSSI.....	468
D.2.1.5	Channel plan that minimizes channel overlap	468
D.2.1.6	Lower transmit power	469
D.2.1.7	Transmit power control	469
D.2.1.8	Neighbor piconet capability	469
D.2.2	Coexistence with IEEE 802.15.1 and IEEE 802.11 FHSS	470
D.2.2.1	PHY collaborative coexistence enhancements.....	470
D.2.2.2	MAC coexistence via collaboration	470
D.2.2.3	Other techniques.....	471
D.3	Coexistence performance.....	471
D.3.1	Allowed operation.....	471
D.3.2	Assumptions for coexistence calculations	471
D.3.3	Performance impact on IEEE 802.15.3 piconets	473
D.3.3.1	IEEE 802.11b overlapping with IEEE 802.15.3	473
D.3.3.2	IEEE 802.15.1 and IEEE 802.11 FHSS overlapping with IEEE 802.15.3	476
D.4	Notes on the calculations	477

Annex E (normative) Protocol implementation conformance statement (PICS) proforma.....	478
E.1 Introduction.....	478
E.1.1 Scope.....	478
E.1.2 Purpose.....	478
E.2 Abbreviations and special symbols.....	478
E.3 Instructions for completing the PICS proforma.....	479
E.4 Identification of the implementation.....	479
E.5 Identification of the protocol	480
E.6 Global statement of conformance	481
E.7 PICS proforma—IEEE Std 802.15.3-2016.....	481
E.7.1 Major roles for IEEE 802.15.3 DEVs.....	481
E.7.2 PHY functions.....	482
E.7.3 Major capabilities for the MAC sublayer	483
E.7.3.1 MAC frames	483
E.7.3.2 MAC sublayer functions	486
E.7.4 Security support	490
Annex F (informative) Implementation considerations.....	491
F.1 Channel time requests.....	491
F.1.1 Types of CTAs.....	491
F.1.2 Interpretation of channel time requests.....	492
F.1.3 Determining CTA Rate Factor from stream requirements	493
F.1.3.1 Example 1.....	494
F.1.3.2 Example 2.....	495
F.1.4 PNC interpretation of CTA rate factor	495
F.1.5 Creating channel time requests from MLME requests	496
F.1.6 Interpreting channel time requests.....	498
F.2 Sample frames.....	500
F.3 Use of multiple mmWave PHYs in a single piconet	502
Annex G (normative) Optional OOK/DAMI modes	504
G.1 Introduction.....	504
G.2 Child piconet operation.....	504
G.3 DAMI	504
G.4 PHY preamble.....	506
G.5 PHY frame format	507
G.6 Interframe space.....	507
G.7 Eye opening for OOK	508
G.8 EVM for DAMI	508

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017

IEEE Standard for High Data Rate Wireless Multi-Media Networks

IMPORTANT NOTICE: IEEE Standards documents are not intended to ensure safety, security, health, or environmental protection, or ensure against interference with or from other devices or networks. Implementers of IEEE Standards documents are responsible for determining and complying with all appropriate safety, security, environmental, health, and interference protection practices and all applicable laws and regulations.

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

1. Overview

1.1 Scope

This standard defines PHY and MAC specifications for high data rate wireless connectivity (typically over 200 Mbps) with fixed, portable, and moving devices. Data rates are high enough to satisfy a set of consumer multimedia industry needs, as well as to support emerging wireless switched point-to-point and high rate close proximity point-to-point applications.

1.2 Purpose

The purpose of this standard is to provide for low complexity, low cost, low power consumption, and high data rate wireless connectivity among devices that support a variety of applications such as a set of consumer multimedia industry needs, wireless switched point-to-point applications in data centers, wireless backhaul/fronthaul intra-device communications, and a wide variety of additional use cases such as rapid large multimedia data downloads and file exchanges between two devices in close proximity, including between mobile devices and stationary devices (kiosks, ticket gates, etc.), and/or wireless data storage devices.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used; therefore, each referenced document is cited in the text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

ANSI X3.66-1979: Advanced data communication control procedures (ADCCP).¹

IEEE Std 802[®]-2014, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.^{2, 3}

ISO/IEC 646:1991, Information Technology—ISO 7-bit coded character set for information interchange.⁴

NIST FIPS Pub 197: Advanced Encryption Standard (AES), Federal Information Processing Standards Publication 197, US Department of Commerce/N.I.S.T., November 26, 2001.⁵

¹ ANSI publications are available from the American National Standards Institute (<http://www.ansi.org>).

² IEEE publications are available from the Institute of Electrical and Electronics Engineers (<http://standards.ieee.org/>).

³ The IEEE standards referred to in Clause 2 are trademarks belonging to the Institute of Electrical and Electronics Engineers, Inc.

⁴ ISO/IEC publications are available from the International Organization for Standardization (<http://www.iso.ch/>) and the International Electrotechnical Commission (<http://www.iec.ch/>). ISO/IEC publications are also available in the United States from the American National Standards Institute (<http://www.ansi.org>).

⁵ NIST FIPS publications are available from the National Institute for Standards and Technology (<http://www.nist.gov>).

3. Definitions, acronyms, and abbreviations

3.1 Definitions

For this standard, the following terms and definitions apply. *The IEEE Standards Dictionary Online* should be consulted for terms not defined in this clause.⁶

ad hoc network: A network typically created in a spontaneous manner. The principal characteristic of an ad hoc network is its limited temporal and spatial extent.

association: The service used to assign a device identifier to a device to enable communications in a piconet.

channel time allocation: A contiguous period of time in the superframe allocated by the piconet coordinator for communication between specified source and destination.

coexistence: The ability of one system to perform a task in a given shared environment where other systems have an ability to perform their tasks and may or may not be using the same set of rules.

data authentication: Authentication of the sender of the data and provision of data integrity.

data integrity: The assurance that the data has not been modified from its original form.

dBm: Decibels relative to a milliwatt.

device: An entity that implements an IEEE 802.15.3 conformant media access control and physical layer interface to the wireless medium.

device address: The EUI-48 associated with device in an IEEE 802.15.3 piconet.

disassociation: The service that removes an existing association.

frame: Format of aggregated bits that are transmitted together in time.

guard interval: A time period placed at the front of a symbol for inter-symbol interference avoidance.

integrity code: A data string generated using a symmetric key that is typically appended to data in order to provide data integrity and source authentication similar to a digital signature.

interoperable: The ability of two systems to perform a given task using a single set of rules.

medium access control command data unit: The unit of data exchanged between two peer medium access control entities using the services of the physical layer to implement the medium access control management protocol.

medium access control protocol data unit: The unit of data exchanged between two peer medium access control entities using the services of the physical layer.

medium access control service data unit: Information that is delivered as a unit between medium access control service access points.

⁶ The *IEEE Standards Dictionary Online* is available at
http://www.ieee.org/portal/innovate/products/standard/standards_dictionary.html.

payload protection: The generic term for providing security services on the contents of a data message, including confidentiality, integrity, and authentication.

piconet: A collection of one or more logically associated devices that share a single identifier with a common coordinator.

piconet coordinator: An entity that has device functionality and also provides coordination and other services, e.g., quality of service, synchronization, association, via the wireless medium for associated devices.

quality of service: A collective measure of the level of service (e.g., availability, latency, throughput, error rate) delivered between devices.

secure frame: A command or data frame in which cryptographic techniques are applied to provide encryption or integrity.

secure piconet: A piconet in which cryptographic techniques are implemented to provide security services.

stream: A unidirectional, logical data connection between two devices that may or may not have quality of service requirements associated with it.

superframe: The basic time division of an IEEE 802.15.3 piconet containing a beacon, the channel time allocation period and optionally the contention access period.

symmetric key: A secret key shared between two or more parties that may be used for encryption/decryption or integrity protection/integrity verification.

wireless medium: The medium used to implement the transfer of protocol data units between peer physical layer entities of a wireless personal area network.

3.2 Acronyms and abbreviations

AAS	asymmetric antenna system
ACK	acknowledgment
AES	advanced encryption standard
AS	application specific
ATP	association timeout period
AV	audio/visual
AWGN	additive white Gaussian noise
BBIFS	beam-forming beam-level interframe spacing
BCRD	bidirectional channel time allocation relinquish duration
BcstID	broadcast identifier
BER	bit error rate
BIFS	backoff interframe space
Blk-ACK	block acknowledgment
BPSK	binary phase-shift keying
BSID	beacon source identifier
BSIFS	beam-forming sector-level interframe spacing
BST	beam switching/steering and tracking
CAP	contention access period

CBC	cipher block chaining
CBC-MAC	cipher block chaining-message authentication code
CCA	clear channel assessment
CCM	counter mode encryption and cipher block chaining message authentication code
CES	channel estimation sequence
CID	Company ID
CMS	common mode signaling
CP	contention period
CRC	cyclic redundancy check
CSMA/CA	carrier sense multiple access with collision avoidance
CTA	channel time allocation
CTAP	channel time allocation period
CWB	continued wake beacon
DAMI	dual alternate mark inversion
DestID	destination identifier
DEV	device
DEVID	device identifier
Dly-ACK	delayed acknowledgment
DME	device management entity
DQPSK	differential quadrature phase-shift keying
DSPS	device synchronized power save
EEP	equal error protection
EIRP	equivalent isotropic radiated power
EPD	Ethertype protocol discrimination
FCS	frame check sequence
FCSL	frame convergence sublayer
FEC	forward error correction
FER	frame error rate
HCS	header check sequence
HR	high rate
HRP	high-rate physical layer
HRPDU	high-rate protocol data unit
HRS	high resolution
HSI	high speed interface
ID	identifier
IE	information element
IFS	interframe space
Imm-ACK	immediate acknowledgment
Imp-ACK	implied acknowledgment
KO	key originator
LAN	local area network
LDPC	low-density parity check
LFSR	linear feedback shift register
LLC	logical link control
LQI	link quality indication
LRP	low-rate physical layer

LRPDU	low-rate protocol data unit
LSB	least significant bit
MAC	medium access control
MAN	metropolitan area network
MCDU	MAC command data unit
MCS	modulation and coding scheme
McstID	multicast identifier
MCTA	management channel time allocation
MIFS	minimum interframe space
MLME	MAC layer management entity
mmWave	millimeter wave
MPDU	MAC protocol data unit
MPR	mandatory PHY rate
MR	medium rate
MSB	most significant bit
MSC	message sequence chart
MSDU	MAC service data unit
NAK	negative acknowledgment
NbrID	neighbor identifier
NLOS	non-line-of-sight
OFDM	orthogonal frequency-division multiplexing
OOK	on-off keying
OUI	organizationally unique identifier
PAL	protocol adaptation layer
PAN	personal area network
PCES	pilot channel estimation sequence
PCTM	pending channel time map
PDU	protocol data unit
PET	pattern estimation and tracking
PHY	physical layer
PIB	PAN information base
PLME	PHY layer management entity
PN	pseudo noise
PNC	piconet coordinator
PNCID	piconet coordinator identifier
PNID	piconet identifier
PPDU	PHY protocol data unit
PRBS	pseudo-random bit sequence
PRNG	pseudo-random number generator
PS	power save
PSD	power spectral density
PSPS	piconet synchronized power save
PSRC	power source
QAM	quadrature amplitude modulation
QoS	quality of service
QPSK	quadrature phase-shift keying

QT	quasi-omni training
RF	radio frequency
RIFS	retransmission interframe space
RS	Reed Solomon
RSSI	received signal strength indication
RSSIR	received signal strength indication relative to sensitivity
RX	receive or receiver
SAP	service access point
SAS	symmetric antenna system
SC	single carrier
S-CAP	sub-contention access period
SDU	service data unit
SEC	security
SECID	security identifier
SFC	secure frame counter
SIFS	short interframe space
SINR	signal-to-interference-plus-noise ratio
SNR	signal-to-noise ratio
SPS	synchronous power save
SrcID	source identifier
SSB	single side band
ST	sector training
STP	stream timeout period
TCM	trellis coded modulation
TPC	transmit power control
TSD	transmit switched diversity
TU	time unit
TX	transmit or transmitter
UEP	unequal error protection
WAN	wide area network
WLAN	wireless local area network
WPAN	wireless personal area network

4. General description

4.1 What is a piconet?

A piconet is a wireless ad hoc data communications system that allows a number of independent data devices (DEVs) to communicate with each other. A piconet is distinguished from other types of data networks in that communications are normally confined to a small area around person or object that typically covers 10 m in all directions and envelops the person or a thing whether stationary or in motion.

This is in contrast to local area network (LAN), metropolitan area network (MAN), and wide area network (WAN), each of which covers a successively larger geographic area, such as a single building or a campus or that would interconnect facilities in different parts of a country or of the world.

4.2 Components of an IEEE 802.15.3 piconet

An IEEE 802.15.3 piconet consists of several components, as shown in Figure 4-1. The basic component is the DEV. One DEV is required to assume the role of the piconet coordinator (PNC) of the piconet. The PNC provides the basic timing for the piconet with the beacon. Additionally, the PNC manages the quality of service (QoS) requirements, power save modes, and access control to the piconet.




Figure 4-1—IEEE 802.15.3 piconet elements

Because IEEE 802.15.3 piconets form without pre-planning and for only as long as the piconet is needed, this type of operation is referred to as an ad hoc network.

This standard allows a DEV to request the formation of a subordinate piconet. The original piconet is referred to as the parent piconet. The subordinate piconet is referred to as either a child or neighbor piconet, depending on the method the DEV used to associate with the parent PNC. Child and neighbor piconets are also referred to as dependent piconets since they rely on the parent PNC to allocate channel time for the operation of the dependent piconet. An independent piconet is a piconet that does not have any dependent piconets.

4.3 Overview of medium access control (MAC) functionality

IEEE 803.15.3 MAC is designed to support the following goals:

- Fast connection time
- Ad hoc networks
- Data transport with quality of service (QoS)
- Security
- Dynamic membership
- Efficient data transfer

4.3.1 Coordination

A piconet is formed when an IEEE 802.15.3 DEV that is capable of acting as the PNC begins transmitting beacons. Thus, even if there are no associated DEVs, the PNC sending the beacon is considered to be a piconet.

4.3.1.1 Starting a piconet

To start a piconet, a DEV that is capable of acting as the PNC scans the available channels to find one that is not being used, as described in 7.2.1. If it finds one that is clear, it starts the piconet by simply sending the beacon after making sure that the channel has remained empty for a specified period of time, as described in 7.2.2. If no channels are available, the DEV has the option of attempting to start a dependent piconet, as described in 7.2.7 and 7.2.8 and summarized in 4.3.1.3 and 4.3.1.4. While the process of starting a piconet does not ensure that the “most capable” PNC is initially selected based on the criteria in 7.2.3, the association and handover process does allow the “most capable” DEV to eventually become the PNC of the piconet.

While the PNC is allowed to handover to a dependent PNC, this does not imply that the dependent PNC will merge the two piconets. IEEE Std 802.15.3 does not provide a process for merging two piconets into a single piconet.

4.3.1.2 Handing over control of the piconet

When a DEV associates with an existing piconet, as described in 7.3.1 and 4.3.4, the PNC checks the capabilities of the new DEV to see if it is more capable to be the PNC of the piconet based on the criteria defined in 7.2.3. If the new DEV is more capable and the current security policies allow it, then the PNC has the option of handing over control of the piconet to the DEV that has just joined. This handover process, as described in 7.2.3, maintains all existing time allocations so that there is no interruption in the delivery of data in the piconet. If the PNC is shutting down or wants to leave the piconet, it also uses the handover process to give control to another DEV in the piconet. The handover process also supports the handing over of a dependent PNC, as described in 7.2.6.

4.3.1.3 Creating a child piconet

A child piconet is one that is formed under an established piconet. The established piconet then becomes the parent piconet. The child piconet functionality is useful for either extending the area of coverage of the piconet or shifting some computational or memory requirements to another PNC-capable DEV. It is possible for the parent piconet to have more than one child piconet. In addition, it is also possible for a child PNC to allow a child piconet as a part of its own piconet.

The child piconet uses a distinct piconet identifier (PNID) and acts as an autonomous piconet except that it is dependent on a private channel time allocation (CTA) from the parent piconet. Association and security membership for the child piconet are handled within the child piconet and do not involve the parent PNC.

The child PNC is a member of the parent piconet and thus is able to exchange data with any DEV in the parent piconet. The child PNC is also a member of the child piconet and thus is able to exchange data with any DEV in the child piconet. Only DEVs in a child piconet that are also members of the parent piconet are

able to communicate directly with DEVs in the parent piconet who are not also members of that child piconet. The use of the child piconet is described in 7.2.7.

4.3.1.4 Creating a neighbor piconet

A neighbor piconet is formed under an established piconet. The established piconet then becomes the parent piconet. The neighbor piconet functionality is a mechanism for sharing the frequency spectrum between different piconets when there are no vacant PHY channels. It is possible for a single piconet to have more than one neighbor piconet or to have both child and neighbor piconets as a part of the parent piconet. In addition, it is possible for the neighbor PNC to allocate a child or neighbor piconets within its own piconet.

The neighbor piconet uses a distinct PNID and is an autonomous piconet except that it is dependent on a private CTA from the parent piconet. Association and security membership for the neighbor piconet are handled within the neighbor piconet and do not involve the parent PNC.

The neighbor PNC is not a member of the parent piconet and thus does not exchange information with any DEV in the parent piconet. The neighbor piconet mechanism is available to other wireless DEVs as a means of sharing the frequency spectrum. Any entity capable of initiating (i.e., requesting status as) an IEEE 802.15.3 neighbor piconet would also be capable of using this as a coexistence method. The use of the neighbor piconet is described in 7.2.8.

4.3.2 Ending a piconet

If the PNC is going to stop operation and there are no other PNC-capable DEVs in the piconet, the PNC places the PNC Shutdown information element (IE), as described in 6.4.5, into the beacon, as described in 7.2.9.1, to notify the members of the piconet.

In the case that the PNC abruptly leaves the piconet without handing over control to another PNC-capable DEV in the piconet, the piconet stops operation. After the association timeout period (ATP) expires, a PNC-capable DEV from the old piconet will be able to start a new piconet using the normal process, as described in 7.2.2.

In the case of dependent piconets, the parent PNC is able to end the dependent piconet via the Disassociation Request command, described in 6.5.1.3, for neighbor piconets, or by using the stream termination procedure, described in 7.5.1.3, for child piconets, as described in 7.2.9.2.

4.3.3 Ending a piconet with a dependent piconet involved

If the parent piconet ends operation, the parent PNC will indicate in the PNC Shutdown IE, as described in 6.4.5, a dependent piconet that will be able to continue to operate. All other dependent piconets will cease operation when the parent piconet ends operation. The dependent PNC that was selected to remain will remove the Parent Piconet IE, as described in 6.4.3, from its Beacon frame, signifying that it is no longer a dependent piconet. In the case where the parent piconet was temporarily disrupted, the parent PNC is able to attempt to join the dependent piconet and potentially receive a transfer of control via PNC handover.

A child piconet ends its piconet with the shutdown procedure and then uses the Channel Time Request command, as described in 6.5.6.1, to terminate the stream and release the resources in the parent piconet. When the child PNC shuts down its piconet, it is not required to leave the parent piconet.

The neighbor piconet uses the Disassociation Request command, as described in 6.5.1.3, to end its relationship with the parent PNC.

When a dependent piconet ends operation it has no effect on the parent piconet except to release resources.

4.3.4 Association and disassociation

In order to participate in a piconet, a DEV needs to join the piconet using the association process, as described in 7.3.1. Associating with the piconet provides the DEV with an identifier that is unique within that piconet, the DEVID, as described in 6.2.3. The DEVID is used instead of the DEV address to save overhead in the system. The association process optionally provides information about the services available in the piconet as well as the services provided by the DEV, as described in 7.3.2. The association process also provides the PNC with the capabilities of the new DEV to enable the PNC to decide if it wants to hand over control of the piconet to the new DEV, as described in 7.2.3 and 4.3.1.2.

When a new DEV joins the piconet, the PNC broadcasts the information about all of the DEVs in the piconet, as described in 7.3.3, and places information in the beacon about the new DEV. This allows other DEVs in the piconet to become aware of the new DEV as well as giving information to the new DEV about the members of the piconet.

When a DEV wants to leave the piconet or if the PNC wants to remove a DEV from the piconet, the disassociation process, as described in 7.3.4, is used. The DEVID of the disassociated DEV is no longer valid, until reissued by the PNC. However, the PNC is not allowed to reissue the DEVID until a waiting period has expired, as described in 7.3.1.

4.3.5 Security overview

Security is one of the following two modes, as described in 8.2:

- a) Mode 0—Open: Security membership is not required and payload protection (either data integrity or data encryption) is not used by the MAC. The PNC is allowed to use a list of DEV addresses to admit or deny entry to the piconet.
- b) Mode 1—Secure membership and payload protection: DEVs establish secure membership with the PNC before they have access to the piconet's resources. Data sent in Mode 1 is allowed to use payload protection (data integrity or data encryption with data integrity). Data integrity is required for most of the commands that are sent when in Mode 1.

When security is enabled, i.e., when Mode 1 is being used, DEVs that wish to join the piconet are required to establish secure membership with the PNC. The DEVs are also allowed to establish a secure relationship with other DEVs for secure communications. A DEV has established a secure membership or a secure relationship when it gets a management key for the security relationship. The process of establishing secure membership or a secure relationship is outside of the scope of this standard. The PNC or DEV that generates and distributes the key is called the key originator.

The payload protection protocol, as described in 9.2.2, uses a symmetric key that is generated by the key originator and is securely distributed to DEVs that have established secure membership or a secure relationship with the key originator, as described in 8.4.2. The symmetric key encryption algorithm used is the advanced encryption standard (AES) 128 in counter mode encryption and cipher block chaining-message authentication code (CCM).

4.3.6 The IEEE 802.15.3 superframe

Timing in the IEEE 802.15.3 piconet is based on the superframe, which is illustrated in Figure 4-2. The superframe is composed of the following three parts:

- The beacon, as described in 6.3.1, which is used to set the timing allocations and to communicate management information for the piconet. The beacon consists of the Beacon frame, as described in 6.3.1, as well as any Announce commands sent by the PNC as a beacon extension, as described in 7.6.3.

- The contention access period (CAP), as described in 7.4.2, which is used to communicate commands and/or asynchronous data if it is present in the superframe.
- The channel time allocation period (CTAP), 7.4.3, which is composed of channel time allocations (CTAs), including management CTAs (MCTAs). CTAs are used for commands, isochronous streams, and asynchronous data connections.

In Figure 4-2 the MCTAs are shown first, but the PNC is allowed to place any number of them at any position in the superframe.




Figure 4-2—IEEE 802.15.3 piconet superframe

The length of the CAP is determined by the PNC and communicated to the DEVs in the piconet via the beacon. However, the PNC is able to replace the functionality provided in the CAP with management CTAs (MCTAs), except in the case of the 2.4 GHz PHY where the PNC is required to allow DEVs to use the CAP, as described in 10.2.10. MCTAs are a type of CTA that is used for communications between the DEVs and the PNC.

The CAP uses carrier sense multiple access/collision avoidance (CSMA/CA) for the medium access, as described in 7.4.2. The CTAP, on the other hand, uses a standard TDMA protocol where the DEVs have specified time windows, as described in 7.4.3.1. MCTAs, as described in 7.4.3.3, are either assigned to a specific source/destination pair and use TDMA for access or they are shared CTAs that are accessed using the slotted aloha protocol, as described in 7.4.3.4.

4.3.7 Channel time management

All data in the IEEE 802.15.3 piconet is exchanged in a peer-to-peer manner. There are three methods for communicating data between DEVs in the piconet, as follows:

- Sending asynchronous data in the CAP, if present, as described in 7.4.2.
- Allocating channel time for isochronous streams in the CTAP, as described in 7.5.1.
- Allocating asynchronous channel time in the CTAP, as described in 7.5.2.

If the CAP is present in the superframe and the PNC allows data in the CAP, DEVs in the piconet are able to use the CAP to send small amounts of data without having to allocate channel time.

If the DEV needs channel time on a regular basis, it makes a request from the PNC for isochronous channel time, as described in 7.5.1.1. If the resources are available, the PNC allocates time in a CTA for the DEV. If the requirements for the data change, then the DEV is able to request a change to the allocation, as described in 7.5.1.2. The source DEV, destination DEV, or the PNC are all allowed to terminate the stream, as described in 7.5.1.3.

For regular CTAs, the PNC is able to change their position within the superframe every superframe. If a DEV misses a beacon, it is unable to use the allocation for a regular CTA. To avoid lost throughput due to

missed beacons, DEVs are allowed to request a special type of CTA called a pseudo-static CTA, as described in 7.4.3.1. If the DEV is allocated a pseudo-static CTA, it is allowed to use the CTA for up to *mMaxLostBeacons* missed beacons. The PNC is allowed to move the locations of these CTAs, but needs to maintain the time for the old allocation for *mMaxLostBeacons* superframes to avoid collisions, as described in 7.4.3.1.

Asynchronous allocation is slightly different. Rather than requesting recurring channel time, an asynchronous channel time request, as described in 7.5.2.1, is a request for a total amount of time to be used to transfer its data. The PNC is then able to schedule time for this request when available based on the channel time requirements. Unlike an isochronous allocation, only the source DEV or PNC are allowed to terminate an asynchronous allocation, as described in 7.5.2.2.

4.3.8 Data communications between DEVs

In order to handle large data frames from layers above the MAC sublayer, this standard supports the fragmentation and defragmentation of these data frames, as described in 7.7. The ability to fragment data frames is also useful to reduce the frame error rate (FER) of a marginal link by decreasing the frame size. The fragments are numbered with a sequence number for the upper layer frame as well as a sequence number for the fragment itself. The total number of fragments of that data frame is also sent to enable the receiving DEV to allocate the correct amount of internal memory to hold the incoming frame.

If the source DEV wishes to verify the delivery of a frame, then one of the acknowledgment (ACK) policies is used, as described in 7.9. This standard provides for three types of ACKs to enable different applications. The no-ACK policy, as described in 7.9.1, is appropriate for frames that do not require guaranteed delivery, where the retransmitted frame would arrive too late or where an upper layer protocol is handling the ACK and retransmission protocol. The immediate-ACK (Imm-ACK) policy, as described in 7.9.2, provides an ACK process in which each frame is individually ACKed following the reception of the frame. The delayed-ACK (Dly-ACK) policy, as described in 7.9.3, lets the source send multiple frames without the intervening ACKs. Instead, the ACKs of the individual frames are grouped into a single response frame that is sent when requested by the source DEV. The Dly-ACK process decreases the overhead in the Imm-ACK process while allowing the source DEV to verify the delivery of frames to the destination.

If the source DEV does not receive the requested ACK, then it has the option of retransmitting the frame, as described in 7.9.6, or dropping the frame. The decision to retransmit or drop the frame depends on the type of data or command that is being sent, the number of times that the source DEV has attempted sending the frame, the length of time it has spent attempting to send the frame, or other implementation-dependent factors.

4.3.9 Information discovery in the piconet

Since IEEE 802.15.3 piconets are ad hoc in nature, it is important for the DEVs in the piconet to be able to find out information about the services and capabilities of the other DEVs in the piconet at any instant in time. This standard supports four methods for discovering information about other DEVs in the piconet: the PNC Information Request command, as described in 6.5.4.1; the Probe Request command, as described in 6.5.4.5; the Announce command, as described in 6.5.5.2; and the Piconet Services command, as described in 6.5.5.1. In addition, the PNC is able to ask a DEV in the piconet to evaluate the channel conditions in either the current channel or in an alternate channel with the Remote Scan Request command, as described in 6.5.7.3. Any DEV in the piconet is able to ask another DEV about the status of the current channel with the Channel Status Request command, as described in 6.5.7.1.

The PNC Information Request command, as described in 7.10.1, is used to obtain information from the PNC about either a specific DEV in the piconet or all of the DEVs in the piconet. The PNC responds to the request with the PNC Information command, as described in 6.5.4.2, which contains information about the DEV (or DEVs) that was requested by the originator of the command.

A DEV uses the Probe Request command, 7.10.2, to find out more detailed information about other DEVs in the piconet. This command allows the originating DEV to retrieve some IEs, as defined in 7.10.2, from a target DEV in the piconet.

One of the goals of connecting DEVs in a piconet is to enable them to share services. However, to do this a DEV needs to be able to discover the services that are available in the piconet as well as to advertise its services to other DEVs in the piconet. This standard enables this with the Piconet Services IE, as described in 6.4.16, which optionally is exchanged in the association process, as described in 7.3.1.

The PNC is responsible for determining the channel in which the piconet will operate. Because the DEVs in the piconet are typically at different geographic locations, they provide additional information about interference or other piconets in the area. The Remote Scan Request command, as described in 7.10.5, is used by the PNC to request information about the channel from other DEVs. In response to the PNC's request, the DEV either scans the channels as requested or rejects the request from the PNC if it is not going to perform the scan.

Any DEV in the piconet is able to request information about the quality of the link between itself and another DEV with the Channel Status Request command, as described in 6.5.7.1. This command is used for two purposes. The first is to allow the DEVs to change transmit power, the data rate, or the requested channel time based on the quality of the data connection between the DEVs. The other use is for the PNC to determine if the DEVs in the piconet are having trouble with the channel. This information along with PNC's scanning of the channel and optional remote scan requests assists the PNC in determining if it needs to change the channel that the piconet is currently using.

4.3.10 Dynamic channel selection

The piconet operates in a dynamic environment and under unlicensed operation rules. Thus, it is subject to interference from licensed users, other IEEE 802.15.3 piconets as well as other unlicensed wireless entities in its channels. To enable the piconet to continue operation in this type of environment, the PNC has the capability to dynamically change the channel that the piconet is using without requiring either user intervention or the disruption of services in the piconet. To evaluate the status of the current channel as well as other channels, the PNC is able to use many methods including the following:

- Gathering information about the current channel from other DEVs in the piconet using the Channel Status Request command, as described in 7.10.4.
- Performing a passive scan of the channels, as described in 7.12.1.
- Requesting other DEVs to perform a channel scan using the Remote Scan Request command, as described in 7.10.5.

If the PNC determines that the current channel is unsuitable, it uses the dynamic channel selection procedure, as described in 7.12.1, to move the piconet to the new channel. The configuration of the piconet and the CTAs do not change with a channel change so that the services provided by the piconet are not interrupted by the change.

4.3.11 Power management

This standard provides techniques to allow DEVs to reduce power usage by enabling the DEVs to turn off for one or more superframes. These techniques are as follows:

- Device synchronized power save (DSPS) mode
- Piconet synchronized power save (PSPS) mode
- Asynchronous power save (APS) mode

DSPS and PSPS are collectively known as synchronized power save (SPS) modes. In the piconet, DEVs operate in one of four power management (PM) modes; ACTIVE mode, DSPS mode, PSPS mode, or APS mode.

PSPS mode, as described in 7.14.1, allows DEVs to sleep at intervals defined by the PNC. The DEV sends a request to the PNC when it wants to enter the PSPS mode. The PNC informs the piconet by setting the DEV's bit in its power save (PS) Status IE in the beacon. The PNC then selects beacons that will be the system wake beacons and indicates the next one in the PS Status IE for the PSPS set. All DEVs in PSPS mode are required to listen to the system wake beacons.

DSPS mode, as described in 7.14.2, is designed to enable groups of DEVs to sleep for multiple superframes but still be able to wake up during the same superframe. DEVs synchronize their sleep patterns by joining a DSPS set that specifies the interval between wake periods for the DEVs and the next time the DEVs will be awake. Besides allowing the DEVs to wake up and exchange traffic at the same time, the use of DSPS sets makes it easy for other DEVs in the piconet to determine exactly when a DSPS DEV will be available to receive traffic.

APS mode, as described in 7.14.3, allows a DEV to conserve power for extended periods until the DEV chooses to listen for a beacon. The only responsibility of a DEV in APS mode is to communicate with the PNC before the end of its ATP in order to preserve its membership in the piconet.

The PNC allocates asynchronous CTAs to a destination DEV that is in either PSPS mode or DSPS mode in the wake superframes for that DEV.

Regardless of the DEV's power management mode, every DEV in the piconet is allowed to power down during parts of the superframe when the DEV is not scheduled to transmit or receive data.

4.3.12 Controlling transmit power in the piconet

The ability to control transmit power in the piconet enables DEVs to minimize interference with other wireless networks that share the same channel as well as to decrease the power usage in some PHY implementations. Two methods are provided by this standard for controlling transmitter power. The first method allows the PNC to set a maximum transmit power for the CAP, beacon, and MCTAs, excluding association MCTAs, as described in 7.12.2.1. Since the link between the PNC and the DEV defines the size of the piconet, controlling the power during these times allows the PNC to reduce transmit power without adversely affecting operation of the piconet.

The second method allows DEVs using a CTA to request either an increase or a decrease in the transmit power of the remote DEV, as described in 7.12.2.2. Thus if two DEVs have a "good" link in a CTA, they are able to reduce their transmitter power to decrease the power usage in some PHY implementations, and to reduce interference to other networks.

4.3.13 Superframe structure using quasi-omni mode

A mmWave piconet operates in either omni mode or quasi-omni mode. The superframe structure for omni mode is illustrated in Figure 4-2. The superframe structure for quasi-omni mode is illustrated in Figure 4-3. In quasi-omni mode, the same Beacon frame is transmitted in different quasi-omni directions in a round-robin way in order to allow mmWave DEVs located in different directional coverage to join the piconet, as described in 7.6.5.




Figure 4-3—Piconet superframe structure in quasi-omni mode

4.3.14 Frame aggregation

Frame aggregation, as described in 7.8, is supported for the purpose of high throughput. Two aggregation methods are provided for mmWave DEVs, one is suitable for normal data and A/V streaming and the other is optimized for low-latency communications.

4.3.15 Beam forming

The beam forming procedure allows a mmWave DEV that is capable of beam forming to increase the antenna gain for supporting high data rate transmission. The beam forming procedures are described in Clause 12.

4.3.16 Channel probing

Channel probing, as described in 7.10.7, is an optional capability that is used by a DEV to determine the best modulation and coding scheme (MCS) to use with the target DEV under current channel condition.

4.3.17 Unequal error protection (UEP)

For flexible error protection that accounts for data that has more emphasis on the MSBs than the LSBs, e.g., video, UEP is provided as an option only for mmWave DEVs. Three types of UEP are defined, as described in 7.17, for use with the three mmWave PHYs.

4.4 Characteristics of the 2.4 GHz PHY

4.4.1 General characteristics

The 2.4 GHz PHY, Clause 10, uses the 2.4 GHz to 2.4835 GHz band that is available for unlicensed use in much of the world, as described in 10.1. Two channel plans are defined, one with 4 channels for high-density applications and one with 3 channels to enable better coexistence with IEEE Std 802.11b™-1999 [B3]⁷

⁷The numbers in brackets correspond to those of the bibliography in Annex A.

networks, as described in 10.2.3. The PHY also supports 5 data rates, ranging from 11 Mb/s to 55 Mb/s, utilizing, quadrature phase-shift keying (QPSK), differential quadrature phase shift keying (DQPSK), 16-quadrature amplitude modulation (16-QAM), 32-QAM and 64-QAM. The base rate of 22 Mb/s is uncoded, while the 11 Mb/s, 33 Mb/s, 44Mb/s, and 55 Mb/s use trellis coded modulation, 10.3.

For efficiency, the PHY calculates the header check sequence over both the MAC and PHY headers and appends this to the MAC header. The header for all frames is sent at the base rate, 22 Mb/s, to allow all DEVs in the piconet to detect traffic. In the 11 Mb/s mode, the entire MAC and PHY header is repeated at the lower modulation rate to increase the probability of receiving the entire header correctly.

The PHY uses a constant-amplitude, zero-autocorrelation (CAZAC) sequence for the preamble, as described in 10.4.2. This sequence has good properties for obtaining synchronization, timing information, and frequency offset.

The on-air bandwidth is limited to 15 MHz in order to allow more channels as well as to decrease the interference to other systems and to decrease the susceptibility to interference from other systems. The transmitter power is constrained by the limitations of the appropriate regulatory bodies.

The receiver of a compliant system reports both the signal level and, if the higher order modulations are used, an indication of the signal quality. This allows a DEV to determine if errors in the channel are due to poor signal quality or due to interference from other systems.

4.4.2 Coexistence and interoperability

While this standard does not require interoperability with other standards, there were choices made with the 2.4 GHz PHY specification that make it easier to design dual-mode radios. The commonalities that allow interoperable radios are discussed in D.1.

Because the 2.4 GHz PHY operates as an unlicensed system, it needs to share the medium with both licensed and unlicensed users in the band. The IEEE 802.15.3 MAC and the 2.4 GHz PHY offer a variety of techniques to enhance the coexistence with other users in the band. The methods provided by this standard include the following:

- Passive scanning
- Dynamic channel selection
- The ability to request channel quality information
- Link quality and RSSI
- A channel plan that minimizes channel overlap
- Lower transmit power
- Transmit power control
- Neighbor piconet capability

The use of these methods to improve coexistence is described in D.2.

4.5 Characteristics of the mmWave PHY

4.5.1 mmWave PHY characteristics

The mmWave PHY, as described in Clause 11, is defined for the frequency band of 57.0–66.0 GHz, as allocated by the regulatory agencies in Europe, Japan, Canada, and the United States as well as any other areas where the regulatory bodies have allocated this band. While four channels are defined for the PHY, regulatory requirements allow fewer in some regions.

A total of three PHYs are defined for the mmWave PHY. They are as follows:

- a) Single Carrier mode in mmWave PHY (SC PHY), as described in 11.2.
- b) High Speed Interface mode in mmWave PHY (HSI PHY), as described in 11.3.
- c) Audio/Visual mode in mmWave PHY (AV PHY), as described in 11.4.

For DEVs that implement the mmWave PHY, at least one of the three PHYs is required. In addition, to promote coexistence and interoperability, a common mode signaling (CMS) is defined based on a low data rate SC PHY mode, as described in 11.1.12.

The SC PHY supports data rates up to 5 Gb/s. The SC PHY supports a wide range of modulations, $\pi/2$ BPSK, $\pi/2$ QPSK, $\pi/2$ 8-phase shift keying (8-PSK), $\pi/2$ 16-QAM, on-off keying (OOK), and dual alternate mark inversion (DAMI). The coding schemes included are Reed-Solomon (RS) coding with low-complexity implementation and low-density parity check (LDPC) coding with high error-correcting capability. Code spreading using either linear feedback shift register (LFSR) code or Golay sequence is also applied to increase robustness of the system.

The HSI PHY, as described in 11.3, is designed for non-line-of-sight (NLOS) operation and uses OFDM with an FEC based on LDPC.

The AV PHY, as described in 11.4, is designed for NLOS operation and the transport of uncompressed, high definition video and audio. It uses OFDM modulation with convolutional inner code and a Reed Solomon(RS) outer code. The AV mode supports omni-directional coverage via the low-rate PHY (LRP) for the purpose of setting up high-throughput connections using the high-rate PHY (HRP).

Different PHYs are a result of demands of different market segments, which were based on the development of the usage models for this standard. For example one usage model is for kiosk applications. This usage model requires 1.5 Gb/s at a 1 m range. The SC-PHY can provide such a data rate at that short range with less complexity thus lower cost than an OFDM PHY. Another usage model required the streaming of uncompressed video. Due to the nature of uncompressed video signals a special PHY, the AV PHY, was selected to provide high throughput. A third usage model involves an ad hoc system to connect computers and devices around a conference table. In this usage model, all of the devices in the wireless personal area network (WPAN) will have bidirectional, NLOS high-speed, low-latency communication, which is provided for by the HSI PHY. Mandatory data rates of all those PHYs are selected according to specific usage models. In addition, higher data rates are provided to give options to the implementers so that they can best address the different market segments.

4.5.2 Piconets using mmWave PHY modes

When a mmWave PHY PNC-capable DEV starts a piconet, the type of piconet it starts depends on the PHY modes that are supported. For example, if the PNC-capable DEV supports only the AV PHY mode, it would start an AV piconet in which the beacon is sent with the AV PHY mode and the contention periods (CPs) would use the AV PHY mode. The AV PNC would also send the Sync frame to improve coexistence with other piconets. DEVs that support only the AV PHY mode are able to find and join the piconet using the AV PHY mode.

The same process is used for a mmWave PHY PNC-capable DEV that supports only HSI PHY mode or SC PHY mode, with the exception that an SC PHY mode PNC does not need to send a Sync frame since its beacon is sent in CMS mode already.

If a PNC-capable DEV supports more than one mmWave PHY mode, then it is able to select the type of piconet it starts, potentially starting more than one piconet, each with a different PHY mode, with the additional piconets as dependent piconets. Alternatively, the multi-mode PNC is able to have a single piconet in which DEVs use any MCS in a CTA that is supported by both of the DEVs.

The complete rules are described in 11.1.

4.6 Conventions

Unless otherwise specified, the term increment or incremented means to increase a value by one.

The editorial conventions in this standard are as follows:

- Frame names (e.g., Beacon frame, MAC Command frame) are capitalized and include the word *frame*.
- Command names (e.g., Association Request command) are capitalized and include the word *command*.
- Field names are capitalized and include the word *field*, which is not capitalized, e.g., Destination PAN Identifier field. The only exception is in a figure that shows the field, in which the word “field” is not included.
- IE names are capitalized and include the acronym IE (e.g., Capability IE).
- MAC and PHY PIB entries are preceded by a lower case prefix, *mac* or *phy*, respectively, and are capitalized at word boundaries, e.g., “*macDevServicesBroadcast*” and “*phyCurrentChannel*”.
- Constants begin with a lower case, italicized “*m*” for MAC related constants and a lower case italicized “*p*” for PHY related constants. Both use the same capitalization and use rules as PIB attributes.
- Service primitives names (e.g., MLME or MCPS) are ALL CAPS, but the .request, .indication, .response, and .confirm in the primitive names are always lower case.
- Primitive parameters are capitalized at word boundaries, e.g., NumberOfPiconets.
- In a PIB attribute, constant or service primitive, acronyms are not capitalized after the first letter, e.g., “*pMinTpcLevel*”, “*macDevServicesBroadcast*”, or “ScanForBsid”.

4.7 Concept of primitives

The services of a layer are the capabilities it offers to the user in the next higher layer or sublayer by building its functions on the services of the next lower layer. This concept is illustrated in Figure 4-4, showing the service hierarchy and the relationship of the two correspondent users and their associated layer (or sublayer) peer protocol entities.




Figure 4-4—Service primitives

The services are specified by describing the information flow between the user and the layer. This information flow is modeled by discrete, instantaneous events, which characterize the provision of a service. Each event consists of passing a service primitive from one layer to the other through a layer SAP associated with an user. Service primitives convey the required information by providing a particular service. These service primitives are an abstraction because they specify only the provided service rather than the means by which it is provided. This definition is independent of any other interface implementation.

A service is specified by describing the service primitives and parameters that characterize it.

A primitive can be one of the following four generic types:

- *Request*: The request primitive is used to request that a service is initiated.
- *Indication*: The indication primitive is used to indicate to the user an internal event.
- *Response*: The response primitive is used to complete a procedure previously invoked by an indication primitive.
- *Confirm*: The confirm primitive is used to convey the results of one or more associated previous service requests.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017

5. Layer management

5.1 Overview of management model

Both MAC and PHY layers conceptually include management entities, called the MAC sublayer management entity and PHY layer management entity (MLME and PLME, respectively). These entities provide the layer management service interfaces for the layer management functions. Figure 5-1 depicts the relationship among the management entities. The protocol adaptation layer (PAL) includes both the FCSL and DME. The reference model depicts the interactions of a single PAL in a single piconet; support for multiple piconets and/or PALs is implementation dependent.




Figure 5-1—The reference model used in this standard

The various entities within this model interact in various ways. Certain of these interactions are defined explicitly within this standard, via a service access point (SAP) across which defined primitives are exchanged. Other interactions are not defined explicitly within this standard, such as the interface between the MAC and the MLME or the interface between the PHY and the PLME. The specific manner in which these MAC and PHY interfaces are integrated into the overall MAC and PHY layers are not specified within this standard.

The various SAPs within this model are the following:

- a) PAL SAP (not shown)
- b) MAC SAP
- c) MLME SAP

The PHY SAP and PLME SAP are not defined in this standard as they are rarely, if ever, exposed in a typical implementation. The PHY management objects and attributes are accessed through the MLME SAP with the generic management primitives defined in 5.2.

The MAC SAP is described further in Annex B.

If the SAP interfaces are not exposed in an IEEE 802.15.3 DEV, then these interfaces do not have to be implemented as described here. If the interfaces are exposed, then they should support the primitives described in this clause.

5.2 Generic management primitives

The management information specific to each layer is represented as a personal area network (PAN) information base (PIB) for that layer. In a LAN/MAN the corresponding information is in the management information base (MIB) and is often associated with a management protocol such as simple network management protocol (SNMP) (see IETF RFC 1157:1990 [B10]). However, piconets are not intended to be managed across a network but rather use the management information to ascertain the characteristics of the layer or sublayer.

The MLME and PLME are viewed as “containing” the PIB for that layer or sublayer. The generic model of PIB-related management primitives exchanged across the management SAPs is to allow the SAP user entity to either “GET” the value of a PIB attribute, or to “SET” the value of a PIB attribute. The invocation of a SET.request primitive may require the layer entity to perform certain defined actions. An attempt to “GET” a PIB attribute identified as a “write only” attribute is not a valid operation. An attempt to “SET” a PIB attribute identified as a “read only” attribute is not a valid operation.

The GET and SET primitives are represented as requests with associated confirm primitives. The DME uses the services provided by the MLME through the MLME SAP to access the management objects in either the MAC PIB or PHY PIB. The primitives are summarized in Table 5-1.

Table 5-1—Summary of generic management primitives

Name	Request	Confirm
MLME-GET	5.2.1	5.2.2
MLME-SET	5.2.3	5.2.4

The parameters used for these primitives are defined in Table 5-2.

Table 5-2—MLME generic management primitive parameters

Name	Type	Valid range	Description
PIBAttribute	Octet string	Any PIB attribute, as defined in 5.4 or 10.7	The name of the PIB attribute
PIBvalue	Variable	As defined in 5.4 or 10.7	The PIB value
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request
ReasonCode	Enumeration	INVALID_PIB_ATTRIBUTE_NAME, INVALID_PIB_ATTRIBUTE_VALUE, READ_ONLY_PIB_ATTRIBUTE, WRITE_ONLY_PIB_ATTRIBUTE	Indicates the reason for a ResultCode of FAILURE

Other SAP-specific primitives are identified in 5.3.

5.2.1 MLME-GET.request

This primitive requests information about a given MAC or PHY PIB attribute. The semantics of this primitive are as follows:

```
MLME-GET.request      (
                        PIBAttribute
                        )
```

The primitive parameter is defined in Table 5-2.

5.2.2 MLME-GET.confirm

This primitive reports the results of an information request about either the MAC or PHY PIB. The semantics of this primitive are as follows:

```
MLME-GET.confirm      (
                        PIBAttribute,
                        PIBvalue,
                        ResultCode,
                        ReasonCode
                        )
```

The primitive parameters are defined in Table 5-2.

5.2.3 MLME-SET.request

This primitive attempts to set the indicated MAC or PHY PIB attribute to the given value. The semantics of this primitive are as follows:

```
MLME-SET.request      (
                        PIBAttribute,
                        PIBvalue
                        )
```

The primitive parameters are defined in Table 5-2.

5.2.4 MLME-SET.confirm

This primitive reports the results of an attempt to set the value of an attribute in either the MAC or PHY PIB. The semantics of this primitive are as follows:

```
MLME-SET.confirm      (
                        PIBAttribute,
                        ResultCode,
                        ReasonCode
                        )
```

The primitive parameters are defined in Table 5-2.

5.3 MLME SAP interface

The services provided by the MLME to the DME are specified in this subclause. These services are described in an abstract way and do not imply any particular implementation. MLME SAP primitives are of the general form ACTION.request followed by ACTION.confirm. An ACTION.indication provides information to the DME that originated either in the local MAC or from another DEV. The DME optionally responds to the indication by issuing an ACTION.response. The DME uses the services provided by the MLME through the MLME SAP. The MLME interface models a single piconet environment; support for multiple piconets is implementation dependent. The primitives are summarized in Table 5-3.

Table 5-3—Summary of MLME primitives

Name	Request	Confirm	Indication	Response
MLME-RESET	5.3.1.1	5.3.1.2	—	—
MLME-SCAN	5.3.2.1	5.3.2.2	5.3.2.3	—
MLME-START	5.3.3.1	5.3.3.2	—	—
MLME-STOP	5.3.4.1	5.3.4.2	—	—
MLME-ASSOCIATE	5.3.5.1	5.3.5.2	5.3.5.3	—
MLME-DISASSOCIATE	5.3.6.1	5.3.6.2	5.3.6.3	—
MLME-MEMBERSHIP-UPDATE	5.3.7.1	5.3.7.2	—	—
MLME-SECURITY-ERROR	—	—	5.3.7.3	—
MLME-SECURITY-MESSAGE	5.3.7.4	5.3.7.5	5.3.7.6	—
MLME-PNC-HANDOVER	—	—	5.3.8.1	5.3.8.2
MLME-NEW-PNC	—	—	5.3.8.3	—
MLME-DEV-INFO	5.3.9.1	5.3.9.2	5.3.9.3	—
MLME-SECURITY-INFO	5.3.10.1	5.3.10.2	5.3.10.3	5.3.10.4
MLME-APPLICATION-SPECIFIC	5.3.11.1	5.3.11.2	5.3.11.3	—
MLME-ANNOUNCE-SERVICE	5.3.12.1	5.3.12.2	—	—
MLME-PICONET-SERVICES	5.3.12.3	5.3.12.4	5.3.12.5	—
MLME-CREATE-STREAM	5.3.13.1	5.3.13.2	5.3.13.3	—
MLME-MODIFY-STREAM	5.3.13.4	5.3.13.5	—	—
MLME-TERMINATE-STREAM	5.3.13.6	5.3.13.7	5.3.13.8	—
MLME-BSID-CHANGE	5.3.14.1	5.3.14.2	—	—
MLME-PICONET-PARM-CHANGE	—	—	5.3.14.3	—
MLME-PS-SET-INFORMATION	5.3.15.1	5.3.15.2	—	—
MLME-SPS-CONFIGURE	5.3.15.3	5.3.15.4	—	—
MLME-PM-MODE-CHANGE	5.3.15.5	5.3.15.6	5.3.15.7	—
MLME-MONITOR-PM-MODE	5.3.15.8	5.3.15.9	5.3.15.10	—
MLME-MULTICAST-CONFIGURATION	5.3.16.1	5.3.16.2	—	—

Table 5-3—Summary of MLME primitives (continued)

Name	Request	Confirm	Indication	Response
MLME-MULTICAST-RX-SETUP	5.3.16.3	5.3.16.4	—	—
MLME-BEACON-EVENT	5.3.17.1	5.3.17.2	5.3.17.3	—
MLME-TXDIV	5.3.18.1	5.3.18.2	—	—

5.3.1 Resetting the MAC/MLME

These primitives support the process of resetting the MAC/MLME, which also resets the PHY/PLME. The parameters used for these primitives are defined in Table 5-4.

Table 5-4—MLME-RESET primitive parameters

Name	Type	Valid range	Description
SetDefaultPIB	Boolean	TRUE, FALSE	If TRUE, all PIB attributes are set to their default values. The default values are implementation dependent. If FALSE, the MAC is reset, but all PIB attributes retain the values that were in place prior to the generation of the MLME-RESET.request primitive.
ResultCode	Enumeration	READY, ERROR	If READY, the MLME has successfully completed initialization and is ready to receive requests. If ERROR, the MLME was unable to successfully complete initialization and is unable to receive requests.

5.3.1.1 MLME-RESET.request

This primitive requests that the MAC entity be reset to its initial conditions. The semantics of this primitive are as follows:

```
MLME-RESET.request          (
                               SetDefaultPIB
                              )
```

The primitive parameter is defined in Table 5-4.

The MLME restores the MAC and its internal variables (other than PIB attributes) to their initial state and values. The SetDefaultPIB parameter governs whether PIB attributes are reinitialized or left unchanged.

5.3.1.2 MLME-RESET.confirm

This primitive is generated by the MLME when it completes its initialization process. The semantics of this primitive are as follows:

```
MLME-RESET.confirm          (
                               ResultCode
                              )
```

The primitive parameter is defined in Table 5-4.

5.3.2 Scanning for piconets

These primitives support the process of determining the presence or absence of piconets, as described in 7.2.1. The parameters used for these primitives are defined in Table 5-5.

Table 5-5—MLME-SCAN primitive parameters

Name	Type	Valid range	Description
ScanForBsid	Boolean	TRUE, FALSE	Indicates if the scan process should search for a specific BSID, as described in 7.2.1.
Bsid	Octet string	As defined in 6.4.2	The text string of a specific piconet for which to scan. This parameter is not used if ScanForBsid is FALSE.
ScanForPnid	Boolean	TRUE, FALSE	Indicates if the scan process should search for a specific PNID, as described in 7.2.1.
Pnid	Integer	0–65535	The identifier (ID) of a specific piconet for which to scan. This parameter is not used if ScanForPnid is FALSE.
ScanForPncAddress	Boolean	TRUE, FALSE	Indicates if the scan process should search for a PNC with a specific MAC address, as described in 7.2.1.
PncAddress	MAC address	Any valid individual MAC address, as described in 6.1	The MAC address of a specific PNC for which to scan. This parameter is not used if ScanForPncAddress is FALSE.
NumberOfPiconets	Integer	0–255	The number of piconets found during the scanning process.
PiconetDescriptionSet	Set of piconet descriptions, as defined in Table 5-6	A set containing zero or more instances of a PiconetDescription	The PiconetDescriptionSet is returned to indicate the results of the scan request.
NumberOfChannels	Integer	0 to the maximum number of PHY-dependent channels, as defined in 10.2.3	Indicates the number of channels scanned.
ChannelRatingList	Ordered list of integers.	0 to the maximum number of PHY-dependent channels, as defined in 10.2.3	Specifies a list of the channels scanned ordered from the best to the worst in terms of interference.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	OTHER	Indicates the reason for a ResultCode of FAILURE.

In Table 5-5, a PiconetDescriptionSet is a set of PiconetDescriptions. Each PiconetDescription consists of the elements shown in Table 5-6.

Table 5-6—Elements of PiconetDescription

Name	Type	Valid range	Description
Bsid	Octet string	As defined in 6.4.2	The text string identifier of a discovered piconet.
Pnid	Integer	0–65535	The PNID of a discovered piconet.
PncAddress	MAC address	Any valid individual MAC address	The MAC address of the PNC of the piconet that was found.
ChannelIndex	Integer	0 to the maximum number of PHY-dependent channels, as defined in 10.2.3	A PHY-dependent channel number on which the piconet was found.
SECmode	Enumeration	MODE_0, MODE_1	The security mode of the piconet that was found, as described in 6.3.1.
SignalQuality	Integer	0–15	Indicates the quality of the received frame or beacon for this piconet. The value is implementation dependent with 0 indicating the lowest quality and 15 the highest quality.
NumApplicationSpecificData	Integer	0–255	Indicates the number of ApplicationSpecificDataSets advertised in the beacon of a discovered piconet.
ApplicationSpecificDataSet	Set of application specific data, as defined in Table 5-7	A set containing zero or more instances of ApplicationSpecificData	The ApplicationSpecificDataSet is returned to indicate the ApplicationSpecificData offered by the discovered piconet.
PhyMode	Enumeration	2.4_GHZ, SC_MMWAVE, HSI_MMWAVE, AV_MMWAVE	The PHY mode that is being used in the piconet that was found.

In Table 5-5, the ChannelRatingList is a set of N integer values, where N equals the number of channel numbers provided by the PHY. The elements of the set are channel numbers and they are ordered from best (least interference) at the lowest set index to worst (most interference) at the highest set index.

In Table 5-6, an ApplicationSpecificDataSet is a set of ApplicationSpecificData. Each ApplicationSpecificData consists of the elements shown in Table 5-7. The ApplicationSpecificDataSets are only the current values advertised in the beacon and potentially are different in subsequent beacons.

Any security features of an existing piconet are ignored during the scan process.

It is not possible to obtain piconet services information, as described in 7.3.2, during the scan process.

Table 5-7—Elements of ApplicationSpecificData

Name	Type	Valid range	Description
UniqueId	Octet string	Any valid organizationally unique identifier (OUI) or Company ID (CID), as defined in 6.4.7	The unique identifier for the entity that defines the format of the data, as described in 6.4.7.
ApplicationData	Octet string	Any valid octet string of length up to ApplicationDataLength	Application-dependent data, as described in 6.4.7.

5.3.2.1 MLME-SCAN.request

This primitive is used to initiate the passive scan procedure to search for either a specific piconet or any piconet. The semantics of this primitive are as follows:

```
MLME-SCAN.request
(
    ScanForBsid,
    Bsid,
    ScanForPnid,
    Pnid,
    ScanForPncAddress,
    PncAddress,
    Timeout
)
```

The primitive parameters are defined in Table 5-5.

5.3.2.2 MLME-SCAN.confirm

This primitive is used to report the result of the request to initiate the passive scan procedure to search for either a specific piconet or any piconet. The semantics of this primitive are as follows:

```
MLME-SCAN.confirm
(
    NumberOfPiconets,
    PiconetDescriptionSet,
    NumberOfChannels,
    ChannelRatingList,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-5. All of the piconets found during the scan will be reported in separate elements of the PiconetDescriptionSet, even if more than one piconet is found on a given channel.

5.3.2.3 MLME-SCAN.indication

This primitive is used to report the result of a passive scan that was initiated by the MAC. The semantics of this primitive are as follows:

```

MLME-SCAN.indication
(
    NumberOfPiconets,
    PiconetDescriptionSet,
    NumberOfChannels,
    ChannelRatingList
)

```

The primitive parameters are defined in Table 5-5.

5.3.3 Starting a piconet

These primitives support the process of creating a new piconet with the DEV acting as PNC, as described in 7.2.2. The parameters used for these primitives are defined in Table 5-8.

Table 5-8—MLME-START primitive parameters

Name	Type	Valid range	Description
Bsid	Octet string	As defined in 6.4.2	The BSID of the new piconet.
SecMode	Enumeration	MODE_0, MODE_1	The security mode of the piconet, as described in 6.3.1.
DevId	Integer	Any valid DEVID, as defined in 6.2.3	The assigned DEVID for the DEV that is acting as the PNC, as described in 7.2.2.
MinDepSuperframePercent	Integer	1–100	The minimum percent of the superframe requested as a CTA for the dependent piconet, as described in 7.2.7 and 7.2.8.
DesiredDepSuperframePercent	Integer	1–100	The desired percent of the superframe requested as a CTA for the dependent piconet, as described in 7.2.7 and 7.2.8.
AllocatedSuperframePercent	Integer	0–100	The percent of the superframe allocated to the new dependent piconet. If the channel time request was rejected, the value shall be set to zero. This parameter is ignored if the DEV is starting an independent piconet.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.

Table 5-8—MLME-START primitive parameters (continued)

Name	Type	Valid range	Description
ReasonCode	Enumeration	NOT_PNC_CAPABLE, NO_CHANNELS_AVAILABLE, ALREADY_PNC, OTHER	Indicates the reason for a ResultCode of FAILURE.
PhyMode	Enumeration	2.4_GHZ, SC_MMWAVE, HSI_MMWAVE, AV_MMWAVE	The PHY that will be used for the beacons and the CP(s) in the piconet that will be started.

5.3.3.1 MLME-START.request

This primitive is used to start a piconet. If the DEV is not a member of the piconet, this primitive causes the DEV to start an independent piconet. If the DEV is a member of the piconet, this primitive causes the DEV to start a child piconet. If the DEV is associated as a neighbor member of a piconet, this primitive causes the DEV to start a neighbor piconet. The semantics of this primitive are as follows:

```
MLME-START.request
(
    Bsid,
    SecMode,
    MinDepSuperframePercent,
    DesiredDepSuperframePercent,
    PhyMode
)
```

The primitive parameters are defined in Table 5-8.

5.3.3.2 MLME-START.confirm

This primitive reports the results of a piconet creation procedure. The semantics of this primitive are as follows:

```
MLME-START.confirm
(
    DevId,
    AllocatedSuperframePercent,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-8.

5.3.4 Stopping a piconet

These primitives support the process of stopping operations as a PNC. The process may result in the shutdown of piconet operations, as described in 7.2.9, or the handover of PNC operations to another DEV in the piconet, as described in 7.2.3 and 7.2.6. The parameters used for these primitives are defined in Table 5-9.

In Table 5-9, the HandoverTargetList is a set of N DEVIDs, where N equals the number of target DEVIDs for a handover provided by the DME. The elements of the set are DEVIDs and they are ordered from preferred target at the lowest set index to the least preferred target at the highest set index. If the

HandoverTargetList contains only the BcstID, as defined in 6.2.3, the DME is not indicating any preference for the handover target.

Table 5-9—MLME-STOP primitive parameters

Name	Type	Valid range	Description
RequestType	Enumeration	SHUTDOWN, HANDOVER	If SHUTDOWN, the current piconet operations will be stopped. If HANDOVER, an attempt to handover PNC operations will be made.
AllowedHandoverTime	Duration	0–65535	If RequestType is HANDOVER, the time in milliseconds in which a handover attempt must be completed.
NumHandoverTargetDev	Integer	0– <i>mMaxNumValidDEVs</i>	The number of DEVs in the HandoverTargetList.
HandoverTargetList	List of DEVIDs	0 to maximum number of DEVIDs, as defined in 6.2.3	If RequestType is HANDOVER, specifies a list of Target DEVIDs for a handover attempt.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	NOT_A_PNC, HANDOVER_FAILED, OTHER	Indicates the reason for a ResultCode.

5.3.4.1 MLME-STOP.request

This primitive initiates the piconet shutdown procedure or the piconet handover procedure. The semantics of this primitive are as follows:

```
MLME-STOP.request
(
    RequestType,
    AllowedHandoverTime,
    NumHandoverTargetDev,
    HandoverTargetList
)
```

The primitive parameters are defined in Table 5-9.

5.3.4.2 MLME-STOP.confirm

This primitive reports the results of the request to stop operations as a PNC. The semantics of this primitive are as follows:

```
MLME-STOP.confirm
(
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-9.

5.3.5 Associating with a piconet

The following primitives support the process of a DEV associating with a PNC, as defined in 7.3.1. The parameters used for these primitives are defined in Table 5-10.

Table 5-10—MLME-ASSOCIATE primitive parameters

Name	Type	Valid range	Description
Bsid	Octet string	As defined in 6.4.2	The BSID of the target PNC for the association.
Pnid	Integer	0–65535	The PNID of the target PNC for the association, as defined in 6.2.2.
PncAddress	MAC address	Any valid individual MAC address	The MAC address of the target PNC for the association.
ChannelIndex	Integer	0–255	A PHY-dependent channel number to search for the target PNC for the association.
PiconetServicesInquiry	Boolean	TRUE, FALSE	Requests that the PNC send the services information about the piconet, as described in 7.3.2.
DevId	Integer	Any valid DEVID, as defined in 6.2.3	The DEVID assigned to a DEV as the result of an association or the DEVID of a DEV that has joined the piconet.
DevAddress	MAC address	Any valid individual MAC address	The MAC address of a DEV that has joined the piconet.
NeighborPiconetRequest	Boolean	TRUE, FALSE	Indicates that the DEV will join as a neighbor PNC rather than as a member of the piconet.
VendorSpecificIE	Octet string	Any valid Vendor Defined IE, as defined in 6.4.17	The Vendor Defined IE, if present, in the Association Response command, as described in 6.5.1.2.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, PNC_NOT_FOUND, PNC_DENIED, PNC_BUSY, ALREADY_ASSOCIATED, NEIGHBOR_REFUSED, OTHER	Indicates the reason for a ResultCode of FAILURE.

5.3.5.1 MLME-ASSOCIATE.request

This primitive initiates the association procedure. The semantics of this primitive are as follows:

```
MLME-ASSOCIATE.request      (
                               Bsid,
                               Pnid,
                               PncAddress,
                               ChannelIndex,
                               NeighborPiconetRequest,
                               PiconetServicesInquiry,
                               Timeout
                              )
```

The primitive parameters are defined in Table 5-10.

5.3.5.2 MLME-ASSOCIATE.confirm

This primitive reports the result of the association procedure. The semantics of this primitive are as follows:

```
MLME-ASSOCIATE.confirm      (
                               DevId,
                               VendorSpecificIE,
                               ResultCode,
                               ReasonCode
                              )
```

The primitive parameters are defined in Table 5-10.

5.3.5.3 MLME-ASSOCIATE.indication

This primitive is used to indicate that a new DEV has associated with the same piconet as this DEV. The semantics of this primitive are as follows:

```
MLME-ASSOCIATE.indication   (
                               DevId,
                               DevAddress
                              )
```

The primitive parameters are defined in Table 5-10.

5.3.6 Disassociation from a piconet

The following primitives are used when a DEV disassociates from a PNC and when the PNC disassociates a DEV from the piconet, as described in 7.3.4. The parameters used for these primitives are defined in Table 5-11.

Table 5-11—MLME-DISASSOCIATE primitive parameters

Name	Type	Valid range	Description
DevId	Integer	Any valid DEVID, as defined in 6.2.3	The DEVID of the disassociated DEV.
DevAddress	MAC address	Any valid individual MAC address	The MAC address of the disassociated DEV.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, CURRENTLY_PNC, DEV_ATP_EXPIRED, PNC_ATP_EXPIRED, DEV_DISASSOCIATED, OTHER_PNC_ACTION, UNKNOWN	Indicates the reason for the disassociation of a DEV from a piconet.

5.3.6.1 MLME-DISASSOCIATE.request

This primitive initiates the procedure for a DEV to disassociate from a piconet. The semantics of this primitive are as follows:

```
MLME-DISASSOCIATE.request
(
    Timeout
)
```

The primitive parameter is defined in Table 5-11.

5.3.6.2 MLME-DISASSOCIATE.confirm

This primitive is used to confirm the result of the disassociation procedure. The semantics of this primitive are as follows:

```
MLME-DISASSOCIATE.confirm
(
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-11.

5.3.6.3 MLME-DISASSOCIATE.indication

This primitive is used to indicate that either this DEV or another DEV has been disassociated from the piconet. The semantics of this primitive are as follows:

```
MLME-DISASSOCIATE.indication    (
                                DevId,
                                DevAddress,
                                ResultCode,
                                ReasonCode
                                )
```

The primitive parameters are defined in Table 5-11.

5.3.7 Security management

These primitives are used to initialize, update, or delete the security information as a result of a membership or key change process, as described in 8.3.4, or as the result of a security event, as described in 8.3.5 and 8.3.7. Primitives are also provided to transfer security messages. These primitives are suitable for use in an authentication process.

The parameters used for the MLME-MEMBERSHIP-UPDATE and MLME-SECURITY-ERROR primitives are defined in Table 5-12.

**Table 5-12—MLME-MEMBERSHIP-UPDATE and
MLME-SECURITY-ERROR primitive parameters**

Name	Type	Valid range	Description
SECID	2 octets	As defined in 6.2.7.2	The identifier for the key.
OrigId	Integer	Any valid DEVID, as defined in 6.2.3, except for the BcstID, the McstID or the UnassocID	Either the PNCID, if this key is for the DEV's PNC personality, or the DEV's DEVID.
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3, except for the BcstID, the McstID or the UnassocID	The DEVID of the target DEV for this relationship.
MembershipStatus	Enumeration	MEMBER, NON-MEMBER	Indicates the membership status for the provided SECID. If NON-MEMBER, KeyInfo is zero length.
KeyOriginator	Boolean	TRUE, FALSE	Indicates if the DEV is the key originator for this relationship. This is always true when the OrigId is the PNCID.
KeyInfo	Octet string	Any valid symmetric key for the symmetric key security operations, as defined in 9.3	The key used for protecting frames between this DEV and the TrgtId DEV.
SrcID	Integer	Any valid DEVID, as defined in 6.2.3, except for the BcstID, the McstID or the UnassocID	The DEVID of the DEV that is the source of a security error.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.

**Table 5-12—MLME-MEMBERSHIP-UPDATE and
MLME-SECURITY-ERROR primitive parameters (continued)**

Name	Type	Valid range	Description
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	NOT_ASSOCIATED, TARGET_UNAVAILABLE, UNAVAILABLE_KEY, FAILED_SECURITY_CHECK, BAD_TIME_TOKEN, INVALID_SEC_VALUE, OTHER	The reason for a security error.

The parameters used for the MLME-SECURITY-MESSAGE primitive are defined in Table 5-13.

Table 5-13—MLME-SECURITY-MESSAGE primitive parameters

Name	Type	Valid range	Description
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the target of the MLME request.
OrigId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the originator of the MLME request.
UniqueId	Octet string	Any valid OUI or CID, as defined in 6.4.7	A unique identifier for the entity that defines the format of the security information, as described in 6.4.7.
SecurityInformation	Octet string	Any valid octet string	Security information that will be passed from one DEV to another peer DEV in the piconet.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, TARGET_UNAVAILABLE, OTHER	The reason for a security error.

5.3.7.1 MLME-MEMBERSHIP-UPDATE.request

This primitive initiates the membership update procedure by either installing or removing a management key. Data keys are generated by the key originator of the relationship and are exchanged between peer MLMEs, as described in 8.3.4. The semantics of this primitive are as follows:

```
MLME-MEMBERSHIP-UPDATE.request    (
                                     OrigId,
                                     TrgtId,
                                     MembershipStatus,
                                     SECID,
                                     KeyOriginator,
                                     KeyInfo,
                                     Timeout
                                     )
```

The primitive parameters are defined in Table 5-12.

5.3.7.2 MLME-MEMBERSHIP-UPDATE.confirm

This primitive indicates the result of a membership update request. The semantics of this primitive are as follows:

```
MLME-MEMBERSHIP-UPDATE.confirm    (
                                     ResultCode,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-12.

5.3.7.3 MLME-SECURITY-ERROR.indication

This primitive allows the MLME of any DEV to indicate a failed security processing operation. The semantics of this primitive are as follows:

```
MLME-SECURITY-ERROR.indication    (
                                     SrcID,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-12.

5.3.7.4 MLME-SECURITY-MESSAGE.request

This primitive initiates the sending of a Security Message command, as described in 6.5.9.1, to the target DEV in the piconet. The semantics of this primitive are as follows:

```
MLME-SECURITY-MESSAGE.request    (
                                     TrgtId,
                                     UniqueId,
                                     Timeout
                                     )
```

The primitive parameters are defined in Table 5-13.

5.3.7.5 MLME-SECURITY-MESSAGE.confirm

This primitive indicates the result of a request to send a Security Message command. The semantics of this primitive are as follows:

```
MLME-SECURITY-MESSAGE.confirm      (
                                     ResultCode,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-13.

5.3.7.6 MLME-SECURITY-MESSAGE.indication

This primitive reports the reception of a Security Message command, as described in 6.5.9.1, from a DEV in the piconet. The semantics of this primitive are as follows:

```
MLME-SECURITY-MESSAGE.indication    (
                                     OrigId,
                                     UniqueId,
                                     SecurityInformation
                                     )
```

The primitive parameters are defined in Table 5-13.

5.3.8 PNC handover

These primitives are used as part of the handover process, as described in 7.2.3 and 7.2.6, where the current PNC's responsibilities are transferred to another DEV in the piconet. The parameters used for these primitives are defined in Table 5-14.

Table 5-14—MLME-PNC-HANDOVER and MLME-NEW-PNC primitive parameters

Name	Type	Valid range	Description
HandoverStatus	Enumeration	STARTED, IMPLICIT, PRELIMINARY, CANCELLED	Indicates if the PNC is beginning or canceling a handover to the DEV.
NewPncDevAddress	MAC address	Any valid individual MAC address	The DEV address of the DEV assuming responsibilities as PNC.
SecMode	Enumeration	MODE_0, MODE_1	The security mode of the piconet, as defined in 6.3.1.

5.3.8.1 MLME-PNC-HANDOVER.indication

This primitive indicates the reception of a directed PNC Handover Request command, 6.5.3.1. The semantics of this primitive are as follows:

```
MLME-PNC-HANDOVER.indication        (
                                     HandoverStatus
                                     )
```

5.3.8.2 MLME-PNC-HANDOVER.response

MLME-PNC-HANDOVER.response ()

5.3.8.3 MLME-NEW-PNC.indication

```
MLME-NEW-PNC.indication (
    NewPncDevAddress,
    SecMode
)
```

5.3.9 Requesting DEV information from the PNC

Table 5-15—MLME-DEV-INFO primitive parameters

Name	Type	Valid range	Description
QueriedDevId	Integer	Any valid DEVID, as defined in 6.2.3, except for the McstID or the UnassocID	The DEVID of the DEV for which information is being requested from the PNC. If it is the BcstID, the request is for information for all of the DEVs in the piconet.
NumDEVInfo	Integer	1– <i>mMaxNumValidDEVs</i>	Number of entries in the DEVInfoSet.
DEVInfoSet	A set of DEV Info fields, as defined in 6.5.4.2	A set containing 1 to <i>mMaxNumValidDEVs</i> instances of fixed-length DEV Info fields	The DEVInfoSet is returned to indicate the results of a PNC Information Request command.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, TARGET_NOT_ASSOCIATED, OTHER	Indicates the reason for a ResultCode of FAILURE.

5.3.9.1 MLME-DEV-INFO.request

This primitive initiates a request to the PNC for information regarding either a single DEV or all of the DEVs in the piconet. The semantics of this primitive are as follows:

```
MLME-DEV-INFO.request      (
                             QueriedDevId,
                             Timeout
                             )
```

The primitive parameters are defined in Table 5-15.

5.3.9.2 MLME-DEV-INFO.confirm

This primitive provides the result of the request to the PNC for information regarding either a single DEV or all of the DEVs in the piconet. The semantics of this primitive are as follows:

```
MLME-DEV-INFO.confirm      (
                             NumDEVInfo,
                             DEVInfoSet,
                             ResultCode,
                             ReasonCode
                             )
```

The primitive parameters are defined in Table 5-15.

5.3.9.3 MLME-DEV-INFO.indication

This primitive indicates the reception of DEV information that was not requested using the MLME-DEV-INFO.request primitive. The semantics of this primitive are as follows:

```
MLME-DEV-INFO.indication    (
                             NumDEVInfo,
                             DEVInfoSet
                             )
```

The primitive parameters are defined in Table 5-15.

5.3.10 Security information retrieval

These primitives are used to request security information about other DEVs in the piconet, as described in 8.4.1. The parameters used for the MLME-SECURITY-INFO primitives are defined in Table 5-16.

5.3.10.1 MLME-SECURITY-INFO.request

This primitive initiates a request to a DEV for security information regarding either a single DEV or all of the DEVs in the piconet. The semantics of the primitive are as follows:

```
MLME-SECURITY-INFO.request  (
                             TrgtId,
                             QueriedDevId,
                             Timeout
                             )
```

Table 5-16—MLME-SECURITY-INFO primitive parameters

Name	Type	Valid range	Description
QueriedDevId	Integer	Any valid DEVID, as defined in 6.2.3, except for the McstID or the UnassocID	The DEVID of the DEV for which information is being requested. If it is set to the BcstID, then the information is being requested for all DEVs.
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3	The DEVID of the DEV for which the security information request is intended.
OrigId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the DEV that initiated the MLME request.
NumSecurityRecords	Integer	0–65535	Number of entries in the SecurityRecordSet.
SecurityRecordSet	A set of SecurityRecord fields, as defined in 6.5.4.4	A set containing 0 or more instances of variable-length SecurityRecord field. The maximum number of instances depends on the size of the records, <i>pMaxFrameBodySize</i> and the length of the secure command security fields, as defined in 6.3.3.2.	The SecurityRecordSet is returned to indicate the results of a Security Information Request command.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, TARGET_NOT_ASSOCIATED, OTHER	Indicates the reason for a ResultCode of FAILURE.

The primitive parameters are defined in Table 5-16.

5.3.10.2 MLME-SECURITY-INFO.confirm

This primitive reports the result of the request to a DEV for security information regarding either a single DEV or all of the DEVs in the piconet. The semantics of the primitive are as follows:

```
MLME-SECURITY-INFO.confirm
(
    TrgtId,
    NumSecurityRecords,
    SecurityRecordSet,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-16.

5.3.10.3 MLME-SECURITY-INFO.indication

This primitive indicates the reception of a request by a DEV for security information it manages regarding either a specific DEV or all of the DEVs in the piconet. The semantics of the primitive are as follows:

```
MLME-SECURITY-INFO.indication      (
                                     OrigId,
                                     QueriedDevId
                                     )
```

The primitive parameters are defined in Table 5-16.

5.3.10.4 MLME-SECURITY-INFO.response

This primitive is used by a DEV to respond to an MLME-SECURITY-INFO.indication. The semantics of the primitive are as follows:

```
MLME-SECURITY-INFO.response      (
                                     OrigId,
                                     NumSecurityRecords,
                                     SecurityRecordSet
                                     )
```

The primitive parameters are defined in Table 5-16.

5.3.11 Application specific data management

These primitives are used to request that the PNC add, modify, or remove application-specific data in the beacon and to report the reception of application specific data in a beacon, as described in 7.15. The parameters used for these primitives are defined in Table 5-17. If the RequestType is "ADD," then the MAC ignores the AsleIndex.

Table 5-17—MLME-APPLICATION-SPECIFIC primitive parameters

Name	Type	Valid range	Description
RequestType	Enumeration	ADD, MODIFY, REMOVE	If ADD, a request that a new AS IE be placed in the beacon. If MODIFY, a request to change the contents of an existing AS IE. If REMOVE, a request that a previously added AS IE be removed from the beacon.
AsleIndex	Octet	0–255	An ID assigned by the PNC to each AS IE successfully added to the beacon. If RequestType is MODIFY or REMOVE, it specifies the AS IE to modify or remove.
RequestID	Octet	0–255	A unique number assigned by the requesting DEV to identify the request.

Table 5-17—MLME-APPLICATION-SPECIFIC primitive parameters (continued)

Name	Type	Valid range	Description
UniqueId	Octet string	Any valid OUI or CID, as defined in 6.4.7	If RequestType is ADD or MODIFY, the unique identifier of the entity that defines the format of ApplicationData.
ApplicationData	Octet string	Octet string	If the RequestType is ADD or MODIFY, the application specific data to add to the beacon, as described in 7.15.
NumApplicationSpecificData	Integer	0–255	Indicates the number of ApplicationSpecificData in the piconet beacon.
ApplicationSpecificDataSet	Set of application specific data, as defined in Table 5-7	A set containing zero or more instances of an ApplicationSpecificData	The ApplicationSpecificDataSet is returned to indicate the ApplicationSpecificData in the piconet beacon.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, UNKNOWN_ID, PNC_DENIED, OTHER	Indicates the reason for a ResultCode of FAILURE.

5.3.11.1 MLME-APPLICATION-SPECIFIC.request

This primitive is used to request to add, modify, or remove application specific data in the piconet beacon. The semantics of this primitive are as follows:

```
MLME-APPLICATION-SPECIFIC.request (
    RequestType,
    RequestID,
    AsIeIndex,
    UniqueId,
    ApplicationData,
    Timeout
)
```

The primitive parameters are defined in Table 5-17.

5.3.11.2 MLME-APPLICATION-SPECIFIC.confirm

This primitive is used to report the result of the request to add, modify, or remove application specific data in the beacon. The semantics of this primitive are as follows:

```
MLME-APPLICATION-SPECIFIC.confirm (
    AsIeIndex,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-17.

5.3.11.3 MLME-APPLICATION-SPECIFIC.indication

This primitive is used to indicate that application specific data was in a successfully received beacon. The MLME passes all of the AS IEs found in every beacon to the DME. The semantics of this primitive are as follows:

```
MLME-APPLICATION-SPECIFIC.indication(
    NumApplicationSpecificData,
    ApplicationSpecificDataSet
)
```

The primitive parameters are defined in Table 5-17.

5.3.12 Piconet services management

These primitives are used to transfer information regarding the services offered by DEVs in a piconet, as described in 7.3.2. The parameters used for these primitives are defined in Table 5-18.

Table 5-18—MLME-ANNOUNCE-SERVICE and MLME-PICONET-SERVICES primitive parameters

Name	Type	Valid range	Description
Uniqueld	Octet string	Any valid OUI or CID, as defined in 6.4.7	The unique identifier for the entity that defines the format of the ServiceData.
ServiceData	Octet string	Octet string	The ServiceData, as described in 7.3.2.
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the DEV to request service information.
NumberOfServices	Integer	0–255	The number of PiconetServices in the PiconetServicesSet.
PiconetServicesSet	Set of piconet services, as defined in Table 5-19	A set containing zero or more instances of a PiconetService	The PiconetServicesSet is returned to indicate the PiconetServices offered.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.

**Table 5-18—MLME-ANNOUNCE-SERVICE and MLME-PICONET-SERVICES
primitive parameters (continued)**

Name	Type	Valid range	Description
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, TARGET_NOT_ASSOCIATED, OTHER	Indicates the reason for a ResultCode of FAILURE.

In Table 5-18, a PiconetServicesSet is a set of PiconetServices. Each PiconetService consists of the elements shown in Table 5-19.

Table 5-19—Elements of PiconetService

Name	Type	Valid range	Description
DevId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the DEV offering the service.
UniqueId	Octet string	Any valid OUI or CID, as defined in 6.4.7	The unique identifier of the entity that defines the format of the ServiceData.
ServiceData	Octet string	Octet string	The ServiceData, as described in 7.3.2.

5.3.12.1 MLME-ANNOUNCE-SERVICE.request

This primitive is used to request the announcement of the availability of a service offered by this DEV. The semantics of this primitive are as follows:

```
MLME-ANNOUNCE-SERVICE.request      (
                                     UniqueId,
                                     ServiceData,
                                     Timeout
                                     )
```

The primitive parameters are defined in Table 5-18.

5.3.12.2 MLME-ANNOUNCE-SERVICE.confirm

This primitive is used to report the result of the request to announce the availability of a service offered by this DEV. The semantics of this primitive are as follows:

```
MLME-ANNOUNCE-SERVICE.confirm      (
                                     ResultCode,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-18.

5.3.12.3 MLME-PICONET-SERVICES.request

This primitive is used to request piconet service information from a DEV in the piconet. If the target DEV is the PNC, then it is a request for all piconet service information previously announced by DEVs and stored in the PNC. The semantics of this primitive are as follows:

```
MLME-PICONET-SERVICES.request      (
                                     TrgtId,
                                     Timeout
                                     )
```

The primitive parameters are defined in Table 5-18.

5.3.12.4 MLME-PICONET-SERVICES.confirm

This primitive is used to report the result of the request for piconet service information from a DEV in the piconet. The semantics of this primitive are as follows:

```
MLME-PICONET-SERVICES.confirm      (
                                     NumberOfServices,
                                     PiconetServicesSet,
                                     ResultCode,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-18.

5.3.12.5 MLME-PICONET-SERVICES.indication

This primitive indicates the reception of piconet services information that was not requested with the MLME-PICONET-SERVICES.request primitive. The semantics of this primitive are as follows:

```
MLME-PICONET-SERVICES.indication   (
                                     NumberOfServices,
                                     PiconetServicesSet
                                     )
```

The primitive parameters are defined in Table 5-18.

5.3.13 Stream management

This mechanism supports the creation, modification, and termination of isochronous streams, as described in 7.5.1. The parameters used for the MLME-CREATE-STREAM, MLME-MODIFY-STREAM, and MLME-TERMINATE-STREAM primitives are defined in Table 5-20.

**Table 5-20—MLME-CREATE-STREAM, MLME-MODIFY-STREAM,
and MLME-TERMINATE-STREAM primitive parameters**

Name	Type	Valid range	Description
RequestID	Integer	0–255	A unique value created by the originating DME to match the request primitive with the response primitive it receives from the PNC MLME.
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the target DEV for an isochronous stream.
SourceDataRate	Integer	$1-(2^{64} - 1)$	The minimum required data rate of the stream source at the MAC SAP in bits per second, as described in F.1.5.
DesiredDataRate	Integer	$1-(2^{64} - 1)$	The desired data rate of the stream source at the MAC SAP in bits per second, as described in F.1.5.
MaxRetries	Integer	0–255	Specifies the maximum number of retries to attempt per transmitted frame.
AckRequested	Boolean	TRUE, FALSE	Indicates if acknowledgments will be used for the stream.
MaxTransmitDelay	Duration	$1-(2^{64} - 1)$	Maximum allowed delay in microseconds for transmitting a MAC service data unit (MSDU) once it is presented to MAC SAP.
UserPriority	Integer	As defined in Table B-1	User priority of the stream, as described in Table B-1.
TypicalDataFrameSize	Integer	$0-pMaxFrameBodySize$	The typical size in octets of an MSDU to be presented to the MAC SAP.
SecMode	Boolean	TRUE, FALSE	Indicates if security is to be applied to the stream.
StreamGrpID	Integer	0–255	A nonzero value specifies that channel time associated with this stream may be shared by other streams associated with the same stream group ID, as described in 7.5.1.
StreamIndex	Integer	Any valid stream index, as defined in 6.2.5	The index of a stream created or the index of a stream to modify or terminate.
OrigId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the source DEV for an isochronous stream.
AvailableDataRate	Integer	$1-(2^{64} - 1)$	If the request was successful, the data rate available for the stream. If the request was unsuccessful, the data rate that the PNC would have been able to allocate.

**Table 5-20—MLME-CREATE-STREAM, MLME-MODIFY-STREAM,
and MLME-TERMINATE-STREAM primitive parameters (continued)**

Name	Type	Valid range	Description
ReliabilityExponent	Integer	0–31	The negative power of 10 that is the maximum frame error ratio (FER) desired including retries and frames lost due to MaxTransmitDelay. For example, a value of 4 corresponds to an FER < 10 ⁻⁴ . If the value of the parameter is zero, then the parameter is ignored.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, TARGET_UNAVAILABLE, RESOURCES_UNAVAILABLE, TERMINATED_BY_PNC, TERMINATED_BY_DEST, TRANSMIT_DELAY_UNSUPPORTED, PNC_BUSY, DEV_IN_PS_MODE, NOT_ASSOCIATED, UNKNOWN_STREAM, OTHER	The reason for a ResultCode of FAILURE.

5.3.13.1 MLME-CREATE-STREAM.request

This primitive is used to request the creation of an isochronous stream. The semantics of this primitive are as follows:

```

MLME-CREATE-STREAM.request
(
    RequestID,
    TrgtId,
    SourceDataRate,
    DesiredDataRate,
    MaxRetries,
    AckRequested,
    MaxTransmitDelay,
    UserPriority,
    TypicalDataFrameSize,
    SecMode,
    StreamGrpID,
    ReliabilityExponent,
    Timeout
)

```

The primitive parameters are defined in Table 5-20.

5.3.13.2 MLME-CREATE-STREAM.confirm

This primitive is used to report the result of the request for the creation of an isochronous stream. The semantics of this primitive are as follows:

```
MLME-CREATE-STREAM.confirm      (
    RequestID,
    StreamIndex,
    AvailableDataRate,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-20.

5.3.13.3 MLME-CREATE-STREAM.indication

This primitive is used to inform a DEV that it is the target of an isochronous stream sourced from another DEV in the piconet. The semantics of this primitive are as follows:

```
MLME-CREATE-STREAM.indication  (
    StreamIndex,
    OrigId
)
```

The primitive parameters are defined in Table 5-20.

5.3.13.4 MLME-MODIFY-STREAM.request

This primitive is used to request a modification to the parameters defining an existing isochronous stream. The semantics of this primitive are as follows:

```
MLME-MODIFY-STREAM.request     (
    StreamIndex,
    SourceDataRate,
    DesiredDataRate,
    MaxRetries,
    AckRequested,
    MaxTransmitDelay,
    UserPriority,
    TypicalDataFrameSize,
    SecMode,
    StreamGrpID,
    ReliabilityExponent,
    Timeout
)
```

The primitive parameters are defined in Table 5-20.

5.3.13.5 MLME-MODIFY-STREAM.confirm

This primitive is used to report the result of the request to modify the parameters defining an existing isochronous stream. The semantics of this primitive are as follows:

```
MLME-MODIFY-STREAM.confirm      (
                                StreamIndex,
                                AvailableDataRate,
                                ResultCode,
                                ReasonCode
                                )
```

The primitive parameters are defined in Table 5-20.

5.3.13.6 MLME-TERMINATE-STREAM.request

This primitive is used to request the termination of an existing isochronous stream. The semantics of this primitive are as follows:

```
MLME-TERMINATE-STREAM.request   (
                                StreamIndex,
                                Timeout
                                )
```

The primitive parameters are defined in Table 5-20.

5.3.13.7 MLME-TERMINATE-STREAM.confirm

This primitive is used to report the result of the request to terminate an existing isochronous stream. The semantics of this primitive are as follows:

```
MLME-TERMINATE-STREAM.confirm   (
                                StreamIndex,
                                ResultCode,
                                ReasonCode
                                )
```

The primitive parameters are defined in Table 5-20.

5.3.13.8 MLME-TERMINATE-STREAM.indication

This primitive is used to report to a DEV other than the DEV terminating the stream that an isochronous stream has been terminated. The semantics of this primitive are as follows:

```
MLME-TERMINATE-STREAM.indication (
                                StreamIndex,
                                ReasonCode
                                )
```

The primitive parameter is defined in Table 5-20.

5.3.14 Piconet parameter management

These primitives allow a DEV acting as PNC to change the BSID of the piconet and allow all DEVs to be informed of changes to piconet characteristics, as defined in 7.11 and 7.12. The parameters used for these primitives are defined in Table 5-21.

Table 5-21—MLME-BSID-CHANGE and MLME-PICONET-PARM-CHANGE primitives parameters

Name	Type	Valid range	Description
IndicationChangeType	Enumeration	BSID, PNID, CHANNEL	Indicates the parameter of the piconet that was changed.
Bsid	Octet string	As defined in 6.4.2	If a request or if the IndicationChangeType is BSID, a new BSID for the piconet, as described in 6.4.2.
Pnid	Integer	0–65535	If the IndicationChangeType is PNID, the new ID for the piconet.
ChannelIndex	Integer	0–255	If the IndicationChangeType is CHANNEL, the new PHY-dependent channel number on which the piconet is operating.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	NOT_A_PNC, OTHER	The reason for a ResultCode of FAILURE.

5.3.14.1 MLME-BSID-CHANGE.request

This primitive initiates changing the BSID. The semantics of this primitive are as follows:

```

MLME-BSID-CHANGE.request
(
    Bsid
)

```

The primitive parameter is defined in Table 5-21.

5.3.14.2 MLME-BSID-CHANGE.confirm

This primitive reports the result of a request to change the BSID. The semantics of this primitive are as follows:

```

MLME-BSID-CHANGE.confirm
(
    ResultCode,
    ReasonCode
)

```

The primitive parameters are defined in Table 5-21.

5.3.14.3 MLME-PICONET-PARM-CHANGE.indication

This primitive reports to a DEV a change to a piconet characteristic that was not requested by the DME. The semantics of this primitive are as follows:

```
MLME-PICONET-PARM-CHANGE.indication
(
    IndicationChangeType,
    Bsid,
    Pnid,
    ChannelIndex
)
```

The primitive parameters are defined in Table 5-21.

5.3.15 Power management

This mechanism supports the process of establishment and maintenance of power management (PM) modes of a DEV, as described in 7.14. The parameters used for these primitives are defined in Table 5-22.

**Table 5-22—MLME-PS-SET-INFORMATION, MLME-SPS-CONFIGURE,
MLME-PM-MODE-CHANGE, and MLME-MONITOR-PM-MODE
primitive parameters**

Name	Type	Valid range	Description
PMMode	Enumeration	ACTIVE, APS, SPS	The PM mode requested by the DEV, as described in 6.5.8.5.
MaxSupportedPSSets	Integer	As defined in 6.5.8.2 and 7.14	The total number of PS sets currently supported by the PNC of this piconet.
NumCurrentPSSets	Integer	As defined in 6.5.8.2 and 7.14	Indicates the number of currently active PS sets in the piconet.
PSSetStructureSet	Set of PSSetStructures	As defined in Table 5-23	The PSSetStructureSet returns the information about the PS sets currently active in the PNC.
SetOperationType	Enumeration	CREATE, JOIN, LEAVE	The requested SPS set operation.
SPSSetIndex	Integer	As defined in 6.5.8.3	If the SetOperationType is JOIN or LEAVE, the SPS set index of the SPS set to join or leave. If the SetOperationType is CREATE, the SPS set index created by the PNC for the new SPS set.
DesiredWakeInterval	Integer	$0-(2^{32}-1)$	If the SetOperationType is CREATE, the period in microseconds at which the requesting DEV would desire to transition from the SLEEP state to the AWAKE state.
WakeInterval	Integer	$0-(2^{32}-1)$	The time period in microseconds at which DEVs in a power save set transition from the SLEEP state to the AWAKE state.

Table 5-22—MLME-PS-SET-INFORMATION, MLME-SPS-CONFIGURE, MLME-PM-MODE-CHANGE, and MLME-MONITOR-PM-MODE primitive parameters (continued)

Name	Type	Valid range	Description
PMActiveEvent	Enumeration	DATA_PENDING, MAX_SLEEP	If PMMode is ACTIVE, an event that causes the MLME to change the PM mode of operation to ACTIVE.
MonitorOperationType	Enumeration	ENABLE, DISABLE	The PM monitor operation requested.
TrgtId	Integer	Any valid DEVID, 6.2.3	Specifies the DEVID of the target DEV for a PM monitor operation.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, TARGET_NOT_ASSOCIATED, UNKNOWN_SET_INDEX, PNC_DENIED, PNC_BUSY, OTHER	The reason for a ResultCode of FAILURE.

In Table 5-22, PSDataSet is a set of PSDataSetStructures. Each PSDataSetStructure consists of the elements shown in Table 5-23.

In Table 5-23, MemberSet is a set of DEVIDs indicating the DEVIDs that are members of the PS set.

Table 5-23—Elements of PSDataSetStructure

Name	Type	Valid range	Description
PSDataSetIndex	Integer	As defined in 6.5.8.2	The identifier for the PS set.
WakeInterval	Integer	$0-(2^{32}-1)$	The period in microseconds at which DEVs in this power save set transition from the SLEEP state to the AWAKE state. The MAC divides this number by the superframe duration and rounds down to the nearest power of two.
NumMembers	Integer	0–255	The number of DEVs in this power save set.
MemberSet	Set of DEVIDs	Any valid DEVID, as defined in 6.2.3	The DEVIDs of the DEVs in this power save set.

5.3.15.1 MLME-PS-SET-INFORMATION.request

This primitive requests the current PS set information from the PNC. The semantics of this primitive are as follows:

```
MLME-PS-SET-INFORMATION.request      (
                                         Timeout
                                         )
```

The primitive parameter is defined in Figure 5-22.

5.3.15.2 MLME-PS-SET-INFORMATION.confirm

This primitive reports the result of the request to obtain the PS set information from the PNC. The semantics of this primitive are as follows:

```
MLME-PS-SET-INFORMATION.confirm      (
                                         MaxSupportedPSSets,
                                         NumCurrentPSSets,
                                         PSDataSetStructureSet,
                                         ResultCode,
                                         ReasonCode
                                         )
```

The primitive parameters are defined in Table 5-22.

5.3.15.3 MLME-SPS-CONFIGURE.request

This primitive requests a change to the current SPS set information maintained by the PNC. Possible requests include: create a new DSPS set and add the current DEVID, add the current DEVID to an existing SPS set, or remove the current DEVID from an existing set. The semantics of this primitive are as follows:

```
MLME-SPS-CONFIGURE.request           (
                                         SetOperationType,
                                         SPSSetIndex,
                                         DesiredWakeInterval,
                                         Timeout
                                         )
```

The primitive parameters are defined in Table 5-22.

5.3.15.4 MLME-SPS-CONFIGURE.confirm

This primitive reports the result of a request to change the current SPS set information. The semantics of this primitive are as follows:

```
MLME-SPS-CONFIGURE.confirm
(
    SetOperationType,
    SPSSetIndex,
    WakeInterval,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-22.

5.3.15.5 MLME-PM-MODE-CHANGE.request

This primitive requests a change to the DEV's PM mode of operation. The semantics of this primitive are as follows:

```
MLME-PM-MODE-CHANGE.request
(
    PMMode,
    Timeout
)
```

The primitive parameters are defined in Table 5-22.

5.3.15.6 MLME-PM-MODE-CHANGE.confirm

This primitive reports the result of the request to change the DEV's PM mode of operation. The semantics of this primitive are as follows:

```
MLME-PM-MODE-CHANGE.confirm
(
    PMMode,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-22.

5.3.15.7 MLME-PM-MODE-CHANGE.indication

This primitive is used to report a change to the DEVs PM mode of operation that was not requested with an MLME-PM-MODE-CHANGE.request but rather was initiated by the MAC. One reason for the MAC to change PM modes is in response to a CTA allocated with the DEV as the destination while the DEV is in a power save mode, as described in 7.14.2.2. This primitive reports only changes in the DEV's PM mode.

Changes in the PM mode of other DEVs in the piconet are reported with the MLME-MONITOR-PM-MODE.indication primitive. The semantics of this primitive are as follows:

```
MLME-PM-MODE-CHANGE.indication
(
    PMMode,
    PMActiveEvent
)
```

The primitive parameter is defined in Table 5-22.

5.3.15.8 MLME-MONITOR-PM-MODE.request

This primitive requests that the MLME enable or disable the monitoring of the PM mode for a DEV in the piconet. The semantics of this primitive are as follows:

```
MLME-MONITOR-PM-MODE.request
(
    MonitorOperationType,
    TrgtId
)
```

The primitive parameters are defined in Table 5-22.

5.3.15.9 MLME-MONITOR-PM-MODE.confirm

This primitive reports the result of a request that the MLME enable or disable the monitoring of the PM mode for a DEV in the piconet. The semantics of this primitive are as follows:

```
MLME-MONITOR-PM-MODE.confirm
(
    MonitorOperationType,
    TrgtId,
    PMMode,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-22.

5.3.15.10 MLME-MONITOR-PM-MODE.indication

This primitive reports a change in the PM mode of a DEV in the piconet, for which PM monitoring was previously enabled using an MLME-MONITOR-PM-MODE.request. This primitive only reports changes to the PM modes of other DEVs in the piconet. Changes to the DEV's PM mode are indicated with the MLME-PM-MODE-CHANGE.indication primitive. The semantics of this primitive are as follows:

```
MLME-MONITOR-PM-MODE.indication
(
    TrgtId,
    PMMode
)
```

The primitive parameters are defined in Table 5-22.

5.3.16 Multicast operations

These primitives support multicast operations. The MLME-MULTICAST-CONFIGURATION primitives are used for multicast traffic that uses a Group Address and a GrpID assigned by the PNC, as described in 7.5.3. The parameters used for these primitives are defined in Table 5-24.

Table 5-24—MLME-MULTICAST-CONFIGURATION primitive parameters

Name	Type	Valid range	Description
RequestType	Enumeration	JOIN, LEAVE	Indicates if this is a request to join a multicast group or leave a multicast group.
GroupAddress	MAC address	Any valid group address, as defined in 6.4.18	A group address representing a specific multicast group.
GrpID	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID assigned by the PNC to a multicast group associated with a specific GroupAddress.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	REQUEST_TIMEOUT, NOT_ASSOCIATED, PNC_DENIED, PNC_BUSY, OTHER	The reason for a ResultCode of FAILURE.

The MLME-MULTICAST-RX-SETUP.request primitive controls the reception of multicast traffic that uses the McstID, as defined in 6.2.3. The parameters for this primitive are defined in Table 5-25.

Table 5-25—MLME-MULTICAST-RX-SETUP.request parameters

Name	Type	Valid range	Description
MulticastStatus	Enumeration	ENABLE, DISABLE, ALL, NONE	If ENABLE or DISABLE, indicates whether the MAC will pass multicast traffic defined by the stream index to the DME. If ALL, then all multicast traffic is passed to the DME. If NONE, then no multicast traffic will be passed to the DME. These restrictions apply only to frames received with the DestID set to the McstID.
SrcID	Integer	Any valid DEVID, as defined in 6.2.3	The DEVID of the source of a multicast stream.
StreamIndex	Integer	Any valid stream index, as defined in 6.2.5	The stream index of a multicast stream.

Table 5-25—MLME-MULTICAST-RX-SETUP.request parameters (continued)

Name	Type	Valid range	Description
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	NOT_ASSOCIATED, UNKNOWN_STREAM, SOURCE_UNAVAILABLE, RESOURCES_UNAVAILABLE	Indicates the reason for a ResultCode of FAILURE.

5.3.16.1 MLME-MULTICAST-CONFIGURATION.request

This primitive is used by a DEV to request to join or leave a multicast group defined by a particular group address. The semantics of this primitive are as follows:

```
MLME-MULTICAST-CONFIGURATION.request (
    RequestType,
    GroupAddress,
    Timeout
)
```

The primitive parameters are defined in Table 5-24.

5.3.16.2 MLME-MULTICAST-CONFIGURATION.confirm

This primitive is used to report the result of a request by a DEV to join or leave a multicast group defined by a particular group address. The semantics of this primitive are as follows:

```
MLME-MULTICAST-CONFIGURATION.confirm (
    GroupAddress,
    GrpID,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-24.

5.3.16.3 MLME-MULTICAST-RX-SETUP.request

This primitive allows the DME to control multicast reception for frames with the DestID set to the McstID, as defined in 6.2.3, and to allow filtering for a particular stream index associated with the McstID. The semantics of this primitive are as follows:

```
MLME-MULTICAST-RX-SETUP.request (
    MulticastStatus,
    SrcID,
    StreamIndex
)
```

The primitive parameters are defined in Table 5-25.

5.3.16.4 MLME-MULTICAST-RX-SETUP.confirm

This primitive indicates the result of configuring the MAC for filtering multicast traffic. The semantics of this primitive are as follows:

```
MLME-MULTICAST-RX-SETUP.confirm    (
                                     ResultCode,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-25.

5.3.17 Timing synchronization

These primitives support the synchronization of upper layer functions with the beacon timing in the piconet. The beginning of the data preamble is observed on the air by all DEVs within a piconet while the delay between the observation and the delivery of the indication is known within a MAC by design (and communicated to the application by implementation-dependent means). The parameters used for these primitives are defined in Table 5-26.

Table 5-26—MLME-BEACON-EVENT primitive parameters

Name	Type	Valid range	Description
BeaconNumber	Integer	0–65535	The beacon number of the beacon that was received.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	NOT_SUPPORTED, OTHER	The reason for a ResultCode of FAILURE.

5.3.17.1 MLME-BEACON-EVENT.request

This primitive requests activation of the MAC synchronization support facility. The semantics of this primitive are as follows:

```
MLME-BEACON-EVENT.request          ( )
```

The primitive has no parameters.

5.3.17.2 MLME-BEACON-EVENT.confirm

This primitive confirms the activation of the MAC synchronization support facility. If the MAC does not support synchronization or if it encounters some other error, the response will have a ResultCode of FAILURE and the ReasonCode set to indicate the type of error. The semantics of this primitive are as follows:

```
MLME-BEACON-EVENT.confirm          (
                                     ResultCode,
                                     ReasonCode
                                     )
```

The primitive parameters are defined in Table 5-26.

5.3.17.3 MLME-BEACON-EVENT.indication

This primitive indicates the beginning of a beacon's data preamble, whether the beacon is transmitted or correctly received by the MAC. The semantics of this primitive are as follows:

```
MLME-BEACON-EVENT.indication      (
                                     BeaconNumber
                                     )
```

The primitive parameters are defined in Table 5-26.

5.3.18 Transmit switched diversity (TSD)

These primitives are used to exchange the TSD capability between a source and a destination, as described in 12.8. The parameters used for these primitives are defined in Table 5-27.

Table 5-27—MLME-TXDIV primitive parameters

Name	Type	Valid range	Description
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the DEVID of the target DEV for exchanging TSD information.
OrigId	Integer	Any valid DEVID, as defined in 6.2.3	The DEVID of the DEV that initiated the MLME request.
SupportedNumAntennas	Integer	0–15	The number of antennas supported by the DEV.
Timeout	Integer	0–65535	The time in milliseconds allowed for the primitive to complete.
TXDiversityThresholdType	Enumeration	LQI type, as defined in 6.4.11	Specifies the type of LQI measurement to use.
TXDiversityTheshold	Integer	Any valid LQI value, as defined in 6.4.27	Specifies the value of the threshold for antenna switching.
AntennaIndex	Integer	0–15	Specifies antenna index to be used.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	NOT_SUPPORTED, OTHER	Indicates the reason for a ResultCode of FAILURE.

5.3.18.1 MLME-TXDIV.request

This primitive is used to request to use TSD. The semantics of this primitive are as follows:

```
MLME-TXDIV.request      (
                          TrgtId,
                          TXDiversityThresholdType,
                          TXDiversityTheshold,
                          Timeout
                          )
```

The primitive parameters are defined in Table 5-27.

5.3.18.2 MLME-TXDIV.confirm

This primitive is used to report the result of the request for TSD. The semantics of this primitive are as follows:

```
MLME-TXDIV.confirm
(
    AntennaIndex,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-27.

5.4 MAC management

The MAC PIB comprises the managed objects, attributes, actions, and notifications required to manage the MAC sublayer of a DEV. The MAC PIB is divided into two groups, PNC characteristics and DEV characteristics. In the Access column of the Table 5-28 and Table 5-29, read only indicates that the parameter is only allowed to be read by the DME while read/write indicates that the DME is able to change it using the MLME-SET.request primitive.

5.4.1 MAC PIB PNC group

The MAC PIB PNC group, Table 5-28, describes both the DEV's PNC capabilities as well as the characteristics of the current piconet.

Table 5-28—MAC PIB PNC group parameters

Managed object	Octets	Definition	Access
<i>macCapEndTime</i>	2	The end time of the CAP interval in the superframes, 7.6.	Read only
<i>macSuperframeDuration</i>	2	Duration of the superframe.	Read only
<i>macPncCapable</i>	1 bit	1 if the DEV has the capability to become the PNC, 0 otherwise.	Read only
<i>macPncDesMode</i>	1 bit	1 if it is desired that the DEV be the PNC.	Read/write
<i>macMaxPsSets</i>	1	The maximum number of PS sets supported by the PNC.	Read only
<i>macBsid</i>	6–32	Identifies the piconet.	Read only
<i>macMaxAssociatedDevs</i>	1	As defined in 6.4.11.	Read only
<i>macMaxChannelTimeRequests</i>	1	As defined in 6.4.11.	Read only
<i>macSec</i>	1 bit	Indicates if the DEV is capable of operating a secure piconet as the PNC.	Read only
<i>macPncServicesBroadcast</i>	1	0x00 = PNC sends information about its services 0x01 = PNC will not send information about its services	Read/write

Table 5-28—MAC PIB PNC group parameters (continued)

Managed object	Octets	Definition	Access
<i>macAllowedChannelSet</i>	Variable	A set of channel indices, one for each channel that the MAC is allowed to use for scanning and starting piconets.	Read/write
<i>macAssocVendorSpecificIe</i>	Variable	A Vendor Defined IE, as defined in 6.4.17, that is sent in the Association Response command, as described in 6.5.1.2, when the DEV is acting as the PNC.	Read/write
<i>macDesiredAtp</i>	2	The ATP value to send in an Association Request command.	Read/write
<i>macPnid</i>	2	If associated with a piconet, the PNID of that piconet.	Read only
<i>macCapData</i>	1	Indicates if the initial setting of the CAP Data Allowed field in the beacon, as described in 6.3.1.1.	Read/write
<i>macCapCommand</i>	1	Indicates if the initial setting of the CAP Commands Allowed field in the beacon, as described in 6.3.1.1.	Read/write
<i>macCapAssociation</i>	1	Indicates if the initial setting of the CAP Data Allowed field in the beacon, as described in 6.3.1.1.	Read/write
<i>macPiconetMaxTxPower</i>	1	The maximum power allowed for transmission during certain times in the superframe as described in 6.3.1.1.	Read/write
<i>macMctaUsed</i>	1	The initial setting of the MCTA Used field, as described in 6.3.1.1.	Read/write
<i>macNextPncCapable</i>	1 bit	Indicates if the DEV is capable of participating in the Next PNC procedure, as described in 7.2.5. 0 = DEV is not capable. 1 = DEV is capable.	Read only
<i>macCapStartTime</i>	2	The time at which the CAP starts when quasi-omni beacons are used.	Read only

5.4.2 MAC PIB characteristic group

The MAC PIB characteristics group, Table 5-29, contains information about the capabilities and characteristics of the DEV.

Table 5-29—MAC PIB characteristic group parameters

Managed object	Octets	Definition	Access
<i>macDevAddress</i>	6	The MAC address of the DEV.	Read only
<i>macDevId</i>	1	The ID of the DEV.	Read only

Table 5-29—MAC PIB characteristic group parameters (continued)

Managed object	Octets	Definition	Access
<i>macPowerManagementMode</i>	1	The current power management mode of the DEV. 0x00 = ACTIVE 0x01 = PSPS 0x02 = DSPS 0x03 = DSPS and PSPS 0x04 = APS	Read only
<i>macPspSupported</i>	1	0x00 = DEV does not support PSPS mode. 0x01 = DEV supports PSPS mode.	Read only
<i>macDspSupported</i>	1	0x00 = DEV does not support DSPS mode. 0x01 = DEV supports DSPS mode.	Read only
<i>macApsSupported</i>	1	0x00 = DEV does not support APS mode. 0x01 = DEV supports APS mode.	Read only
<i>macMaxStreams</i>	1	Maximum number of streams that the DEV is able to handle.	Read only
<i>macPowerSource</i>	1	0x00 = Battery power. 0x01 = Mains power.	Read/write
<i>macSecurityOptionImplemented</i>	1	0x00 = Mode 0. 0x01 = Mode 1.	Read only
<i>macDevServicesBroadcast</i>	1	0x00 = DEV sends information about its services. 0x01 = DEV will not send information about its services.	Read/write
<i>macAlwaysAwake</i>	1 bit	As defined in 6.4.11.	Read only
<i>macListenToSource</i>	1 bit	As defined in 6.4.11.	Read only
<i>macListenToMulticast</i>	1 bit	As defined in 6.4.11.	Read only
<i>macCtaRelinquishCapable</i>	1 bit	Indicates if the DEV is capable of using time relinquished in a CTA by another DEV. 0x00 = DEV does not support. 0x01 = DEV does support.	Read only
<i>macDlyAckCapable</i>	1 bit	Indicates if the DEV is capable of using Dly-ACK as the source. 0 = DEV does not support. 1 = DEV does support.	Read only
<i>macImpAckCapable</i>	1 bit	Indicates if the DEV is capable of using implied ACK (Imp-ACK) as the source. 0 = DEV does not support. 1 = DEV does support.	Read only
<i>macStpCapable</i>	1 bit	Indicates if the DEV is capable of using the stream timeout period (STP). 0 = DEV does not support. 1 = DEV does support.	Read only

5.5 MAC SAP

The MAC SAP defines the logical interface between the MAC and the FCSL above it. This logical interface description includes a list of primitives and their definitions. Although these primitives and their definitions

are informative, they provide a context in which to understand the parameters that need to be passed between the MAC and the FCSL so that each sublayer may fulfill its specified functions.

The IEEE 802.15.3 MAC SAP primitives are summarized in Table 5-30.

Table 5-30—Summary of MAC SAP primitives

Name	Request	Confirm	Indication	Response
MAC-ASYNC-DATA	5.5.1	5.5.2	5.5.3	—
MAC-ISOCH-DATA	5.5.4	5.5.5	5.5.6	—

The parameters used for these primitives are defined in Table 5-31.

Table 5-31—MAC-ISOCH-DATA and MAC-ASYNC-DATA primitive parameters

Name	Type	Valid range	Description
RequestID	Integer	0–511	An identifier, which is used to correlate a request to a response. It is unique per stream index for outstanding isochronous requests. It is unique among all outstanding asynchronous requests.
StreamIndex	Integer	Any valid stream index, as defined in 6.2.5	The stream with which the data is associated.
TransmitTimeout	Duration	$1-(2^{32}-1)$	Maximum allowed delay in microseconds from when the MSDU is been presented to the MAC SAP until the frame has finished transmission and the acknowledgment, if required is successfully received.
MaxRetries	Integer	0–255	Specifies the maximum number of retries to attempt per transmitted frame with a maximum value no greater than the MaxRetries value supplied in the original stream creation request.
SecMode	Boolean	TRUE, FALSE	Indicates if security is to be applied to the MSDU or if the MSDU was received securely.
UserPriority	Integer	As defined in Table B-1	User priority of the stream, as described in Table B-1.
AckRequested	Boolean	TRUE, FALSE	Indicates if the request requires an acknowledgment, as described in 7.9, of the MSDU at the MAC layer.
ConfirmRequested	Enumeration	NEVER, ALWAYS, ON_ERROR	Indicates when a confirm primitive is required for the request.
Data	Variable number of octets	Any octet string	MSDU portion of the primitive.

Table 5-31—MAC-ISOCH-DATA and MAC-ASYNC-DATA primitive parameters (continued)

Name	Type	Valid range	Description
TransmitDelay	Duration	$1-(2^{32} - 1)$	Delay in microseconds from when the MSDU is presented to the MAC SAP until the frame has finished transmission and the acknowledgment, if required, has been successfully received. If the transmission fails due to timeout, this field shall be set to the TransmitTimeout value for this frame.
TrgtId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the target DEVID of an MSDU.
OrigId	Integer	Any valid DEVID, as defined in 6.2.3	Specifies the originator DEVID of an MSDU.
ResultCode	Enumeration	SUCCESS, FAILURE	Indicates the result of the MLME request.
ReasonCode	Enumeration	TRANSMIT_TIMEOUT, MAX_RETRIES, NOT_ASSOCIATED, OTHER	The reason for a ResultCode of FAILURE.
UncompressedVideo	Boolean	TRUE, FALSE	Set to TRUE if the Data parameter contains uncompressed video data.
UepAllowed	Boolean	TRUE, FALSE	Indicates if UEP is allowed for the transmission of the data.
ErrorsFreeData	Boolean	TRUE, FALSE	Indicates if the Data parameter that was received contains errors or is free from errors.
InterlacedFieldIndication	Enumeration	TOP_FIELD, BOTTOM_FIELD, NOT_INTERLACED	If UncompressedVideo is TRUE, the parameter indicates if the pixels belong to the top or bottom field of an interlaced image or that the image is not interlaced.
VideoFrameNumber	Integer	As defined in 6.2.9.1.4	If UncompressedVideo is TRUE, the parameter contains a sequential numbering of the video frames that are being transferred.
HPosition	Integer	As defined in 6.2.9.1.4	If UncompressedVideo is TRUE, the parameter contains the horizontal position of the first pixel in the MSDU.
VPosition	Integer	As defined in 6.2.9.1.4	If UncompressedVideo is TRUE, the parameter contains the vertical position of the first pixel in the MSDU.

The parameters InterlacedFieldIndication, VideoFrameNumber, HPosition, and VPosition are used by the AV PHY in the Video Subheader field, as described in 6.2.9.1.4. These fields are sent in the MAC header because it has lower error rate than the subframes carrying video data. For uncompressed video applications, video data that is in error can still be displayed, but only if the position of the pixel data is known.

5.5.1 MAC-ASYNC-DATA.request

This primitive is used to initiate the transfer of an MSDU from one MAC entity to another MAC entity or entities. All asynchronous data should use Multi-protocol Data frames to avoid interoperability problems in piconets where different FCSs are being used. The semantics of this primitive are as follows:

```
MAC-ASYNC-DATA.request      (
                                RequestID,
                                TrgtId,
                                TransmitTimeout,
                                MaxRetries,
                                UserPriority,
                                AckRequested,
                                ConfirmRequested,
                                Data
                                )
```

The primitive parameters are defined in Table 5-31.

5.5.2 MAC-ASYNC-DATA.confirm

This primitive is used to report the result of a request to transfer an asynchronous MSDU from one MAC entity to another MAC entity or entities. This primitive is only generated if the ConfirmRequested parameter in the MAC-ASYNC-DATA.request with the same RequestID value is ALWAYS or is ON_ERROR and the ResultCode is FAILURE. The semantics of this primitive are as follows:

```
MAC-ASYNC-DATA.confirm      (
                                RequestID,
                                TransmitDelay,
                                ResultCode,
                                ReasonCode
                                )
```

The primitive parameters are defined in Table 5-31.

5.5.3 MAC-ASYNC-DATA.indication

This primitive is used to indicate the reception of an asynchronous MSDU. The semantics of this primitive are as follows:

```
MAC-ASYNC-DATA.indication    (
                                TrgtId,
                                OrigId,
                                Data
                                )
```

The primitive parameters are defined in Table 5-31.

5.5.4 MAC-ISOCH-DATA.request

This primitive is used to initiate the transfer of an isochronous MSDU from one MAC entity to another MAC entity or entities. The semantics of this primitive are as follows:

```
MAC-ISOCH-DATA.request      (
    RequestID,
    StreamIndex,
    TransmitTimeout,
    MaxRetries,
    AckRequested,
    ConfirmRequested,
    Data,
    DataType,
    UepAllowed,
    InterlacedFieldIndication,
    VideoFrameNumber,
    HPosition,
    VPosition
)
```

The primitive parameters are defined in Table 5-31.

5.5.5 MAC-ISOCH-DATA.confirm

This primitive is used to report the result of a request to transfer an isochronous MSDU from one MAC entity to another MAC entity or entities. This primitive is only generated if the ConfirmRequested parameter in the MAC-ISOCH-DATA.request with the same RequestID value is ALWAYS or is ON_ERROR and the ResultCode is FAILURE. The semantics of this primitive are as follows:

```
MAC-ISOCH-DATA.confirm      (
    RequestID,
    StreamIndex,
    TransmitDelay,
    ResultCode,
    ReasonCode
)
```

The primitive parameters are defined in Table 5-31.

5.5.6 MAC-ISOCH-DATA.indication

This primitive is used to indicate the reception of an isochronous MSDU. The semantics of this primitive are as follows:

```
MAC-ISOCH-DATA.indication      (  
                                TrgtId,  
                                OrigId,  
                                StreamIndex,  
                                Data,  
                                DataType,  
                                ErrorFreeData,  
                                InterlacedFieldIndication,  
                                VideoFrameNumber,  
                                HPosition,  
                                VPosition  
                                )
```

Data identified as video by the source DEV may be passed by the destination DEV to the higher layers even if the data was received in error. The error status of the data passed up by this primitive is indicated in the ErrorFreeData parameter.

The primitive parameters are defined in Table 5-31.

6. MAC frame formats

6.1 Frame format conventions

The MAC in all DEVs shall be able to validate the error free reception of every frame received from the PHY using the frame check sequence (FCS). Note that the PHY only passes frames to the MAC that have passed the header check sequence (HCS) test.

The MAC frames in the MAC sublayer are described as a sequence of fields in a specific order. Each figure in Clause 6 depicts the fields as they appear in the MAC frame and in the order in which they are transmitted in the wireless medium, from left to right where the leftmost bit is transmitted first in time.

In the figures, all bits within fields are numbered from 0 (left) to k (right) where the length of the field is $k+1$ bits. The octet boundaries within a field are obtained by taking the bit-numbers of the field modulo 8. Octets within numeric fields that are longer than a single octet are depicted in increasing order of significance, from lowest numbered bit on the left to the highest numbered bit on the right. The octets in fields longer than a single octet are sent to the PHY in order from the octet containing the lowest numbered bits to the octet containing the highest numbered bits. For any text fields, the first character is in the first octet of the field with other characters following sequentially. An example of the bit and octet ordering is illustrated in Figure 6-1.




Figure 6-1—Example of bit and octet ordering

The payload in the data frame is sent with the lowest numbered octet first, least significant bit (LSB) first, over the air, as illustrated in Figure 6-2. The FCS is an exception to this convention and is transmitted with the most significant bit (MSB) first. In the case of the HCS, the PHY determines the bit order in which the HCS is sent.

Values specified in decimal are coded in unsigned binary unless otherwise stated.

Without further qualification, “reception” by the MAC sublayer implies that the frame contents are valid and that the protocol version is supported. However, reception implies nothing about frame addressing, nor whether the frame type or other fields in the MAC header are meaningful to the MAC entity that has received the frame.




Figure 6-2—Example of payload transmission order

Unless otherwise stated, any reserved field or subfield shall be set to zero upon transmission and shall be ignored on reception.

Reserved values in nonreserved fields shall not be transmitted by conformant DEVs. However, a DEV may receive frames with values that it considers to be reserved values in nonreserved fields. These fields, along with other fields in the same frame that rely on the interpretation of these fields, shall be ignored on reception.

All DEVs shall be assigned a DEV address, which is the 48-bit extended universal identifier (EUI-48), as defined by IEEE Std 802-2014⁸ and assigned by the IEEE Registration Authority.⁹ An individual MAC address is a MAC address with the group bit set to zero, as defined in IEEE Std 802-2014. All MAC addresses shall be sent in the canonical form defined in IEEE Std 802-2014.

6.2 General frame format

The MAC frame format, illustrated in Figure 6-3, comprises a set of fields that occur in a fixed order in all frames. The figures in this subclause are a representation of the MAC Header field and MAC Frame Body field. The HCS is not shown since this is calculated and verified by the PHY. The MAC frame shall be formatted as illustrated in Figure 6-3. The maximum size of the MAC Frame Body field, *pMaxFrameBodySize*, is a PHY-dependent parameter that includes the frame payload and FCS fields, but not the PHY preamble, PHY header, MAC header, MAC subheader, or MAC header validation. The parameter *pMaxFrameBodySize* is defined in the following subclauses:

- 10.2.8.1 for the 2.4 GHz PHY
- 11.2.7.1 for the SC PHY mode
- 11.3.6.3 for the HSI PHY mode
- 11.4.1.3.1 for the AV PHY mode

The number of octets in the MAC Frame Body field shall range from zero to *pMaxFrameBodySize*, inclusive. The maximum length of the MAC Frame Body field includes the length of the security fields, if present.

⁸Information on references can be found in Clause 2.

⁹Interested applicants should contact the IEEE Registration Authority, <http://standards.ieee.org/develop/regauth/>.

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks

Octets: 2	2	1	1	3	1	variable	0 or 4
Frame Control	PNID	DestID	SrcID	Fragmentation Control	Stream Index	Frame Payload	FCS
MAC Header						MAC Frame Body	

Figure 6-3—MAC Header field and MAC Frame Body field formats

The Non-secure MAC Frame Body field shall be formatted as illustrated in Figure 6-4 when the security (SEC) field is set to zero in the Frame Control field.

Octets: variable	4
Frame Payload	FCS
Non-secure MAC Frame Body	

Figure 6-4—Non-secure MAC Frame Body field format

The Secure MAC Frame Body field shall be formatted as illustrated in Figure 6-5 when the SEC field shall be set to one in the Frame Control field. The Secure Payload field in the Secure MAC Frame Body field contains the Frame Payload protected as indicated in 9.2.2.

Octets: 2	2	variable	8	4
SECID	Secure Frame Counter	Secure Payload	Integrity Code	FCS
Frame payload				
Secure MAC Frame Body				

Figure 6-5—Secure MAC Frame Body field format

6.2.1 Frame Control field

The Frame Control field shall be formatted as illustrated in Figure 6-6.

Bits: b0–b2	b3–b5	b6	b7–b8	b9	b10	b11	b12	b13	b14	b15
Protocol Version	Frame Type	SEC	ACK Policy	Retry	More Data	Imp-ACK Request	Imp-ACK NAK	CTA Relinquish	Blk-ACK	Reserved

Figure 6-6—Frame Control field format

6.2.1.1 Protocol Version field

The Protocol Version field is invariant in size and placement across all revisions of IEEE Std 802.15.3. For this revision of the standard the value of the protocol version is 0b000. All other values are reserved. The

revision level will be incremented only when a fundamental incompatibility exists between a new revision and the prior revision of the standard. A DEV that receives a frame with a higher revision level than it supports may discard the frame without indication to the sending DEV.

6.2.1.2 Frame Type field

The Frame Type field indicates the type of frame that is being sent. Table 6-1 lists the valid frame type values and their description. The format and the usage of each of the individual frame types is defined in 6.3.

Table 6-1—Valid frame type values

Type value b5 b4 b3	Frame type description	Subclause
0b000	Beacon frame	6.3.1
0b001	Imm-ACK frame	6.3.2.1
0b010	Delayed ACK (Dly-ACK) frame	6.3.2.2
0b011	Command frame	6.3.3
0b100	Data frame	6.3.4
0b101	Multi-protocol Data frame	6.3.5
0b110	Sync frame	6.3.6
0b111	Reserved	—

6.2.1.3 SEC field

The SEC field shall be set to one when the Frame Payload field is protected using the key specified by the security ID (SECID). The SEC field shall be set to zero otherwise. Frames with the SEC field set to one shall use the secure frame format for that frame type, as described in 6.3.

6.2.1.4 ACK Policy field, Imp-ACK Request field, and Blk-ACK field

The ACK Policy field, Imp-ACK Request field, and Blk-ACK field are used to indicate the type of acknowledgment procedure that the addressed recipient is required or allowed to perform. The use of the ACK procedures is described in 7.9. The allowed values for the ACK Policy field, Imp-ACK Request field, and Blk-ACK field are defined in Table 6-2. The ACK policy of a frame is determined by the combination of the ACK Policy field, Imp-ACK Request field, and Blk-ACK field.

An ACK policy of Dly-ACK or Dly-ACK Request is valid only in the data frames of a stream that is currently employing the Dly-ACK mechanism. It is not valid for frames using the asynchronous stream index or the MCTA stream index.

6.2.1.5 Retry field

The Retry field shall be set to one in any data or command frame that is a retransmission of an earlier frame. It shall be set to zero in all other frames.

Table 6-2—Valid ACK policy field type values

Blk-ACK field	Imp-ACK Request field	ACK Policy field	ACK policy type	Description
0	0	00	No ACK	The recipient(s) does not acknowledge the transmission, and the sender treats the transmission as successful without regard for the result, as described 7.9.1.
0	0	01	Imm-ACK	The addressed recipient returns an Imm-ACK frame after successful reception, as described in 7.9.2.
0	0	10	Dly-ACK	The addressed recipient keeps track of the frames received with this policy until requested to respond with a Dly-ACK frame, as described in 7.9.3.
0	0	11	Dly-ACK request	The addressed recipient returns either an Imm-ACK or a Dly-ACK frame after successful reception, as described in 7.9.3.
0	1	01	Imp-ACK	The addressed recipient returns an Imm-ACK frame, a data frame, or a command frame after successful reception, as described in 7.9.6.
1	0	01	Blk-ACK	The addressed recipient returns a Data frame with the MAC subheader including the Blk-ACK Bitmap field after successful reception of the MAC header, as described in 7.9.5.

6.2.1.6 More Data field

The More Data field shall be set to zero if the DEV will not use the rest of the channel time in that CTA, as described in 7.4.3.1. The More Data field shall be set to zero in the last frame of an extended beacon and in a Beacon frame that is not part of an extended beacon, as described in 7.6.2. In all other cases the More Data field shall be set to one. This includes frames, other than the last one, that are part of an extended beacon.

6.2.1.7 Imp-ACK Negative Acknowledgment (NAK) field

The Imp-ACK NAK field shall be set to one by a DEV when all of the following conditions are true:

- The DEV is responding to a frame addressed to the DEV for which it has successfully received the MAC header with an ACK policy of Imp-ACK.
- The FCS check, 6.2.7.6, for the Frame Payload field failed.

The Imp-ACK NAK field shall be set to zero otherwise. The use of the Imp-ACK NAK field is described in 7.9.6.

6.2.1.8 CTA Relinquish field

The CTA Relinquish field shall be set to one when the DEV relinquishes CTA ownership to another DEV. It shall be set to zero otherwise. The use of the CTA Relinquish field is described in 7.4.3.8.

6.2.2 Piconet ID (PNID) field

The PNID field contains the unique identifier for the piconet, as described in 7.11.3. The PNID normally remains constant during the current instantiation of the piconet and may be persistent for multiple sequential instantiations of the piconet by the same PNC. The PNID field shall be set to the current PNID for the piconet and is used to identify frames from DEVs in the piconet.

6.2.3 SrcID and DestID fields

There are two DEVID fields in the MAC frame format. These fields are used to indicate the source DEVID (SrcID) and destination DEVID (DestID). A DEVID for a DEV is assigned by the PNC during the association of the DEV. The DEVID is unique to an associated DEV within a piconet. The following DEVIDs are assigned or reserved:

- 0x00 is the PNCID
- 0xED-0xF6 are reserved
- 0xF7, 0xF8, 0xF9, 0xFA, 0xFB, or 0xFC are NbrIDs
- 0xFD is the McstID
- 0xFE is the UnassocID
- 0xFF is the BcstID

The maximum number of valid DEVs, *mMaxNumValidDEVs*, is the maximum number of DEVIDs that the PNC is able to allocate in a piconet. This includes all of the regular DEVIDs, the PNCID, the GrplIDs, and the NbrIDs, but not the reserved IDs, the BcstID, McstID, or the UnassocID.

6.2.4 Fragmentation Control field

The Fragmentation Control field is used to aid in the fragmentation and a reassembly of MSDUs and command frames. The Fragmentation Control field shall be formatted as illustrated in Figure 6-7.

Bits: b0–b8	b9–b15	b16–b22	b23
MSDU Number	Fragment Number	Last Fragment Number	Reserved

Figure 6-7—Fragmentation Control field format

The three octets that compose the Fragmentation Control field may be used for reporting PHY-dependent receive status information to the transmitting DEV in an Imm-ACK frame, a Dly-ACK frame, or a Data frame sent in response to a Blk-ACK request that has a zero length Data Payload field. If the source DEV is not reporting PHY-dependent receive status information in an Imm-ACK, Dly-ACK frame, or a Data frame sent in response to a Blk-ACK request that has a zero length Data Payload field, it shall set the fragmentation field of the frame to all zeros, i.e., 0x000000. All other values are PHY dependent. The receive status for the 2.4 GHz PHY is defined in 10.7. The receive status for the mmWave PHYs is defined in 11.1.8.3.

For Data frames sent in response to a Blk-ACK request that have a non-zero length Data Payload field, the Fragmentation Control field shall contain the fragmentation information.

6.2.4.1 MSDU Number field

The MSDU Number field indicates the sequence number of the current MSDU or command frame.

For data frames, each DEV shall maintain one modulo-512 counter for each of its isochronous streams, and one for its asynchronous data traffic. The MSDU numbers for all command frames shall be assigned from a single modulo-512 counter.

Each MSDU number counter shall be set to zero when the DEV is initialized. The MSDU number counter for an isochronous stream shall be set to zero when the stream index for the isochronous stream is first assigned by the PNC, as described in 7.5.1.1.

6.2.4.2 Fragment Number field

The Fragment Number field indicates the order of the current fragment within the current MSDU. The Fragment Number field shall be set to zero in all unfragmented frames.

6.2.4.3 Last Fragment Number field

The Last Fragment Number field indicates the total number of fragments within the current MSDU. The value of this field is equal to one less than the number of fragments. The Last Fragment Number field shall be the same for every fragment of a fragmented MSDU and shall be set to zero for all unfragmented MSDUs.

6.2.5 Stream Index field

The Stream Index field assigned values are as follows:

- 0x00 is assigned for asynchronous data
- 0xFB is assigned for beam forming starting from sector-level training
- 0xFC is assigned for beam forming starting from beam-level training
- 0xFD is assigned for MCTA traffic
- 0xFE is assigned for unassigned streams
- 0xFF is reserved

DEVs use other values of the stream index as dynamically assigned by the PNC during the setup of the data stream, as described in 6.5.6.1. The PNC allocates a unique stream index value for each isochronous stream in the piconet. A DEV shall support all stream index values that are not reserved.

6.2.6 MAC header validation

When the PHY receives a frame, it validates the received frame's MAC header before passing the MAC Header field and its associated Frame Payload field to the MAC. The protection mechanism used to validate the MAC Header field is PHY dependent. In addition, the bit order and the length of the protection mechanism, *pLengthHcs*, are also PHY dependent. The MAC Header field protection mechanism is defined in the following subclauses:

- 10.2.9 for the 2.4 GHz PHY
- 11.2.3.2.2 for the SC PHY mode
- 11.3.3.4 for the HSI PHY mode
- 11.4.1.4 for the AV PHY mode

6.2.7 MAC Frame Body field

6.2.7.1 Frame Payload field

The Frame Payload field is a variable-length field that carries the information that is to be transferred to a DEV or group of DEVs in the piconet. In the case of a secure frame, it also includes the required security information and the secure payload, Figure 6-5.

6.2.7.2 Secure session ID (SECID) field

The SECID field contains an identifier for the key that is being used to protect the frame. The SECID field shall be formatted as illustrated in Figure 6-8.

Octets: 1	1
Key Originator	Key ID

Figure 6-8—SECID field format

The Key Originator field for all keys except the piconet group data key shall be set to the DEVID of the key originator in the relationship. The Key Originator field for the piconet group data key shall be set to the BcstID.

Key ID field shall have a unique value for the key associated with the security relationship. The Key ID field value for a given key is selected by the key originator in a security relationship, as described in 8.3.8.

6.2.7.3 Secure Frame Counter (SFC) field

The Secure Frame Counter (SFC) field contains a counter that is used to ensure the uniqueness of the nonce in a secure frame. A DEV shall not reuse a frame counter with the same time token, as described in 6.3.1.1, and key, as described in 8.3.5. The DEV shall initialize the SFC to zero for the first frame sent and increment it for each successive secure frame sent. When the time token, as described in 6.3.1, is updated, the DEV shall reset the SFC to zero. In the case where the DEV receives a new key, the DEV shall set the SFC to zero.

6.2.7.4 Secure Payload field

The Secure Payload field is a variable-length field that contains the information, protected by the symmetric key security operations, as defined in 9.3, that is to be transferred to a DEV or group of DEVs in the piconet. As illustrated in Figure 6-5, the Secure Payload field is a part of the Frame Payload field and does not include the SECID, SFC, or Integrity Code fields.

6.2.7.5 Integrity Code field

The Integrity Code field contains an encrypted integrity code that is used to cryptographically protect the integrity of the MAC header and Frame Payload. The integrity code is computed as specified in 9.3.

6.2.7.6 Frame Check Sequence (FCS) field

The Frame Check Sequence (FCS) field contains a 32-bit CRC. The CRC described here is equivalent to ANSI X3.66-1979. The MSB of the FCS is the coefficient of the highest order term, and the field is sent

over the wireless medium, as indicated in 6.1. The FCS shall be calculated over the entire Frame Payload field, which is referred to here as the Calculation field.

If the Frame Payload field has zero length (as in an immediate ACK frame) the FCS shall not be sent.

The FCS is calculated using the following standard generator polynomial of degree 32:

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

The FCS is the one's complement of the modulo 2 sum of the remainders in item a) and item b) below:

- The remainder resulting from $((x^k \times (x^{31} + x^{30} + \dots))$ divided (modulo 2) by $G(x)$. The value k is the number of bits in the calculation field.
- The remainder resulting from the calculation field contents, treated as a polynomial, is multiplied by x^{32} and then divided by $G(x)$.

The FCS field shall be transmitted in the order specified in 6.1.

At the transmitter, the initial remainder of the division shall be preset to all ones and is then modified via division of the calculation fields by the generator polynomial $G(x)$. The one's complement of this remainder is the FCS field.

At the receiver, the initial remainder shall be preset to all ones. The serial incoming bits of the calculation fields and FCS, when divided by $G(x)$ in the absence of transmission errors, results in a unique non-zero remainder value. The unique remainder value is the following polynomial:

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

6.2.8 SC and HSI aggregated frame format

Figure 6-9 illustrates the SC and HSI aggregated frame format.



Figure 6-9—SC and HSI aggregated frame format

6.2.8.1 Standard aggregation format

6.2.8.1.1 Non-secure standard aggregation format

To use standard aggregation, the Aggregation field and Low-latency Mode field in PHY header shall be set as described in 11.2.3.2.1.

The MAC Subheader field for standard aggregation shall be formatted as illustrated in Figure 6-10.

Octets: 2	5	...	5	1
Blk-ACK Bitmap	Subheader 1	...	Subheader 8	RX Buffer Size

Figure 6-10—MAC Subheader field format for standard aggregation

The Blk-ACK Bitmap field shall be formatted as illustrated in Figure 6-11.

Octets: 1	1
ACK or MSB ACK	ACK or LSB ACK

Figure 6-11—Blk-ACK bitmap field format

In equal error protection (EEP) mode, the ACK or MSB ACK field is a bitmap that indicates if a subframe was correctly received. The bit position zero, which is the first bit from right in Figure 6-11, corresponds to the first subframe of the frame that is being ACKed. If a subframe was correctly received, then the bit for that subframe shall be set to one and shall be set to zero otherwise. The ACK or LSB ACK field shall be set to the same value as the ACK or MSB ACK field.

To use UEP mode, the UEP field in PHY header shall be set as described in 11.2.3.2.4. In UEP mode, the ACK or MSB ACK field is a bitmap that indicates if the MSBs of a subframe were correctly received. The bit position zero, which is the first bit from right in Figure 6-11, corresponds to the first subframe of the frame that is being ACKed. If the MSBs of the subframe were correctly received, then the bit for that subframe shall be set to one and shall be set to zero otherwise. The ACK or LSB ACK field is a bitmap that indicates if the LSBs of a subframe were correctly received. The bit position zero of the ACK or LSB ACK field corresponds to the first subframe of the frame that is being ACKed. If the LSBs of the subframe were correctly received, then the bit for that subframe is set to one and shall be set to zero otherwise.

The Subheader field shall be formatted as illustrated in Figure 6-12.

Bits: b0–b4	b5	b6	b7	b8–b18	b19–b20	b21	b22–b30	b31–b37	b38	b39
MCS information	FCS present	Retry	Resolution indication	Subframe length	Subframe information	Skewed constellation	MSDU number	Fragment number	Last fragment	Reserved

Figure 6-12—Subheader field format

In UEP mode, the MCS information field shall be set to the UEP MCS (as defined in 11.2.2.9 for SC PHY and 11.3.2.1 for HSI PHY) that is used for the subframe.

In EEP mode, the MCS Information field shall be set to zero. All the subframes use the MCS indicated in the PHY header.

The FCS Present field shall be set to one if the subframe uses an FCS and shall be set to zero otherwise. The FCS is optional for subframes for the cases in which the upper layer will handle checking the data integrity. If the FCS present field is set to zero, then the subframe is considered correctly received only if the PHY header, MAC header, and MAC subheader were all correctly received.

The Retry field shall be set to one if the subframe is a retransmission and shall be set to zero otherwise.

The Resolution Indication field shall be set to zero if the resolution of the Subframe Length field is 1 octet and shall be set to one if the resolution is 512 octets.

The Subframe Length field is used to determine the length of the subframe before coding, not including the FCS. If the resolution of the Subframe Length field is 1 octet, then this field contains the length of the subframe in octets. If the resolution of the Subframe Length field is 512 octets, then the length of the subframe in octets is 512 times the value of this field. If the Subframe length field is set to zero, the corresponding subframe is not present in the aggregated data frame.

The Subframe Information field indicates if the subframe contains MSB, LSB, or MSB and LSB combined data. Valid values for the Subframe Information field are as follows:

- 0 → The subframe contains only MSB data
- 1 → The subframe contains only LSB data
- 2 → The subframe contains MSB and LSB data
- 3 → Reserved

The Skewed Constellation field shall be set to one if a skewed constellation is used in the subframe and shall be set to zero otherwise.

The MSDU Number field indicates the MSDU number of the subframe. Each subframe that contains a fragment from the same MSDU shall have the same MSDU number.

The Fragment Number field indicates the fragment sequence number of the subframe within the current MSDU, if the subframe contains a MSDU fragment. This field shall be set to zero if the corresponding subframe contains an unfragmented MSDU.

The Last Fragment field shall be set to one if the subframe contains the last fragment of the MSDU and shall be set to zero otherwise.

The RX Buffer Size field indicates the free buffer space at the target DEV as a multiple of the preferred fragment size, as defined in 6.4.11.

The MAC Frame Body field for standard aggregation shall be formatted as illustrated in Figure 6-13.

Octets: variable	4 or 8	...	variable	4 or 8
Subframe Payload 1	FCS or Combined FCS	...	Subframe Payload <i>n</i>	FCS or Combined FCS

Figure 6-13—MAC Frame Body field format for standard aggregation

The maximum number of subframes that are aggregated in one frame shall be *mMaxSubframeSize*, as defined in 7.16.

In EEP mode, the subframe contains the FCS field, as defined in 6.2.7.6. In UEP mode, the subframe contains the Combined FCS field, as defined in Figure 6-41. The Subframe Payload field is formatted as defined in 6.3.5 for the Multi-protocol Data frame.

6.2.8.1.2 Secure standard aggregation format

The secure MAC Subheader field for standard aggregation shall be formatted as illustrated in Figure 6-14.

Octets: 5	2	5	...	5	1
Security Header	Blk-ACK bitmap	Subheader 1	...	Subheader 8	RX Buffer Size

Figure 6-14—Secure MAC Subheader field format for standard aggregation

The Security Header field shall be formatted as illustrated in Figure 6-15.

Octets: 2	2	1
SECID	SFC	Subframe Security

Figure 6-15—Security Header field format

The SECID field is used to identify the key set that is used to encrypt and/or authenticate the data in the frame, as defined in 6.2.7.2.

The SFC field contains a counter that is used to ensure the uniqueness of the nonce of a secure frame, as defined in 6.2.7.3.

The Subframe Security field contains a bitmap that indicate if a subframe applies security or not. The bit position zero, which is the first bit from right in Figure 6-15, corresponds to the first subframe. The bit shall be set to one if the subframe applies security and shall be set to zero otherwise.

The secure MAC Frame Body field for standard aggregation shall be formatted as illustrated in Figure 6-16.

Octets: variable	8	4 or 8	...	variable	8	4 or 8
Subframe payload 1	Integrity code	FCS or Combined FCS	...	Subframe payload <i>n</i>	Integrity code	FCS or Combined FCS

Figure 6-16—Secure MAC Frame Body field format for standard aggregation

The Integrity Code field is used to cryptographically protect the integrity of the header and payload, as defined in 6.2.7.5.

6.2.8.2 Low-latency aggregation formats**6.2.8.2.1 Non-secure low-latency aggregation format**

To use low-latency aggregation mode, the Aggregation field and the Low-latency Mode field in PHY header shall be set as described in 11.2.3.2.1.

The MAC Subheader field for low-latency aggregation shall be formatted as illustrated in Figure 6-17.

Octets: 1	1	2	2	64
UEP	RX Buffer Size	MSDU Request Number	MSDU Response Number	Blk-ACK bitmap

Figure 6-17—MAC Subheader field format for low-latency aggregation

The UEP field shall be formatted as illustrated in Figure 6-18.

Bits: b0–b4	b5–b7
UEP MCS	Reserved

Figure 6-18—UEP field format

In EEP mode, the UEP MCS field shall be set to zero and all subframes use the MCS indicated in the PHY header.

In UEP mode, the UEP MCS field indicates the MCS used for UEP mode for the subframes, as defined in 11.2.2.9.

The RX Buffer size field indicates the free buffer space at the destination DEV as a multiple of the preferred fragment size, as defined in 6.4.11.

The MSDU Request Number field and the MSDU Response Number field shall be formatted as illustrated in Figure 6-19.

Bits: b0–b8	b9–b15
MSDU Sequence Number	Reserved

Figure 6-19—MSDU Request/Response Number field format

The MSDU Sequence Number field is used as the basis for the Blk-ACK bitmap.

The MSDU Request Number field indicates the most recent MSDU sequence number acknowledged at the transmitter.

The MSDU Response Number field indicates the first MSDU sequence number of the transmitted Blk-ACK bitmap-field. The MSDU Response Number field is used to generate the offset for the Blk-ACK bitmap field. The MSDU Response Number field shall be a copy from the received MSDU Request Number, upon reception of a valid MAC Subheader.

For EEP mode the Blk-ACK Bitmap field shall be formatted as illustrated in Figure 6-20.

Bits: b0	b1	...	b255	b256–b511
Subframe ACK #1	Subframe ACK #2	...	Subframe ACK #256	Reserved

Figure 6-20—EEP mode Blk-ACK Bitmap field format

The Subframe ACK field is an offset from the MSDU sequence number of the received MSDU Response Number field. It shall be set to one if the corresponding subframe was correctly received and shall be set to zero otherwise.

For UEP mode, the Blk-ACK Bitmap field shall be formatted as illustrated in Figure 6-21.

Bits: b0	b1	b2	b3	...	b510	b511
Subframe MSB ACK #1	Subframe LSB ACK #1	Subframe MSB ACK #2	Subframe LSB ACK #2	...	Subframe MSB ACK #256	Subframe LSB ACK #256

Figure 6-21—UEP mode Blk-ACK Bitmap field format

The Subframe MSB ACK field shall be set to one if the MSBs of the corresponding subframe were correctly received and shall be set to zero otherwise.

The Subframe LSB ACK field shall be set to one if the LSBs of the corresponding subframe were correctly received and shall be set to zero otherwise.

The MAC Frame Body field for low-latency aggregation shall be formatted as illustrated in Figure 6-22.

Octets: 3	variable	4 or 8	...	3	variable	4 or 8
MSDU Subheader n	Subframe Payload 1	FCS or Combined FCS	...	MSDU Subheader n	Subframe Payload n	FCS or Combined FCS

Figure 6-22—MAC Frame Body field format for low-latency aggregation

The maximum number of subframes that are aggregated in one frame shall be $mMaxSubframeSize$, as defined in 7.16:

The MSDU Subheader field shall be formatted as illustrated in Figure 6-23.

Bits: b0–b8	b9–b14	b15	b16–b23
MSDU Number	Subframe Length	Reserved	MSDU Subheader HCS

Figure 6-23—MSDU Subheader field format

The MSDU Number field indicates the sequence number of the MSDU that is aggregated in the subframe. The MSDU number is unique in a window of 256 MSDUs, allowing retransmission of an out-of-order MSDU over several frames.

Subframe Length field contains the length of each subframe in unit of four octets. This field may be set to zero to allow idle transmission of data if data is not present at the FCSL. The HCS field shall be inverted in the case of a zero-length MSDU, so that the MSDU HCS check fails.

The MSDU Subheader HCS field contain a CRC defined as the one's complement of the remainder of the division of the first 16 bits of the header by the polynomial $x^8 + x^2 + x + 1$. A serial implementation is illustrated in Figure 6-24.




Figure 6-24—HCS implementation example for MSDU subheader

In EEP mode, the subframe contains the FCS field, as defined in 6.2.7.6. In UEP mode, the subframe contains the Combined FCS field, as defined in Figure 6-41.

6.2.8.2.2 Secure low-latency aggregation format

The Secure MAC Subheader field for low-latency aggregation shall be formatted as shown in Figure 6-25.

Octets: 2	2	1	1	2	2	32/64
SECID	SFC	UEP	RX Buffer Size	MSDU Request Number	MSDU Response Number	Blk-ACK Bitmap

Figure 6-25—Secure MAC Subheader field format for low-latency aggregation

The SECID field is used to identify the key set that is used to encrypt and/or authenticate the data in the frame, as defined in 6.2.7.2.

The SFC field contains a counter that is used to ensure the uniqueness of the nonce of a secure frame, as defined in 6.2.7.3.

The secure MAC Frame Body field for low-latency aggregation shall be formatted as shown in Figure 6-26.

Octets: 3	variable	8	4 or 8	...	3	variable	8	4 or 8
Secure MSDU Subheader 1	Subframe Payload 1	Integrity Code	FCS or Combined FCS	...	Secure MSDU Subheader <i>n</i>	Subframe Payload <i>n</i>	Integrity Code	FCS or Combined FCS

Figure 6-26—Secure MAC Frame Body field format for low-latency aggregation

The Secure MSDU Subheader field shall be formatted as illustrated in Figure 6-27.

Bits: b0–b8	b9–b14	b15	b16–b23
MSDU Number	Subframe Length	Subframe Security	MSDU Subheader HCS

Figure 6-27—Secure MSDU Subheader field format

The Subframe Security field shall be set to one if the subframe applies security and shall be set to zero otherwise.

The Integrity Code field is used to cryptographically protect the integrity of the header and payload, as defined in 6.2.7.5.

6.2.9 AV aggregated frame format

The AV aggregated frame format is optimized to carry uncompressed audio and video in an efficient manner. The AV aggregated frame format is used instead of the standard aggregation or low-latency aggregation formats.

Figure 6-28 illustrates the AV aggregated MAC Frame Body field format.

MAC Frame Body				
Subframe 1	Subframe 2	...	Subframe 6	Subframe 7

Figure 6-28—AV aggregated MAC Frame Body field format

The collection of subframes in a single frame is referred to as the MAC Frame Body field. The MAC Frame Body field of a MAC frame may have one to seven, inclusive, subframes of varying sizes.

6.2.9.1 Extended MAC header format

The HRP Extended MAC Header field is used to describe the contents of an AV aggregated frame, which typically is used to transport uncompressed audio and video.

The HRP Extended MAC Header field shall be formatted as illustrated in Figure 6-29.

Octets: 10	2	5	6	24	16
MAC header	Extended Control Header	MAC Extension Header	Security Header	Video Header	Reserved

Figure 6-29—HRP Extended MAC Header field format

The LRP Extended MAC Header field shall be formatted as illustrated in Figure 6-30.

Octets: 10	2	5	5
MAC Header	Extended Control Header	MAC Extension Header	Security Header

Figure 6-30—LRP Extended MAC Header field format

The MAC Extension Header field and Security Header field are present in LRP frames only if their corresponding fields are set to one in the Extended Control Header.

6.2.9.1.1 Extended Control Header field

The Extended Control Header field shall be formatted as illustrated in Figure 6-31.

Bits: b0–b2	b3	b4	b5	b6–b8	b9–b15
Frame Class	MAC Extension Header Present	Security Header Present	Video Header Present	Frame Type	Reserved

Figure 6-31—Extended Control Header field format

Valid values for the Frame Class field are given in Table 6-3.

Table 6-3—Frame Class field values

Field value	Frame class
0b000	Regular
0b0001	AV aggregated
0b010	Omni-ACK
0b011	Beacon
0b100–0b111	Reserved

Regular frames have a single payload field without MAC level aggregation while AV aggregated frames have one or more subframes as a part of the frame.

The MAC Extension Header Present field shall be set to one if the MAC Extension Header field is in the Extended MAC Header field and shall be set to zero otherwise.

The Security Header Present field shall be set to one if the Security Header field is in the Extended MAC Header field and shall be set to zero otherwise.

The Video Header Present field shall be set to one if the Video Header field is in the Extended MAC Header field and shall be set to zero otherwise.

For HRP frames, the MAC Extension Header Present field, Security Header Present field, and Video Header Present field indicate the validity of the associated field. Because the HRP header is a fixed length, these fields are always present in the Extended MAC Header field. If a header in the HRP MAC header is present but not valid, it may be set to any value and shall be ignored upon reception. For LRP frames, unused headers, as indicated by the appropriate header present field, are not present in the MAC header.

The Frame Type field is only defined for regular frame class; it shall be set to zero in other frames. AV aggregated frames use the Type field in the MAC Extension Header field while the Omni-ACK and Beacon frames do not require the Frame Type field. Valid values for the Frame Type field are as follows:

- 0 → MAC commands
- 1 → Data
- 2 → Audio
- 3–7 → Reserved

6.2.9.1.2 MAC Extension Header field

The MAC Extension Header field is illustrated in Figure 6-32.

Bits: b0–b3	...	b24–b27	b28–b31	b32–b39
Type 1	...	Type 7	Reserved	ACK Groups

Figure 6-32—MAC Extension Header field format

The Type field indicates the type of data that is contained in the subframe. Valid values for the Type field are as follows:

- 0x0 → MAC commands
- 0x1 → Data
- 0x2 → Audio
- 0x3 → Video
- 0x4–0xF → Reserved

The ACK Groups field shall be formatted as illustrated in Figure 6-33.

Bits: b0	b1	...	b6	b7
Subframe 1	Subframe 2	...	Subframe 7	LSB FCS

Figure 6-33—ACK Groups field format

The bit for a subframe shall be set to one if the subframe is in the same ACK group as the previous (i.e., lower numbered subframe). Otherwise, it is the first subframe in an ACK group and its bit shall be set to zero. The first bit, corresponding to subframe 1, shall always be set to zero as it is always the start of an ACK group. No more than 5 ACK groups shall be defined; therefore, the number of bits set to zero among the subframe bits shall not exceed five.

The LSB FCS field shall be set to one if the LSB FCS is part of the calculation to determine if a subframe was correctly received, as defined in 7.1. It shall be set to zero if the LSB FCS is ignored in determining if a

subframe was correctly received. The setting of this field applies only to those subframes that are sent with a UEP HRP mode. For all other subframes, the LSB FCS field shall be set to zero.

6.2.9.1.3 Security Header field

The Security Header field shall be formatted as illustrated in Figure 6-34.

Octets: 4	2
Security Control	SFC

Figure 6-34—Security Header field format

The Security Control field shall be formatted as illustrated in Figure 6-35.

Bits: b0–b15	b16–b17	b18–b19	...	b30–b31
SECID	Reserved	Subframe 1 Security	...	Subframe 7 Security

Figure 6-35—Security Control field format

The SECID field is used to identify the key set that is used to encrypt and/or authenticate the data in the frame, as defined in 6.2.7.2.

The Subframe Security field indicates the type of security that is applied to a subframe. Valid values are as follows:

- 0b00 → No security applied
- 0b01 → Encryption and integrity code
- 0b10–0b11 → Reserved

The SFC field is defined in 6.2.7.3. The SFC shall be incremented for each subframe in the frame, even if security is not applied to the subframe.

6.2.9.1.4 Video Header field

The Video Header field shall be formatted as illustrated in Figure 6-36.

Octets: 5	5	5	5	4
Video Control 1	Video Control 2	Video Control 3	Video Control 4	Reserved

Figure 6-36—Video Header field format

Unless otherwise stated, the numbers for all of the fields begin with zero, e.g., the first video frame is number zero, then number one, and so on. One of the first video subframes in the first frame sent in a stream should have the Video Frame Number field, the H-Position field, and V-Position field set to zero.

The Video Subheader field is included as part of the MAC header because the error rate for the MAC header is lower than that of the subframes.

The Video Control field shall be formatted as illustrated in Figure 6-37.

Bits: b0–b3	b4–b6	b7	b8–b23	b24–b39
Reserved	Video Frame Number	Interlaced Field Indication	H-Position	V-Position

Figure 6-37—Video Control field format

The Interlaced Field Indication field shall be set to one if the video subframes carry pixels for the bottom field. It shall be set to zero if the video subframes carry pixels for the top field or if the video subframes carry pixels for non-interlaced video modes.

The Video Frame Number field contains a counter that keeps track of the video frame to which the pixels in the subframe belong. The video frame number is calculated follows:

- For progressive video, the Video Frame Number field shall be incremented sequentially. After reaching the max value of 0x7, the next value shall be zero. All frames belonging to the same video frame have identical Video Frame Number values.
- For interlaced video, the Video Frame Number field shall be incremented in a step of two. Thus, each video frame has two frame numbers. All frames belonging to the first field have even Video Frame Numbers and all frames belonging to the second field have odd Video Frame Numbers. For example, for the first uncompressed video frame, the frames belonging to the first field have a Video Frame Number set to zero, and the frames belonging to the second field have a Video Frame Number set to one. Therefore, the same video frame has two Video Frame Numbers.

The H-Position field contains the horizontal position of the first pixel in the subframe where zero is on the left side of the screen.

The V-Position field contains the vertical position of the first pixel in the subframe where zero corresponds to the top of the screen. For interlaced formats, the V-positions of the lines range from 0 to 539 independent of status as even or odd frame number.

6.2.9.2 Subframe format

The subframes with Type Video in the MAC frame shall be formatted as illustrated in Figure 6-38.

Octets: variable	8
Subframe Payload	Combined FCS

Figure 6-38—Subframe field format for type video

Subframes other than those with Type Video shall be formatted as illustrated in Figure 6-39.

Octets: variable	4
Subframe Payload	FCS

Figure 6-39—Subframe field format for type other than video

The Subframe Payload field shall be formatted as defined in 6.3.5 for the Multi-protocol Data frame. If the subframe has an integrity code, as indicated by the Subframe Security field, then Subframe Payload field shall be formatted as illustrated in Figure 6-40, where the Payload Data field shall be formatted as defined in 6.3.5 for the Multi-protocol Data frame.

Octets: 8	variable
Integrity Code	Payload Data

Figure 6-40—Subframe Payload with integrity code

The Integrity Code field is defined in 6.2.7.5.

The Combined FCS field shall be formatted as illustrated in Figure 6-41.

Bits: b0–b3	b4–b7	b8–b11	...	b28–b31	b32–b35	b36–b39	b40–b43	...	b52–b55	b56–b59	b60–b63
LSB FCS 5	MSB FCS 5	LSB FCS 6	...	MSB FCS 8	LSB FCS 1	MSB FCS 1	LSB FCS 2	...	MSB FCS 3	LSB FCS 4	MSB FCS 4

Figure 6-41—Combined FCS field format

The LSB FCS field contains the FCS, as defined in 6.2.7.6, calculated over the four LSBs of each octet in the Subframe Payload field. The LSB FCS field occupies the LSBs of the Combined FCS field.

The MSB FCS field contains the FCS, as defined in 6.2.7.6, calculated over the four MSBs of each, if present, and Subframe Payload field. The MSB FCS field occupies the MSBs of the Combined FCS field.

The FCS field contains the FCS, as defined in 6.2.7.6, calculated over Subframe Payload field.

Subframes with frame type other than data do not support aggregation or fragmentation.

Subframes with frame type data support the aggregation of MSDUs or the fragmentation of MSDUs in their payloads. The format of a Subframe Payload field for a subframe with Frame Type data is illustrated in Figure 6-42.

Octets: variable	variable	...	variable
Sub-payload 1	Sub-payload 2	...	Sub-payload <i>n</i>

Figure 6-42—Subframe Payload field with Frame Type data

Each Sub-payload field shall be formatted as shown in Figure 6-43.

Bits: b0–b19	b20–b29	b30	b31	variable
Length	Sequence Number	Last Fragment	First Fragment	MSDU

Figure 6-43—Sub-payload field format

The Length field indicate the length of the MSDU field in octets.

The Sequence Number field is incremented for each fragment, regardless if it is of the same or different MSDU.

The Last Fragment field and First Fragment field shall be set as indicated in Table 6-4.

Table 6-4—Fragment field settings

Fragment type	Last Fragment field	First Fragment field
First fragment	0	1
Last fragment	1	0
Middle fragment	0	0
Complete MSDU	1	1

The MSDU field for subframes with frame type data shall be formatted, as defined in 6.3.5.

6.2.9.3 AV aggregated

The AV aggregated frame shall be formatted as illustrated in Figure 6-44. The AV aggregated frame shall not be used for LRP frames.




Figure 6-44—AV aggregated frame format

The AV aggregated frame has one to seven, inclusive, subframes. The subframe format is defined in 6.2.9.2.

6.2.9.4 Regular

The HRP regular frame shall be formatted as illustrated in Figure 6-45.



Figure 6-45—HRP regular frame format

The LRP regular frame shall be formatted as illustrated in Figure 6-46.




Figure 6-46—LRP regular frame format

The subframe format is defined in 6.2.9.2.

6.2.9.5 Directional ACK

The Directional ACK frame format is defined in 11.4.3.8. The Directional ACK frame is used to acknowledge HRP frames and beam formed LRP frames. The ACK policy for a directional ACK shall be Imm-ACK. The directional ACK shall be only used with the AV PHY HRP frames.

6.3 Format of individual frame types

6.3.1 Beacon frame

6.3.1.1 Non-secure Beacon frame

The Non-secure Beacon frame Frame Payload field shall be formatted as illustrated in Figure 6-47.

Octets: 19	variable	...	variable	4
Piconet synchronization parameters	IE-1	...	IE- <i>n</i>	FCS

Figure 6-47—Non-secure Beacon frame Frame Payload field format

The individual IEs in the Beacon frame Frame Payload field are listed in Table 6-13. These IEs are encoded in type, length, value format and are defined in 6.4. The IEs in the beacon payload may appear in any order except for the channel time allocation (CTA) IEs, which shall be the first IEs of the beacon payload following the Piconet Synchronization Parameters field.

The Piconet Synchronization Parameter field shall be formatted as illustrated in Figure 6-48.

Octets: 6	2	2	1	1	1	6
Time Token	Superframe Duration	CAP End Time	Max TX Power Level	Piconet Mode	PNC Response	PNC Address

Figure 6-48—Piconet Synchronization Parameters field format

The Time Token field contains a strictly increasing counter, which is incremented in each beacon. The Time Token counter shall be set to zero by the PNC when a piconet is started, as defined in 7.2.2. The beacon number is defined to be the 16 LSBs of the time token.

NOTE—For a 1 ms superframe, the time token would roll over in just under 9000 years. At that point the PNC would shutdown the piconet and start a new one.

The Superframe Duration field contains the duration of the current superframe. The resolution of this field is 1 μ s and therefore has a range of [0–65535] μ s. However, the valid range of this field is [*mMinSuperframeDuration*, *mMaxSuperframeDuration*].

The CAP End Time field specifies the end of the CAP interval for the current superframe, as defined in 7.6. The resolution of this field is 1 μ s which gives a range of [0–65535] μ s. The CAP begins a SIFS after the end of the beacon and continues until the CAP end time.

The Max TX Power Level field is used to indicate the maximum TX power level allowed in the current superframe by the PNC in the piconet, as described in 7.12.2.1. The value is in dBm encoded in two's complement format. For example, a +2 dBm TX power level is encoded as 0x02 while a –2 dBm TX power level is encoded as 0xFE. However, if the PNC does not want to limit the TX power, then it shall set the field to 0x7F.

The Piconet Mode field defines certain characteristics about the piconet and the superframe. The encoding of this octet shall be formatted as illustrated in Figure 6-49.

Bits: b0	b1	b2	b3	b4	b5	b6	b7
CAP Data	CAP Commands	CAP Association	MCTA Used	SEC Mode	Reserved	Other IEs Unchanged	CTA IEs Unchanged

Figure 6-49—Piconet Mode field

If the CAP Data field, CAP Commands field, or CAP Association field value is set to one, then that type of data or command is allowed to be sent in the CAP of the current superframe. The CAP Commands field applies to all commands except for the Association Request command, which is covered by the CAP Association field. Otherwise, that type of frame is not allowed to be sent in the CAP. The use of these fields is described in 7.4.2 and 10.2.10.

The MCTA Used field shall be set to one if the PNC will be using open or association MCTAs in the superframe.

The SEC Mode field indicates the current security settings in the piconet, as defined in 8.2. The field is encoded as illustrated in Table 6-5.

Table 6-5—SEC Mode field encodings

Type value	SEC mode
0	Mode 0
1	Mode 1

The Other IEs Unchanged field may be set to one if all the IEs (other than CTA IEs) in the beacon payload are identical to the IEs (other than CTA IEs) contained in the previous beacon, as described in 7.6.2.

The CTA IEs Unchanged field may be set to one if the CTA IEs in the beacon are identical to the CTA IEs contained in the previous beacon, as described in 7.6.2.

The PNC Response field shall be formatted as illustrated in Figure 6-50.

Bits: b0–b3	b4–b7
MCTA Allocation Rate	Reserved

Figure 6-50—PNC Response field format

The MCTA Allocation Rate field indicates the frequency with which the PNC will be allocating either open MCTAs or directed uplink MCTAs for each DEV. For example, if the MCTA Allocation Rate field is set to a value of 8, the PNC is indicating that it will be providing either an open MCTA or a directed uplink MCTA for each DEV in the piconet at least once out of every 8 superframes. A value of 15 means the PNC is not giving any guarantees about when it will allocate MCTAs. A value of 0 indicates that the PNC is using only the CAP to provide access to the PNC.

The PNC Address field contains the DEV address of the PNC, as described in 6.1.

The MAC header settings for a Non-secure Beacon frame shall be set and interpreted as described in Table 6-6.

Table 6-6—MAC header settings for a Non-secure Beacon frame

Header field	Setting on transmission	Interpretation on reception
Frame type	Beacon value in Table 6-1	Decoded
SEC	0	Decoded
ACK policy	No-ACK value in Table 6-2	May be ignored
Retry	0	May be ignored
More data	As required, 7.6.2	Decoded
DestID	BcstID	Decoded
SrcID	PNCID	Decoded
Fragmentation control	0x000000	May be ignored
Stream index	0x00	May be ignored

6.3.1.2 Secure Beacon frame

The Secure Beacon frame Secure Payload field shall be formatted as illustrated in Figure 6-51. The Secure Beacon frame format is used when the piconet is operating in a secure mode.

Octets: 2	2	19	variable	...	variable	8	4
SECID	SFC	Piconet synchronization parameters	Information element-1	...	Information element- <i>n</i>	Integrity code	FCS

Figure 6-51—Secure Beacon frame Secure Payload field format

The SECID field is defined in 6.2.7.2.

The SFC field is used by the DEV for this frame to ensure uniqueness of the nonce, as defined in 6.2.7.3.

The Piconet Synchronization Parameters field is defined in 6.3.1.1.

The Integrity Code is defined in 6.2.7.5.

The MAC header settings for a Secure Beacon frame shall be set and interpreted as described in Table 6-7.

Table 6-7—MAC header settings for a Secure Beacon frame

Header field	Setting on transmission	Interpretation on reception
Frame type	Beacon value in Table 6-1	Decoded
SEC	1	Decoded
ACK policy	No-ACK value in Table 6-2	May be ignored
Retry	0	May be ignored
More data	As required, 7.6.2	Decoded
DestID	BcstID	Decoded
SrcID	PNCID	Decoded
Fragmentation control	0x000000	May be ignored
Stream index	0x00	May be ignored

6.3.2 Acknowledgment frames

6.3.2.1 Imm-ACK frame

The Imm-ACK frame has neither a Frame Payload field nor an FCS field.

The MAC header settings for an Imm-ACK frame shall be set and interpreted as described in Table 6-8.

Table 6-8—MAC header settings of an Imm-ACK frame

Header field	Setting on transmission	Interpretation on reception
Frame type	Imm-ACK value in Table 6-1	Decoded
SEC	0	May be ignored
ACK policy	No-ACK value in Table 6-2	May be ignored
Retry	0	May be ignored
More data	0	May be ignored
DestID	SrcID of the received frame	Decoded
SrcID	DestID of the received frame	Decoded

Table 6-8—MAC header settings of an Imm-ACK frame (continued)

Header field	Setting on transmission	Interpretation on reception
Fragmentation control	0x000000 or Receive Status value, 6.2.4	May be ignored
Stream index	0x00	May be ignored

6.3.2.2 Delayed ACK (Dly-ACK) frame

The Dly-ACK frame is only used in response to an isochronous stream data frame with the ACK Policy field set to Dly-ACK Request. The Dly-ACK frame Frame Payload field shall be formatted as illustrated in Figure 6-52.

Octets: 1	1	1	2	...	2	4
Max Burst	Max Frames	MPDUs ACKed	MPDU ID Block-1	...	MPDU ID Block- <i>n</i>	FCS

Figure 6-52—Dly-ACK frame Frame Payload field format

The Max Burst field indicates the number of frames of *pMaxFrameBodySize* that may be sent in one burst. A burst is the collection of the frames that are pending acknowledgment via a Dly-ACK frame.

The Max Frames field indicates the maximum number of frames, regardless of size, that may be sent before requesting a Dly-ACK from the DEV receiving the frames.

Any burst shall meet the restrictions of both the Max Frames field and the Max Burst field, as described in 7.9.3.

The MPDUs ACKed field shall contain the number of MPDUs that are being ACKed with this frame.

The MPDU ID Block fields shall be sent in the same order as the data frames were received. The MPDU ID Block field shall be formatted as illustrated in Figure 6-53.

Bits: b0–b8	b9–b15
MSDU number	Fragment number

Figure 6-53—MPDU ID Block field format

The MAC header settings for a Dly-ACK frame shall be set and interpreted as described in Table 6-9.

Table 6-9—MAC header settings of a Dly-ACK frame

Header field	Setting on transmission	Interpretation on reception
Frame type	Dly-ACK value in Table 6-1	Decoded
SEC	0	May be ignored

Table 6-9—MAC header settings of a Dly-ACK frame (continued)

Header field	Setting on transmission	Interpretation on reception
ACK policy	No-ACK value in Table 6-2	May be ignored
Retry	0	May be ignored
More data	0	May be ignored
DestID	SrcID of the received frame	Decoded
SrcID	DestID of the received frame	Decoded
Fragmentation control	0x000000 or Receive Status value, 6.2.4	May be ignored
Stream index	0x00	May be ignored

6.3.3 Command frame

When sending command frames the ACK Policy field in the Frame Control field shall be set to either Imm-ACK or no-ACK.

Only certain commands may be fragmented, as indicated in 6.5. For commands that are not allowed to be fragmented, the Fragmentation Control field shall be set to 0x000000.

6.3.3.1 Non-secure Command frame

The Non-secure Command frame Frame Payload field shall be formatted as shown in Figure 6-54. The command types are described in 6.5.

Octets: variable	4
Command block	FCS

Figure 6-54—Non-secure Command frame Frame Payload field format

The Command Block field shall be formatted as shown in Figure 6-55.

Octets: 2	2	variable
Command Type	Length	Payload

Figure 6-55—Command Block field format

The Command Type field indicates the type of command and is defined in Table 6-22.

The Length field contains the length of the Payload field in octets.

The Payload field contains information specific to the MAC command. The Payload field for each of the MAC commands are defined in 6.5.

The MAC header settings for a Non-secure Command frame shall be set and interpreted as described in Table 6-10.

Table 6-10—MAC header settings of a Non-secure Command frame

Header field	Setting on transmission	Interpretation on reception
Frame type	Command value in Table 6-1	Decoded
SEC	0	Decoded
ACK policy	As required for command protocol	Decoded
Retry	As appropriate	Decoded
More data	As appropriate	Decoded
DestID	As appropriate	Decoded
SrcID	As appropriate	Decoded
Stream index	0x00 or 0xFD	May be ignored

6.3.3.2 Secure Command frame

The Secure Command frame Secure Payload field shall be formatted as illustrated in Figure 6-56.

Octets: 2	2	variable	8	4
SECID	SFC	Command Block	Integrity Code	FCS

Figure 6-56—Secure Command frame Secure Payload field format

The SECID field is defined in 6.2.7.2.

The SFC field is defined in 6.2.7.3.

The Command Block field shall be formatted as illustrated in Figure 6-55.

The Integrity Code field is defined in 6.2.7.5.

This frame format is used when the piconet is operating in a secure mode.

The MAC header settings for a Secure Command frame shall be set and interpreted, as described in Table 6-11.

Table 6-11—MAC header settings of a Secure Command frame

Header field	Setting on transmission	Interpretation on reception
Frame type	Command value in Table 6-1	Decoded
SEC	1	Decoded

Table 6-11—MAC header settings of a Secure Command frame (continued)

Header field	Setting on transmission	Interpretation on reception
ACK policy	As required for command protocol	Decoded
Retry	As appropriate	Decoded
More data	As appropriate	Decoded
DestID	As appropriate	Decoded
SrcID	As appropriate	Decoded
Stream index	0x00 or 0xFD	May be ignored

6.3.4 Data frame**6.3.4.1 Non-secure Data frame**

The Non-secure Data frame Frame Payload field shall be formatted as shown in Figure 6-57.

Octets: variable	0 or 4
Data Payload	FCS

Figure 6-57—Non-secure Data frame Frame Payload format

The Frame Type field shall be set to the Data frame value in Table 6-1, and the SEC field shall be set to zero. The other fields in the MAC header take on values that are appropriate for that particular Non-secure Data frame. All fields in the MAC header of a Non-secure Data frame shall be decoded on reception.

Null data frames, which are non-secure data frames with no Data ID field, no Data Header field and a zero-length Data Payload field, are allowed. For example, a null data frame may be used with Dly-ACK negotiation, as described in 7.9.3.

The FCS is not sent for zero-length frames, as described in 6.2.7.6.

6.3.4.2 Secure Data frame

The Secure Data frame Secure Payload field shall be formatted as illustrated in Figure 6-58.

Octets: 2	2	variable	8	4
SECID	SFC	Data Payload	Integrity Code	FCS

Figure 6-58—Secure Data frame Secure Payload field format

The Frame Type field shall be set to the Data frame value in Table 6-1, and the SEC field shall be set to one. The other fields in the MAC header take on values that are appropriate for that particular Secure Data frame. All fields in the MAC header of a Secure Multi-protocol Data frame shall be decoded on reception.

The SECID field is defined in 6.2.7.2.

The SFC field is defined in 6.2.7.3.

The Data ID field is defined in 6.3.4.1.

The Data Header field is defined in 6.3.4.1.

6.3.5 Multi-protocol Data frame

6.3.5.1 Non-secure Multi-protocol Data frame

The Non-secure Multi-protocol Data frame is used to support multiple protocols simultaneously operating in the piconet. Frame Payload field shall be formatted as shown in Figure 6-59.

Octets: 1	variable	variable	0 or 4
Data ID	Data Header	Data Payload	FCS

Figure 6-59—Non-secure Multi-protocol Data frame Frame Payload format

The Frame Type field shall be set to the Multi-protocol Data frame value in Table 6-1, and the SEC field shall be set to zero. The other fields in the MAC header take on values that are appropriate for that particular data frame. All fields in the MAC header of a Non-secure Multi-protocol Data frame shall be decoded on reception.

The Data ID field indicates the contents of the Data Header field. If the Data Payload field is zero length, i.e., a null data frame, the Data ID field and Data Header field are omitted. Valid values of the Data ID field are given in Table 6-12.

Table 6-12—Data ID values

Data ID field value	Data Header field size	Data Header field contents
0x00	As defined in Figure 6-60	As defined in Figure 6-60
0x01	3 octets	OUI or CID as defined in 6.4.7
0x02	2 octets	An Ethertype
0x03–0xFF	—	Reserved

The Data ID value of 0x00 is used to support Ethertype protocol discrimination (EPD), as defined in IEEE Std 802-2014. If the Data ID field is set to 0x00, then the Data Header field shall be formatted as illustrated in Figure 6-60.

Octets: 6	6	2
Destination address	Source address	Ethertype

Figure 6-60—Data Header field format when Data ID field is 0x00

The Destination Address field and Source Address field contain the destination and source MAC addresses for the data payload. Note that these addresses do not necessarily correspond to the DestID and SrcID of the frame in a bridged network.

The Ethertype field contains an Ethertype, as defined in IEEE Std 802-2014.

NOTE—Additional fields related to the Ethertype will follow the Data Header field in the Data Payload field based on value of the Ethertype, e.g., virtual LAN tags.¹⁰

If the Data ID field is set 0x01, then the format of the remainder of the Data Payload field is determined by the entity that has been assigned the OUI or CID.

If the Data ID field is set to 0x02, then the format of the remainder of the Data Payload field is determined by the value of the Ethertype.

6.3.5.2 Secure Multi-protocol Data frame

The Secure Multi-protocol Data frame Secure Payload field shall be formatted as illustrated in Figure 6-61.

Octets: 2	2	1	variable	variable	8	4
SECID	SFC	Data ID	Data Header	Data Payload	Integrity Code	FCS

Figure 6-61—Secure Multi-protocol Data frame Secure Payload field format

The Frame Type field shall be set to the Multi-protocol Data frame value in Table 6-1, and the SEC field shall be set to one. The other fields in the MAC header take on values that are appropriate for that particular Secure Multi-protocol Data frame. All fields in the MAC header of a Secure Data frame shall be decoded on reception.

The SECID field is defined in 6.2.7.2.

The SFC field is defined in 6.2.7.3.

The Data ID field is defined in 6.3.5.1.

The Data Header field is defined in 6.3.5.1.

The Integrity Code field is defined in 6.2.7.5.

6.3.6 Sync frame

The Sync frame Frame Payload field shall be formatted as illustrated in Figure 6-62.

Octets: 4	4	...	4	4
Synchronization Parameters	CTA Block-1	...	CTA Block- <i>n</i>	FCS

Figure 6-62—Sync frame format

¹⁰Notes in text, tables, and figures are given for information only and do not contain requirements needed to implement the standard.

The Synchronization Parameters field shall be formatted as illustrated in Figure 6-63.

Octets: 2	2
Superframe Duration	Frame Start Time

Figure 6-63—Synchronization Parameters field format

The Superframe Duration field indicates the duration of the current superframe, as described in 6.3.1.1.

The Frame Start Time field indicates the time stamp for the Sync frame, which is the start time of the preamble of a Sync frame, measured from the start of the superframe, in units of microseconds.

The CTA Block field shall be formatted as illustrated in Figure 6-64.

Octets: 2	2
CTA Location	CTA Duration

Figure 6-64—CTA Block field format

The CTA Location field indicates the start time of the allocation, measured from the start of the superframe, as described in 6.3.1.1.

The CTA Duration field specifies the duration of the CTA, as described in 6.3.1.1.

6.4 Information elements (IEs)

The IEs are listed in Table 6-13. The required column indicates the type of DEVs that are required to be able to support the IE.

Table 6-13—IEs

Element ID hex value	Element	Subclause	Present in beacon	Required
0x00	CTA IE	6.4.1	As needed	All DEVs
0x01	BSID IE	6.4.2	In every beacon	All DEVs
0x02	Parent Piconet IE	6.4.3	As needed	Optional
0x03	DEV Association IE	6.4.4	As needed	All DEVs
0x04	PNC shutdown IE	6.4.5	As needed	Optional
0x05	Piconet Parameter Change IE	6.4.6	As needed	All DEVs
0x06	AS IE	6.4.7	As needed	Optional
0x07	Pending channel time map (PCTM) IE	6.4.8	As needed	All DEVs
0x08	PNC Handover IE	6.4.9	As needed	All DEVs

Table 6-13—IEs (continued)

Element ID hex value	Element	Subclause	Present in beacon	Required
0x09	CTA Status IE	6.4.10	As needed	All DEVs
0x0A	Capability IE	6.4.11	Non-beacon IE	All DEVs
0x0B	Transmit Power Parameters IE	6.4.12	Non-beacon IE	All DEVs
0x0C	PS Status IE	6.4.13	As needed	All DEVs
0x0D	CWB IE	6.4.14	As needed	All DEVs
0x0E	Overlapping PNID IE	6.4.15	Non-beacon IE	All DEVs
0x0F	Piconet Services IE	6.4.16	Non-beacon IE	Optional
0x10	Group ID IE	6.4.18	Non-beacon IE	Optional
0x11	Stream Renew IE	6.4.19	Non-beacon IE	All DEVs
0x12	Next PNC IE	6.4.20	As needed	All DEVs
0x13	Piconet Channel Status IE	6.4.21	Non-beacon IE	Optional
0x14	Synchronization IE	6.4.22	As needed	mmWave DEVs
0x15	TSD IE	6.4.23	Non-beacon IE	mmWave DEVs
0x16	UEP Specific IE	6.4.24	Non-beacon IE	Optional
0x17	IFS IE	6.4.25	As needed	SC PHY DEVs
0x18	CTA Relinquish Duration IE	6.4.26	Non-beacon IE	Optional
0x19	Feedback IE	6.4.27	Non-beacon IE	mmWave DEVs
0x1A	Mapping IE	6.4.28	Non-beacon IE	mmWave DEVs
0x1B	BST Clustering IE	6.4.29	Non-beacon IE	mmWave DEVs
0x1C	PET Clustering IE	6.4.30	Non-beacon IE	mmWave DEVs
0x1D	Beam PET IE	6.4.31	Non-beacon IE	mmWave DEVs
0x1E	HRS Beam PET IE	6.4.32	Non-beacon IE	Optional
0x1F	PET Amplitude IE	6.4.33	Non-beacon IE	Optional
0x20	PET Phase IE	6.4.34	Non-beacon IE	Optional
0x21	Sync Frame Frequency IE	6.4.35	Non-beacon IE	Optional
0x22	Directional Peer IE	6.4.36	As needed	Optional
0x23–0x7F	Reserved			
0x80–0xFF	Vendor Defined IE	6.4.17	As needed	Optional

The format of an individual IE is shown in Figure 6-65. Unless otherwise specified, these elements may appear in any order in the frames that are allowed to include more than one of these elements.

Octets: 1	1	variable
Element ID	Length	Content

Figure 6-65—IE format

The Element ID field contains the value from Table 6-13 that identifies the IE.

The Length field contains the length of the Content field in octets.

The Content field contains information specific to the IE.

6.4.1 CTA IE

The CTA IE Content field shall be formatted as illustrated in Figure 6-66. Because the length parameter supports only 255 octets of payload in an IE, the PNC may split the CTA information into more than one CTA IE entry in the beacon. The CTA blocks shall be ordered by increasing value of the CTA location with the highest value being the last.

Octets: 7	7	...	7
CTA Block-1	CTA Block-2	...	CTA Block- <i>n</i>

Figure 6-66—CTA IE Content field format

The CTA blocks shall be formatted as illustrated in Figure 6-67.

Octets: 1	1	1	2	2
DestID	SrcID	Stream Index	CTA Location	CTA Duration

Figure 6-67—CTA block

The DestID indicates the DEV to which the source DEV may send the frames.

The SrcID indicates the DEV to which the channel time is being allocated.

If the CTA is for a child piconet, the DestID and SrcID shall both be the DEVID of the DEV that is the child piconet's PNC.

If the CTA is for a neighbor piconet, the DestID and SrcID shall both be the DEVID assigned by the PNC for the neighbor piconet and shall be one of the reserved neighbor piconet IDs, as described in 6.2.3.

The Stream Index field, as described in 6.2.5, indicates the stream corresponding to the channel time allocation.

The CTA Location field indicates the start time of the allocation in microseconds measured from the start of the beacon, as described in 7.6.

The Duration field specifies the duration of the CTA in microseconds.

The end time of each allocation is the start time contained in the CTA Location field plus the CTA duration.

6.4.2 BSID IE

The BSID IE is used to provide a text string to identify the piconet. The BSID IE Content field shall be formatted as illustrated in Figure 6-68.

Octets: 6–32
Piconet BSID

Figure 6-68—BSID IE Content field format

The Piconet BSID field is a set of ISO/IEC 646:1991 encoded characters that is used to identify the piconet. The setting of the piconet BSID is described in 7.11.3. The BSID field shall not be null terminated. The length of the field is determined from the Length field of the IE.

6.4.3 Parent Piconet IE

The Parent Piconet IE is used to provide a text string to identify the parent piconet and the DEV address of the parent PNC. The Parent Piconet IE Content field shall be formatted as illustrated in Figure 6-69.

Octets: 6	6–32
Parent PNC Address	Parent Piconet BSID

Figure 6-69—Parent Piconet IE Content field format

The Parent PNC Address field contains the DEV address, as described in 6.1, of the parent PNC for the piconet.

The Parent Piconet BSID field contains the piconet BSID from the BSID IE, as described in 6.4.2, in the parent PNC's beacon.

6.4.4 DEV Association IE

The DEV Association IE is used to notify current members in the piconet about one or more DEVs that have either just associated or disassociated from the piconet. The DEV Association IE Content field shall be formatted as illustrated in Figure 6-70.

Octets: 11	11	...	11
DEV-1 Association Info	DEV-2 Association Info	...	DEV- <i>n</i> Association Info

Figure 6-70—DEV Association IE Content field format

The DEV Association Info fields shall be formatted as illustrated in Figure 6-71.

Octets: 6	1	1	3
DEV Address	DEVID	DEV status	DEV Capabilities

Figure 6-71—DEV Association Info fields

The DEV Address field contains the address of the DEV, as described in 6.1, that corresponds to the DEVID.

The DEVID is the identifier assigned by the PNC to a DEV.

The DEV Status field shall be formatted as illustrated in Figure 6-72.

Bits: b0	b1–b7
Association Status	Reserved

Figure 6-72—DEV Status field format

The Association Status field shall be encoded as follows:

- 0 → Disassociated
- 1 → Associated

The DEV Capabilities field is defined in 6.4.11.

6.4.5 PNC Shutdown IE

The PNC Shutdown IE is used to indicate that the PNC is shutting down. The PNC Shutdown IE Content field shall be formatted as illustrated in Figure 6-73.

Octets: 1
Remaining DEVID

Figure 6-73—PNC Shutdown IE Content field format

The Remaining DEVID field indicates which dependent piconet PNC is allowed to continue operation, as described in 7.2.9.1. It shall be set to the PNCID if there are no dependent piconets in the current piconet.

6.4.6 Piconet Parameter Change IE

The Piconet Parameter Change IE Content field shall be formatted as illustrated in Figure 6-74.

Octets: 1	2	1	2	2	6-32
Change Type	Change Beacon Number	New Channel Index	Superframe Timing	PNID	BSID

Figure 6-74—Piconet Parameter Change IE Content field format

The Change Type field shall be set as given in Table 6-14.

The New Channel Index, Superframe Timing, PNID and BSID fields are defined in Table 6-14.

Table 6-14—Description of field contents for change type values

Change Type field value	Interpretation	Field to decode	Description of field contents
0	PNID	PNID	The new PNID, 6.2.2, that will take effect beginning with the superframe that has the beacon number equal to the Change Beacon Number field.
1	BSID	BSID	The new BSID, 6.4.2, that will take effect beginning with the superframe that has the beacon number equal to the Change Beacon Number field.
2	MOVE	Superframe timing	The offset in microseconds between the beacon's expected transmission time and the time that it will be sent by the PNC, 7.11.1. The change occurs with the beacon that has the beacon number equal to the Change Beacon Number field.
3	SIZE	Superframe timing	The new superframe duration, 7.11.2, that will be used for the superframe that has the beacon number equal to the Change Beacon Number field.
4	CHANNEL	New channel index	The channel index of the PHY channel that the piconet will begin using the beacon that has the beacon number equal to the Change Beacon Number field. The mapping of the channel number is PHY dependent. For the 2.4 GHz PHY, the mapping is defined in 10.2.3. For the SC PHY and HSI PHY, the mapping is defined in 11.1.5. For the AV PHY, the mapping is defined in 11.4.1.1.
5-255	Reserved	None	

The Change Beacon Number field is the beacon number of the superframe when the change will take effect. The difference between the beacon number of the beacon which first includes this IE and the Change Beacon Number field is defined to be the *NbrOfChangeBeacons*. For a piconet without pseudo-static CTAs, *NbrOfChangeBeacons* shall be at least two. For a piconet that has pseudo-static CTAs, *NbrOfChangeBeacons* shall be at least *mMaxLostBeacons*. For a piconet that has child or neighbor piconets, *NbrOfChangeBeacons* shall be at least eight. However, a child or neighbor PNC may set the *NbrOfChangeBeacons* to a different number based on the Change Beacon Number field in the parent PNC's beacon, as defined in 7.12.1.

6.4.7 AS IE

The purpose of the application specific (AS) IE is to provide information specific to the application for enhanced operation that is outside of the scope of this standard. The AS IE Content field shall be formatted as illustrated in Figure 6-75.

Octets: 3	variable
Unique ID	Application Specific Data

Figure 6-75—AS IE Content field format

The Unique ID field contains an OUI or CID assigned by the IEEE Registration Authority,¹¹ that refers to an entity that defines the format of the field that follows. If a value of the Unique Id field is not understood by the receiving DEV, that DEV shall ignore the remainder of the associated AS IE.

The Application Specific Data field is provided by the PNC. Its use by an application specific capable DEV is outside of the scope of this standard.

More than one AS IE may be placed in any beacon. The negotiation of the application specific capability between the DEV and the PNC is outside of the scope of this standard.

6.4.8 Pending channel time map (PCTM) IE

The PCTM IE is used to request that a DSPS or APS DEV switch to ACTIVE mode. The PCTM IE Content field shall be formatted as illustrated in Figure 6-76.

Octets: 1	1–32
Start DEVID	PCTM

Figure 6-76—PCTM IE Content field format

The Start DEVID field indicates the DEVID that corresponds to the first bit in the PCTM.

The PCTM field contains a bitmap of 1 to 32 octets in length. Each bit of the PCTM field when set to one indicates the PNC is requesting that the DEV whose DEVID is equal to the start DEVID plus the bit position in the PCTM bitmap listen to the next beacon for a CTA, as described in 7.14.2. The bit position 0, i.e., the first bit or LSB of the bitmap, corresponds to the start DEVID.

The bits corresponding to the PNCID, UnassocID, BcstID, McstID, NbrIDs, and the reserved DEVIDs, as described in 6.2.3, shall be set to zero upon transmission by the PNC and shall be ignored upon reception.

¹¹Interested applicants should contact the IEEE Registration Authority, <http://standards.ieee.org/develop/regauth/>.

6.4.9 PNC Handover IE

The PNC Handover IE is included in the last beacons sent by the old PNC just prior to the old PNC relinquishing control of the piconet. The PNC Handover IE Content field shall be formatted as illustrated in Figure 6-77.

Octets: 6	1	2
New PNC Address	New PNC DEVID	Handover Beacon Number

Figure 6-77—PNC Handover IE Content field format

The New PNC Address field contains the DEV address, 6.1, of the DEV that will be taking over as PNC.

The New PNC DEVID field contains the current DEVID of the DEV that will be taking over as PNC.

The Handover Beacon Number field contains the beacon number of the first beacon that will be sent by the new PNC. The last beacon sent by the old PNC will have a beacon number one less than the Handover Beacon Number field.

6.4.10 CTA Status IE

The CTA Status IE is used by the PNC to inform the DEVs of certain characteristics of a CTA. The CTA Status IE Content field shall be formatted as illustrated in Figure 6-78.

Octets: 1	1	1	1	2	2
DestID	SrcID	Stream Index	Channel Time Request Info	CTA Sub-Rate	Start Beacon Number

Figure 6-78—CTA Status IE Content field format

The DestID field contains the DEVID of the destination for this CTA.

The SrcID field contains the DEVID of the source for this CTA.

The Stream Index field indicates the stream for which the PNC is providing information, as described in 6.5.6.1.

The Channel Time Request Info field shall be formatted as illustrated in Figure 6-79.

Bits: b0–b2	b3	b4	b5	b6	b7
Priority	Terminate	PM Channel Time Request Type	CTA Type	CTA Rate Type	Reserved

Figure 6-79—Channel Time Request Info field format

The Priority field is defined in B.1.

The Terminate field shall be set to one if the stream has been terminated. Otherwise it shall be set to zero.

The PM Channel Time Request Type field is defined in 6.5.6.1.

The CTA Type field is defined in 6.5.6.1.

The CTA Rate Type field is defined in 6.5.6.1.

The CTA Sub-Rate field is set to the number of beacons between every CTA, as described in 6.5.6.1. If one or more CTAs are allocated per superframe, this value shall be set to zero.

The Start Beacon Number field is set to the beacon number, as described in 6.3.1.1, of the first beacon where the CTA of the new or modified stream will first appear.

6.4.11 Capability IE

The Capability IE Content field shall be formatted as illustrated in Figure 6-80.

Octets: 7/13
Overall Capabilities

Figure 6-80—Capability IE format

For the 2.4 GHz PHYs, the Overall Capabilities contains the 2.4 GHz PHY Capabilities field. The 2.4 GHz PHY Capability field shall be formatted as illustrated in Figure 6-81.

Octets: 4	3
PNC Capabilities	2.4 GHz DEV Capabilities

Figure 6-81—2.4 GHz PHY Capability field format

For mmWave PHYs, the Overall Capabilities field contains the mmWave PHY Capabilities field. The mmWave PHY Capability field shall be formatted as illustrated in Figure 6-82.

Octets: 4	6	3
PNC Capabilities	mmWave DEV Capabilities	Beam Forming Capabilities

Figure 6-82—mmWave PHY Capability field format

The PNC Capabilities field shall be formatted as illustrated in Figure 6-83.

Octets: 1	1	1	1
Max Associated DEVs	Max Channel Time Requests	Max TX Power	PNC Rating

Figure 6-83—PNC Capabilities field format

The Max Associated DEVs field indicates the number of associated DEVs this DEV is able to manage if it is PNC-capable and becomes the PNC. Non-PNC-capable DEVs shall set this field to zero.

The Max Channel Time Requests field indicates the number of channel time requests the DEV is capable of handling as a PNC. This field shall be set to zero in a non-PNC-capable DEV.

The Max TX Power Level field indicates the maximum transmit power that is possible for the DEV. The power level is in dBm, encoded in twos complement notation. For example, if a DEV was capable of 14 dBm TX power, the field would take on the value 0x0E; while if the DEV was capable of -4 dBm TX power, the field would take on the value 0xFC.

The PNC Rating field shall be formatted as illustrated in Figure 6-84. Bits b4–b7 are arranged in order of preference for PNC selection, with the highest preference (PNC-capable) corresponding to the MSB.

Bits: b0–b2	b3	b4	b5	b6	b7
Reserved	Next PNC Capable	PSRC	SEC	PNC Des-Mode	PNC Capable

Figure 6-84—PNC Rating field format

The Next PNC Capable field shall be set to one if the DEV is capable of performing the next PNC functionality, as described in 7.2.5. Otherwise, the field shall be set to zero.

The PSRC field shall be set to one if the DEV is receiving power from the alternating current mains and shall be set to zero otherwise.

The SEC field shall be set to one if the DEV is capable of acting as a key originator (see 8.4). Otherwise, the SEC field shall be set to zero.

The PNC Des-Mode field is the desired mode of the DEV. This field shall be set to one if it is desired that the DEV be the PNC of the piconet and the PNC Capable field is set to one. Otherwise, this field shall be set to zero.

The PNC Capable field shall be set to one if the DEV is capable of being a PNC in the piconet. Otherwise, the PNC Capable field shall be set to zero.

The 2.4 GHz DEV Capabilities field shall be formatted as illustrated in Figure 6-85.

Bits: b0	b1	b2	b3	b4	b5	b6	7
Supported Data Rates					Preferred Fragment Size		
Bits: b8	b9	b10	b11	b12	b13	b14	b15
Always AWAKE	Listen to Source	Listen to Multicast	Dly-ACK	Imp-ACK	CTA Relinquish	STP	Reserved
Bits: b16	b17	b18	b19	b20	b21	b22	b23
Reserved							

Figure 6-85—2.4 GHz DEV Capabilities field format

The mmWave DEV Capabilities field shall be formatted as illustrated in Figure 6-86.

The Supported Data Rates or Supported MCSs field is a PHY-dependent mapping that indicates the data rates that the DEV is capable of using. For the 2.4 GHz PHY, this field is the Supported Data Rates field, it shall be formatted as illustrated in Figure 6-85, and the mapping of a field value to a set of data rates is defined in Table 10-21. For the mmWave PHYs, this field is the Supported MCSs field, it shall be formatted as illustrated in Figure 6-86, and the mapping of a field value to a set of MCSs is defined in 11.1.8.1.

The Preferred Fragment Size field is a PHY-dependent mapping that indicates the maximum MAC frame size preferred to be received by the DEV when fragmentation is used. For the 2.4 GHz PHY, this field shall be formatted as illustrated in Figure 6-85, and the mapping of a field value to a preferred fragment size is defined in Table 10-22. For the mmWave PHYs, this field shall be formatted as illustrated in Figure 6-86, and the mapping of a field value to a preferred fragment size is defined in 11.1.8.2.

The Always AWAKE field shall be set to one to indicate the DEV is in ACTIVE mode and that it will listen to all CTAs, regardless of the DestID or SrcID. Otherwise, the field shall be set to zero.

The Listen to Source field shall be set to one to indicate the DEV is in ACTIVE mode and that it will listen to all CTAs where the SrcID is equal to the DEVID of a DEV that is currently the source of a stream to that DEV regardless of the DestID of those CTAs. Otherwise, the field shall be set to zero.

The Listen to Multicast field shall be set to one to indicate the DEV is in ACTIVE mode and that it will listen to all multicast CTAs regardless of the SrcID or the Stream Index. Otherwise, the field shall be set to zero.

The Dly-ACK field shall be set to one if the DEV is capable of performing the Dly-ACK procedure, as defined in 7.9.3. Otherwise, the field shall be set to zero.

The Imp-ACK field shall be set to one indicate that the DEV is capable of performing the Imp-ACK procedure, as defined in 7.9.6. Otherwise, the field shall be set to zero.

The CTA Relinquish field shall be set to one if the DEV is capable of participating in the CTA relinquish procedure, as defined in 7.4.3.8. Otherwise, the field shall be set to zero.

Bits: b0	b1	b2	b3	b4	b5	b6	b7
Supported MCSs							
Bits: b8	b9	b10	b11	b12	b13	b14	b15
Preferred Fragment Size			Always AWAKE	Listen to Source	Listen to Multicast	Dly-ACK	Imp-ACK
Bits: b16	b17	b18	b19	b20	b21	b22	b23
CTA Relinquish	STP	SC Capable	HSI Capable	AV Capable	OOK Capable	DAMI Capable	Blk-ACK
Bits: b24	b25	b26	b27	b28	b29	b30	b31
UEP Capable		UEP Type			HRP RX Capable	HRP TX Capable	TSD Support
Bits: b32	b33	b34	b35	b36	b37	b38	b39
Supported IFS				Supported Aggregation		Pilot Word Capability	PCES Capability
Bits: b40	b41	b42	b43	b44	b45	b46	b47
Sync Frame Capable	Reserved						

Figure 6-86—mmWave DEV Capabilities field format for mmWave PHYs

The STP field shall be set to one to indicate that the DEV is capable of renewing its streams within the stream timeout period, as defined in 7.3.4. Otherwise, the field shall be set to zero.

The values of the fields in the PNC Capabilities field and the DEV Capabilities field shall not change while a DEV is associated in a piconet.

For the mmWave PHYs, the SC Capable field shall be set to one if the DEV supports the SC PHY, as defined in 11.2. It shall be set to zero otherwise.

The HSI Capable field shall be set to one if the DEV supports the HSI PHY, as defined in 11.3, and shall be set to zero otherwise.

The AV Capable field shall be set to one if the DEV supports the AV PHY, as defined in 11.4, and shall be set to zero otherwise.

The OOK Capable field shall be set to one if the DEV supports the OOK mode, as defined in Annex G, and shall be set to zero otherwise.

The DAMI Capable field shall be set to one if the DEV supports the DAMI mode, as defined in Annex G, and shall be set to zero otherwise.

The Blk-ACK field shall be set to one if the DEV is capable of performing Blk-ACK, as defined in 7.9.5. Otherwise, the field shall be set to zero.

The UEP Capable field indicates if the DEV supports UEP in the PHY. The valid values of the UEP Capable field are given in Table 6-15.

Table 6-15—UEP Capable field values

Value	UEP support
0	No UEP
1	SC UEP
2	HSI UEP
3	AV UEP

The UEP Type field indicates the type of UEP that is supported based on the value of the UEP Capable field. The valid values of the UEP Type field are given in Table 6-16.

Table 6-16—UEP Type field values

Value	PHY UEP type of mmWave PHYs
0	No UEP support
1	UEP type 1 using different FECs (SC PHY)
2	UEP type 2 using different MCSs (SC PHY)
3	UEP type 3 using different MCSs (All PHYs)
4	UEP type 3 using skewed constellation (All PHYs)
5	UEP type 3 using different MCSs and skewed constellation (SC and HSI PHYs)
6	All UEP types (SC PHY)
7	Reserved

The HRP RX Capable field shall be set to one if the DEV supports HRP RX for the AV PHY, as defined in 11.4, and shall be set to zero otherwise.

The HRP TX Capable field shall be set to one if the DEV supports HRP TX for the AV PHY, as defined in 11.4, and shall be set to zero otherwise.

The TSD Support field shall be set to one if the DEV is capable of TSD, as defined in 12.8; it shall be set to zero otherwise.

The Supported Interframe Spacing (IFS) field indicates the minimum value of the IFS supported by a DEV. The valid values of the field are given in Table 6-17.

Table 6-17—Supported IFS encoding

Field value	SIFS and MIFS duration
0	0.2 μ s
1	0.4 μ s
2	0.6 μ s
3	0.8 μ s
4	1.0 μ s
5	2.0 μ s
6	2.5 μ s
7–15	Reserved

The minimum allowed IFS for either SIFS or MIFS may also be constrained by the PHY mode in use. A source DEV may use the shortest SIFS or MIFS supported by the destination DEV. The SIFS and MIFS used may be different.

The Supported Aggregation field indicates the type of aggregation that is supported by the DEV. The valid values for the Supported Aggregation field are as follows:

- 0 → No aggregation support
- 1 → Standard aggregation support
- 2 → Low-latency aggregation support
- 3 → Standard and low-latency aggregation support

The Pilot Word Capability field shall be set to one if the DEV supports a pilot word of length 8. It shall be set to zero otherwise.

The PCES Capability field shall be set to one if the DEV supports the use of PCES and shall be set to zero otherwise.

The Sync Frame Capable field shall be set to one if the DEV supports Sync frame transmission, as defined in 7.18, and shall be set to zero otherwise.

The Beam Forming Capabilities field shall be formatted as illustrated in Figure 6-87.

Bits: b0–b1	b2–b5	b6–b9	b10–b12	b13	b14–b15	b16–b17	b18–b23
LQI Type	Number TX Sectors	Number RX Sectors	Antenna Type	PET	Number TX Quasi-Omni Directions	Number RX Quasi-Omni Directions	Reserved

Figure 6-87—Beam Forming Capabilities field format

The LQI Type field indicate the type of LQI used in beam forming procedure. Valid values of the LQI type fields are as follows:

- 0 → RSSIR

- 1 → SNR
- 2 → SINR
- 3 → Reserved

The Number TX Sectors field indicates the number of TX sectors supported by the DEV.

The Number RX Sectors field indicates the number of RX sectors supported by the DEV.

The Antenna Type field indicate the supported antenna type of the DEV. Valid values of the Antenna Type field are as follows:

- 0 → No beam forming capability
- 1 → Beam forming antenna capable
- 2–7 → Reserved

The PET field shall be set to one if the DEV supports pattern estimation and tracking and shall be set to zero otherwise.

The Number TX Quasi-omni Directions field indicates the number of TX quasi-omni directions supported by the DEV.

The Number RX Quasi-omni Directions field indicates the number of RX quasi-omni directions supported by the DEV.

6.4.12 Transmit Power Parameters IE

The Transmit Power Parameters IE is used to communicate the transmit power control (TPC) capabilities of a DEV. The Transmit Power Parameters IE Content field shall be formatted as illustrated in Figure 6-88.

Octets: 1	1	1
TX Power Levels	TX Power Step Size	Current TX Power

Figure 6-88—Transmit Power Parameters IE Content field format

The TX Power Levels field indicates the number of levels supported by a DEV.

The TX Step Size field indicates the TX power level step size in 1 dB resolution, e.g., a number 4 in this field means that the DEV has nominally 4 dB steps.

If a DEV does not support TPC, it shall set the TX Power Levels and TX Power Step Size fields to zero.

The Current TX Power field is the DEV's estimate of its transmitter power measured at the antenna interface. The value is in dBm encoded in twos complement format. For example, a +2 dBm TX power level is encoded as 0x02 while a –2 dBm TX power level is encoded as 0xFE.

6.4.13 PS Status IE

The PS Status IE Content field shall be formatted as illustrated in Figure 6-89.

Octets: 2	1	1	1–32
Next Wake Beacon	PS Set Index	Start DEVID	DEVID Bitmap

Figure 6-89—PS Status IE Content field format

The Next Wake Beacon field shall be set to the beacon number, as described in 6.3.1.1, of the next wake beacon for the set indicated in the PS Set Index field, as described in 7.14.1. It shall be set to zero when the PS Set Index field is zero, i.e., for the set of DEVs in APS mode.

The PS Set Index field is set to the index of the power save set as follows:

- 0 → APS
- 1 → PSPS
- 2–253 → DSPS

The Start DEVID field indicates the DEVID that corresponds to the first bit in the DEVID bitmap.

The DEVID Bitmap field is 1 to 32 octets in length. Each bit of the DEVID Bitmap field corresponds to a DEVID that is equal to the start DEVID plus the bit position in the bitmap. The bit position zero, i.e., the first bit or LSB of the bitmap, corresponds to the start DEVID. The bit corresponding to a DEVID shall be set to one when a DEV that is a member of this PS set is in a power save mode. It shall be set to zero otherwise.

The bits corresponding to the PNCID, UnassocID, BcstID, McstID, NbrIDs, and the reserved IDs, as described in 6.2.3, shall be set to zero upon transmission by the PNC and shall be ignored upon reception.

6.4.14 CWB IE

The continued wake beacon (CWB) IE Content field shall be formatted as illustrated in Figure 6-90.

Octets: 1	1–32
Start DEVID	DEVID bitmap

Figure 6-90—CWB IE Content field format

The Start DEVID field indicates the DEVID that corresponds to the first bit in the DEVID bitmap.

The DEVID Bitmap field is 1 to 32 octets in length. Each bit of the DEVID bitmap corresponds to a DEVID that is equal to the start DEVID plus the bit position in the bitmap. The bit position zero, i.e., the first bit or LSB of the bitmap, corresponds to the start DEVID. The bit corresponding to a DEVID shall be set to one when the PNC is requesting that the DEV listen to the next beacon for a CTA block, as described in 7.14.2.3. It shall be set to zero otherwise.

The bits corresponding to the PNCID, UnassocID, BcstID, McstID, NbrIDs, and the reserved DEVIDs, as described in 6.2.3, shall be set to zero upon transmission by the PNC and shall be ignored upon reception.

6.4.15 Overlapping PNID IE

The Overlapping PNID IE is used to communicate the PNIDs that a DEV has detected either in its channel or in other channels. The Overlapping PNID IE Content field shall be formatted as illustrated in Figure 6-91.

Octets: 2	1	...	2	1
PNID 1	Channel Index 1	...	PNID n	Channel Index n

Figure 6-91—Overlapping PNID IE Content field format

The PNID field contains the PNID from a frame that a DEV has received since the last time this IE was sent by the DEV to the PNC.

If the DEV has received a beacon from a different piconet on the current channel with the same PNID but a different PNC address, it will add that PNID and channel index to this IE. Otherwise, the IE will not contain the same PNID/channel index pair as for the current piconet. Thus a DEV will report piconets with the same PNID in other channels but will not erroneously report frames from the current piconet as being from an overlapping piconet.

The Channel Index field contains the channel on which the PNID was found.

6.4.16 Piconet Services IE

The Piconet Services IE is used to provide information about the application layer capabilities of an individual DEV. The Piconet Services IE Content field shall be formatted as illustrated in Figure 6-92.

Octets: 0/1	0/1/3	0–127
DEVID	Services ID	Piconet Services

Figure 6-92—Piconet Services IE Content field format

The DEVID field identifies the DEV corresponding to the Piconet Services field. If the PNC is sending the IE as the aggregate capabilities of the piconet, the DEVID field shall be set to the BcstID.

If the value of the DEVID field is the PNCID, then the Piconet Services field is not present and the Services ID field is 1 octet in length. In this case, the Services ID field contains the following information:

- 0 → Reserved
- 1 → Broadcast of Piconet Services IE not allowed
- 2 → Piconet Services IE not supported
- 3–254 → Reserved
- 255 → Other failure

If the value of the DEVID field is not the PNCID, then the Services ID field contains a Unique ID field, as defined in 6.4.7. In this case, the format of the Piconet Services field is defined by the entity indicated in the Unique ID field.

The Piconet Services field is used to indicate the application layer capabilities of the DEV indicated by the DEVID. The content of the Piconet Services field is outside of the scope of this standard.

6.4.17 Vendor Defined IE

The Vendor Defined IE Content field shall be formatted as illustrated in Figure 6-93.

Octets: 3	variable
Unique ID	Vendor Defined information

Figure 6-93—Vendor Defined IE Content field format

The Unique ID field is defined in 6.4.7.

The format of the Vendor Defined Information field is defined by the entity identified in the Unique ID field. Its use by a DEV is outside of the scope of this standard.

6.4.18 Group ID IE

The Group ID IE is used to list the DEVs that are members of a multicast group. The Group ID IE Content field shall be formatted as illustrated in Figure 6-94.

Octets: 6	1	1	1–32
Group Address	GrpID	Start DEVID	Group IDs

Figure 6-94—Group ID IE Content field format

The Group Address field is a 48-bit MAC address with the I/G bit set to indicate a group MAC address, as defined in IEEE Std 802-2014.

The GrpID field contains the DEVID that has been assigned by the PNC for the address in the Group Address field.

The Start DEVID field indicates the DEVID that corresponds to the first bit in the Group IDs field.

The Group IDs field contains a bitmap of 1 to 32 octets in length. Each bit of the Group IDs field when set to one indicates the DEV whose DEVID is equal to the start DEVID plus the bit position in the Group ID bitmap is a member of the multicast group identified by the Group Address field and GrpID field. The bits in the Group IDs field are set to zero otherwise. The bit position 0, i.e., the first bit or LSB of the bitmap, corresponds to the start DEVID.

The bits corresponding to the PNCID, UnassocID, BcstID, McstID, NbrIDs, and the reserved DEVIDs, 6.2.3, shall be set to zero upon transmission by the PNC and shall be ignored upon reception.

6.4.19 Stream Renew IE

The Stream Renew IE is used by a DEV to renew the STP for streams for which it is the source. The Stream Renew IE Content field shall be formatted as illustrated in Figure 6-76.

Octets: 1	1	...	1
Stream Index 1	Stream Index 2	...	Stream Index n

Figure 6-95—Stream Renew IE Content field format

The Stream Index fields contain the stream indices that the DEV wants to maintain.

6.4.20 Next PNC IE

The Next PNC IE is used to inform the members of the piconet which DEV will become the PNC after an implicit handover, as described in 7.2.5. The Next PNC IE Content field shall be formatted as illustrated in Figure 6-96.

Octets: 1	2
Next PNC	Next PNID

Figure 6-96—Next PNC IE Content field format

The Next PNC field contains the DEVID of the DEV that the PNC has chosen to be the Next PNC.

The Next PNID field shall contain the PNID to be used by the next PNC if it takes over as PNC, as described in 7.2.5.

6.4.21 Piconet Channel Status IE

The Piconet Channel Status IE is used to report those DEVs from which the source DEV has correctly received a MAC header. The Piconet Channel Status IE Content field shall be formatted as illustrated in Figure 6-97.

Octets: 32	1
Bitmap of DEVs in Range	Channel Status Details

Figure 6-97—Piconet Channel Status IE Content field format

The Bitmap of DEVs in Range field is a listing of all of the DEVs in the piconet that were heard by the DEV sending the command. The bit position zero, i.e., the first bit or LSB of the bitmap, corresponds to DEVID zero. The bit corresponding to a DEVID is set to one if a MAC header was correctly received from a DEV with that DEVID since the last Piconet Channel Status IE was sent. Otherwise, the bit shall be set to zero.

The Channels Status Details field shall be set to one if the DEV scanned for power save DEVs. Otherwise, this value shall be set to zero.

6.4.22 Synchronization IE

The Synchronization IE Content field shall be formatted as illustrated in Figure 6-98. The Synchronization IE shall be supported by DEVs that support the SC PHY or HSI PHY.

Octets: 4	1	5	5
Quasi-omni Beacon Info	Section Indications	Association S-CAP Info	Regular S-CAP Info

Figure 6-98—Synchronization IE Content field format

The Quasi-omni Beacon Info field shall be formatted as illustrated in Figure 6-99.

Bits: b0–b15	b16–b18	b19–b21	b22–b25	b26–b31
Beacon Offset Time	Beacon Index	Number Beacon Frames	PNC Directional Antenna Capabilities	Reserved

Figure 6-99—Quasi-omni Beacon Info field format

The Beacon Offset Time field contains the time in microseconds that is the start of this Beacon frame delayed from the start of the superframe.

The Beacon Index field indicates the index of the current Beacon frame.

The Number Beacon Frames field indicates the number of Beacon frames that will be sent as part of the quasi-omni beacon.

The PNC Directional Antenna Capabilities field shall be formatted as illustrated in Figure 6-100.

Bits: b0	b1	b2	b3
Antenna Symmetry	PNC Beam Forming Capable	Sectors Only	BST Support

Figure 6-100—PNC Directional Capabilities field format

The Antenna Symmetry field shall be set to one if the PNC has a symmetric antenna system (SAS) and shall be set to zero otherwise, i.e., the PNC has an asymmetric antenna system (AAS).

The PNC Beam Forming Capable field shall be set to one if the PNC is capable of beam forming and shall be set to zero otherwise.

The Sectors Only field shall be set to one if the PNC supports only sectorized antenna and shall be set to zero otherwise.

The BST Support field shall be set to one if the PNC supports beam switching/steering and tracking, as described in Clause 12, and shall be set to zero otherwise.

The Section Indication field shall be formatted as illustrated in Figure 6-101.

Bits: b0	b1	b2	b3	b4–b7
PNC Quasi-omni Tracking Present	Sector Training Present	Association S-CAP Present	Regular S-CAP Present	Reserved

Figure 6-101—Section Indication field format

The PNC Quasi-omni Tracking Present field shall be set to one if the PNC quasi-omni tracking section exists in the beacon and shall be set to zero otherwise.

The Sector Training Present field shall be set to one if the sector training section exists in the beacon and shall be set to zero otherwise.

The Association S-CAP Present field shall be set to one if association S-CAP sections exist in the CAP and shall be set to zero otherwise.

The Regular S-CAP Present field shall be set to one if regular S-CAP sections exist in the CAP and shall be set to zero otherwise.

The Association S-CAP info field shall be formatted as illustrated in Figure 6-102.

Bits: b0–b15	b16–b31	b32–b33	b34–b35	b36–b37	b38–b39
Association S-CAP Start Time	S-CAP Duration	Total Number Association S-CAPs	Number Association S-CAPs in Superframe	First PNC Quasi-omni RX Index	Reserved

Figure 6-102—Association S-CAP Info field format

The Association S-CAP Start Time field specifies the start time offset of the association S-CAPs from the start of the superframe in microseconds.

The S-CAP Duration field specifies the duration of the association S-CAP, as described in 7.6.5.2, in microseconds.

The Total Number Association S-CAPs is the same as the total number of PNC quasi-omni RX directions.

The Number Association S-CAPs in Superframe field specifies the number of association S-CAPs in the current superframe.

The First PNC Quasi-omni RX Index field specifies the PNC's quasi-omni RX index that will be used in the first association S-CAP in the current superframe.

The Regular S-CAP info field shall be formatted as illustrated in Figure 6-103.

Bits: b0–b15	b16–b31	b32–b33	b34–b35	b36–b37	b38–b39
Regular S-CAP Start Time	S-CAP Duration	Total Number Regular S-CAPs	Number Regular S-CAPs in Superframe	First PNC Quasi-omni RX Index	Reserved

Figure 6-103—Regular S-CAP Info field format

The Regular S-CAP Start Time field specifies the start time offset of the regular S-CAP from the start of the superframe in microseconds.

The S-CAP Duration field specifies the duration of the regular S-CAPs in microseconds.

The Total Number Regular S-CAPs field is set to be the same as the total number of PNC quasi-omni RX directions.

The Number Regular S-CAPs field in Superframe field specifies the number of regular S-CAPs in the current superframe.

The First PNC Quasi-omni RX Index field specifies the PNC's quasi-omni RX index that will be used in the first regular S-CAP in the current superframe.

6.4.23 TSD IE

The TSD IE is used for information exchange in TSD procedure as described 12.8. The TSD IE Content field shall be formatted as illustrated in Figure 6-104.

Octets: 1	1	1	1
Mode	Number Of Transmit Directions	TSD Feedback Period	Transmit Direction Index

Figure 6-104—TSD IE Content field format

The Mode field is encoded as indicated in Table 6-18.

Table 6-18—Mode field encoding

Field value	Mode	Description
0	Announce	Announce TSD IE sent from PNC to DEV
1	Request to switch	DEV requests PNC to switch to the next transmit direction
2	Request to stay	DEV requests PNC to stay at the current transmit direction
3	Response	In response of a request
4–255	Reserved	

The Number Of Transmit Directions field indicates the number of TX directions supported by the PNC.

The TSD Feedback Period field indicates the time interval between the TSD IEs sent from the DEV to the PNC for feedback.

The Transmit Direction Index field indicates the index of the transmit direction of the PNC. A field value of zero indicates that the PNC shall select one of the transmit directions that was not used previously.

6.4.24 UEP Specific IE

The UEP Specific IE is used to indicate the separating position of MSB and LSB data. For example, it is used to handle different components of a video signal, e.g., RGB, YCbCr. This IE shall be only used in UEP type 2. The UEP Specific IE Content field shall be formatted as illustrated in Figure 6-105.

Octets: 4
MSB/LSB Separation

Figure 6-105—UEP Specific IE Content field format

The MSB/LSB Separation field shall be formatted as illustrated in Figure 6-106.

Octets: 1	1	1	1
Component 1	Component 2	Component 3	Reserved

Figure 6-106—MSB/LSB Separation field format

The Component fields contain the information to indicate the separating position of MSB and LSB data. Valid values of the Component 1, 2, and 3 fields are as follows:

- 0 → All bits are MSB
- 1 → b_0 is LSB, b_1 – b_7 are MSB
- 2 → b_0 – b_1 is LSB, b_2 – b_7 are MSB
- 3 → b_0 – b_2 is LSB, b_3 – b_7 are MSB
- 4 → b_0 – b_3 is LSB, b_4 – b_7 are MSB
- 5 → b_0 – b_4 is LSB, b_5 – b_7 are MSB
- 6 → b_0 – b_5 is LSB, b_6 – b_7 are MSB
- 7 → b_0 – b_6 is LSB, b_7 is MSB
- 8 → All bits are LSB
- 9–255 → Reserved

6.4.25 IFS IE

The IFS IE Content field shall be formatted as illustrated in Figure 6-107.

Octets: 1
IFS Value

Figure 6-107—IFS IE Content field format

The IFS Value field shall be formatted as illustrated in Figure 6-108.

Bits: b0–b3	b4–b7
SIFS	MIFS

Figure 6-108—IFS value field format

The mapping of a field value to a specific IFS value is defined in Table 6-17.

6.4.26 CTA Relinquish Duration IE

The CTA Relinquish Duration IE Content field shall be formatted as illustrated in Figure 6-109.

Octets: 2
CTA Relinquish Duration

Figure 6-109—CTA Relinquish Duration IE Content field format

The CTA Relinquish Duration field contains the time in microseconds for which the time in the CTA is relinquished.

6.4.27 Feedback IE

The Feedback IE shall be formatted as illustrated in Figure 6-110.

Bits: b0–b3	b4–b9	b10–b13	b14–b19	b20	b21–b23	b24–b27	b28–b31
Best Sector	Best Sector LQI	Second Best Sector	Second Best Sector LQI	End Of Sector Training	Reserved	Total Number Of Feedbacks	Index of Current Feedback

Figure 6-110—Feedback IE Content field format

The Best Sector field shall be set to the index of the best transmit sector for the originating DEV. In the SAS case, this also represents the index of the best receiver sectors for the target DEV.

The Best Sector LQI field shall be set to the value of LQI measured by the Target DEV for the best sector. The Best Sector LQI field is encoded in dB. The corresponding LQI values in the case of SNR/SINR and RSSIR are given in Table 6-19 and Table 6-20. The LQI can be measured using the CES field in the preamble.

The Second Best Sector field shall be set to the index of the second best transmit sector for the originating DEV. In the SAS case, this also represents the index of the second best receiver sectors for the target DEV. If the second best sector is not supported, then this field shall be set to zero.

The Second Best Sector LQI field shall be set to the value of LQI measured by the Target DEV for the second best sector. The Second Best Sector LQI is encoded in the same manner as the Best Sector LQI. If the second best sector is not supported, then this field shall be set to zero.

The encoding of LQI for SNR/SINR is defined in Table 6-19. The SNR value is equal to E_c/N_0 (dB) + 20 dB.

Table 6-19—SNR/SINR value for LQI field

Field value	SNR/SINR value
0b000000	≤ 0 dB
0b000001	0.5 dB
0b000010	1 dB
...	...
0b111110	31 dB
0b111111	> 31.5 dB

The encoding of LQI for RSSIR is defined in Table 6-20 where 0 dB indicates that the received signal level is equal to receiver sensitivity for the respective base rate.

Table 6-20—RSSIR value for LQI field

Field value	RSSIR value
0b000000	≤ 0 dB
0b000001	1 dB
0b000010	2 dB
...	...
0b111110	62 dB
0b111111	> 63 dB

The End Of Sector Training field is used to indicate that a DEV wants to end the training stage. Any DEV who wants to terminate the training stage at the end of sector-level training may set this field to one during sector-level training period. However, if one of DEVs, DEV1, sets the field to one but the other DEV, DEV2, does not, then DEV1 shall help DEV2 to finish beam-level training by using DEV1's best transmit and receive sectors.

The Total Number Of Feedbacks field shall be set to the number of Announce commands that will be sent as part of the feedback process.

The Index Of Current Feedback field shall be set to number of the current Feedback IE in the series that is being sent. The first Feedback IE that is sent shall have the Index of the Current Feedback field set to zero.

6.4.28 Mapping IE

The Mapping IE Content field shall be formatted as illustrated in Figure 6-111.

Bits: b0–b5	b6–b7	b8–b13	b14–b17	b18–b21	b22–b23
Number TX Beams	Sync Mode	Number Of RX Beams	Best Cluster Tracking Period	Second Best Cluster Tracking Period	Reserved

Figure 6-111—Mapping IE Content field format

When sent at the end of sector-level training, the Number of TX Beams field and Number of RX Beams field shall be set to one less than the number of beams that the sending device will be using during the beam-level training. When sent at the end of the beam-level training, these fields shall be set to the number of high resolution (HRS) beams to be used during tracking.

If the system is SAS, the Number of RX Beams field shall be set to zero.

The Sync Mode field shall encode the length of the sync sequence to be used for subsequent training sequences. The value of the field shall be the same as that used in the PHY header, as described in 11.2.3.1.1 and 11.3.3.2.

The Best Cluster Tracking Period and Second Best Cluster Tracking Period fields shall be set to the tracking period of the best cluster and the second best cluster, respectively. The Tracking Period fields shall be encoded as defined in Table 6-21.

Table 6-21—Tracking period encoding

Field value	Tracking period value
0b0000	Reserved
0b0001	0.004 ms
0b0010	0.008 ms
0b0011	0.016 ms
0b0100	0.032 ms
0b0101	0.064 ms
0b0110	0.128 ms
0b0111	0.256 ms
0b1000	0.512 ms
0b1001	1.024 ms
0b1010	2.048 ms
0b1011	4.096 ms

Table 6-21—Tracking period encoding (continued)

Field value	Tracking period value
0b1100	8.192 ms
0b1101	16.384 ms
0b1110	32.768 ms
0b1111	Reserved

6.4.29 BST Clustering IE

The BST Clustering IE is used for AAS sector- and beam-level training, as defined in 12.5.1.1.1 and 12.5.1.2.1, respectively. It is also used for SAS sector- and beam-level training, as defined in 12.5.1.1.2 and 12.5.1.2.2, respectively. The BST Clustering IE Content field shall be formatted as illustrated in Figure 6-112.

Octets: variable	variable
BST TX Cluster Mapping	BST RX Cluster Mapping

Figure 6-112—BST Clustering IE Content field format

The BST TX Cluster Mapping field and BST RX Cluster mapping field shall be formatted as illustrated in Figure 6-113. In the case of a SAS system, the TX and RX clusters are identical, and only one BST Cluster Mapping field shall be included in the IE.

Bits: b0–b3	b4–b7	b8–b13	...	b(4+N*4)–b(7+N*4)	b(8+N*4)–b(13+N*4)/0
Number Of Clusters	Number Of Beams In Cluster 0	Number Of Beams In Cluster 1	...	Number Of Beams In Cluster N	Reserved

Figure 6-113—BST TX Cluster Mapping field and BST RX Cluster Mapping field format

The Number Of Clusters field contains one less than the number of TX or RX clusters, C_T and C_R , respectively, as defined in 12.2.4.

The Number Of Beams in Cluster field shall be set to one less than the number of beams that make up that cluster.

The Reserved field shall be included if there is an even number of clusters and shall be omitted otherwise.

6.4.30 PET Clustering IE

The PET Clustering IE is used for AAS sector- and beam-level training, as defined in 12.5.1.1.1 and 12.5.1.2.1, respectively. It is also used for SAS sector- and beam-level training, as defined in 12.5.1.1.2 and 12.5.1.2.2, respectively. The PET Clustering IE Content field shall be formatted as illustrated in Figure 6-114.

Octets: variable	variable
PET TX Cluster Mapping	PET RX Cluster Mapping

Figure 6-114—PET Clustering IE Content field format

The PET TX Cluster Mapping field and PET RX Cluster Mapping field shall be formatted as illustrated in Figure 6-115. In the case of a SAS system, the transmit and receive clusters are identical, and only one Cluster Mapping field shall be included in the IE.

Octets: 1	2	2	...	2
Number Of Clusters	Cluster 0 Descriptor	Cluster 1 Descriptor	...	Cluster <i>N</i> Descriptor

Figure 6-115—PET TX Cluster Mapping field and PET RX Cluster Mapping field format

The Number of Clusters field shall be set to one less than the number of TX or RX clusters, C_T and C_R , respectively, as defined in 12.2.4.

The Cluster Descriptor fields shall be for the TX or RX arrays as identified by the position of the field in the PET Clustering IE. The Cluster Descriptor field shall be formatted as illustrated in Figure 6-116.

Octets: 1	1
Center Beam Index	Cluster Encoding

Figure 6-116—Cluster Descriptor field format

The Cluster Encoding field describes the geometry of the cluster and shall be as described in 12.2.4.

The Center Beam Index field shall be the index of the beam or HRS beam around which the cluster is formed.

6.4.31 Beam PET IE

The Beam PET IE Content field shall be formatted as illustrated in Figure 6-117.

Octets: 4	4
PET TX Configuration	PET RX Configuration

Figure 6-117—Beam PET IE Content field format

For the SAS case, the transmit and receive PET configurations are the same, so the PET RX Configuration field shall be omitted.

The PET TX Configuration field and PET RX Configuration field shall be formatted as shown in Figure 6-118.

Bits: b0–b3	b4–b7	b8–b15	b16–b23	b24–b26	b27–b29	b30–b31
Number Of Antennas, z-axis	Number Of Antennas, x-axis	Beam Codebook ID, z-axis	Beam Codebook ID, x-axis	Amplitude Resolution	Phase Resolution	Reserved

Figure 6-118—PET Configuration field format

The Number of Antennas, z-axis and Number of Antennas, x-axis fields shall be set to one less than the number of antennas along the z-axis and x-axis, respectively, as described in 12.2.3. The z-axis and x-axis are only applied here as an example to define an antenna pattern for a linear two-dimensional (2-D) antenna array, as described in 12.2.3. The reference to z-axis and x-axis are changeable according to the implementation requirement. These values shall be for the TX or RX arrays as identified by the position of the field in the PET Clustering IE.

The Beam Codebook ID fields shall identify the codebooks, as described in 12.3.1, to be used for the respective axes.

The Amplitude Resolution field shall indicate the number of discrete values for amplitude that can result from pattern estimation. The field shall be coded as N , where there are 2^N possible resulting amplitude values. The value of this field shall be limited to the range 0 to 4.

Similarly, the Phase Resolution field shall indicate the number of discrete phase values that can result from pattern estimation. This field shall also be coded as N , where there are 2^N possible resulting phase values. The value of this field shall be limited to the range 0 to 4.

6.4.32 HRS Beam PET IE

The HRS Beam PET IE Content field shall be formatted as illustrated in Figure 6-119.

Octets: 2	2
PET TX HRS Configuration	PET RX HRS Configuration

Figure 6-119—HRS Beam PET IE Content field format

For the SAS case, the PET TX HRS Configuration field and PET RX HRS Configuration field are the same, so the PET RX HRS Configuration field shall be omitted.

The PET TX HRS Configuration field and PET TX HRS Configuration field shall be formatted as illustrated in Figure 6-120.

Octets: 1	1
HRS Beam Codebook ID, z-axis	HRS Beam Codebook ID, x-axis

Figure 6-120—PET HRS Configuration field format

The HRS Beam Codebook ID fields shall identify the codebooks, as described in 12.3.1, to be used for their respective axes.

6.4.33 PET Amplitude IE

The PET Amplitude IE Content field shall be formatted as illustrated in Figure 6-121.

Bits: b0–b3	b4–b7	...	b(4+M*4)–b(7+M*4)	b(8+M*4)–b(13+M*4)/0
Element 0 Amplitude	Element 1 Amplitude	...	Element M Amplitude	Reserved

Figure 6-121—PET Amplitude IE Content field format

The Element Amplitude fields are sent in the predefined antenna element order, as defined in 12.2.3. If F represents the amplitude relative to the element with the highest amplitude, the value of the Element Amplitude field shall be the one less than the numerator, N , of a fraction $N/D = F$, where D is the number of possible values as specified by the Amplitude Resolution field of the Beam PET IE. For example, if the Amplitude Resolution field value is 3, then there are $2^3 = 8$ possible values and $D = 8$. In this case, one of the antenna element amplitudes will be $N = 8$, and the others will have values in the range of $N = 1$ to $N = 8$. For this case, one Element Amplitude value shall be seven while the others will have values in the range zero to seven.

The Reserved field shall be present if there is an odd number of antenna elements; it shall be omitted otherwise.

6.4.34 PET Phase IE

The PET Phase IE Content field shall be formatted as illustrated in Figure 6-122.

Bits: b0–b3	b4–b7	...	b(4+M*4)–b(7+M*4)	b(8+M*4)–b(13+M*4)/0
Element 0 Phase	Element 1 Phase	...	Element M Phase	Reserved

Figure 6-122—PET Phase IE Content field format

The Element Phase fields are sent in the predefined antenna element order, as defined in 12.2.3. If F represents the phase as a fraction of a full circle, the Element Phase field shall set to the numerator, N , of a fraction $N/D = F$, where D is the number of possible values as specified by the Phase Resolution field of the Beam PET IE. For example, if the phase resolution field value is 4, then there are $2^4 = 16$ possible values and $D = 16$. In this case, each LSB of the phase value is equal to 22.5° , and the values shall be in the range of $N = 0$ to $N = 15$, which corresponds to phases of 0° to 342.5° .

The Reserved field shall be present if there is an odd number of antenna elements; it shall be omitted otherwise.

6.4.35 Sync Frame Frequency IE

The Sync Frame Frequency IE Content field shall be formatted as illustrated in Figure 6-123.

Bits: b0–b4	b5	b6–b7
Sync Frame Frequency	Sync Frame Direction	Reserved

Figure 6-123—Sync Frame Frequency IE Content field format

The Sync Frame Frequency field is set to the number of superframes between transmission of two Sync frames, as defined in 6.3.6, requested by the PNC. If the Sync Frame Frequency field is set to zero, then the PNC is requesting the DEV cease sending Sync frames.

The Sync Frame Direction field shall be set to zero if the PNC is requesting omni-directional transmission of the Sync frame and shall be set to one if the PNC is requesting directional transmission of the Sync frame. If the request is for directional transmission, then the Sync frame direction is a round robin of the DEV's available directions.

6.4.36 Directional Peer IE

The Directional Peer IE Content field shall be formatted as illustrated in Figure 6-124.

Octets: 1	1	1	2
SrcID	DestID	Configuration	Allocated Superframes

Figure 6-124—Directional Peer IE Content field format

The SrcID field contains the DEVID of the source for the directional peer communication in the regular CAP and/or regular S-CAP.

The DestID field contains the DEVID of the destination for the directional peer communication in the regular CAP and/or regular S-CAP.

The Configuration field shall be formatted as illustrated in Figure 6-125.

Bits: b0	b1	b2–b7
Request/Release	Unidirectional/Bidirectional Allowance	Reserved

Figure 6-125—Configuration field format

The Request/Release field shall be set to one to request directional communication and shall be set to zero to release the directional communication.

The Unidirectional/Bidirectional Allowance field shall be set to one for unidirectional communication and shall be set to zero otherwise.

The Allocated Superframes field indicates the number of superframes allocated for directional communication between the source and destination.

6.5 MAC commands

The MAC commands are listed in Table 6-22. If the column labeled “Associated” in Table 6-22 is marked with an “X,” then that MAC command shall only be sent by a DEV that is associated in the piconet. If the column labeled “Secure membership (if required)” in Table 6-22 is marked with an “X” and secure membership is required for the piconet, then that MAC command shall only be sent by a DEV that has established secure membership with the PNC in the piconet. Because a neighbor PNC is not a member of the piconet, it sends only non-secure MAC commands. The PNC or destination DEV shall ignore any MAC command from a DEV that is not allowed to be sent, as indicated in Table 6-22. The “Required” column indicates the type of DEVs that are required to support the command.

For peer-to-peer communications, if the DEV has established a secure relationship with a peer DEV, and the “Secure membership (if required)” column is marked with an “X,” that command shall be sent to the peer DEV with a secure command using the key specified in Table 8-1.

Table 6-22—Command types

Command type hex value b15–b0	Command name	Subclause	Associated	Secure membership (if required)	Required
0x0000	Association Request command	6.5.1.1			All DEVs
0x0001	Association Response command	6.5.1.2	X		All DEVs
0x0002	Disassociation Request command	6.5.1.3	X		All DEVs
0x0003	Request Key command	6.5.2.1	X	X	Optional
0x0004	Request Key Response command	6.5.2.2	X	X	Optional
0x0005	Distribute Key Request command	6.5.2.3	X	X	Optional
0x0006	Distribute Key Response command	6.5.2.4	X	X	Optional
0x0007	PNC Handover Request command	6.5.3.1	X	X	All DEVs
0x0008	PNC Handover Response command	6.5.3.2	X	X	All DEVs
0x0009	PNC Handover Information command	6.5.3.3	X	X	All DEVs
0x000A	PNC Information Request command	6.5.4.1	X	X	All DEVs
0x000B	PNC Information command	6.5.4.2	X	X	All DEVs
0x000C	Security Information Request command	6.5.4.3	X	X	Optional
0x000D	Security Information command	6.5.4.4	X	X	Optional
0x000E	Probe Request command	6.5.4.5	X		All DEVs
0x000F	Probe Response command	6.5.4.6	X		All DEVs

Table 6-22—Command types (continued)

Command type hex value b15–b0	Command name	Subclause	Associated	Secure membership (if required)	Required
0x0010	Piconet Services command	6.5.5.1	X		All DEVs
0x0011	Announce command	6.5.5.2	X		All DEVs
0x0012	Channel Time Request command	6.5.6.1	X	X	All DEVs
0x0013	Channel Time Response command	6.5.6.2	X	X	All DEVs
0x0014	Channel Status Request command	6.5.7.1	X	X	All DEVs
0x0015	Channel status response	6.5.7.2	X	X	All DEVs
0x0016	Remote Scan Request command	6.5.7.3	X	X	All DEVs
0x0017	Remote Scan Response command	6.5.7.4	X	X	All DEVs
0x0018	Transmit Power Change command	6.5.7.5	X	X	All DEVs
0x0019	PS Set Information Request command	6.5.8.1	X	X	All DEVs
0x001A	PS Set Information Response command	6.5.8.2	X	X	All DEVs
0x001B	SPS Configuration Request command	6.5.8.3	X	X	All DEVs
0x001C	SPS Configuration Response command	6.5.8.4	X	X	All DEVs
0x001D	PM Mode Change command	6.5.8.5	X	X	All DEVs
0x001E	Security Message command	6.5.9.1	X		Optional
0x001F	Announce Response command	6.5.5.3	X		All DEVs
0x0020	PM Mode Change Response command	6.5.8.6	X	X	All DEVs
0x0021	AS IE Request command	6.5.9.3	X	X	Optional
0x0022	AS IE Response command	6.5.9.4	X	X	Optional
0x0023	Multicast Configuration Request command	6.5.10.1	X	X	Optional
0x0024	Multicast Configuration Response command	6.5.10.2	X	X	Optional
0x0025–0x00FF	Reserved				
0x0100–0xFFFF	Vendor Defined command	6.5.9.2	X		Optional

Unless otherwise stated in the command descriptions, in all commands sent between the PNC and a DEV, the SEC field shall be set to zero when the piconet is operating in security mode 0. When the piconet is operating in security mode 1, unless otherwise stated in the command description, the SEC field shall be set

to one and the DEV shall be a secure member of the piconet in order to send that command. The ACK Policy field shall be set to Imm-ACK for all commands unless otherwise stated in the command description.

The only commands that may be fragmented are as follows:

- PNC Information command, as described in 6.5.4.2
- PNC Handover Information command, as described in 6.5.3.3
- PS Set Information Response command, as described in 6.5.8.2

The fragmentation and defragmentation of these commands employ the same method as that used for data frames, as described in 7.7.

6.5.1 Association and disassociation commands

These commands are used by a DEV to join a piconet and by a DEV or the PNC to end a DEV's membership in the piconet.

6.5.1.1 Association Request command

The Association Request command Payload field shall be formatted as illustrated in Figure 6-126. The SEC field in the Frame Control field shall be set to zero. The DestID shall be set to the PNCID. The SrcID shall be set to either the UnassocID, as described in 6.2.3, or the DEV's newly allocated DEVID, as described in 7.3.1.

Octets: 6	As defined in 6.4.11	2	1	variable
DEV Address	Overall Capabilities	ATP	DEV Utility	IEs

Figure 6-126—Association Request command Payload field format

The DEV Address field is the address of the DEV, as described in 6.1, requesting association.

The Overall Capabilities field is defined in 6.4.11.

The Association Timeout Period (ATP) field is maximum amount of time in milliseconds that the association relationship will be maintained in the absence of communication between the PNC and DEV, as described in 7.3.4.

The DEV Utility field shall be formatted as illustrated in Figure 6-127.

Bits: b0	b1	b2–b4	b5–b7
Piconet Services Inquiry	Neighbor PNC	Best PNC TX Quasi-omni Pattern	Reserved

Figure 6-127—DEV Utility field format

The Piconet Services Inquiry field shall be set to one if the associating DEV is requesting that the PNC send the Piconet Services command, as described in 6.5.5.1, and shall be set to zero otherwise.

The Neighbor PNC field shall be set to one if the DEV intends to be a neighbor PNC, as described in 7.2.8, in the current piconet and shall be set to zero otherwise.

The Best PNC TX Quasi-omni Pattern field contains the PNC TX quasi-omni pattern, as determined by the DEV from receiving beacons, that the PNC should use for further quasi-omni transmissions to the DEV.

If the PNC uses AAS, as determined from the received beacon, the DEV may set the Response TX sector field to the index of the beacon that was initially heard by the DEV. It shall be set to zero otherwise.

The IEs field contains zero or more IEs.

6.5.1.2 Association Response command

The Association Response command Payload field shall be formatted as illustrated in Figure 6-128. The ACK Policy field shall be set to no-ACK. The SEC field in the Frame Control field shall be set to zero. The DestID shall be set to the UnassocID, as described in 6.2.3. The SrcID shall be set to the PNCID.

Octets: 6	1	2	1	variable
DEV Address	DEVID	ATP	Reason Code	IEs

Figure 6-128—Association Response command Payload field format

The DEV Address field is the address of the DEV, as described in 6.1, requesting association.

The DEVID field is the identifier allocated to the DEV if the association is successful. If this field contains the UnassocID, the DEV is not allowed to associate for the reason indicated in the reason code. For the successful association of a neighbor PNC, the DEVID shall be one of the reserved NbrIDs, as described in 6.2.3.

The ATP field contains the finalized value for the Association Timeout Period in milliseconds. This value may be different from that requested by the DEV in its Association Request command if the PNC is not able to support the value requested.

The valid values of the Reason Code are as follows:

- 0 → Success
- 1 → Already serving maximum number of DEVs
- 2 → Lack of available channel time to serve the DEV
- 3 → Channel too severe to serve the DEV
- 4 → PNC turning off with no PNC-capable DEV in the piconet
- 5 → Neighbor piconet not allowed
- 6 → Channel change in progress
- 7 → PNC handover in progress
- 8 → Association denied
- 9–254 → Reserved
- 225 → Other failure

The IEs field contains zero or more IEs.

6.5.1.3 Disassociation Request command

The Disassociation Request command Payload field shall be formatted as illustrated in Figure 6-129.

Octets: 1
Reason Code

Figure 6-129—Disassociation Request command Payload field format

The valid reason codes are as follows:

- 0 → ATP expired
- 1 → Channel too severe to serve the DEV
- 2 → PNC unable to service DEV
- 3 → PNC turning off with no PNC-capable DEV in the piconet
- 4 → DEV leaving piconet
- 5–254 → Reserved
- 225 → Other failure

The Disassociation Request command shall use the secure command frame format if the DEV is a secure member of the piconet.

6.5.2 Security commands

This set of commands is used to establish the security and privacy functions between a DEV and the PNC and between DEVs in the piconet.

6.5.2.1 Request Key command

The Request Key command is used to request a payload protection key from the key originator. The SEC field in the Frame Control field shall be set to one.

This command shall be protected using the management key that is shared between the requesting DEV and the key originator. The Request Key command has no Payload field.

6.5.2.2 Request Key Response command

The Request Key Response command is used by a key originator in a security relationship to send the requested key in an encrypted format to the requesting DEV. The SEC field in the Frame Control field shall be set to one. This command shall be protected using the management key that is shared between the requesting DEV and the key originator. The integrity code is generated using the management key that is shared between the requesting DEV and the key originator. The Request Key Response command Payload field shall be formatted as illustrated in Figure 6-130.

Octets: 2	variable
SECID	Encrypted Key

Figure 6-130—Request Key Response command Payload field format

The SECID field is the unique identifier for the security relationship with which the key is associated. The SECID field is used to identify the type of key and the key originator and is defined in 6.2.7.2.

The Encrypted Key field is defined in the symmetric key security operations, as described in 9.3.1.

6.5.2.3 Distribute Key Request command

The Distribute Key Request command is used to transmit a key to another DEV. The SEC field in the Frame Control field shall be set to one. This command may have the ACK Policy field set to no-ACK only if the source ID is the PNCID. This command shall be protected using the management key that is shared between the requesting DEV and the key originator. The Distribute Key Request command Payload field shall be formatted as illustrated in Figure 6-131.

Octets: 2	variable
SECID	Encrypted Key

Figure 6-131—Distribute Key Request command Payload field format

The SECID field is the unique identifier for the security relationship with which the key is associated. The SECID field is used to identify the type of key and the key originator and is defined in 6.2.7.2.

The Encrypted Key field is defined in the symmetric key security operations, as described in 9.3.1.

6.5.2.4 Distribute Key Response command

The Distribute Key Response command is used in a distribute key protocol to inform the key originator whether or not the key was properly received. The SEC field in the Frame Control field shall be set to one. This command shall be protected using the management key that is shared between the requesting DEV and the key originator. The Distribute Key Response command Payload field shall be formatted as illustrated in Figure 6-132.

Octets: 2
SECID

Figure 6-132—Distribute Key Response command Payload field format

The SECID field is the unique identifier for the security relationship with which the key is associated. The SECID field is used to identify the type of key and the key originator and is defined in 6.2.7.2.

6.5.3 PNC handover commands

These commands are used to handover PNC responsibilities.

6.5.3.1 PNC Handover Request command

This command is used by the PNC to hand over its responsibility to another DEV in the piconet that is capable of being a PNC. The PNC Handover Request command Payload field shall be formatted as illustrated in Figure 6-133.

Octets: 1	1	1	1
Number Of DEVs	Number Of Channel Time Request Blocks	Number Of PS Sets	Handover Status

Figure 6-133—PNC Handover Request command Payload field format

The Number Of DEVs field indicates the total number of DEVs that are currently members of the piconet. In addition, this field indicates the number of DEV Information records that will be transferred from the old PNC to the new PNC via the PNC Information command, as described in 6.5.4.2.

The Number Of Channel Time Request Blocks field is the number of Channel Time Request Block fields, excluding requests for asynchronous channel time, currently being served by the PNC that will be transferred from the old PNC to the new PNC via the PNC Handover Information command, as described in 6.5.3.3.

The Number Of PS Sets field indicates the total number of PS sets that will be transferred from the old PNC to the new PNC via the PS Set Information Response command, as described in 6.5.8.2.

The allowed values of the Handover Status field are as follows:

- 0 → The PNC is starting the PNC handover process with the destination DEV.
- 1 → The PNC is canceling the handover process with the destination DEV.
- 2 → The PNC is only transferring information and is not beginning the PNC handover process.
- 3 → The PNC is starting the PNC handover process with information that was sent previously, as described in 7.2.4.
- 4–255 → Reserved.

6.5.3.2 PNC Handover Response command

The format of the PNC Handover Response command Payload field shall be as illustrated in Figure 6-134.

Octets: 1
Reason Code

Figure 6-134—PNC Handover Response command Payload field format

The Reason Code field indicates that the new PNC is either ready to take over as the new PNC or that it will be unable to become the PNC. The valid Reason Code values are as follows:

- 0x00 → Success, ready for handover
- 0x01–0xEC → Success, member of parent piconet with DEVID equal to Reason Code value
- 0xED–0xF6 → Reserved
- 0xF7–0xFC → Success, associated in parent piconet with NbrID equal to Reason Code value
- 0xFE → Handover refused, unable to join parent piconet
- 0xFF → Handover refused, unable to act as PNC for more than one piconet

6.5.3.3 PNC Handover Information command

The PNC Handover Information command Payload field shall be formatted as illustrated in Figure 6-135.

Octets: 1	12	2	...	1	12	2
DEVID	Channel Time Request Block 1	Next Beacon	...	DEVID	Channel Time Request Block n	Next Beacon

Figure 6-135—PNC Handover Information command Payload field format

The DEVID field contains the identifier of the source of the Channel Time Request Block field that follows in the command.

The Channel Time Request Block field is defined in 6.5.6.1. Note that asynchronous Channel Time Request Block fields are not passed in this command, thus the Num Targets field in the Channel Time Request Block field is always one. Consequently, the Channel Time Request Block fields will all be a fixed length.

The Next Beacon field indicates the beacon number, as described in 6.3.1.1, of the next superframe when this channel time request will be allocated.

6.5.4 Information request commands

This set of commands is used to obtain information about another DEV in the piconet. The PNC Information Request and PNC Information commands are used to retrieve data about any or all of the currently associated DEVs in the piconet. The Security Information Request and Security Information commands are used to retrieve security information about any or all of the currently associated DEVs in the piconet. The Probe Request and Probe Response commands are used to retrieve IEs from a specific DEV in the piconet.

6.5.4.1 PNC Information Request command

The DestID for the PNC Information Request command Payload field shall be the PNCID. The PNC Information Request command shall be formatted as illustrated in Figure 6-136.

Octets: 1
Queried DEVID

Figure 6-136—PNC Information Request command Payload field format

The Queried DEVID field contains the DEVID of the DEV whose information is being requested from the PNC. If the value of this field is BcstID, then the DEV is requesting information regarding the entire list of associated DEVs from the PNC.

6.5.4.2 PNC Information command

This command may be sent either as a response to the PNC Information Request command by a DEV or it may be sent unsolicited. In either case the SrcID shall be the PNCID. This command may be sent either in a directed command frame to a DEV or it may be sent in a broadcast command frame meant for all DEVs in

the piconet. If the DestID is BcstID, then the ACK Policy field shall be no-ACK. The PNC Information command Payload field shall be formatted as illustrated in Figure 6-137.

Octets: 18/24	18/24	---	18/24
DEV-1 Info	DEV-2 Info	...	DEV- <i>m</i> Info

Figure 6-137—PNC Information command Payload field format

The DEV Info field shall be formatted as illustrated in Figure 6-138.

Octets: 6	1	1	As defined in 6.4.11	2	1
DEV Address	DEVID	Membership Status	Capability	ATP	System Wake Beacon Interval

Figure 6-138—Format of a DEV Info field in a PNC Information command

The DEV Address field contains the address of the DEV, as described in 6.1, corresponding to the DEVID.

The DEVID field contains the ID assigned to the DEV by the PNC. This field shall not contain the BcstID, the UnassocID, the McstID, or the reserved IDs, as described in 6.2.3.

The DEV Info Utility field shall be formatted as illustrated in Figure 6-139.

Bits: b0	b1–b7
Membership Status	Reserved

Figure 6-139—DEV Info Utility field format

The Membership Status field shall be set to zero if the DEV is associated but is not a secure member of the piconet and shall be set to one if the DEV is associated and a secure member of the piconet.

The Overall Capabilities field shall be formatted as illustrated in Figure 6-81 and is defined in 6.4.11.

The ATP field is defined in 6.5.1.1.

The System Wake Beacon Interval field, as described in 6.5.8.3, is the value that the DEV sent to the PNC via the SPS Configuration Request command, as described in 6.5.8.3.

6.5.4.3 Security Information Request command

The Security Information Request command enables a DEV to request security information regarding a single DEV or all DEVs. The Security Information Request command Payload field shall be formatted as illustrated in Figure 6-140.

Octets: 1
Queried DEVID

Figure 6-140—Security Information Request command Payload field format

The Queried DEVID field indicates the DEV whose security information is being requested. If the value of this field is the BcstID, then the DEV is requesting all of the security information maintained by the target DEV.

6.5.4.4 Security Information command

The Security Information command Payload field shall be formatted as illustrated in Figure 6-141.

Octets: 1	1	variable	variable	...	variable
Total Number Of Frames	Sequence Number	DEV-1 Security Record	DEV-2 Security Record	...	DEV- <i>m</i> Security Record

Figure 6-141—Security Information command Payload field format

The Total Number Of Frames field indicates the number of frames that will be sent to complete this request.

The Sequence Number field specifies the number of frames that have been sent prior to this frame by this DEV in response to the initial request. Thus the first frame has a Sequence Number of 0 while the last frame has a Sequence Number equal to the Total Number Of Frames minus one.

A given Security Record field shall be formatted as illustrated in Figure 6-142.

Octets: 2	6	1	2	variable
Length	DEV Address	DEVID	Verification Info Length	Verification Info

Figure 6-142—Format of an Security Record field in an Security Information command

The DEV Address field contains the address of the DEV, as described in 6.1, corresponding to the DEVID.

The DEVID field contains the ID assigned to the DEV by the PNC. If the DEV is not currently associated in this piconet, the field shall be set to the UnassocID. This field shall not contain the broadcast or multicast DEVIDs.

The Verification Info Length field indicates the length of the verification information that is included in the Security Record field. If this length is zero, no Verification Info field shall be included.

The Verification Info field specifies the security information that may be used to verify the identity of that particular DEV.

6.5.4.5 Probe Request command

The Probe Request command is used either to request information about a DEV or to see if a DEV is still present in the piconet. This command may be exchanged between any two DEVs in the piconet according to the rules outlined in Table 6-23 and Table 6-24. The individual IEs used in this frame are described in 6.4. The Probe Request command Payload field shall be formatted as illustrated in Figure 6-143.

Octets: 4	2
Information Requested	Request Index

Figure 6-143—Probe Request command Payload field format

The Information Requested field shall be formatted as illustrated in Figure 6-144.

Bits: b0	b1–b31
IE Request Type	IEs Requested

Figure 6-144—Information Requested field format

The IE Request Type field indicates the format of the IEs requested field. This field shall be set to zero if the IEs Requested field is a bitmap and shall be set to one if the IEs Requested field is a binary encoding of the IE's ID.

If the IE Request Type field indicates that the IEs Requested field is a bitmap, then the sender shall set a value of one in a bit to request the IE that corresponds to the bit position. Otherwise, the sender shall set the bit to zero. The bit position for an IE is same as the value of the element-ID for that IE. That is, the bit position of n in information request field corresponds with the IE whose element ID, Table 6-13, is n .

If the IE Request Type field indicates that the rest of the bits are binary coded, then the IEs Requested field contains the element ID of the IE that is being requested by the sender of this command from its intended recipient.

Both the IE Request Type field and the IEs Requested field shall be set to zero when the source DEV is not requesting any information from the destination DEV.

If the IEs Requested field indicates that the CTA Status IE, as described in 6.4.10, is being requested from the destination DEV, the first octet of the Request Index field is set to the stream index of the stream for which CTA information is requested. If the Request Index field is set to zero, the DEV is requesting information about all isochronous streams directed to the requesting DEV and to the BcstID and McstID. If the Information Requested field indicates that the CTA Status IE is not being requested from the destination DEV, the Request Index field has no meaning and shall be set to zero.

Table 6-23 lists the rules that shall apply to requesting IEs from another DEV based on the identity of the originator of the request.

Table 6-23—Rules for requesting IEs in a Probe Request command

IE	Subclause	PNC allowed to request?	DEV allowed to request?
CTA IE	6.4.1	Shall not request	Shall not request
BSID IE	6.4.2	Shall not request	May request
Parent Piconet IE	6.4.3	Shall not request	May request
DEV Association IE	6.4.4	Shall not request	Shall not request
PNC shutdown IE	6.4.5	Shall not request	Shall not request
Piconet Parameter Change IE	6.4.6	Shall not request	Shall not request
AS IE	6.4.7	May request	May request
Pending channel time map (PCTM) IE	6.4.8	Shall not request	May request
PNC Handover IE	6.4.9	Shall not request	Shall not request
CTA Status IE	6.4.10	Shall not request	May request
Capability IE	6.4.11	May request	May request
Transmit Power Parameters IE	6.4.12	May request	May request
PS Status IE	6.4.13	Shall not request	Shall not request
CWB IE	6.4.14	Shall not request	Shall not request
Overlapping PNID IE	6.4.15	May request	Shall not request
Piconet Services IE	6.4.16	May request	May request
Vendor Defined IE or reserved	6.4.17	May request	May request
Group ID IE	6.4.18	Shall not request	May request
Stream Renew IE	6.4.19	Shall not request	Shall not request
Next PNC IE	6.4.20	Shall not request	Shall not request
Piconet Channel Status IE	6.4.21	May request	May request
Synchronization IE	6.4.22	Shall not request	Shall not request
TSD IE	6.4.23	Shall not request	Shall not request
UEP Specific IE	6.4.24	May request	May request
IFS IE	6.4.25	Shall not request	Shall not request
CTA Relinquish Duration IE	6.4.26	Shall not request	Shall not request
Feedback IE	6.4.27	Shall not request	Shall not request
Mapping IE	6.4.28	Shall not request	Shall not request
BST Clustering IE	6.4.29	Shall not request	Shall not request
PET Clustering IE	6.4.30	Shall not request	Shall not request
Beam PET IE	6.4.31	Shall not request	Shall not request
HRS Beam PET IE	6.4.32	Shall not request	Shall not request
PET Amplitude IE	6.4.33	Shall not request	Shall not request

Table 6-23—Rules for requesting IEs in a Probe Request command (continued)

IE	Subclause	PNC allowed to request?	DEV allowed to request?
PET Phase IE	6.4.34	Shall not request	Shall not request
Sync Frame Frequency IE	6.4.35	Shall not request	Shall not request
Directional Peer IE	6.4.36	Shall not request	Shall not request

6.5.4.6 Probe Response command

The Probe Response command is used to return information about a DEV to a requesting DEV. The individual IEs used in this frame are described in 6.4. The Probe Response command Payload field shall be formatted as illustrated in Figure 6-145.

Octets: variable
IEs Provided

Figure 6-145—Probe Response command Payload field format

The IEs Provided field contains the IEs, as described in 6.4, that the source DEV of this command is providing to the destination. The elements themselves may be placed in any order.

Table 6-24 lists the rules that shall apply to responding to a request for an IE based on the sender of the request.

Table 6-24—Rules for responding to requests in Probe commands

IE	Subclause	DEV receives request from DEV	DEV receives request from PNC	PNC receives request from DEV
CTA IE	6.4.1	Shall ignore	Shall ignore	Shall ignore
BSID IE	6.4.2	Shall ignore	Shall ignore	Shall respond
Parent Piconet IE	6.4.3	Shall ignore	Shall ignore	Shall respond
DEV Association IE	6.4.4	Shall ignore	Shall ignore	Shall ignore
PNC shutdown IE	6.4.5	Shall ignore	Shall ignore	Shall ignore
Piconet Parameter Change IE	6.4.6	Shall ignore	Shall ignore	Shall ignore
AS IE	6.4.7	May respond	May respond	May respond
Pending channel time map (PCTM) IE	6.4.8	Shall ignore	Shall ignore	Shall respond
PNC Handover IE	6.4.9	Shall ignore	Shall ignore	Shall ignore
CTA Status IE	6.4.10	Shall ignore	Shall ignore	Shall respond
Capability IE	6.4.11	Shall respond	Shall respond	Shall respond
Transmit Power Parameters IE	6.4.12	Shall respond	Shall respond	Shall respond

Table 6-24—Rules for responding to requests in Probe commands (continued)

IE	Subclause	DEV receives request from DEV	DEV receives request from PNC	PNC receives request from DEV
PS Status IE	6.4.13	Shall ignore	Shall ignore	Shall ignore
CWB IE	6.4.14	Shall ignore	Shall ignore	Shall ignore
Overlapping PNID IE	6.4.15	Shall ignore	May respond	Shall ignore
Piconet Services IE	6.4.16	May respond	May respond	May respond
Vendor Defined IE or reserved	6.4.17	May respond	May respond	May respond
Group ID IE	6.4.18	Shall ignore	Shall ignore	Shall respond
Stream Renew IE	6.4.19	Shall ignore	Shall ignore	Shall ignore
Next PNC IE	6.4.20	Shall ignore	Shall ignore	Shall ignore
Piconet Channel Status IE	6.4.21	May respond	May respond	May respond
Synchronization IE	6.4.22	Shall ignore	Shall ignore	Shall ignore
TSD IE	6.4.23	Shall ignore	Shall ignore	Shall ignore
UEP Specific IE	6.4.24	May respond	May respond	May respond
IFS IE	6.4.25	Shall ignore	Shall ignore	Shall ignore
CTA Relinquish Duration IE	6.4.26	Shall ignore	Shall ignore	Shall ignore
Feedback IE	6.4.27	Shall ignore	Shall ignore	Shall ignore
Mapping IE	6.4.28	Shall ignore	Shall ignore	Shall ignore
BST Clustering IE	6.4.29	Shall ignore	Shall ignore	Shall ignore
PET Clustering IE	6.4.30	Shall ignore	Shall ignore	Shall ignore
Beam PET IE	6.4.31	Shall ignore	Shall ignore	Shall ignore
HRS Beam PET IE	6.4.32	Shall ignore	Shall ignore	Shall ignore
PET Amplitude IE	6.4.33	Shall ignore	Shall ignore	Shall ignore
PET Phase IE	6.4.34	Shall ignore	Shall ignore	Shall ignore
Sync Frame Frequency IE	6.4.35	Shall ignore	Shall ignore	Shall ignore
Directional Peer IE	6.4.36	Shall ignore	Shall ignore	Shall ignore

6.5.5 Information announcement commands

6.5.5.1 Piconet Services command

The Piconet Services command is sent by the PNC to provide information about the application layer capabilities of all of the DEVs in a piconet. The Piconet Services command Payload field shall be formatted as illustrated in Figure 6-146.

Octets: variable	variable	...	variable
Piconet Services IE-1	Piconet Services IE-2	...	Piconet Services IE-n

Figure 6-146—Piconet Services command Payload field format

The Piconet Services IE is defined in 6.4.16.

6.5.5.2 Announce command

The Announce command is used to send unrequested information about a DEV to one or more DEVs in the piconet. The individual IEs used in this frame are described in 6.4. The Announce command Payload field shall be formatted as illustrated in Figure 6-147.

Octets: variable
IEs Provided

Figure 6-147—Announce command Payload field format

The IEs Provided field contains the IEs, as described in 6.4, that the source DEV of this command is providing to the destination. The elements themselves may be placed in any order.

Table 6-25 lists the rules that shall apply to sending an unrequested IE based on the sender of the request.

Table 6-25—Rules for sending IEs in an Announce command

IE	Subclause	PNC allowed to send?	DEV allowed to send?
CTA IE	6.4.1	Shall not send	Shall not send
BSID IE	6.4.2	Shall not send	Shall not send
Parent Piconet IE	6.4.3	Shall not send	Shall not send
DEV Association IE	6.4.4	May send	Shall not send
PNC shutdown IE	6.4.5	May send	Shall not send
Piconet Parameter Change IE	6.4.6	May send	Shall not send
AS IE	6.4.7	May send	May send
Pending channel time map (PCTM) IE	6.4.8	May send	Shall not send
PNC Handover IE	6.4.9	May send	Shall not send

Table 6-25—Rules for sending IEs in an Announce command (continued)

IE	Subclause	PNC allowed to send?	DEV allowed to send?
CTA Status IE	6.4.10	May send	Shall not send
Capability IE	6.4.11	May send	May send
Transmit Power Parameters IE	6.4.12	May send	May send
PS Status IE	6.4.13	May send	Shall not send
CWB IE	6.4.14	Shall not send	Shall not send
Overlapping PNID IE	6.4.15	Shall not send	May send
Piconet Services IE	6.4.16	May send	May send
Vendor Defined IE or reserved	6.4.17	May send	May send
Group ID IE	6.4.18	May send	Shall not send
Stream Renew IE	6.4.19	Shall not send	May send
Next PNC IE	6.4.20	May send	Shall not send
Piconet Channel Status IE	6.4.21	May send	May send
Synchronization IE	6.4.22	May send	Shall not send
TSD IE	6.4.23	May send	May send
UEP Specific IE	6.4.24	May send	May send
IFS IE	6.4.25	May send	Shall not send
CTA Relinquish Duration IE	6.4.26	May send	May send
Feedback IE	6.4.27	May send	May send
Mapping IE	6.4.28	May send	May send
BST Clustering IE	6.4.29	May send	May send
PET Clustering IE	6.4.30	May send	May send
Beam PET IE	6.4.31	May send	May send
HRS Beam PET IE	6.4.32	May send	May send
PET Amplitude IE	6.4.33	May send	May send
PET Phase IE	6.4.34	May send	May send
Sync Frame Frequency IE	6.4.35	May send	Shall not send
Directional Peer IE	6.4.36	May send	May send

6.5.5.3 Announce Response command

The Announce Response command is used to confirm the receipt of IEs from another DEV. The individual IEs used in this frame are described in 6.4. The Announce command Payload field shall be formatted as illustrated in Figure 6-148.

Octets: 2	2	...	2
Announce Response Block 1	Announce Response Block 2	...	Announce Response Block <i>n</i>

Figure 6-148—Announce Response command Payload field format

The Announce Response Block fields shall be formatted as illustrated in Figure 6-149.

Octets: 1	1
IE Element ID	Reason Code

Figure 6-149—Announce Response Block field format

The IE Element ID field contains the element ID of the IE for which the response is being sent.

The Reason Code field indicates the result of the Announce command for the IE specified in the IE Element ID field. The valid values of the Reason Code field are as follows:

- 0 → Success
- 1 → Unsupported IE
- 2–254 → Reserved
- 255 → Other failure

6.5.6 Channel time allocation request, modification, and termination commands

This group of commands is used for the request, modification, termination and grant of channel time within the CTAP.

6.5.6.1 Channel Time Request command

The Channel Time Request command may be used to request, modify, or terminate CTAs corresponding with either isochronous streams or asynchronous data traffic. The Channel Time Request command structure shall be formatted as illustrated in Figure 6-150. The DEV that sends this command is the originator and is seeking from the PNC channel time allocations during which to communicate with a target DEV or DEVs.

Octets: 12–138	12–138	...	12–138
Channel Time Request Block 1	Channel Time Request Block 2	...	Channel Time Request Block <i>n</i>

Figure 6-150—Channel Time Request command Payload field format

Each Channel Time Request Block field corresponds to a channel time request. If the DEV is making a request for asynchronous channel time where the destinations share CTAs, then there shall be only one

asynchronous Channel Time Request Block field in the command, and it shall be the last Channel Time Request Block field in the Channel Time Request command. The Channel Time Request Block field for a given Channel Time Request command shall be formatted as illustrated in Figure 6-151.

Octets: 1	1–127	1	1	1	1	2	2	1	1
Num Targets	Target ID List	DSPS Set Index	Stream Request ID	Stream index	Channel Time Request control	CTA rate factor	Channel Time Request TU	Minimum number of TUs	Desired Number Of TUs

Figure 6-151—Channel Time Request Block field format

The Num Targets field indicates the number of target DEVIDs in the target ID list. For isochronous requests, i.e., stream index not equal to the asynchronous stream index, the num targets field shall be set to one. For asynchronous requests, the num targets fields shall take on values from 1 to 127.

The Target ID List field is a series of DEVIDs with which the originating DEV seeks to establish communications by requesting channel time allocations from the PNC.

The DSPS Set Index field is used to identify the DSPS set with which the channel time request corresponds, if the channel time request is for a DSPS allocation. Only valid DSPS set indices, as described in 6.5.8.3, are allowed for a DSPS allocation request. Otherwise, the field shall be set to zero and shall be ignored on reception.

The Stream Request ID field is used to uniquely identify the DEV's request before it receives a stream index from the PNC. If the channel time request is for a new isochronous stream, then the stream request ID is a non-zero identifier generated by the originating DEV that is unique among the DEV's channel time requests. The stream request ID shall remain constant during the entire frame exchange sequence for establishing a new stream. If the channel time request is to modify or terminate an existing stream or the request is for an asynchronous allocation, the stream request ID shall be set to zero and shall be ignored on reception.

The Stream Index field is defined in 6.2.5. In the case where the DEV is requesting the creation of an isochronous stream, it is set to the unassigned stream value, as described in 6.2.5, by the originating DEV. In the case where the DEV is requesting the reservation or termination of an asynchronous channel time, it is set to the asynchronous stream value, as described in 6.2.5. When the stream index is other than the unassigned stream index or asynchronous stream index value, this channel time request is a request to modify or terminate an existing CTA. In the case where the DEV is requesting a specific MCTA interval, as described in 7.4.3.3, the stream index shall be set to the MCTA stream value, as described in 6.2.5, and the Target ID List field shall contain only the PNCID. In the case where the DEV is requesting an allocation for beam forming, the stream index shall be set to beam forming stream value, as described in 6.2.5.

The Channel Time Request Control field shall be formatted as illustrated in Figure 6-152.

Bits: b0–b2	b3	b4	b5	b6	b7
User Priority	Reserved	PM Channel Time Request Type	CTA Type	CTA Rate Type	Target ID List Type

Figure 6-152—Channel Time Request Control field format

The User Priority field is defined in Table B-1.

The PM Channel Time Request Type field indicates the type of request. It shall be set to zero to request an ACTIVE channel time allocation and shall be set to one to request a DSPS channel time allocation. For sub-rate allocations, an ACTIVE allocation puts no restriction on the superframe of the first CTA. A DSPS allocation synchronizes all CTAs with the DSPS set awake superframes of the DSPS set specified by the DSPS index. The value of the CTA Rate Factor shall be no smaller than the DSPS set's wake beacon interval.

The CTA Type field indicates whether a pseudo-static CTA is being requested. The CTA Type field shall be set to one if the channel time request is for a pseudo-static CTA and shall be set to zero otherwise.

The CTA Rate Type field indicates whether a super-rate CTA or a sub-rate CTA is being requested. The CTA Rate Type field shall be set to one for a sub-rate CTA and zero for a super-rate CTA.

The CTA Rate Factor field in conjunction with the CTA Rate Type field specifies the frequency at which the requesting DEV would like the PNC to allocate channel time. In the case of a super-rate request, the PNC will interpret the CTA Rate Factor as the maximum spacing allowed between CTAs allocated in a superframe, as described in F.1.3.

For instance, in the case where the CTA Rate Type field is set to zero, a value indicating a super-rate CTA request, and the CTA Rate Factor field contains a value N greater than zero, the requesting DEV is requesting super-rate CTAs from the PNC. If these super-rate CTAs are allocated by the PNC, they will appear N times per superframe. A PNC shall support at least 8 CTAs per stream in the same superframe. The CTA Rate Type field set to zero and the CTA Rate Factor field set to zero shall be reserved.

NOTE—A sub-rate request always has a CTA Rate Factor greater than one. Thus a CTA Rate Type equal to one and a CTA Rate Factor equal to one is not allowed.

In the case where the CTA Rate Type field is set to one, a value indicating a sub-rate CTA request, and the CTA Rate Factor field contains a non-zero value N , the requesting DEV is requesting sub-rate CTAs from the PNC. If these sub-rate CTAs are allocated by the PNC, they will appear in the beacon once every N superframes. The CTA Rate Factor in this case shall be limited to powers of 2 (i.e., 2, 4, 8, ...), up to and including the value of 65536, which shall be represented by a CTA Rate Factor equal to zero.

If the Channel Time Request Block field is a request for an MCTA interval, only the CTA Rate Factor field and stream index shall be interpreted by the PNC. All other fields except the stream index and num targets fields shall be set to zero.

The Target ID List Type field shall be set to zero for asynchronous group channel time requests and shall be set to one for individual asynchronous channel time requests, as described in 7.5.2.1.

The Channel Time Request Time Unit (TU) field indicates the unit of time that the DEV is using for the CTA(s) it is requesting. This allows the PNC to know the units of channel time the DEV is able to make use of so that the PNC will efficiently allocate channel time. The resolution of this field is 1 μ s and therefore has a range of [0–65535] μ s.

For an isochronous request, the Minimum Number of TUs field indicates the minimum number of Channel Time Request TUs required by the originating DEV to support the stream.

For an isochronous request, the Desired Number of TUs field indicates the number of Channel Time Request TUs that is desired by the requesting DEV. The Desired Number of TUs field shall be greater than or equal to the Minimum Number of TUs field.

In the case of an isochronous super-rate allocation, the Minimum Number of TUs and the Desired Number of TUs are the number of TUs requested in each superframe. In the case of an isochronous sub-rate

allocation, the fields contain the number of TUs requested in each of the superframes containing the sub-rate CTA. For example, a request for a Minimum Number of TUs of 4 with a sub-rate CTA Rate Factor of 4 indicates that the DEV is requesting 4 TUs every fourth superframe. Likewise, a request for a Minimum Number of TUs of 11 with a super-rate CTA Rate Factor of 4 indicates that the DEV is requesting at least 11 TUs per superframe, spread into 4 allocations that are evenly spaced in the superframe.

For an asynchronous request, the concatenation of the Minimum Number of TUs field and the Desired Number of TUs field indicates the total number of TUs that are requested for this allocation, i.e., it is interpreted as a single, 2-octet field. Note that this is a request for a total amount of time rather than a recurring use of time in the superframe. The use of this field is defined in 7.5.2.

6.5.6.2 Channel Time Response command

The Channel Time Response command Payload field shall be formatted as illustrated in Figure 6-153.

Octets: 1	1	1	1
Stream Request ID	Stream Index	Available Number Of TUs	Reason Code

Figure 6-153—Channel Time Response command Payload field format

The Stream Request ID field is defined in 6.5.6.1.

The Stream Index field is defined in 6.5.6.1.

The Available Number of TUs field is used by the PNC to indicate to the requesting DEV the number of TUs per CTA Rate Factor it has assigned to the requested isochronous stream. In the case of a super-rate allocation, it is the number of TUs assigned in each superframe. In the case of a sub-rate allocation it is the number of TUs assigned in each of the sub-rate superframes.

For isochronous channel time requests, if the Available Number of TUs is greater than or equal to the Minimum Number of TUs requested and less than or equal to the Desired Number of TUs requested, then the requesting DEV is informed that there is channel time available. If, however, the Available Number of TUs field is less than the Minimum Number of TUs requested, then the requesting DEV is informed that the PNC is unable to fulfill the DEV's request for channel time. In this case, the Available Number of TUs will be set by the PNC to the number of TUs that the PNC would have been able to allocate for this request, as described in 7.5.1.1.

For asynchronous stream requests, the response frame is sent only if the PNC is unable to fulfill the request, in which case the available number of TUs field is set to zero.

The Reason Code field indicates whether a channel time request was successful or unsuccessful. The codes assignable to this field are as follows:

- 0 → Success
- 1 → Success, DEV in PS mode
- 2 → Target DEV unassociated
- 3 → Target DEV not a member
- 4 → Priority unsupported
- 5 → Stream terminated by PNC
- 6 → Stream terminated by target DEV
- 7 → Channel time unavailable

- 8 → Destination DEV in power save mode
- 9 → Unable to allocate as pseudo-static CTA
- 10 → Superframe overloading
- 11 → Requested super-rate or sub-rate unsupported
- 12 → Request denied
- 13 → PNC handover in progress
- 14 → STP expired
- 15–254 → Reserved
- 255 → Other failure

6.5.7 Channel status commands

This group of commands is used to request and provide information about the remote DEV's view of the channel and to change the transmitter power based on the current channel conditions.

6.5.7.1 Channel Status Request command

The Channel Status Request command may be sent by any DEV in the piconet to any other DEV in the piconet, including the PNC, to request the current channel condition as experienced at the target DEV. This command may also be sent by the PNC as a broadcast frame, i.e., the DestID set to the BcstID. The Channel Status Request command has no Payload field.

6.5.7.2 Channel status response

The Channel Status Response command is sent by the target DEV in response to the originating DEV's request to let the originating DEV know the current channel condition at the target DEV. When the DestID of this command is the PNCID, the values in the command shall correspond to all frames exchanged by the DEV with other DEVs in the piconet. When the DestID of this command is a non-PNC DEVID, the values in the command shall correspond to the frames exchanged between the requesting DEV and the target DEV. The Channel Status Response command Payload field shall be formatted as illustrated in Figure 6-154.

Octets: 2	2	2	2	2
Measurement Window Size	TX Frame Count	RX Frame Count	RX Frame Error Count	RX Frame Loss Count

Figure 6-154—Channel Status Response command Payload field format

The Measurement Window Size field is the number of superframes during which the measurements were taken. The minimum Measurement Window Size for a valid measurement for this command shall be 2 superframes. A Measurement Window Size of zero indicates that the responding DEV does not provide channel status statistics.

The TX Frame Count field contains the total number of frames, not including Imm-ACK frames, that were transmitted by the sender of this command to the destination of this command. This count includes all transmission attempts, including retransmissions of the same frame.

The RX Frame Count field contains the total number of frames, not including Imm-ACK frames, that were correctly received by the sender of this command, as described in 7.1. Only the directed frames transmitted by the destination of this command intended for the sender of this command are included.

The RX Frame Error Count field contains the total number of frames, not including Imm-ACK frames, that were received in error by the sender of this command from the destination of this command. A frame is

considered to have been received in error if the header is correctly received but the frame is not correctly received, as described in 7.1.

The RX Frame Loss Count field contains the number of frames in streams with the ACK Policy field set to no-ACK, not including Imm-ACK frames, that were determined by the originator of the command to have been lost. The originating DEV determines this for a particular stream index by observing gaps in the Fragmentation Control field of received frames. These numbers are accumulated for all streams between the originating DEV and the target DEV.

6.5.7.3 Remote Scan Request command

The SrcID for Remote Scan Request command shall be the PNCID. The Remote Scan Request command Payload field shall be formatted as illustrated in Figure 6-155.

Octets: 1	...	1
Channel 1	...	Channel <i>n</i>

Figure 6-155—Remote Scan Request command Payload field format

The Channel Number field indicates the channels that are to be scanned. The mapping of the channel number is PHY dependent. For the 2.4 GHz PHY, the mapping is defined in 10.2.3.

6.5.7.4 Remote Scan Response command

The DestID for the Remote Scan Response command shall be the PNCID. The Remote Scan Response command Payload field shall be formatted as illustrated in Figure 6-156.

Octets: 1	1	variable	1	variable	variable
Reason Code	Number Channels	Channel Rating List	Number Piconets	Remote Piconet Description Set	IE

Figure 6-156—Remote Scan Response command Payload field format

The allowed Reason Code field values are as follows:

- 0 → Success
- 1 → Request denied
- 2 → Invalid channel requested
- 3–254 → Reserved
- 255 → Other failure

If the request is denied, then the Remote Scan Response command shall include only the Command Type, Length, and Reason Code fields.

The Number Channels field indicates the number of channels that were scanned by the remote DEV.

The Channel Rating List field contains a list of channel indices ordered from best (least interference) to worst (most interference). The Channel Rating List field shall be formatted as shown in Figure 6-157.

Octets: 1	...	1
Best Channel Index	...	Worst Channel Index

Figure 6-157—Channel Rating List field format

The Number Piconets field indicates the number of piconets that were found. If the DEV did not find any piconets, then the number shall be set to zero and there shall be no remote piconet description sets in the command.

The Remote Piconet Description Set is a collection of one or more Remote Piconet Description fields. Each Remote Piconet Description field shall be formatted as shown in Figure 6-158.

Octets: 2	1	1	1	6	2–34	2–40
PNID	Scanned Frame Type	Channel Index	Piconet Type	PNC Address	BSID IE	Parent Piconet IE

Figure 6-158—Remote Piconet Description field format

The PNID is the identifier of the piconet that was found by the DEV.

The Scanned Frame Type field indicates the type of frame that was received by the DEV with the piconet information. The allowed values are as follows:

- 0 → The DEV found the PNID in a beacon.
- 1 → The DEV found the PNID only in a non-Beacon frame.
- 2–255 → Reserved.

The Channel Index field indicates the PHY channel where the information was found. The mapping of the channel number is PHY dependent. For the 2.4 GHz PHY, the mapping is defined in 10.2.3.

The Piconet Type field indicates the type of piconet that was found. The allowed values are as follows:

- 0 → Non-dependent piconet
- 1 → Dependent piconet
- 2–255 → Reserved

The PNC Address field is defined in 6.3.1. If the DEV found a beacon, it shall put the PNC Address from the beacon into the command. Otherwise, the PNC Address field shall be set to zero.

The BSID IE is defined in 6.4.2. If the DEV found a beacon, it shall put this IE into the command. Otherwise, it shall include a BSID IE with zero length, i.e., only the element ID and length fields.

The Parent Piconet IE is defined in 6.4.3. If the DEV found a beacon from a dependent piconet, it shall put this IE into the command. Otherwise, it shall include a Parent Piconet IE with zero length, i.e., only the element ID and length fields.

The IE field is provided for future expansion of the standard. The IE is not defined in this revision of the standard, and so the source DEV may omit the IE in the Remote Scan Response command. The IE shall be ignored upon reception by the destination DEV of this command.

6.5.7.5 Transmit Power Change command

The Transmit Power Change command is used to request a change in the transmit power of a DEV. The Transmit Power Change command Payload field shall be formatted as illustrated in Figure 6-159.

Octets: 1
TX Power Change

Figure 6-159—Transmit Power Change command Payload field format

The TX Power Change field contains the requested TX power level change in dB at the destination DEV in twos complement format. For example, a +2 dB change in the TX power level is 0x02 while a –2 dB TX power level change is encoded as 0xFE.

6.5.8 Power management commands

These commands are used to enable DEVs to conserve power as well as by other DEVs that want to know when DEVs using power management will be available for communication.

6.5.8.1 PS Set Information Request command

The PS Set Information Request command is used to acquire information from the PNC regarding the number of PS sets and their structure. The PS Set Information Request command has no Payload field.

6.5.8.2 PS Set Information Response command

The PS Set Information Response command Payload field shall be formatted as illustrated in Figure 6-160.

Octets: 1	1	8–39	...	8–39
Max Supported PS Sets	Number Current PS Sets	PS Set Structure 1	...	PS Set Structure <i>n</i>

Figure 6-160—PS Set Information Response command Payload field format

The Max Supported PS Sets field indicates the number of PS sets supported by the PNC of this piconet.

The Number Current PS Sets field is a count of the number of PS Set Structures in this command as well as the number of currently active PS sets in the piconet.

Each PS Set Structure shall be formatted as illustrated in Figure 6-161.

Octets: 1	2	2	1	1	1–32
PS Set Index	Wake Beacon Interval	Next Wake Beacon	Bitmap Length	Start DEVID	DEVID Bitmap

Figure 6-161—PS Set Structure field format

When the PS Set Index field is zero, the DEVID Bitmap field lists the DEVs currently in APS mode, if any. When the PS Set Index field is one, the DEVID Bitmap field indicates the DEVs currently in PSPS mode, if any. When the PS Set Index field is any value between 0x02 and 0xFD, inclusive, the DEVID Bitmap field indicates the DEVs currently in this particular DSPS set. The PS set indices are defined as follows:

0x00 → APS set
 0x01 → PSPS set
 0x02–0xFD → DSPS sets
 0xFE → Unallocated DSPS set
 0xFF → Reserved

The Wake Beacon Interval field is defined in 6.5.8.3. This field is set to the system wake beacon interval for PS sets 0 and 1. For all other PS sets it is set to the wake beacon interval of that DSPS set. Note that the wake beacon interval has no interpretation for PS set 0, as described in 7.14.3.

The Next Wake Beacon field is defined in 6.5.8.4. This field is set to the next system wake beacon for PS sets 0 and 1. For all other PS sets it is set to the next wake beacon of that DSPS set. Note that the Next Wake Beacon field has no interpretation for PS set 0, as described in 7.14.3.

The Bitmap Length field contains the number of octets in the DEVID bitmap. This field shall take on values from 1 to 32, inclusive.

The Start DEVID field indicates the DEVID corresponding to the LSB in the DEVID bitmap.

The DEVID Bitmap field is a bitmap of the DEVIDs in a specific PS set. A value of 0 in a bitmap position indicates that the DEV corresponding to that DEVID is not part of the PS set. A value of 1 in a bitmap position indicates that the DEV corresponding to that DEVID is in the PS set.

6.5.8.3 SPS Configuration Request command

The SPS Configuration Request command is used to set up and manage SPS set memberships for SPS DEVs currently participating or requesting to participate in one or more SPS modes. The SPS Configuration Request command Payload field shall be formatted as illustrated in Figure 6-162.

Octets: 1	1	2
Operation Type	SPS Set Index	Wake Beacon Interval

Figure 6-162—SPS Configuration Request command Payload field format

The Operation Type field indicates whether a DEV is requesting either to join or to leave an existing SPS set. The valid operation types are as follows:

- 0 → Join
 1 → Leave
 2–255 → Reserved

The SPS Set Index field is used to identify the SPS set the requesting DEV wants to create/join, join, configure, or leave. The SPS Set Index field shall not be set to zero (APS) in this command.

Table 6-26 lists the interpretation of the fields for various combinations of the Operation Type field and SPS Set Index field.

Table 6-26—SPS Configuration Request command parameter entries

Operation type	SPS set index value	Wake beacon interval	Comments
0 (create/join)	Unallocated DSPS set (0xFE)	Any valid DSPS wake beacon interval value. This value is decoded by the PNC upon reception.	The unallocated DSPS Set Index (0xFE) shall be used to request the PNC to establish a new DSPS set.
0 (join)	0x01	Any valid PSPS system wake beacon interval value. This value is decoded by the PNC upon reception.	The PSPS mode is permanently associated with PS set index 0x01.
0 (join)	0x02–0xFD	Shall be set to zero and ignored by the PNC upon reception.	The DEV is requesting to join an existing DSPS set.
1 (leave)	0x01–0xFD	Shall be set to zero and ignored by the PNC upon reception.	The DEV is requesting to leave the indicated SPS set.

The Wake Beacon Interval field contains the number of superframes, including the current one, between wake beacons, as described in 7.14. For example, a wake beacon interval of 8 indicates that the DEV is requesting a wake beacon every 8th beacon, Figure 7-63. Valid values for the wake beacon interval for either the PSPS or DSPS ranges are in powers of 2 (e.g., 2, 4, 8, ...). Furthermore, the wake beacon interval shall have a value between 2 and 256 for PSPS and between 2 and 65536 for DSPS. Because the value 65536 cannot be represented with 2 octets, a wake beacon interval of 0 shall represent the interval value 65536.

6.5.8.4 SPS Configuration Response command

The SPS Configuration Response command is sent by the PNC as a response to an SPS Configuration Request command received from a DEV. The SPS Configuration Response command Payload field shall be formatted as illustrated in Figure 6-163.

Octets: 1	1	2
Reason Code	SPS Set Index	Next Wake Beacon

Figure 6-163—SPS Configuration Response command Payload field format

The Reason Code field contains the result of the SPS Configuration Request command. The valid reason codes are as follows:

- 0 → Success

- 1 → Already member
- 2 → Invalid SPS set (i.e., attempting either to configure PS set 0 or to join a non-existing set)
- 3 → Set creation failed
- 4 → Unique Wake Beacon Interval required
- 5 → DSPS set deleted by PNC
- 6 → PNC handover in progress
- 7–254 → Reserved
- 255 → Other failure

The Reason Code field is set to zero (Success), if the create or join operation is successful. If the SPS Set Index field has been set to zero or any value not representing an SPS set, the reason code shall be set to "Invalid SPS Set." If a DEV requests to join a DSPS set where it is already a member, the reason code shall be set to "Already member." However, a DEV is allowed to make multiple requests to join the DSPS set. This has the effect of updating the DEV's desired system wake beacon interval value. If a DSPS set creation fails for any other reason than listed above, the reason code shall be set to "Set creation failed."

The SPS Set Index is defined in 6.5.8.3.

The Next Wake Beacon field indicates the beacon number, as described in 6.3.1.1, of either the next system wake beacon for the PSPS mode when the SPS set index is set to one, or the next DSPS wake beacon when the SPS set index is greater than or equal to 0x02 and less than or equal to 0xFD.

6.5.8.5 PM Mode Change command

The Power Management (PM) Mode Change command Payload field shall be formatted as illustrated in Figure 6-164.

Octets: 1
PM Mode

Figure 6-164—PM Mode Change command Payload field format

The PM Mode field shall be set as follows:

- 0 → ACTIVE mode
- 1 → APS mode
- 2 → SPS mode
- 3–255 → Reserved

6.5.8.6 PM Mode Change Response command

The PM Mode Change Response command is sent by the PNC as a response to a PM Mode Change Request command received from a DEV. The PM Mode Change Response command Payload field shall be formatted as illustrated in Figure 6-165.

Octets: 1	1
PM mode	Reason code

Figure 6-165—PM Mode Change Response command Payload field format

The PM Mode field is defined in 6.5.8.5 and is set by the PNC to indicate current PM mode of the DEV.

The Reason Code field indicates the result of the attempt by the DEV to change PM modes. The valid reason codes are as follows:

- 0 → Success
- 1 → Not a member of any existing SPS sets
- 2–254 → Reserved
- 255 → Other failure

6.5.9 Special commands

6.5.9.1 Security Message command

The Security Message command is used to send security-related information to another DEV in the piconet. The SEC field in the Frame Control field shall be set to zero. The Security Message command Payload field shall be formatted as illustrated in Figure 6-166.

Octets: 3	variable
Unique ID	Security Information

Figure 6-166—Security Message command Payload field format

The Unique ID is defined in 6.4.7.

The Security Information field contains security-related information whose format is defined by the entity indicated in the Unique ID field. Its use by a DEV is outside of the scope of this standard.

6.5.9.2 Vendor Defined command

The Vendor Defined command Payload field shall be formatted as illustrated in Figure 6-167.

Octets: 3	variable
Unique ID	Vendor Defined Data

Figure 6-167—Vendor Defined command Payload field format

The Unique ID field is defined in 6.4.7.

The format of the Vendor Defined Data field is defined by the entity indicated in the Unique ID field. Its use by a DEV is outside of the scope of this standard.

6.5.9.3 AS IE Request command

The AS IE Request command is used to send an AS IE to the PNC to be put in the beacon. The DestID shall be set to the PNCID. The AS IE Request command Payload field shall be formatted as illustrated in Figure 6-168.

Octets: 1	1	1	variable
Request ID	Request type	AS IE Index	AS IE

Figure 6-168—AS IE Request command Payload field format

The Request ID field is an identifier generated by the originating DEV that is unique among the DEV's AS IE requests.

The Request Type field indicates the type request. The valid values are as follows:

- 0 → Add
- 1 → Modify
- 2 → Remove
- 3–255 → Reserved

The AS IE Index field is assigned by the PNC and is used to identify the AS IE for the request. If the Request Type field is "add," then the AS IE Index field shall be ignored by the PNC.

The AS IE field is defined in 6.4.7. If the Request Type field is "remove," then the length of Application Specific Data field in the AS IE shall be zero.

6.5.9.4 AS IE Response command

The AS IE Request command is used to respond to a request to put an AS IE in the beacon. The SrcID shall be set to the PNCID. The AS IE Request command Payload field shall be formatted as illustrated in Figure 6-169.

Octets: 1	1	1
Request ID	AS IE Index	Reason Code

Figure 6-169—AS IE Response command Payload field format

The Request ID field is defined in 6.5.9.3.

The AS IE Index field is defined in 6.5.9.3.

The valid values of the Reason Code field are as follows:

- 0 → Success
- 1 → Request rejected
- 2 → Unknown AS IE index
- 2–254 → Reserved
- 255 → Other failure

6.5.10 Multicast configuration commands

6.5.10.1 Multicast Configuration Request command

The Multicast Configuration Request command is used to request a GrpID, 7.5.3. The DestID shall be set to the PNCID. The Multicast Configuration Request command Payload field shall be formatted as illustrated in Figure 6-126.

Octets: 1	6
Action	Group Address

Figure 6-170—Multicast Configuration Request command Payload field format

The valid values of the Action field are as follows:

- 0 → Join
- 1 → Leave
- 2–255 → Reserved

The Group Address field is defined in 6.4.17.

6.5.10.2 Multicast Configuration Response command

The Multicast Configuration Response command is used by the PNC to respond to a request for a GrpID, as described in 7.5.3. The SrcID shall be set to the PNCID. The Multicast Configuration Response command Payload field shall be formatted as illustrated in Figure 6-171.

Octets: 6	1	1
Group Address	GrpID	Reason code

Figure 6-171—Multicast Configuration Response command Payload field format

The Group Address field is defined in 6.4.17.

If the request for a GrpID was successful, the GrpID field is the DEVID, as defined in 6.2.3, that has been assigned by the PNC for the address in the Group Address field. Otherwise, the GrpID field shall be set to zero.

The valid values of the Reason Code are as follows:

- 0 → Success
- 1 → Failure, lack of DEVIDs
- 2 → Failure, handover in progress
- 3 → Failure, resources unavailable
- 4 → Failure, not a valid group address
- 5–254 → Reserved
- 255 → Other failure

7. MAC functional description

7.1 Introduction

This clause provides a description of the MAC functionality.

In this clause, unless otherwise indicated, receiving a frame means that the PHY has successfully received a frame over the medium and both the FCS and HCS calculations match their respective data, as defined in 6.2.6 and 10.2.9.

Asynchronous MSDUs shall be delivered to the FC SL in the order of reception.

An example message sequence chart (MSC) is shown in Figure 7-1, which illustrates three MLME requests and the associated timeouts. In the first case, the request completes before the timeout expires, and so the confirm returns with the ResultCode set equal to SUCCESS. In the second case, the requested action completes unsuccessfully before the timer expires; therefore, the confirm primitive is returned with the ResultCode set equal to FAILURE, and the ReasonCode indicates the reason for the failure, if known. In the third case, the requested action does not complete before the timeout expires, and so the confirm primitive is returned with the ResultCode set equal to FAILURE and the ReasonCode set equal to TIMEOUT.




Figure 7-1—MSC showing examples of primitive timers

Throughout this clause, some of the procedures and MSCs are written as though the optional MLME SAP, 5.3, is exposed and is supported by the MLME. Unless stated otherwise, the procedures initiated by an MLME primitive may also be initiated internally in the MAC. For procedures not shown to be initiated by an MLME primitive, the procedure is executed by the MAC without interaction with the DME.

For a MAC header to be correctly received by the MAC, the PHY first verifies the HCS. In addition, the MAC header shall have a protocol revision supported by the MAC; have a DestID equal to a DEVID, BcstID, McstID, or when applicable the PNCID or UnassocID; and have a PNID equal to the PNID of the

piconet with which the DEV is synchronized. Because the FCS validation is not required, it is possible for the MAC header to be correctly received even if the frame is not correctly received.

For a frame to be correctly received by the MAC, the MAC header shall be correctly received and the frame payload shall pass the FCS validation. If a DEV correctly receives a frame from an unassociated DEV, it may ignore the frame and may choose not to respond to the frame. If secure membership is required in the piconet and a DEV correctly receives a frame from a DEV that is not a member of the piconet, it shall ignore the frame and shall not respond to the frame, except for the ACK, if the ACK Policy field is set to Imm-ACK, Imp-ACK, or Dly-ACK Request.

7.2 Starting, maintaining, and stopping piconets

An IEEE 802.15.3 piconet begins when a PNC-capable DEV takes on the responsibility of being the PNC. This subclause describes the processes involved in starting and maintaining the piconet. The types of piconets defined in this standard are as follows:

- Independent piconet: A piconet with no dependent piconets and no parent piconets.
- Parent piconet: A piconet that has one or more dependent piconets.
- Dependent piconet: A piconet that requires a time allocation in another piconet, called the parent piconet, and is synchronized with the parent piconet's timing.

There are two types of dependent piconets. They are as follows:

- Child piconet: A dependent piconet where the PNC is a member of the parent piconet.
- Neighbor piconet: A dependent piconet where the PNC is not a member of the parent piconet.

7.2.1 Scanning through channels

All DEVs shall use passive scanning to detect an active piconet. That is, DEVs shall be in receive mode for a period of time in a channel no less than *mMinChannelScan* to look for Beacon frames or, if supported, Sync frames from a PNC. If a particular BSID, PNID, or PNC address to scan for is not specified with an MLME-SCAN.request, the DEV searches for any Beacon frame or, if supported, a Sync frame. If a particular BSID, PNID, or PNC address to scan for is specified, the DEV shall ignore all received frames not matching the parameter or parameters contained in the request.

In addition, the searching DEV shall collect statistics on each channel scanned and save them in the ChannelRatingList, as described in 5.3.2.

DEVs search for piconets by traversing through all available PHY channels. A DEV may search the channels in any order as long as all valid channels are contained in the search pattern. The result of a scan shall include information on any parent, child, as described in 7.2.7, or IEEE 802.15.3 neighbor, as described in 7.2.8, piconets that were detected. This provides a complete inventory of each channel.

While searching, if any frame type other than a Beacon frame is received, the searching DEV shall stay in the channel for a minimum of *mMinChannelScan* from the time of reception of first frame and look for a beacon from the PNC. The DEV shall scan all indicated channels to find piconets before returning the scan information via the MLME-SCAN.confirm primitive. The DEV shall report only piconets found due to the reception of a Beacon frame or, if supported, a Sync frame as a part of the MLME-SCAN.confirm primitive.

Figure 7-2 illustrates the message flow for a successful scan operation.




Figure 7-2—MSC for scan operation

7.2.2 Starting a piconet

A DEV that is instructed to start a piconet through MLME-START.request, as described in 5.3.3.1, shall try only to start its own piconet and shall not attempt to associate with an existing piconet.

The MAC should choose the channel with the least amount of interference to start the piconet based on the results of a recent scan (either DME initiated or MAC initiated).

Once the MAC has received the MLME-START.request primitive, it shall listen to the channel for *mMinChannelScan* duration to determine if the channel is still clear. If, at the end of this listening period, the MAC determines that the channel is clear, the DEV, now the PNC, shall commence broadcasting its beacon once every superframe duration. When the piconet starts, the PNC allocates an additional DEVID to itself for the purposes of exchanging data with other DEVs that become members of the established piconet.

Once a PNC has established a piconet, the PNC should periodically allocate channel time in the CTAP so that there is quiet time for it to scan the channel for other piconets. If the PNC detects another piconet in the same channel that is not one of its own child or neighbor piconets, it may take action to improve coexistence with the other piconet. Some of the actions the PNC may take include the following:

- Changing to a different channel, as described in 7.12.1.
- Becoming a child or neighbor piconet of the other piconet, as described in 7.2.7 and 7.2.8.
- Reducing the piconet's transmit power, as described in 7.12.2.

The PNC shall continue in its role for as long as it desires or until the PNC determines that an associating DEV is more capable, at which time the PNC may decide to initiate the PNC handover procedure, as described in 7.2.3.

If the MAC determines that no channels are available, it will respond with an MLME-START.confirm with a ResultCode of FAILURE and ReasonCode of NO_CHANNELS_AVAILABLE.

Figure 7-3 illustrates the message flow for a successful start operation.




Figure 7-3—MSC for starting a piconet

7.2.3 PNC handover

When the PNC leaves the piconet or when it transfers its PNC functionality to another DEV, the PNC may choose a DEV that is capable of being a PNC as its successor. PNC handover is optional for a PNC using a mmWave PHY. The PNC Capable field in the PNC Capabilities field, as described in 6.4.11, is used to indicate that a DEV is capable of being a PNC. The PNC shall use the information in the PNC Capabilities field of the other DEVs in the piconet with the evaluation criteria defined in Table 7-1 to select the most qualified PNC-capable DEV that is currently a member of the piconet to be the new PNC. The PNC shall send a PNC Handover Request command, as described in 6.5.3.1, to its chosen DEV with the parameters specified in 6.5.3.1. If the piconet is not a dependent piconet, the DEV shall accept the nomination and be prepared to receive the piconet information records. If the DEV is currently the PNC of a dependent piconet, it may refuse the request by sending a PNC Handover Response command to the PNC with the Reason Code field set to “Handover refused, unable to act as PNC for more than one piconet.” If both the current and the new PNC are members of the same dependent piconet, then the DEV shall accept the handover request unless it is unable to join the parent piconet as either a regular DEV or a neighbor PNC. In the case where the DEV is unable to join the parent piconet, the DEV sends the PNC Handover Response command to the PNC with the Reason Code field set to “Handover refused, unable to join parent piconet.”

Table 7-1—Comparison order of fields for PNC handover

Order	Information	Note
1	PNC Des-mode field in PNC capabilities field	PNC Des-mode=1 is preferred
2	SEC field in PNC capabilities field	SEC=1 is preferred
3	PSRC field in PNC capabilities field	PSRC=1 is preferred
4	Maximum associated DEVs	Higher value is preferred
5	Maximum channel time requests	Higher value is preferred
6	Transmitter power level (PHY dependent)	Higher value is preferred
7	Maximum PHY rate (PHY dependent)	Higher value is preferred
8	DEV address	Higher value is preferred

When the handover is initiated, the HandoverStatus is STARTED. If the handover timer expires, a PNC Handover Request command shall be sent to the DEV with a HandoverStatus of CANCELLED. In addition, if the DEV sees a PNC Shutdown IE from the PNC during the handover process, it knows that the handover was canceled.

The PNC may allocate channel time with the chosen PNC-capable DEV as the destination for the purpose of transferring information about the DEVs in the piconet and their current channel time requests. The PNC shall first send a PNC Information command, as described in 6.5.4.2, to the chosen PNC-capable DEV. In the PNC Information command, the PNC shall include all DEVs that are associated in the piconet, including any associated neighbor PNCs, the DEV personality of the PNC, and an entry for the PNCID. Once the PNC has successfully sent this command, it shall then begin sending all of the current channel time requests to the chosen PNC-capable DEV using a PNC Handover Information command, as described in 6.5.3.3. Once the PNC has successfully sent the PNC Handover Information command, it shall send a PS Set Information Response command, as described in 6.5.8.2, to the new PNC. The PNC may fragment the PNC Information, PNC Handover Information, and PS Set Information Response commands using the process described in 7.7.

The PNC Handover Information command shall not be sent if the PNC has indicated in the PNC Handover Request command that it does not have any Channel Time Request Block fields to transfer. The PS Set Information Response command shall not be sent if the PNC has indicated in the PNC Handover Request command that it does not have any PS sets to transfer.

The handover procedure will transfer all information necessary for the new PNC to take over except for the following:

- Asynchronous Channel Time Request Block fields will not be transferred. All DEVs with asynchronous data to send need to send a new Channel Time Request command, as described in 6.5.6.1, to the new PNC after it has sent its first beacon.
- CTA locations are not transferred, except in the preceding beacons.

The current PNC shall not accept any new commands that would change any of the information that will be transferred to the new PNC once the PNC handover process has been initiated. The PNC shall refuse these requests with the Reason Code field set to "Handover in progress" for commands that have this reason code defined.

Once the chosen PNC-capable DEV has received the required information from the current PNC, it shall respond to the current PNC with a PNC Handover Response command, as described in 6.5.3.2. This will signal to the current PNC that the chosen PNC-capable DEV is ready to take over as the new PNC.

After the PNC receives the PNC Handover Response command with Reason Code field indicating success, it shall place a PNC Handover IE, as described in 6.4.9, in the beacon with the Handover Beacon Number field set to the beacon number of the superframe in which the new PNC will send its first beacon. Upon receiving the PNC Handover IE, the chosen PNC-capable DEV will prepare to broadcast its first beacon as the new PNC. After sending the last beacon, the old PNC relinquishes control of the piconet and stops generating beacons. The new PNC shall broadcast its first beacon at the time the beacon would have been sent by the old PNC. This time may vary from the actual time due to clock inaccuracies of the old and new PNCs. The new PNC shall start sending beacons with the time token counter set to one more than the time token of the last beacon that will be sent by the old PNC. The new PNC shall begin using the PNCID as the SrcID for all beacon or command frames transmitted. The new PNC shall use the PNCID or its previously assigned DEVID as the SrcID for all data frames transmitted. When the PNC handover is successful, the association of the remaining DEVs with the piconet is unaffected and hence they are not required to reassociate with the new PNC.

The PNC shall ensure that the handover announcement complies with the rules for beacon announcements in 7.6.4.

Figure 7-4 illustrates the message sequence of a PNC handover to a PNC-capable DEV that is currently a member of the piconet.




Figure 7-4—PNC handover MSC

The mandatory and optional information transfer for PNC handover is illustrated in Figure 7-5.

The optional security information handover referenced in Figure 7-4 is shown in Figure 8-2.

The DME initiates the handover process using MLME-STOP.request with RequestType set to HANDOVER. This process is illustrated in Figure 7-6. The DME may choose the target DEV or DEVs for the handover or allow the PNC to determine the target DEV. If the handover completes successfully, the MLME-STOP.confirm primitive is generated with a ResultCode set to SUCCESS. If the handover does not successfully complete within the time period specified by the DME, the PNC shall perform the PNC shutdown process defined by 7.2.9. After completion of the shutdown process, the MLME-STOP.confirm primitive is generated with the ResultCode set to FAILURE and the ReasonCode set to HANDOVER_FAILED.

Note that the PNC handover process should not stop any of the isochronous data connections. Figure 7-4 illustrates only the handover process and hence does not show other traffic. However, since the PNC needs to allocate sufficient channel time to transfer the DEV and channel time request data, some of the data traffic may be affected depending on the traffic conditions within the piconet.




Figure 7-5—Handover information transfer MSC

When a DEV joins a piconet, the PNC shall compare the PNC Capabilities field of the new DEV to its own. If the PNC Des-Mode field is set in the new DEV but not in the current PNC, the current PNC shall perform PNC handover. If the new DEV is more qualified to be the PNC, based on the PNC selection criteria in Table 7-1, the PNC may perform PNC handover.

As Table 7-1 shows, PNC Des-Mode is the top priority field in the PNC selection criteria. Since the PNC Des-Mode is the highest priority, a DEV with this field set is more likely to become the PNC of the piconet. Thus, this field should be set if it is desirable for the DEV to be the PNC of the piconet. If only one DEV has the PNC Des-Mode field set, then that DEV would become the PNC.

If the piconet is using mode 1 security, then the new PNC and DEVs need to follow the security procedures in 8.3.1 in addition to the handover process described here.

A dependent PNC receiving a parent beacon with a PNC Handover IE may immediately insert the Piconet Parameter Change IE into its beacons with the Change Type field set to MOVE, as described in 6.4.6, and the Superframe Timing field set to zero. A member of a piconet that receives this Piconet Parameter Change IE in the beacon from the PNC shall not transmit after the superframe that has a beacon number equal one less than the Change Beacon Number field in the Piconet Parameter Change IE until it has correctly received a beacon from its PNC, as described in 7.1. This requirement applies to DEVs in independent as well as dependent piconets.

7.2.4 Preliminary handover

The Preliminary Handover procedure shall not be used by a dependent PNC.

If a DEV has been chosen as next PNC, the PNC shall periodically send handover information to that DEV so that it has the information it needs to take over as PNC. This sending of handover information is known as preliminary handover. In order to send the handover information, the PNC shall send the Handover Request

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 7-6—DME initiated PNC handover

command with the Handover Status field set to Preliminary Handover. The DEV receiving this frame shall store the information for use at a later time, but not take over as PNC at this time.

Information transfer shall take place as shown in Figure 7-5. In addition, any security information may be transferred, as defined in 8.4.1.

When all the information has been received, the next PNC shall send the Handover Response command with Reason Code set to “Success, Ready for Handover.”

The current PNC may issue a handover based on previously transferred information by sending the Handover Request command with the Handover Status field set to “Handover Based on Previous Information.” Once the chosen PNC-capable DEV is prepared to start operating as PNC, it shall respond to the current PNC with a PNC Handover Response command with Reason Code set to “Success, Ready for Handover.”

The handover process now proceeds, as described in 7.2.3, from the point where the Handover Response command is sent.

7.2.5 Next PNC

The Next PNC procedure shall not be used for dependent piconets.

Although the PNC will normally attempt to perform a handover when it shuts down or leaves the piconet, this is not always possible. The PNC could be powered off in such a way that handover is not performed. The PNC could also move out of range of the rest of the DEVs in the piconet, or vice versa. In order to avoid the interruption in communication that would occur when the PNC disappears without handover, the PNC may choose a DEV in the piconet to be the next PNC.

The PNC should periodically evaluate the PNC capabilities of the members of the piconet and other factors to select a DEV to be the next PNC. In addition to the PNC capabilities field, the PNC should use information provided by the DEVs in the Piconet Channel Status IE, as described in 6.4.21, to determine which potential PNC could best be heard by the other DEVs in the piconet. In the context of the Piconet Channel Status IE, "heard" means correctly receiving a MAC header.

A DEV should continuously determine the DEVs that it can hear while it is in the AWAKE state, as defined in 7.14. If MCTAs are used in the piconet, DEVs that are not extremely power sensitive should listen to all of the uplink MCTAs.

The Next PNC IE, as described in 6.4.20, should be periodically announced in the beacon according to the beacon information announcement procedure, as defined in 7.6.4.

The PNC also selects the Next PNID for the Next PNC IE. The PNC shall select a Next PNID such that it is different from any other PNID that has been detected, as described in 7.11.3.

If no other DEV in the piconet is PNC capable, the PNC shall set the Next PNC field to zero and the Next PNID field to the current PNID in the Next PNC IE.

If a DEV has been chosen as next PNC, the PNC shall periodically perform preliminary handover, as defined in 7.2.4.

When the DEV identified as the next PNC fails to detect $mMaxLostBeacons+1$ consecutive beacons, it shall scan the channel for any frames. Since no frames are transmitted in dynamic CTAs or the CAP if the beacon was not correctly received, the next PNC shall use the presence of any frames from the current piconet as an indication that the beacon is still being transmitted but the next PNC has lost contact with the PNC. In this case the next PNC shall not take over the role of PNC for that piconet. If no MAC headers are correctly received during the 2 superframe duration scan, the next PNC shall begin sending out beacons using the same superframe duration and channel. The process of the next PNC taking over as PNC is known as implicit handover. There is a possibility that the next PNC has left the range of all of the other DEVs in the piconet as opposed to the original PNC going out of range or powering off.

The next PNC shall start its beacon after the time it would have expected the beacon from the previous PNC so that if they are both in range of some DEVs, their beacons will not collide. The next PNC should choose a position for the beacon that was unused based on the previous CTAs, to minimize the probability of collision. In order to reduce the possibility of two piconets in close range using the same PNID, the next PNC shall use the announced Next PNID as the PNID of the piconet.

DEVs that are not identified as the next PNC should store the Next PNID from the beacon. When the DEVs that are not identified as the next PNC fail to detect $mMaxLostBeacons$, they should start scanning for beacons. If they do not see beacons with the current PNID, but see beacons with the PNID field set to the Next PNID and the PNC Address field set to the DEV address that corresponds to the DEVID of the Next PNC, they are automatically associated with that piconet.

7.2.6 Dependent PNC handover

The dependent PNC handover process begins in the same manner as a regular PNC handover, as described in 7.2.3, with the current PNC sending a PNC Handover Request command to the target DEV that it has

selected to become the new PNC, as shown in Figure 7-7. In this and the two subsequent figures, the identities PNC, DEV-2 and DEV-3 are all relative to the dependent piconet and not the parent piconet. If the target DEV is not a member of the parent piconet, then that DEV shall begin the association process to join the parent piconet and, if required, become a secure member of the parent piconet. The target DEV may request to associate with the parent piconet as either a neighbor PNC or a member of the piconet. While the target DEV is attempting to join the parent piconet, the current dependent PNC shall send the target DEV the information about all of the DEVs with a PNC Information command, all of the current channel time requests with a PNC Handover Information command and the power save information, if any, using a PS Set Information Response command. The target DEV may also request the transfer of any security information at this point using a Security Information Request command. Note that the transfer of this information will not interfere with the target DEV's association because the former occurs only during the time reserved for the dependent piconet while the latter occurs only during the time reserved for the parent piconet.




Figure 7-7—Successful PNC handover in a dependent piconet

Figure 7-7 references three processes not included in the figure. The MSC for the optional security information handover is shown in Figure 8-2. The association process that is used to become member of the parent piconet is shown in Figure 7-15 while the stream modification request to handover control of a CTA is shown in Figure 7-34.

Once the transfer of the information is complete and the target DEV has joined the parent piconet, the target DEV shall send a PNC Handover Response command to the dependent PNC with a Result Code set to the DEVID that was assigned to it by the parent PNC. This informs the dependent PNC that the target DEV is

ready to take over control of the piconet. At this point, the dependent PNC shall send a Channel Time Request command to the parent PNC to handover the control of the dependent piconet CTA to the new dependent PNC, as described in 7.5.1.2. Once the parent PNC changes the SrcID and DestID of the dependent piconet CTA, the current dependent PNC shall either complete the handover process to the new PNC or it shall shutdown the dependent piconet because it will not be able to regain control of the CTA.

After the dependent PNC receives a beacon from the parent PNC with the change in the SrcID of the dependent piconet CTA, the current dependent PNC shall begin placing the PNC Handover IE in its beacon, using the procedure indicated in 7.2.3, with the Handover Beacon Number field set to indicate the first beacon that will be sent by the new PNC. The last superframe controlled by the current dependent PNC will be the one in which the beacon number is one less than the Handover Beacon Number field. The following superframe will begin when the target DEV, now the new dependent PNC, sends its first beacon.

There are multiple points in the handover process where it is possible for the handover to fail. The current dependent PNC may cancel the handover process up until the time when it requests that the parent PNC handover control of the dependent piconet CTA to the new dependent PNC. The dependent PNC cancels the process by sending a PNC Handover Request command to the target DEV with the Handover Status field set to one to indicate that the process has been canceled.

The handover process will also fail if the target DEV fails to join the parent piconet. If the target DEV attempts to join as a neighbor PNC but the parent PNC does not support neighbor PNCs or does not wish to allow any more neighbor PNCs, then the association request by the new dependent PNC will be rejected. In that case, the DEV may also try to join as a regular DEV, in which case the dependent piconet would become a child piconet after the handover process.

If the target DEV fails to join the parent piconet as either a regular DEV or a neighbor PNC, it shall send a PNC Handover Response command to the dependent PNC with the Reason Code set to “Handover refused, unable to join parent piconet” as illustrated in Figure 7-8. The target DEV may refuse the handover at any time while the dependent PNC is sending the information about the piconet.




Figure 7-8—Failed dependent PNC handover when target DEV fails to join parent piconet

If the parent PNC rejects the request to handover control of the CTA to the new dependent PNC, the dependent PNC shall send a PNC Handover Request command to the target DEV with the Handover Status field set to one to indicate that the handover process is being canceled, as illustrated in Figure 7-9.

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 7-9—Failed dependent PNC handover when control for the dependent piconet CTA is handed over in the parent piconet

If the dependent PNC cancels the handover process, the target DEV may disassociate from the parent piconet. If DEV-2 joined the parent piconet as a neighbor PNC, it should disassociate from the parent piconet if the handover process is canceled to free up that resource for other DEVs that need to form a neighbor piconet.

7.2.7 Child piconet

When a PNC-capable DEV that is a member of an existing piconet wants to form a child piconet, the DEV shall use the Channel Time Request command, defined in 6.5.6.1 to request a pseudo-static private CTA. A private CTA is a CTA for which the SrcID and DestID are identical. The DEV shall set the Target ID List field in the Channel Time Request command to contain only the DEVID of the originating DEV, the Stream Index field to zero, and the PM Channel Time Request Type field to ACTIVE. The PNC will recognize this as a request for a child piconet. The PNC may allocate a private CTA for the child piconet depending on the availability of network resources, its capabilities and security policy.

If the PNC rejects the formation of a child PNC for any other reason than insufficient channel time or unable to allocate as pseudo-static, it shall send a Channel Time Response command with the Reason Code field set to “request denied.”

The DEV, now the child PNC, shall start sending its beacon in its allocated private CTA. The child PNC shall use a PNID that is distinct from the parent PNID. The child piconet beacon contains a Parent Piconet IE, as described in 6.4.3. Also included in the child piconet beacon is a private CTA for the parent piconet, using the PNCID for both the SrcID and DestID. This is provided to reserve the time, not to convey any information to the parent PNC.

It is possible for more than a single child piconet to be created from a common parent piconet. It is also possible for another dependent piconet to be formed in a child or neighbor piconet. There is no restriction in this standard on the number of levels that may be created. However, there is a practical limitation to the number of dependent piconets and the levels that are able to be supported.

The standard does not provide for the direct frame transfer between a member of a child piconet and a member of a parent piconet. Furthermore, this includes any other child piconets that are dependents of the parent. However, the child PNC DEV is a member of the parent piconet and thus may exchange data with any DEV in the parent piconet. The child PNC DEV is also a member of the child piconet and thus may exchange data with any DEV in the child piconet.

If the child PNC misses *mMaxLostBeacons* parent PNC beacons, the child PNC shall stop transmitting beacons to its piconet. When the child PNC hears the parent's beacons again, it shall resume sending its beacon as long as its ATP has not expired.

The FCSL initiates the formation of a child piconet using an MLME-START.request primitive while a DEV is currently a member of a piconet.

Figure 7-10 illustrates the relationship between the parent piconet superframe and the child piconet superframe. In the figure the superframe periodicity is the same for both the child and the parent piconets. Note that the CTA positions CTA1, CTA2, etc., are not to scale in Figure 7-10 and so are illustrative only.




Figure 7-10—Parent piconet and child piconet superframe relationship

The MSC for creating a child piconet is illustrated in Figure 7-11. The association and channel time request processes are defined in 7.3.1 and 7.5.1.1, respectively. The MSC for the association process is shown in Figure 7-15 while the MSC for a channel time request is shown in Figure 7-30.




Figure 7-11—MSC for creating a child piconet

The child piconet is an autonomous piconet except that it is dependent on a private CTA from the parent piconet. Association, security, etc. shall be handled within the child piconet and do not involve the parent PNC.

7.2.8 Neighbor piconet

If after following the scan procedure in 7.2.1, no free channels are available, then a neighbor PNC-capable DEV (i.e., a PNC-capable DEV from a different system), may attempt to start a neighbor piconet on the same channel as the existing piconet. To start a neighbor piconet, the neighbor PNC-capable DEV shall send an Association Request command, as described in 6.5.1.1, to the PNC. The Neighbor PNC field in the DEV Utility field shall be set as indicated in 6.5.1.1 when the Association Request command is sent. A neighbor PNC is not required to establish a secure relationship with the parent PNC, and so a PNC operating in mode 1 may reject the request for the neighbor piconet.

If the neighbor association request is accepted, then the PNC shall set the DEVID in the Association Response command to be one of the unused NbrIDs, as described in 6.2.3. If the request was rejected, as described in 6.5.1.2, depending on the reason code, the neighbor PNC-capable DEV may retry the request at a later time. If the reason code in the rejection indicates that neighbor piconets are not supported, then the neighbor PNC-capable DEV should not retry the request while that DEV is the PNC of the parent piconet.

After the association request is accepted, the neighbor PNC-capable DEV then sends a Channel Time Request command, as described in 6.5.6.1, to obtain a private pseudo-static CTA for the neighbor piconet. The Channel Time Request command shall have the Target ID List field to contain only the NbrID that was assigned to the neighbor PNC-capable DEV by the PNC in the Association Response command.

If the PNC permits the formation of a neighbor piconet and there is sufficient channel time available, the PNC shall allocate a private CTA using the NbrID as both the source and destination DEVID.

The neighbor PNC-capable DEV, now the neighbor PNC, shall start sending its beacon in its private CTA. The neighbor PNC shall use a PNID that is distinct from the parent PNID. If the neighbor piconet is operating an IEEE 802.15.3 piconet, its beacon shall contain a Parent Piconet IE, as described in 6.4.3.

If the neighbor PNC is operating an IEEE 802.15.3 piconet, a private CTA for the parent piconet is included in its beacon, using the PNCID for both the SrcID and DestID. This is provided to reserve the time for the parent piconet, not to convey any information to the parent PNC.

If the network operated by the neighbor PNC is not an IEEE 802.15.3 piconet, the neighbor PNC shall allow communications in its network only during the time allocated by the parent piconet using methods appropriate to its protocol. It shall ensure that its network does not have transmission outside of its allocated CTA.

If the neighbor PNC misses *mMaxLostBeacons* parent PNC beacons, the neighbor PNC shall stop its own transmissions and the transmissions of the DEVs in its piconet. When the neighbor PNC receives the parent PNC's beacon again, it may return to normal operation as long as its ATP has not expired.

The DME initiates the formation of a neighbor piconet by first using an MLME-ASSOCIATE.request with the NeighborPiconetRequest field set to TRUE. After successfully associating as a neighbor device, the DME uses an MLME-START.request to start neighbor PNC operations.

Figure 7-12 illustrates the relationship between the parent piconet superframe and the neighbor piconet superframe.




Figure 7-12—Parent piconet and neighbor piconet superframe relationship

The MSC for the initiation of the neighbor piconet is illustrated in Figure 7-13. The association and channel time request processes are defined in 7.3.1 and 7.5.1.1, respectively. The MSC for the association process is shown in Figure 7-15 while the MSC for a channel time request is shown in Figure 7-30.




Figure 7-13—MSC for initiating a neighbor piconet

The neighbor piconet is an autonomous piconet except that it is dependent on a private CTA from the parent piconet. Association, security, etc. shall be handled within the neighbor piconet and do not involve the parent PNC.

The neighbor PNC is not a member of the parent piconet and shall only send the following commands to the parent PNC:

- Association Request command
- Disassociation Request command
- Channel Time Request command
- Vendor Defined commands
- Security Message command
- Any Probe Request, Probe Response or Announce commands
- Any required Imm-ACK frames

The parent PNC is not a member of the neighbor piconet.

7.2.9 Stopping piconet operations

If the PNC is going to leave the piconet and there is no DEV capable of taking over as PNC or there is not sufficient time for a handover, the PNC will shut down the piconet operations. The DME may initiate the shutdown piconet operations using an MLME-STOP.request primitive with the RequestType set to SHUTDOWN.

7.2.9.1 Stopping an independent or parent piconet

If the PNC is going to remove itself from the piconet and no other DEVs are capable of taking over as the PNC, the PNC shall place the PNC Shutdown IE, as described in 6.4.5, in the beacon. The PNC shall ensure that the shutdown announcement complies with the rules for beacon announcements in 7.6.4. The only exception to this requirement is if the PNC will be shutting down and does not have enough time to wait for the next system wake beacon to complete the handover process. The process of stopping a piconet without handing over is illustrated in Figure 7-14.




Figure 7-14—MSC for stopping a piconet

If the parent PNC is not itself a dependent PNC and it is currently supporting one or more dependent piconets, the parent PNC shall select the dependent piconet PNC with the lowest DEVID to continue operation without interruption. The parent PNC shall notify the dependent PNC that it may continue operation by including the dependent DEVID in the PNC Shutdown IE, as described in 6.4.5, in the beacon. The other dependent PNCs, not seeing their DEVID in the PNC Shutdown IE, shall either cease operations, change channels or join another piconet as a dependent piconet by the time of the last beacon sent by the parent PNC. If there is time, the dependent PNC should perform the shutdown procedure for its own piconet.

If the dependent PNC whose DEVID was listed in the parent PNC's PNC Shutdown IE is coordinating an IEEE 802.15.3 piconet, it shall remove the Parent Piconet IE from its Beacon frame, signifying that it is no longer a dependent piconet.

7.2.9.2 Parent PNC stopping a dependent piconet

If the parent PNC wishes to stop a child piconet, it shall terminate the stream allocated to the child piconet using the isochronous stream termination procedure, as described in 7.5.1.3. If the parent PNC wishes to stop a neighbor piconet, it shall either send a Disassociation Request command, as described in 6.5.1.3, to the neighbor PNC or terminate the stream allocated to the neighbor piconet using the isochronous stream termination procedure, as described in 7.5.1.3. In either case, the dependent PNC shall change channels, join another piconet as a dependent piconet, or immediately initiate its shutdown procedure, as described in 7.2.9.1. The parent PNC shall listen for the dependent PNC shutdown beacon sequence to determine when the dependent piconet CTA should be removed. The parent PNC may set a maximum time for the completion of the dependent shutdown sequence, after which the CTA will be removed regardless of the completion of the dependent shutdown procedure. If the dependent PNC is a neighbor that is not operating an IEEE 802.15.3 piconet, the parent PNC shall provide the same time as it would allow for its own

shutdown sequence for the neighbor PNC to cease operations as a dependent piconet of the parent piconet before removing its private CTA.

7.2.9.3 Dependent PNC termination of a dependent piconet

After stopping piconet operations for its own piconet, as described in 7.2.9.1, a child PNC shall inform its parent PNC that it no longer requires channel time for child piconet operations by sending the parent PNC a Channel Time Request command terminating the CTA used for the child piconet.

After stopping piconet operations for its own piconet, as described in 7.2.9.1, a neighbor PNC shall inform its parent PNC that it no longer requires channel time for neighbor piconet operations by sending a Disassociation Request command to the parent PNC. Upon receiving a Disassociation Request command from a neighbor PNC, a parent PNC shall remove the CTA used by the neighbor piconet.

7.2.10 Non-PNC-capable DEVs

Simple DEVs may be implemented without providing support for the PNC role in a piconet. The implication of this is that these DEVs would be unable to form a piconet by themselves. Therefore, these DEVs should be of a type that are normally used only in conjunction with a DEV that provides PNC capability.

7.3 Association and disassociation with a piconet

Membership in a piconet depends on the security mode of the piconet. For a piconet that does not implement security, as described in 8.2.1, membership occurs immediately upon completion of the association process, as described in 7.3.1. For a piconet that implements security, as described in 8.2.2, membership occurs immediately upon receipt of the MLME-MEMBERSHIP-UPDATE.request primitive with the MembershipStatus parameter set to MEMBER. A DEV is removed from membership in a piconet via the disassociation process.

7.3.1 Association

Before a DEV has completed the association process, all frames sent to the PNC by the DEV shall be exchanged either in the CAP of the superframe or in an association MCTA.

An unassociated DEV initiates the association process by sending an Association Request command, as described in 6.5.1.1, to the PNC. When the PNC receives an Association Request command, it shall send an Association Response command, indicating that the DEV has been associated and the DEVID it has been assigned or that the request has been rejected with the reason for the rejection, as defined in 6.5.1.2. For association using MCTAs, the Association Response command, as described in 6.5.1.2, is sent in an MCTA with PNCID as source and UnassocID as destination.

The PNC shall acknowledge all correctly received Association Request commands, as described in 7.1, by sending an Imm-ACK. The ACK to an Association Request command does not mean that the DEV is associated. The PNC needs some time to ensure that the DEV should be allowed in the piconet, to ensure that there are enough resources available to support another DEV in the piconet, and to allocate a DEVID. The PNC may maintain a list of DEV addresses that are allowed to join the piconet. If the list is in use, when the PNC receives an Association Request command, the PNC shall consult the list to determine if the DEV address in the request is included in the list. If the DEV address is not in the list, the PNC shall send an Association Response command with the reason code set to "association denied," as described in 6.5.1.2, indicating that the association failed. If the PNC determines that there are not enough resources available to support the new DEV, the PNC shall send an Association Response command with the reason code set to the appropriate value in 6.5.1.2. If the PNC determines that the DEV will be associated, the PNC shall send an Association Response command with the reason code set to "success." The time difference between when

the PNC sends the Imm-ACK to the Association Request command from a DEV and when it sends an Association Response command meant for the same DEV shall not exceed *mAssocRespConfirmTime*.

The Association Response command is not a directed frame. If an Imm-ACK was required for this command, when there were multiple DEVs trying to associate during the same time interval, all of them would try to ACK and collide. Therefore, the ACK Policy field for the Association Response command is set to no-ACK. Instead each DEV trying to associate shall compare its DEV address with the DEV Address field in the Association Response command and if there is a match, accept the DEVID for all future communications.

An unassociated DEV that receives the Association Response command with the DEV address matching its own shall send during the CAP or an association MCTA a second Association Request command with the SrcID field, as described in 6.2.3, set to its newly assigned DEVID. The PNC upon receiving this Association Request command shall respond with an Imm-ACK with the DestID set to the SrcID of the Association Request command. The PNC after acknowledging this second request shall then initialize the DEV Association IE with the requesting DEV's DEVID, DEV address, DEV Capabilities field, and Association Status field set to "associated." The requesting DEV upon receiving the Imm-ACK to its second Association Request command shall consider itself associated. All other DEVs that are members of the piconet receiving the beacon containing the DEV Association IE may use the DEV Association IE to update their internal list of associated DEVs in the piconet. The PNC shall ensure that the DEV Association IE announcement for a newly associated DEV complies with the rules for beacon announcements in 7.6.4.

The PNC starts the ATP timer once it has sent the Association Response command for the new DEV. The associating DEV needs to send the second Association Request command before the ATP timer expires. If the PNC receives the second association request command after the ATP timer expires, the PNC shall send the Disassociation Request command, as described in 6.5.1.3, to the DEV requesting association to indicate that it has failed the association process.

Figure 7-15 illustrates the message flow for a successful association process.

When a PNC allows a DEV to associate, it shall assign a DEVID that is unique within the piconet. The first DEVID assigned should be randomly selected by the PNC. The PNC should assign DEVIDs so as to minimize the likelihood of collision with duplicate DEVID assignments within an interfering piconet and should assign the least recently used DEVIDs available.

NOTE—A simple method to accomplish these two goals is for the PNC to construct a list of available DEVIDs whenever it starts a new piconet. The first DEVID in the list is selected at random, followed by DEVIDs in monotonically increasing order that wrap around to one after DEVID 0xEC. This minimizes the chance that DEVIDs are duplicated if piconets with identical PNIDs interfere with each other. As DEVIDs become available for reassignment, they are added to the end of the list, which is thereby maintained in the least recently used order.

With the exception of the device that contains the PNC, there shall be a one-to-one correspondence between a DEV's device address and its assigned DEVID. The DEV that contains the PNC shall be assigned two DEVIDs: the PNCID shall be assigned to the PNC function within the DEV, and the other value of the DEVID shall be for use for all of the non-PNC traffic. When there is a coordination handover, as described in 7.2.3, the new PNC assumes the PNCID. The former PNC shall continue to use its non-PNCID DEVID for its non-PNC traffic. Hence the PNC is seen as two logical operational entities within the same DEV.

A DEVID released as a result of disassociation (see 7.3.4) shall not be reassigned until twice the disassociated DEV's ATP has elapsed.

After the association process is complete, the PNC broadcasts the PNC Information command, as described in 7.3.3.

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 7-15—MSC of DEV-2 associating

7.3.2 Piconet services

Piconet services information is provided by the PNC, if supported, to an associating DEV upon request. In order to request the piconet services information, the associating DEV sets the Piconet Services Inquiry field in the Association Request command with the SrcID set to the newly assigned DEVID, as described in 6.5.1.1. If the DEV sets the Piconet Services Inquiry field, the PNC shall send a Piconet Services command, as described in 6.5.5.1, with DestID set to the newly assigned DEVID after that PNC has received the Association Request command with the SrcID set to the newly assigned DEVID from the DEV, as described in 7.3.1. This process is illustrated in Figure 7-16.




Figure 7-16—PNC sending the Piconet Services command to a newly associated DEV in response to a request in the association process

An associating DEV may inspect the Piconet Services IEs in the Piconet Services command, as described in 6.5.5.1, returned by the PNC to determine information about other DEVs.

DEVs that are members of the piconet may place their own Piconet Services IE in the PNC's record of piconet services by sending the Piconet Services IE to the PNC using the Announce command. The PNC then sends an Announce command with DestID set to the BcstID containing the Piconet Services IE that it has added to its internal record of piconet services. If the PNC supports this capability, it retains the Piconet Services IEs of DEVs that have been sent to the PNC via the Announce command. The PNC will only save Piconet Services IEs for which it has space. Thus it is possible that the PNC would not retain a DEV's Piconet Services IE. After a DEV disassociates from the piconet, the PNC shall delete the DEV's Piconet Services IE from its own record. The process of sending a Piconet Services IE to the PNC is illustrated in Figure 7-17.




Figure 7-17—DEV sending the Piconet Services IE to the PNC

If a DEV sends a Probe Request command to the PNC requesting the Piconet Services IE, the PNC responds with Probe Response commands that contain all of the Piconet Services IEs that it has in its internal record. If a DEV has not provided a Piconet Services IE to the PNC, the PNC sends the Piconet Services IE in the Probe Response command with a zero length Content field. If the PNC did not have enough space to save the Piconet Services IE that a DEV provided, it shall send in the Probe Response command a Piconet Services IE with length 1, i.e., it contains only the DEVID. If the PNC does not support the Piconet Services IE or if its policy is not to broadcast the Piconet Services IE, it shall respond to the request with a Piconet Services IE with length 2, the DEVID field set to the PNCID, and the Services ID field set to the appropriate value, 6.4.16. The process of requesting the Piconet Services IEs from the PNC is illustrated Figure 7-18.




Figure 7-18—DEV requesting Piconet Services IEs from the PNC

If a DEV has a need for privacy, it is not required to provide information that would be available outside of the security operations of the piconet. The MAC PIB element *macDevServicesBroadcast*, as described in 5.4.2, indicates if the DEV will send the Piconet Services field. Likewise, the PNC is not required to furnish this information if it violates the security policy as set in the MAC PIB element *macPncServicesBroadcast*, as described in 5.4.1.

A DEV may use the Probe Request command, as described in 6.5.4.5, to request the Piconet Services IE from another DEV, the target DEV, in the piconet. If the target DEV supports sending the Piconet Services IE and its internal policy allows sending the information, the target DEV shall respond with a Probe Response command with the requested IE. If the target DEV does not support the piconet services IE or if its policy does not allow sending the piconet services IE, the target DEV shall respond with a Probe Response command with a zero-length Piconet Services IE, that is an IE with only the Element ID and Length fields.

It is outside of the scope of this standard to define the content or use of the Piconet Services field.

7.3.3 Broadcasting piconet information

The PNC shall broadcast the piconet information using the PNC Information command, as described in 6.5.4.2, after a DEV becomes a member of the piconet. This means that if security is required for the piconet, the PNC will wait until after secure membership has been established to broadcast the piconet information with the new DEV. In addition, the PNC shall send the piconet information for each of the DEVs that are a member of the piconet and any neighbor PNCs that are associated in the piconet at least once every *mBroadcastDevInfoDuration* via a PNC Information command. When the PNC broadcasts this command, the PNC shall include an entry for the DEV personality of the PNC as well as an entry for the

PNCID. The DEV Info field for the PNCID shall contain the same information as the DEV Info field for the PNC's DEV personality with the exception of the DEVID field, which shall be set to the PNCID.

7.3.4 Disassociation

When a PNC wants to remove a DEV from the piconet, the PNC shall send a Disassociation Request command, as described in 6.5.1.3, to that DEV with an appropriate reason code. Similarly when a DEV wants to leave the piconet, the DEV shall send a Disassociation Request command to the PNC with an appropriate reason code.

All Disassociation Request commands, when received correctly, shall be acknowledged by the intended recipient.

All DEVs in the piconet shall send frames to the PNC often enough to assure that the ATP is not reached. If the PNC does not receive any frame originating from an associated DEV within this timeout duration, the PNC shall disassociate the DEV. The DEV may send a Probe Request command without requesting any information to cause the PNC to reset the ATP if the DEV does not have any other traffic that it needs to send to the PNC.

If the DEV is STP capable, as indicated by the STP field in its Capability IE, and is the originator of any allocated streams, it shall send an Announce command that includes the Stream Renew IE to reset the STP of all its streams. This command will also reset the ATP of the DEV.

If the DEV supports reporting the Piconet Channel Status, as defined in 7.2.5, it should include the Piconet Channel Status IE in the Announce command sent to the PNC that is used to reset the ATP.

If the beacons from the PNC are not received by the DEV for longer than the ATP, the DEV shall consider itself disassociated from the piconet and may try to associate again. The DEV notifies the DME that the ATP expired using the MLME-DISASSOCIATE.ind primitive with the ReasonCode set to DEV_ATP_EXPIRED. In addition, if an associated DEV receives a broadcast PNC Information command from the PNC that is missing its DEV Info field, i.e., none of the DEV Info fields (6.5.4.2) contains its DEV address, the DEV shall consider itself disassociated from the piconet.

The PNC shall send a Disassociation Request command to a DEV that sends a frame after its ATP has expired.

The PNC upon receiving a Disassociate Request command or an ATP expiration shall include in the beacon a DEV Association IE. The PNC shall ensure that the DEV Association IE announcement for a disassociated DEV complies with the rules for beacon announcements in 7.6.4. The PNC shall perform the stream termination procedures for each of the assigned CTAs with the disassociated DEV as the SrcID or DestID. The other DEVs that are members of the piconet shall use this information to update their internal DEV association table and to determine whether they will discontinue listening or transmitting in a CTA. Note that when a DEV is disassociated, it loses its DEVID, and so the PNC will reset the bits that refer to this DEVID in all of the relevant bitmaps, e.g., PS Status IE, PCTM IE, CWB IE. The PNC will also remove the disassociated DEV from any PS sets and multicast groups that it has joined and shall delete the Piconet Services IE, if any, for that DEV from its internal storage.

If the DEV is disassociated, the PNC shall terminate all of the streams allocated to the disassociated DEV with its DEVID as either the SrcID or DestID. The PNC follows the process described in 7.5.1.3 except that it does not send any of the commands that would have had the disassociated DEV as the destination.

[illegible]

Figure 7-20 illustrates the message flow for a disassociation initiated by the PNC.




Figure 7-20—PNC-initiated disassociation MSC

The channel time is divided into superframes, with each superframe beginning with a beacon. The superframe is composed of three major parts: the beacon, the optional CAP, and the CTAP, as shown in Figure 7-21. The CTAP is used for asynchronous and isochronous data streams as well as commands while the CAP may be used for commands and non-stream data, as regulated by the PNC. During the CAP, the DEVs access the channel in a distributed style using CSMA/CA and a backoff procedure. During the CTAP, the PNC controls the channel access by assigning CTAs to an individual DEV or to a group of DEVs with each CTA having a fixed start time and duration.

210
Copyright © 2016 IEEE. All rights reserved.




Figure 7-21—Superframe structure

7.4.1 Interframe space (IFS)

There are four IFSs that are defined: the minimum interframe space (MIFS), the short interframe space (SIFS), the backoff interframe space (BIFS), and the retransmission interframe space (RIFS). The actual values of the MIFS, SIFS, BIFS, and RIFS are PHY dependent and are defined as follows:

- In 10.2.7.1 for the 2.4 GHz PHY
- In 11.2.6 for the SC PHY
- In 11.3.5.5 for the HSI PHY
- In 11.4.1.2 for the AV PHY

All Imm-ACK frames, frames sent as a response frame for Imp-ACK, and Dly-ACK frames shall start transmission over the medium a SIFS after the end of the transmission of the previous frame that requested the ACK. The IFS between all received Imm-ACK frames and Dly-ACK frames and the next frame transmitted over the medium shall be no less than a SIFS. The IFS in a CTA between a frame and the next frame transmitted over the medium by the same DEV if the first frame had the ACK Policy field set to either no-ACK or Dly-ACK shall be no less than a MIFS.

During the CTAP, all DEVs shall use an IFS no less than a RIFS for retransmissions. During a CP, however, the retransmissions shall follow the CAP rules described in 7.4.2. The rules for acknowledgment and retransmissions are described in 7.9. The interframe space requirement for the beacon is ensured by the location of the CTAs, which is determined by the PNC, as described in 7.4.3.6.

7.4.2 Carrier sense multiple access with collision avoidance (CSMA/CA)

The basic medium access mechanism during the CAP is carrier sense multiple access with collision avoidance (CSMA/CA). The CSMA/CA contention method is also used for open CTAs and association CTAs. The PNC controls the type of data or commands that may be sent in the CAP via the CAP Control field in the Piconet Mode field of the Piconet Synchronization Parameters field, as described in 6.3.1, in the beacon. A DEV shall only send frames of the type indicated by the Piconet Mode field in the beacon for the current superframe. The CAP Control field in the Piconet Mode field may be changed by the PNC from superframe to superframe.

Association CTAs are CTAs that have the SrcID set to the UnassocID and the DestID set to the PNCID. Association CTAs shall be used only to send Association Request commands to the PNC. Open CTAs are CTAs that have the SrcID set to the BcstID and the DestID set to a valid DEVID including GrpIDs, the McstID, and the BcstID. Open CTAs may be used to send either commands or data to the DEV or DEVs indicated by the DestID.

To minimize collisions, a transmitting DEV is required to first sense that the medium is idle for a random length of time. The MAC shall use the CCA capabilities of the PHY to detect whether the channel is busy or idle. Only if the medium is idle after that time shall the DEV start its transmission. This process of waiting before transmission is termed “backoff.” The backoff procedure shall not be applied to the transmission of the beacon that is transmitted by the PNC at the beginning of the superframe.

During a CP, a DEV is allowed to transmit one frame at a time with backoff being applied to every frame attempted during the CP, except for the Imm-ACK frame. The PNC may send a command a SIFS following the Imm-ACK of a frame in a CP or following a frame with ACK Policy field set to no-ACK in a CP. In this case, the PNC is not required to perform the backoff procedure before sending its frame.

In no case shall a DEV or the PNC extend its transmissions that started during a CP past the end of that CAP. If an Imm-ACK is expected for that frame, the remaining time in the CP needs to be long enough to accommodate the current frame, 2 SIFS times, and the Imm-ACK frame at the same PHY rate as the transmitted frame. If there is insufficient time remaining in the CP for the entire frame exchange sequence, then the DEV or the PNC shall not commence transmission of the frame.

The following backoff procedure shall be performed when sending frames (other than Imm-ACK) during a CP.

The backoff algorithm uses the following information:

- *retry_count*: An integer that takes on values in the range 0 to 3, inclusive.
- *backoff_window(retry_count)*: A table that has values [7, 15, 31, 63].
- *pBackoffSlot*: A PHY-dependent parameter that is based on the amount of time it takes to sense the channel. For the 2.4 GHz PHY, this is defined in 10.2.7.1. For the SC PHY, this is defined in 11.2.6.1. For the HSI PHY, this is defined in 11.3.5.4. For the AV PHY, this is defined in 11.4.1.2.
- *bw_random(retry_count)*: A random integer drawn from a uniform distribution over the interval [0, *backoff_window(retry_count)*]. The random number generated for a DEV should be statistically uncorrelated with random numbers generated by other DEVs. If the DEV does not possess a random number source, the random integer should be generated using its unique DEV address (and any other information that the implementer wishes to use) and a pseudo-random number generator (PRNG) such as MGF1, as defined in IEEE Std 1363™-2000 [B6]. Note that the current state of the PRNG should be maintained and subsequent backoffs should use subsequent integers in the pseudo-random sequence. It is important that designers recognize the need for statistical independence among the random number streams among DEVs.

The backoff time in a CP is measured at the air interface and indicates when a DEV may begin transmitting data. The DEV first waits a BIFS duration, as described in 7.4.1, from when the medium is determined to be idle before beginning the backoff algorithm. At the beginning of the CAP, the DEV may begin the backoff algorithm a SIFS after the end of the beacon transmission. If the PNC indicates that it is using an extended beacon, as described in 7.6.2, then the DEV shall wait until a SIFS after the last Announce command sent by the PNC as a part of the extended beacon before beginning the backoff procedure. At the beginning of a contention access CTA, the DEV may begin the backoff algorithm at the start time of the CTA.

The DEV shall then choose *backoff_count* = *bw_random(retry_count)* and shall maintain a counter for *backoff_count* that is decremented only when the medium is idle for the entire duration of *pBackoffSlot*. The *retry_count* shall be set to zero for the first transmission attempt of a frame. Whenever the channel is busy, the *backoff_count* shall be suspended. The channel shall be determined to be idle for the duration of a BIFS period before the backoff slot countdown is resumed. When the backoff counter reaches zero, the DEV may transmit a frame. When a backoff count of zero is drawn, the DEV can transmit immediately following the channel having been idle for a BIFS.

The backoff counter shall also be suspended outside of a CP duration. The backoff counter shall also be suspended if there is not enough time remaining in a CP for the DEV to send the frame. Note that the backoff counter is not maintained across superframes. A DEV shall choose a new *backoff_count* at the start of every CP. If the total time elapsed since the frame was queued for transmission has exceeded the transmission timeout specified for the frame, the backoff counter shall be reset, and the attempted transmission shall be canceled.

When a directed frame is transmitted and the expected ACK is not correctly received by the DEV, as described in 7.1, the `retry_count` shall be incremented but shall not be set to more than 3. The `backoff_count` shall then be set to `bw_random(retry_count)`. If the maximum number of retries for that frame has not been exceeded, the backoff procedure is again resumed.

For DEVs supporting mmWave PHY, Blk-ACK is allowed in a CP. In a CP, the Data frame sent in response to a Blk-ACK request shall have a zero-length MAC Frame Body field. The directional use of CP is defined in 7.6.5.2.

7.4.3 Channel time allocation period channel access

Channel access in the CTAP is based on a TDMA method in which all CTAs have guaranteed start time and duration. The guaranteed start times enable both power saving and good QoS characteristics. All the CTAs for the current superframe are broadcast in the beacon.

7.4.3.1 Channel time allocations (CTAs)

The PNC divides the CTAP into channel time allocations (CTAs). A DEV that is given a directed CTA is guaranteed that no other DEVs will compete for the channel during the indicated time duration of the CTA. A DEV with a CTA may or may not make use of all the allocated time duration within the CTA. The selection of a stream, command, or asynchronous data for transmission during a CTA is determined locally by the DEV depending on the number of pending frames and the value of their User Priority fields. See B.3.1 for more information on priority management.

There are two types of CTAs: dynamic CTA and pseudo-static CTA. The type of a CTA requested is indicated in the Channel Time Request command, as specified in 6.5.6.1.

The PNC may move dynamic CTAs within the superframe on a superframe-by-superframe basis. This allows the PNC the flexibility to rearrange CTA assignments to optimize the utilization of the assignments. The PNC moves a dynamic CTA by simply changing the CTA parameters in the beacon. Dynamic CTAs may be used for both asynchronous and isochronous streams.

If multiple CTAs per superframe were requested by the DEV in the Channel Time Request command, as described in 6.5.6.1, the PNC shall attempt to spread the CTAs out evenly within the superframe. If the PNC is unable to spread out the allocations, as described in F.1.4, it should deny the channel time request.

The PNC should attempt to allocate the CTAs of all SPS DEVs first in the superframe. Exceptions to placing these allocations first are as follows and in order of priority:

- QoS streams that need multiple CTAs within a superframe and require a location immediately following the beacon, if the CAP is used.
- If CAP is not used, a single CTA that follows the beacon without the *mFirstCtaGap* restriction and is one of the following:
 - An MCTA with the PNCID as the SrcID—directed, broadcast, or multicast
 - A single pseudo-static CTA

Pseudo-static CTAs shall be allocated only for isochronous streams and shall not be sub-rate allocations, as described in 6.5.6.1. If the PNC needs to change the duration or location of a pseudo-static CTA within the superframe, it shall change the corresponding CTA blocks in the beacon. The PNC shall not create any new CTAs for other stream indices that overlap with the old time interval of the pseudo-static CTAs for *mMaxLostBeacons* number of superframes. However, the PNC may overlap the old and new time intervals of the same pseudo-static CTA within a superframe as it does not create the possibility of frame collisions. If the PNC sees the transmission of a PDU during the new allocation by the source of the old allocation before the expiration of *mMaxLostBeacons* number of superframes, the PNC may reuse the old allocation for

another pair of DEVs. When the source DEV of a pseudo-static CTA receives a beacon with the new CTA, it shall cease using the old CTA and begin using the new CTA. When the destination DEV of a pseudo-static CTA receives a beacon with the new CTA, it shall begin receiving during the new CTA and may also receive during the old CTA.

If the PNC needs to simultaneously change the positions of one or more of the CTAs in the superframe, including pseudo-static CTAs, the PNC may place the Piconet Parameter Change IE in the beacon with the Change Type field set to MOVE and the Superframe Duration field set to zero. After the superframe with the beacon number equal to the value in the Change Beacon Number field in the Piconet Parameter Change IE, a DEV shall not transmit until the DEV successfully receives a beacon with a beacon number greater than or equal to the value in the Change Beacon Number field in the Piconet Parameter Change IE.

While the PNC is changing the time interval of the pseudo-static CTA, it is possible for the destination DEV to miss traffic for up to $mMaxLostBeacons$ superframes. If the destination wants to avoid this, it would need to listen for the entire superframe duration whenever it misses a beacon.

A private CTA is a CTA where the same DEV is both the source and the destination. A private CTA is not used for communication in the piconet. Instead, it is used to reserve channel time for some other use. For example, a private CTA would be used for a dependent piconet, as described in 7.2.7 and 7.2.8. Private CTAs shall be pseudo-static CTAs, so that its position and duration remain relatively constant for the other use. A DEV requests a private CTA by using a Channel Time Request command, as described in 6.5.6.1, with Target ID List field containing only its own DEVID and the Stream Index field set to the unassigned stream values, as defined in 6.2.5.

The PNC may allocate CTAs with the SrcID set to the BcstID. These CTAs are referred to as either open CTAs or open MCTAs, depending on the stream index, and use contention based channel access instead of TDMA. The type of contention access that is used is indicated by the Stream Index field in the CTA block. A stream index set to the asynchronous stream index indicates that CSMA/CA is used while the MCTA stream index indicates that slotted aloha is used. A DEV may request that the PNC modify the frequency and duration of the open CTAs or open MCTAs by sending a Channel Time Request command, as described in 6.5.6.1, to the PNC with the DestID set to the BcstID and the stream index set to the MCTA index. The CTA Rate Factor field, CTA Rate Type field, Channel Time Request TU field, and Minimum Number of TUs field shall be set by the DEV to the desired duration and frequency requested for the open CTA allocation. The PNC is not required to allocate open CTAs or open MCTAs in the manner requested by a DEV; rather it may use this information to determine the frequency and duration of open CTAs or open MCTAs that it allocates. A DEV may modify its request for open CTAs or open MCTAs by sending another Channel Time Request command to the PNC with new parameters. A DEV is not allowed to terminate either an open CTA or an open MCTA. However, a DEV requests zero time for open CTAs or open MCTA by sending a Channel Time Request command to the PNC with the Minimum Number of TUs field set to zero.

The More Data field, as described in 6.2.1.6, in the Frame Control field is set to one to indicate that the source DEV could be sending more frames in the CTA. In order to save power at the destination DEV, a source DEV may indicate that it will not use the remaining time in the current CTA by setting the More Data field to zero. The source DEV may retransmit a frame with More Data set to zero for which an ACK was expected but was not received. If the destination DEV receives a frame with the More Data field set to zero with an ACK Policy other than no-ACK, it should continue to listen for an implementation-dependent time after sending an acknowledgment to make sure that the source DEV is not going to retransmit the frame because it did not receive the ACK. The source DEV may choose to send a zero-length frame with the More Data field set to zero when it has no more frames to send in a CTA.

The More Data field shall be ignored by the destination for all frames sent in the CAP, with the exception of any Announce commands used for the extended beacon.

7.4.3.2 Channel time allocation (CTA) and channel time usage

The DEVs that are members of the piconet shall use the Channel Time Request command, as described in 6.5.6.1, whenever they wish to make a change in their CTAs. Once a Channel Time Request command is received from a DEV, the PNC shall remember that as the outstanding request for that stream for every superframe until another Channel Time Request command for that stream is received from the DEV. The CTAs within the CTAP are based on the current pending requests from all the DEVs and the currently available channel time within the CTAP. The start time of each CTA is referenced to the start of the Beacon frame, as described in 7.6. The algorithm used to allocate the channel time and assign CTAs is outside of the scope of this standard.

The PNC shall not allocate any MCTAs or dynamic CTAs within *mFirstCtaGap* following the end of the beacon except with the PNC as the source.

When a source DEV has a frame of any type for a destination DEV, the source DEV may send it during any CTA for that source DEV and destination DEV pair or to use the CAP to communicate that frame. The source DEV may also send a frame to a destination DEV in any CTA assigned to that source even if the destination DEV is different from that indicated in the CTA block, provided the source DEV has determined that the destination DEV will be receiving in that CTA, as described in 6.4.11. The stream index in a transmitted frame shall be one of the following:

- The asynchronous stream index, as defined in 6.2.5
- The MCTA stream index, as defined in 6.2.5
- An established stream from the source DEV to the destination DEV

If the DestID of the CTA is the McstID or the BcstID, the source DEV may still send directed frames to any associated DEV. However, it is possible that the target DEV will not be receiving during the CTA if it is in a power save mode, as described in 7.14, or if it is not receiving multicast traffic, as described in 7.5.3.

In any superframe there may be one or more DEVs in the piconet that receives the beacon in error. This may not happen to the same DEV all the time but may happen to different DEVs at different times depending upon their location and type of interference to which they are subjected. If a DEV did not receive the beacon, it shall not transmit during the CAP or during any MCTA or dynamic CTA, except to ACK a directed frame sent to the DEV with the ACK Policy field set to one of Imm-ACK, Imp-ACK, or Dly-ACK Request. DEVs with pseudo-static CTAs are allowed to transmit during these CTAs as long as the number of consecutive lost beacons is less than or equal to *mMaxLostBeacons*. A DEV shall stop transmitting in its pseudo-static CTA when the number of consecutive lost beacons exceeds *mMaxLostBeacons*. If a DEV that is the destination of a pseudo-static CTA misses a beacon, it should listen for the entire duration of the superframe in case the pseudo-static CTA is in the process of being moved. Any DEV that misses a beacon may also listen for the entire duration of the superframe to receive frames for which it is the destination.

In no case shall a DEV extend its transmissions that started during a CTA beyond the end of that CTA. Hence, the source DEV shall check whether there is enough time remaining in the CTA for the transmission of current frame and SIFS. If an Imm-ACK or Dly-ACK is expected for that frame, the DEV shall check whether there is enough time remaining in the time slot to accommodate the current frame, 2 SIFS periods, and the Imm-ACK or Dly-ACK frame at the same PHY rate as the transmitted frame. If there is not enough time remaining for this entire frame exchange sequence, then the DEV shall abort the transmission and not use the remainder of the CTA.

The PNC may compute more than one superframe time allocation at a time and keep them repeating over time until the situation changes. If the PNC wants to increase the allocated channel time allocation on a regular basis, it shall allocate more time up to the maximum requested by sending a Channel Time Response command, as described in 6.5.6.2, to the source DEV and change the allocation(s) in the beacon. If the source DEV requires additional channel time it will need to use the stream modification procedure, as

described in 7.5.1.2. The PNC may also reduce the channel time allocation for a stream by sending the Channel Time Response command to the source DEV and changing the allocations in the beacon.

In any individual superframe, the PNC may allocate more time for a dynamic CTA than the amount indicated in the Channel Time Response command.

7.4.3.3 Management CTAs

Management CTAs (MCTAs) are CTAs that have the stream index set to the MCTA stream index, as described in 6.2.5. A PNC may choose to use MCTAs instead of the CAP for sending command frames, unless otherwise restricted by the PHY, as described in 10.2.10. When MCTAs are used, the PNC shall ensure that sufficient MCTAs are allocated to allow for the transmission of commands to and from the PNC. There may not be any MCTAs in a superframe or there may be as few as a single MCTA in a superframe where the ownership of the MCTA changes from superframe to superframe. At the other extreme, there may be one or more uplink and downlink MCTAs per member DEV per superframe plus MCTAs for association. The PNC is responsible for determining the appropriate number of MCTAs in a superframe in the same way that the PNC is responsible for choosing the CAP size if a CAP is used. The PNC determines which DEVs will be allocated MCTAs and the frequency of the allocations. The PNC shall allocate at least one association MCTA every $mMctaAssocPeriod$ if the CAP is not used.

An open MCTA is one where the SrcID is the BcstID, as described in 6.2.3, and the DestID is set to a valid DEVID including GrpIDs, the McstID, and the BcstID. Any DEV that is associated in the piconet may attempt to send a command frame to the DEV or DEVs that are indicated in the DestID in an open MCTA. An MCTA with the UnassocID as the SrcID is an association MCTA. Any DEV not currently associated in the piconet may attempt to send an Association Request command to the PNC in an association MCTA. Association Request commands shall not be sent in open MCTAs. Likewise, only Association Request commands shall be sent in association MCTAs. Open MCTAs with the DestID set to the PNCID enable the PNC to service a large number of DEVs with low MCTA requirements by using a minimum number of MCTAs. When there are few DEVs in a piconet, it might be more efficient to use MCTAs assigned to a DEV instead of using an open MCTA. Open MCTAs in which the DestID is not the PNCID are used by DEVs to send frames to other DEVs without having to request channel time.

It is the PNC's responsibility to determine the number and type of MCTAs to use for each superframe. A DEV may request the frequency of MCTA allocations by sending a Channel Time Request command, as described in 6.5.6.1, to the PNC with the stream index set to the MCTA stream index, as described in 6.2.5, and the CTA Rate Factor, as described in 6.5.6.1, set to the DEV's desired interval for uplink MCTAs, the Num Targets field set to one and the Target ID field set to the PNCID. All other parameters of the Channel Time Request Block field shall be set to zero and may be ignored by the PNC upon reception.

If commands are not allowed in the CAP, the PNC shall assign an MCTA with the new DEV's DEVID as the SrcID as soon as possible after a successful association, 7.3.1, preferably in the next superframe, in order to support fast connections.

The access mechanism for regular MCTAs, i.e., neither open nor association MCTAs, is TDMA, as described in 7.4.3.1.

7.4.3.4 Slotted aloha access

Slotted aloha is used for contention access in an open MCTA or an association MCTA. The access to an open or association MCTA shall be controlled by a contention window CW_a maintained by each DEV. The contention window shall be derived from the number a , where a is the number of retransmission attempts made by the DEV. For the first access attempt, a shall be set to zero. The size of the contention window, CW_a , is defined as follows:

$$CW_a = \begin{cases} 256 & \text{for } 2^{a+1} \geq 256 \\ 2^{a+1} & \text{for } 2^{a+1} < 256 \end{cases}$$

The open or association MCTA used for the a^{th} retransmission attempt shall be chosen by a uniformly distributed random integer value, r_a , within the interval $[1, CW_a]$. The random number generated for a DEV should be statistically uncorrelated with random numbers generated by other DEVs. If the DEV does not possess a random number source, the random integer should be generated using its unique DEV address (and any other information that the implementer wishes to use) and a pseudo-random number generator (PRNG) such as MGF1, as defined in IEEE Std 1363-2000 [B6]. Note that the subsequent retransmission attempts should use backoff counters drawn from subsequent entries in the pseudo-random list of integers. It is important that designers recognize the need for statistical independence among the random number streams among DEVs.

The DEV shall start counting r_a beginning with the open or association MCTA(s) in the current superframe and continue across superframes. The lack of an Imm-ACK indicates the failure of the previous access attempt.

The first open or association MCTA after the DEV begins the access process is specified by number $r = 1$. The open or association MCTA with number equal to r_a is the MCTA that the DEV shall access. The DEV shall not access the MCTA before its counter has reached the open or association MCTA with the number $r = r_a$. After receiving an ACK, a shall be reset to zero.

7.4.3.5 Allocation of MCTAs

The PNC shall indicate in every beacon, as described in 6.3.1, the rate at which it will be allocating either open MCTAs or directed uplink MCTAs in the MCTA Allocation Rate field in the Piconet Synchronization Parameters field. If the PNC is not using either open MCTAs or directed uplink MCTAs it shall indicate this with the appropriate value of the MCTA Allocation Rate field, as described in 6.3.1.1. Likewise, if the PNC will not be guaranteeing the rate at which MCTAs will be allocated, it shall also indicate this in the MCTA Allocation Rate field.

The intent of the MCTA Allocation Rate field is to enable the DEVs in the piconet to approximately determine the length of time required to send a command to the PNC. This information might be used to set the timeout parameters for the MLME primitives, as described in 5.3.

7.4.3.6 Guard time

In a TDMA system, guard times are required to keep transmissions in adjacent CTAs from colliding. In addition, a SIFS time is required to ensure sufficient turnaround time between transmissions. A CTA is defined by the start time and the duration, as specified in the CTA IE. Guard time is the time between the end of one CTA and the start of the next CTA. Including SIFS as part of CTAs and allocating guard time between CTAs ensures that transmissions are spaced by at least a SIFS. Figure 7-22 is an illustration of the allocation of the guard time such that the transmissions are separated by at least a SIFS if the owners of adjacent CTAs drift towards the other CTA. The PNC shall allocate sufficient guard time between CTAs to ensure that transmissions in adjacent CTAs do not overlap.

The required guard time depends on the maximum drift between a DEV's local time and the ideal time. This drift is a function of the time elapsed since a synchronizing reference event. In an IEEE 802.15.3 piconet, the synchronizing event is the start of the preamble of a beacon. The maximum drift, $MaxDrift$, is calculated as follows:

$$MaxDrift = \text{Clock accuracy} \times \text{interval}$$




Figure 7-22—Guard time

Propagation delay will also affect timing uncertainty, but in a piconet, the 10 m range limits propagation delay to around 33 ns, or even 66 ns for DEVs 20 m apart at opposite ends of a piconet. This is much lower than the resolution of the CTA timing, and it is ignored when calculating the guard time.

The PNC may calculate a single worst case guard time for all CTAs in the superframe, or it may calculate and assign guard time based on the type of CTA (dynamic or pseudo-static) and the position of the CTA in the superframe.

Pseudo-static CTAs require longer guard times than dynamic CTAs because pseudo-static CTAs allow transmission even when up to $mMaxLostBeacons$ are missed by the transmitting DEV. Guard times are calculated based on the worst-case drift in a superframe and the maximum allowed number of lost beacons for each of the adjacent CTAs. Guard time may be calculated by the PNC as follows:

$$GuardTime = (MaxLostBeacons_{CTA_n} + MaxLostBeacons_{CTA_{n+1}} + 2) \times MaxDrift$$

MaxLostBeacons for each CTA depends on whether the CTA is pseudo-static or dynamic. MaxLostBeacons is zero for a dynamic CTA and is equal to $mMaxLostBeacons$ if the CTA is pseudo-static. The PNC calculates the MaxDrift using the superframe duration and the clock accuracy, $pClockAccuracy$. The PNC then calculates the start time and duration of each CTA such that there is sufficient guard time between the end of one CTA and the start of the next CTA.

A DEV transmitting in a CTA starts transmission of the preamble for the first frame at the point which it calculates is the start of the CTA based on its local clock. In the case of no-ACK or delayed-ACK, the transmitting DEV shall ensure that there is enough time remaining in the CTA to transmit the frame and allow for a SIFS before the end of the CTA as calculated by that DEV, as shown in Figure 7-23.




Figure 7-23—SIFS and Guard time at the end of a CTA—no-ACK

If Imm-ACK is used, the transmitting DEV shall also ensure there is enough time for the ACK and another SIFS as shown in Figure 7-24.




Figure 7-24—SIFS, ACK, SIFS, and Guard time at the end of a CTA—Imm-ACK

As with any CTA, the PNC shall include sufficient guard time between the last CTA in the Superframe and the beacon as shown in Figure 7-25.




Figure 7-25—Guard time at the end of the superframe

The PNC shall begin transmission of the beacon preamble at the point in time that it calculates is the start of the superframe based on its local clock. All DEVs will resynchronize their clock based on the beacon arrival.

Because the clock in one DEV may be fast and another may be slow relative to the ideal time, a DEV that is expecting to receive either the beacon or a frame during the CAP or in a CTA shall begin receiving before the time that it calculates to be the start of the beacon, CAP or CTA and shall continue receiving after the time that it calculates to be within one SIFS of the end of the CTA. The amount of time that the DEV listens before the start of the CTA and after the end of the CTA is up to the implementer. It may be calculated based on the type of CTA, the superframe duration and *pClockAccuracy*. The DEV shall be able to receive a frame that is transmitted within the bounds of allowable transmission for the CTA, accounting for the worst-case drift.

7.4.3.7 Calculating channel time requests

Each DEV sends channel time requests to the PNC to indicate the amount of channel time required for transmission.

The requesting DEV shall include the frame transmission time, if known *a priori*, and the ACK transmission time, if used, and one MIFS or SIFS as appropriate per frame or ACK when calculating channel time requests. Figure 7-26 shows an example of channel time being requested for a CTA where Imm-ACKs are used.




Figure 7-26—Channel time request for frames with immediate ACKs

When No-ACK is used, the channel time request is calculated differently because there is a MIFS in between each frame in the CTA instead of a SIFS. However, there is a SIFS at the end of the CTA to allow time for the DEVs to switch from transmit to receive and from receive to transmit. Figure 7-27 shows an example of a channel time request when no-ACK is used.




Figure 7-27—Channel time request with no ACKs

A Channel Time Request TU in the CTA may cover more than one frame as shown in Figure 7-28.




Figure 7-28—Channel Time Request TU covering multiple frames

If the frame size is not known a priori, it is up to the requesting DEV to decide the amount of channel time to request for the CTA. The calculation method in Figure 7-28 is preferred for better CTA efficiency.

7.4.3.8 Relinquishing CTA time to another DEV

The PNC gives transmit control to the DEV that is the SrcID of a CTA for the duration of the CTA. The DEV that has transmit control in a CTA may, subject to the restrictions in this subclause, relinquish the remaining time in a CTA to another DEV. The DEV that relinquishes the channel time is referred to as the

originating DEV while the DEV that is given the transmit control of the time in the CTA is referred to as the target DEV. The DEV that is the SrcID of the CTA begins the CTA with transmit control for the CTA.

The originating DEV relinquishes the remaining time in the CTA to a target DEV by setting the CTA Relinquish field in the header of a frame which has the DestID set to the DEVID of the target DEV. Transmit control of the CTA can only be exchanged between the source DEV of the CTA and the destination DEV(s). If there are multiple destination DEVs for the CTA, transmit control may be given to any one of them.

The originating DEV should not initiate a CTA relinquish procedure unless it has determined that the receiving DEV has set the CTA Relinquish Supported field to one in the DEV Capabilities, as defined in 6.4.11.

A DEV that receives transmit control of a CTA keeps control until the CTA end time or until it relinquishes the transmit control back to source DEV of the CTA.

The target DEV, after listening for an *mCtaRelinquishTimeout* following reception of a frame with the CTA Relinquish field set, may send a data frame to any potential DestID of the CTA. The target DEV listens to verify that the originating DEV is not retrying the frame that had the Relinquish CTA field set.

If the DEV relinquishing time in a CTA does not correctly receive a header, as described in 7.1, from the DEV to which it relinquished the CTA within *mCtaRelinquishTimeout* time, it shall maintain transmission control of the channel.

The CTA Relinquish procedure shall not be used in the CAP or in any contention CTA.

Any DEV that is responding to frames in a CTA may set the More Data field in the header of an Imm-ACK frame or data frame to indicate to the current CTA transmit control owner that they would like to get a transmission opportunity. The current CTA transmit control owner may choose to relinquish the CTA, send a frame with an Implied ACK request, or to ignore the request.

If the DEV that has transmit control of the CTA wishes to relinquish the CTA and has no data to send, it may set the CTA Relinquish field in an Imm-ACK frame. The same rules apply as when the CTA Relinquish field is used in a data frame.

The target of a CTA relinquish shall relinquish the CTA only to DEV that is SrcID of the CTA. The current CTA transmit control owner shall always initiate the relinquish frame exchange.

Figure 7-29 illustrates the process of relinquishing transmit control to another DEV in the piconet.




Figure 7-29—MSC for relinquishing CTA time when the target DEV has data to send

7.5 Channel time management

Channel time management in an IEEE 802.15.3 piconet involves the following:

- The creation, modification, and termination of isochronous data streams between two or more DEVs in the piconet.
- The reservation and termination of asynchronous channel time for the exchange of asynchronous data between two or more DEVs in the piconet.

A DEV may support one or more streams depending upon the application it is designed to support. A PNC needs to support as many isochronous streams as it desires to source and sink.

7.5.1 Isochronous stream management

Creating, modifying, and terminating isochronous streams between two or more DEVs in a piconet is accomplished via negotiation between the originating DEV and the PNC using the Channel Time Request and Channel Time Response commands, as described in 6.5.6.1 and 6.5.6.2. Once a stream index and its CTA are established, the CTA may be modified or terminated.

Only a DEV that is either a member of the piconet or associated as a neighbor PNC shall send a Channel Time Request command to the PNC.

There is no absolute guarantee of the length of delay between the time of the request and the reception of a beacon containing the requested CTA.

The StreamGrpID parameter of the MLME-CREATE-STREAM.request and MLME-MODIFY-STREAM.request primitives allows a DME to define associations between streams for the purpose of sharing channel time. When a StreamGrpID other than zero is assigned, a CTA should be shared fairly among all streams within the group associated with the StreamGrpID. The DME is responsible for assigning unique StreamGrpIDs for stream groups.

The PNC may split a super-rate allocation into more CTAs than required to satisfy the CTA Rate Factor and may also split a sub-rate allocation into more than one CTA in the same superframe.

A DEV may request the creation (and subsequent modification and termination) of allocations, isochronous and asynchronous, in which the DestID is the PNCID. Likewise, the PNC may create allocations, both isochronous and asynchronous, in which the SrcID is the PNCID. However, while these allocations are transferred in the PNC handover process, the purpose for the allocations might be lost when the DEV that is the PNC changes. The PNC shall only allocate multiple CTAs per superframe if the time allocated for each CTA is no more than one Channel Time Request TU different from any of the other CTAs allocated for that stream.

7.5.1.1 Isochronous stream creation

In the case where the originating DEV is going to request a new isochronous stream with a target DEV, the originating DEV shall send a Channel Time Request command, as described in 6.5.6.1, to the PNC with the following parameter values:

- Target ID List field is set to the DEVID with which the originating DEV is requesting a new stream.
- Stream Index field is set to the unassigned stream value, as described in 6.2.5.
- Stream Request ID field is set to a unique value between 1 and 255 for the duration of the negotiation.
- Priority field is set to a value between 0b011 and 0b110, as defined in Figure 6-152.
- All the other Channel Time Request command parameters are set to appropriate values, as defined in 6.5.6.1.

The PNC upon receiving the Channel Time Request command from the originating DEV shall respond with a Channel Time Response command, as described in 6.5.6.2, to the originating DEV with the following Channel Time Response command field values if the requested channel time is available:

- The Stream Index field is set to an unused value other than the asynchronous stream index. This indicates that the isochronous stream has been allocated channel time. The PNC should assign the least recently used stream index available.
- The Available Number of TUs field is set to a value greater than or equal to the Minimum Number of TUs and less than or equal to the Desired Number of TUs requested.
- The Reason Code field is set to "success."

The PNC may update the beacon with the newly assigned isochronous stream CTAs before it receives an ACK to the Channel Time Response command from the originating DEV.

The PNC shall announce the creation of all streams with the CTA Status IE, as described in 6.4.10, using the beacon information announcement mechanism, as described in 7.6.4. The PNC shall issue the initial CTA for the stream in the superframe indicated in the CTA Status IE.

The CTA Status IE shall have the stream index of the new allocation, the Channel Time Request Control field from the corresponding Channel Time Request command, and the CTA Sub-Rate field set appropriately, 6.4.10. In addition, the PNC shall allocate the first CTA of the stream in the superframe with the beacon number, as described in 6.3.1.1, indicated by the Start Beacon Number field of the CTA Status IE.

If, however, either the requested channel time is not available or the PNC is not able to support the requested priority, as described in B.3.1, the PNC shall respond to the requesting DEV with the following parameter values:

- The Stream Index field shall be set to the unassigned stream value, as described in 6.2.5.
- The Available Number of TUs field shall be set to the number of TUs that the PNC had available for allocation to this request.
- The Reason Code field shall be set to “priority unsupported,” “channel time unavailable,” or “unable to allocate as pseudo-static CTA” value.

If the request is for a private pseudo-static CTA, and the PNC will not support the creation of a child piconet, it shall respond with the reason code set to “request denied.”

If the stream index is set to beam forming stream value then the request is for a CTA for beam forming. If the PNC grants this request, it shall assign the stream index value to be the beam forming stream index.

The requesting DEV upon receiving this Channel Time Response command and indicated parameters may accept its denied request as final, it may resend its original request, or it may modify its original request with new parameters.

If the target DEVID is not a member of the piconet, the PNC shall respond to the requesting DEV with the following parameter values:

- The Stream Index field shall be set to the unassigned stream value.
- The Available Number of TUs field set to zero.
- The Reason Code field shall be set to either “target DEV not a member” or “target DEV unassociated” depending on the status of the target DEVID.

DEVs perform multicast negotiations at a higher layer. A DEV may set up a multicast stream either by using a GrpID assigned by the PNC, as described in 7.5.3, or by using the McstID. A DEV enables reception of a multicast stream that has the DestID set to the McstID with the MLME-MULTICAST-RX-SETUP.request. This primitive tells the MAC to receive frames from a particular source DEV with the DestID set to the McstID and with the stream index specified in the primitive. A DEV enables the reception of multicast traffic addressed to the GrpID by joining a multicast group, as described in 7.5.3.

If the target DEV is in DSPS and is part of the same SPS Set as the channel time request, and the PNC grants the allocation, the PNC shall respond with a Channel Time Response command with a Reason Code of SUCCESS. Neither the originator nor the target DEV need to change PM modes for the stream to be allocated.

However, if the target DEV is either in DSPS mode and is not part of the same SPS Set as the channel time request or in APS mode, and the PNC grants the channel time request, the PNC shall set the Reason Code in the Channel Time Response command to “Success, DEV in PS mode.” The PNC shall place the PCTM IE in the beacon with a bit set for the target DEV, as described in 6.4.8.

When the Target DEV in DSPS or APS mode receives a beacon with its bit set in the PCTM IE, it shall send a PM Mode Change command to the PNC. If the DEV is going to remain in a power save mode it shall set

the PM Mode field in the PM Mode Change command to the appropriate value, either “SPS” or “APS.” The PNC shall then terminate the stream, as described in 7.5.1.3.

If the power save DEV is going to listen to the new allocation, it shall set the PM Mode field in the PM Mode Change command to ACTIVE. The PNC shall then begin allocating the channel time in the beacon for the stream. The PNC shall no longer set the bits for the DEV in the PS Status IEs.

If the PNC does not receive the PM Mode Change command from the power save DEV within a timeout determined by the PNC, the PNC shall terminate the channel time request, as described in 7.5.1.3, and unset the PS DEV’s bit in the PCTM IE.

An STP capable DEV shall renew its allocated streams at least once every ATP by sending an Announce command with the Stream Renew IE to the PNC containing the stream index of every stream that the originator wants to keep. Streams that have been modified within the current ATP may be omitted from the Stream Renew IE sent during the same ATP.

If the Target DEV is DSPS mode, after the PNC sets the DSPS DEV’s bit in the PCTM IE the PNC shall provide in the DSPS DEV’s next wake superframe an MCTA with the DSPS DEV as the source and the PNC as the destination that is long enough to handle a PM Mode Change command, a Channel Time Request command with four isochronous Channel Time Request Block fields, and the associated Imm-ACKs and SIFs. This allows the DSPS DEV to request a change to one of the current channel time allocations, to request new channel time, or to request that a channel time allocation be terminated.

Figure 7-30 illustrates the sequence of messages involved in successfully establishing a DEV-2 to DEV-3 stream in a piconet.




Figure 7-30—MSC for creating a DEV-2 to DEV-3 stream

Figure 7-31 illustrates the sequence of messages involved in an unsuccessful attempt to establish a DEV-2 to DEV-3 stream in a piconet.




Figure 7-31—MSC for a denied DEV-2 to DEV-3 stream

7.5.1.2 Isochronous stream modification

Only the originating DEV or the PNC may modify an established isochronous stream. The originating DEV that is requesting a modification of the channel time allocated to one of its streams shall send the PNC a Channel Time Request command with the following parameter values:

- Target ID List field is set to the DEVID with which the originating DEV has an established stream.
- Stream Request ID field set to zero.
- Stream Index field set to the index of the stream to be modified.
- The CTA Type field shall be set to the same value as in the original request for that stream index.
- All the other Channel Time Request command parameters are set to appropriate values, as defined in 6.5.6.1.

The PNC upon receiving the Channel Time Request command shall check to see if the requested resources are available. If the requested channel time is not available, the PNC shall:

- Send a Channel Time Response command with:
 - The Available Number of TUs equal to the previously assigned Available Number of TUs.
 - The Stream Request ID field set to zero.
 - The Reason Code field set to “channel time unavailable.”
- Make no modification to the existing beacon CTA blocks for the DEV requesting the modification.

If the requested channel time is available, the PNC shall:

- Reserve the requested channel time.
- Send a Channel Time Response command where:
 - The Stream Request ID field is set to zero.
 - The Reason Code field is set to “success.”
- and
 - The “new” Available Number of TUs field is greater than the “previous” Available Number of TUs field and is less than or equal to the Desired Number of TUs (for a requested an increase in channel time),
- or
 - The “new” Available Number of TUs field less than or equal to the Desired Number of TUs field (for a requested decrease in channel time).
- Build a new beacon with the modified CTA.

NOTE—If a request is rejected, the Available Number of TUs is calculated using the Channel Time Request TU associated with the last accepted request. Only if the request is accepted will the Channel Time Request TU size for an exiting stream be modified.

The PNC shall announce the modification of all streams with the CTA Status IE, as described in 6.4.10, using the beacon information announcement mechanism, as described in 7.6.4. The PNC shall issue the first modified CTA for the stream in the superframe indicated in that IE. If the target DEV is in DSPPS mode, the PNC shall also allocate an uplink MCTA in the same superframe as when the CTA is first allocated with the DSPPS DEV as the source and the PNC as the destination that is long enough to handle a PM Mode Change command, a Channel Time Request command with four isochronous Channel Time Request Block fields, and the associated Imm-ACKs and SIFs.

A dependent PNC, the originator DEV, may handover control of the dependent piconet's CTA to another DEV, the target DEV, in the parent piconet. The target DEV shall be either a member of the piconet or a DEV that has associated as a neighbor PNC, as described in 7.2.8. To handover control of the dependent piconet's CTA, the originator DEV shall send a Channel Time Request command to the parent PNC with the following parameters:

- The Num Targets field set to one.
- The Target ID List field containing the DEVID of the target DEV that is to receive control of the CTA.
- The Stream Request ID field set to zero.
- The Stream Index field set to the stream index of a CTA that has already been allocated to the dependent PNC as a private, pseudo-static CTA.
- All other fields set to the same values as in the last successful Channel Time Request for this Stream Index.

If the target DEV indicated in the Target ID List is either a member of the parent piconet or is an associated neighbor PNC and the Channel Time Request command has the correct entries as indicated above, the parent PNC shall grant the request to change the source and destination for the stream and shall send a Channel Time Response command to the originator with the Reason Code set to "Success." The PNC shall continue to place the CTA block for the allocation in the beacon but shall change the SrcID and DestID to be equal to the target DEV's DEVID. Once the PNC has changed the SrcID and DestID in the CTA block, the target DEV will have gained control of the CTA and will be allowed to request modification or termination of the allocation.

If the target DEV is not a member of the piconet and it is not an associated neighbor PNC, the parent PNC shall reject the request and shall send a Channel Time Response command to the originator with the Reason Code set to either "target DEV unassociated" or "target DEV not a member" depending on the status of the target DEV.

If the Channel Time Request command has improper entries, e.g., the Stream Index does not exist or the Stream Index is not associated with a private, pseudo-static CTA, then PNC shall reject the request and shall send a Channel Time Response command to the originator DEV with the Reason Code set to "request denied."

If the DEV is STP capable, a request to modify a stream will reset the STP of that stream.

Figure 7-32 illustrates the message sequence involved in requesting a modification to an existing stream.




Figure 7-32—MSC for modifying a stream

Figure 7-33 illustrates the message sequence involved when a requested stream modification for an existing stream is denied.




Figure 7-33—MSC for a denied stream modification

The MSC for the handover of the control of a private, pseudo-static CTA is illustrated in Figure 7-34.




Figure 7-34—MSC for handing over control of a private, pseudo-static CTA

7.5.1.3 Isochronous stream termination

Only the PNC, the originating DEV, or the target DEV may terminate an established stream. In the case of multicast or broadcast streams, only the originating DEV or the PNC may terminate the stream. In the case where either the originating DEV or the target DEV desires to terminate a specific stream, it shall send the PNC the Channel Time Request command with the following parameter values:

- Target DEVID List field is set to the DEVID of the DEV to which the originating DEV has an established stream.
- Stream Index field shall be set to the value of the stream to be terminated.
- All other fields shall be set to zero.

The PNC, upon receiving a Channel Time Request command from a DEV requesting stream termination, shall respond with an Imm-ACK. In the case where the originating DEV is requesting a stream termination, the PNC shall then notify the target DEV of the termination via CTA Status IE in the beacon. The CTA Status IE shall have the stream index set to the value of the terminated stream and the Terminate field set to indicate that the stream was terminated, as described in 6.4.10. The PNC shall ensure that the CTA Status IE announcements comply with the rules for beacon announcements in 7.6.4. Figure 7-35 illustrates the MSC for termination of a stream by a source DEV.

In the case where the target DEV is requesting a stream termination, the PNC shall then notify the originating DEV of the termination via a Channel Time Response command. Figure 7-36 illustrates the MSC for termination of a stream by a target DEV.

In the case where the PNC decides to terminate an originating DEV's stream, the PNC shall notify the source DEV via a Channel Time Response command and the target DEV via a CTA Status IE in the beacon. The CTA Status IE shall have the stream index set to the value of the terminated stream and the Terminate field set to indicate that the stream was terminated, as described in 6.4.10. The PNC shall ensure that the CTA Status IE announcements comply with the rules for beacon announcements in 7.6.4. Figure 7-37 illustrates the termination of a source DEV's stream by the PNC.

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 7-35—MSC of source DEV-2 requesting termination of its stream




Figure 7-36—MSC of target DEV-3 requesting termination of source DEV-2's stream




Figure 7-37—MSC of PNC terminating a stream

If the originator of a stream is STP capable, the PNC may terminate any stream that has not been renewed within the ATP of the originator of the stream. If the PNC terminates the stream, it shall use the stream termination procedure described in this subclause for the PNC terminating a stream with the Reason Code in the Channel Time Response command set to “STP expired.”

7.5.2 Asynchronous channel time reservation and termination

This subclause describes the process used to allocate asynchronous channel time with other DEVs in the piconet.

7.5.2.1 Asynchronous channel time reservation

There are two methods for requesting asynchronous channel time. They are as follows:

- Request a single CTA for multiple target DEVs, i.e., group asynchronous channel time requests.
- Request individual CTAs for each of the target DEVs, i.e., a individual asynchronous channel time request.

The DEV requesting asynchronous channel time shall only use one of the two methods at a time. The DEV switches between the two methods by sending a Channel Time Request command that utilizes the new method. If the DEV changes methods, the PNC shall drop previously received asynchronous channel time requests from that DEV. A DEV shall not send a Channel Time Request that requests both types of asynchronous allocations and the PNC shall reject any request received from a DEV that requests both types of asynchronous allocations.

When a DEV is requesting the creation or modification of channel time for asynchronous data transmissions with a target DEV or DEVs, the originating DEV shall send a Channel Time Request command, 6.5.6.1, to the PNC with the following parameter values:

- The target ID list shall contain either:
 - A list of all of the target DEVs. Only one Channel Time Request Block field is used for all destinations with the same TU for all of the target DEVs.
- or

- Only one DEV in the destination list. In this case the originating DEV may send multiple Channel Time Request Block fields in the command and the TU may be different in each of the Channel Time Request Block fields.
- Stream Index field shall be set to zero.
- Priority field shall be set to a value of either 0b000 or 0b001, as defined in B.3.1.
- The DSPS Set Index, PM Channel Time Request Type, CTA Type, and CTA Rate Type fields shall be set to zero and may be ignored upon reception.
- All the other Channel Time Request command parameters are set to appropriate values, as defined in 6.5.6.1.

The PNC upon receiving the Channel Time Request command from the originating DEV shall respond with an Imm-ACK to the requesting DEV. If the requested channel time is available, the PNC places the CTA block(s) in a beacon with the source and target DEVID fields appropriately set.

In the case of a group asynchronous allocation, the PNC shall place multiple CTA blocks in the beacon, one for each of the destinations. Each CTA block shall have the asynchronous stream index and the same SrcID, start time, and duration but different DestIDs. The PNC may also split a group asynchronous allocation into several CTAs in a single superframe, with any such CTA again announced by multiple CTA blocks that overlap in time but have different DestIDs. Such splits shall only be done on the TU boundaries.

For an individual asynchronous allocation, if any of the DestIDs is currently in a power save mode, the PNC shall allocate the CTAs for the power save DEV(s) in the DEV's next wake beacon, either system or DSPS. It is the responsibility of the source DEV to be able to handle different CTAs for destination DEVs in power save mode as opposed to DEVs in ACTIVE mode.

For group asynchronous allocations, if some of the DEVs are in a power save mode, then the PNC shall allocate as much as possible of the requested channel time in their wake beacons. The rest of the channel time may be allocated in non-wake beacons.

There is no guarantee of the delay between the time of the request and the reception of a beacon containing the requested CTA. If a frame's timeout interval expires while waiting for its requested CTA in the beacon, an MAC-ASYNC-DATA.confirm shall be sent with the ReasonCode set to TX_TIMEOUT.

When the PNC allocates an asynchronous CTA it decrements its count of unallocated TUs by the number of TUs allocated in the CTA. When the count of unallocated TUs reaches zero, the PNC no longer allocates the CTA and drops the channel time request.

If the request is rejected, the PNC shall send a Channel Time Response command indicating the rejection. Note that if an asynchronous request is queued by the PNC but that channel time is not immediately available due to resource constraints, that does not constitute a rejection of the request.

A new asynchronous channel time request to a target DEV replaces the previous request for that target DEV and unallocated TUs from the previous request shall be replaced by the current request. A new group asynchronous channel time request replaces all previous asynchronous requests and unallocated TUs from all previous asynchronous requests shall be replaced by the current request. The originating DEV may change the request method, TU size, destinations, and desired TUs between subsequent asynchronous channel time requests.

The PNC may time out the request for an asynchronous CTA and purge them after *mAsyncRequestLifetime*. The requesting DEV should resend a new request after *mAsyncRequestLifetime* if it desires more channel time.

Figure 7-38 illustrates the sequence of messages involved in reserving channel time for the exchange of asynchronous data between DEV-2 and DEV-3 in a piconet.




Figure 7-38—MSC for reserving asynchronous data channel time

If the target DEV is in APS mode and the PNC grants the channel time request, the PNC shall set the Reason Code in the Channel Time Response command to “Success, DEV in PS mode.” The PNC shall place the PCTM IE in the beacon with a bit set for the target DEV, as described in 6.4.8.

When the Target DEV in APS mode receives a beacon with its bit set in the PCTM IE, it shall send a PM Mode Change command to the PNC. If the DEV is going to remain in APS mode it shall set the PM Mode field in the PM Mode Change command to “APS.” The PNC shall then terminate the asynchronous channel time, as described in 7.5.2.2.

If the power save DEV is going to listen to the new allocation, it shall set the PM Mode field in the PM Mode Change command to ACTIVE. The PNC shall then begin allocating the channel time in the beacon for the asynchronous allocation. The PNC shall no longer set the bits for the DEV in the PS Status IEs.

If the PNC does not receive the PM Mode Change command from the APS DEV within a timeout determined by the PNC, the PNC shall terminate the channel time request, 7.5.2.2, and unset the DEV’s bit in the PCTM IE.

7.5.2.2 Asynchronous channel time termination

Only the PNC or the originating DEV shall be allowed to terminate an asynchronous CTA. In the case where the originating DEV is going to terminate a specific asynchronous CTA, it shall send to the PNC the Channel Time Request command with the following parameter values:

- TargetID List field shall be set to the DEVIDs of the DEVs with whom the originating DEV is going to terminate the asynchronous connection.
- Stream Index field shall be set to the asynchronous stream index, as described in 6.2.5.
- All other fields shall be set to zero.

The PNC, upon receiving the Channel Time Request command from a DEV requesting termination of the asynchronous channel time, shall respond with an Imm-ACK and shall cease allocating the channel time.

In the case where the PNC terminates the asynchronous channel time, the PNC shall notify the source DEV via a directed Channel Time Response command, as described in 6.5.5.2, with the reason code set to the appropriate value.

7.5.3 Multicast group configuration

Group MAC addresses are defined in IEEE Std 802-2014. Because IEEE Std 802.15.3 uses DEVIDs for addressing, the PNC may assign a DEVID to be used for a group address. The PNC also keeps track of all of the DEVs that request the use of a particular group address by maintaining a list of their DEVIDs and the associated group address. A group of DEVs that have been registered with the PNC using a particular group address is called a multicast group.

A DEV requests a DEVID for a group address, called a GrpID, from the PNC using the Multicast Configuration Request command, 6.5.10.1, with the Group Address field set to the desired group address and the Action field set to "Join." If a GrpID is not currently assigned as a DEVID for that Group Address and the PNC has the resources available, the PNC should assign a GrpID for the Group Address and respond to the originating DEV with the Multicast Configuration Response command, 6.5.10.2. If the request was successful, the PNC adds the originating DEV to the multicast group associated with the GrpID.

If the PNC has already assigned a GrpID for the address in the Group Address field and the PNC has the resources available, it shall add the originating DEV's DEVID to the multicast group.

If the originating DEV's request is granted, the PNC shall send the Multicast Configuration Response command to the originating DEV with the GrpID field set to the value assigned to that group address and the Reason Code set to "Success."

If the address in the Group Address field does not correspond to a valid group address, the PNC shall not assign a GrpID and shall send the Multicast Configuration Response command to the originating DEV with the GrpID set to zero and the Reason Code field set to "Failure, not a valid group address."

If the PNC is unable to fulfill the originating DEV's request for a GrpID, the PNC shall send the Multicast Configuration Response command to the originating DEV with the GrpID set to zero and the Reason Code field set to the appropriate value.

When a DEV no longer needs to use the group address, it shall send the Multicast Configuration Request command to the PNC with the Group Address field set to the address and the Action field set to "Leave." When the PNC receives this command, it shall remove the DEV from the multicast group and respond with the Multicast Configuration Response command with the GrpID field set to zero, the Group Address field set to the same value as in the request command, and the Reason Code field set to "Success." The PNC shall always respond to a properly formatted Multicast Configuration Request command with the Action field set to "Leave" with a Multicast Configuration Response command with the Reason Code set to "Success." If the address in the Group Address field corresponds to a multicast group that has the originating DEV as a member, the PNC shall remove the DEV from the multicast group.

If the PNC is unable to support an existing multicast group, it shall terminate all CTAs with the GrpID as either the SrcID or DestID using the stream termination procedure in 7.5.1.3. The PNC shall also place a DEV Association IE in the beacon, following the repetition rules described in 7.6.4, with the DEVID set to the GrpID and the Association Status field set to disassociated. When a DEV receives such an IE, it will know that the multicast group no longer exists.

If a multicast group no longer has any members, either due to disassociation or requests from the DEVs to leave the group, the PNC shall free the GrpID. A GrpID shall be allocated and reused according to the rules for assigning DEVIDs described in 7.3.1. A GrpID shall not be reported in the PNC Information command.

During PNC handover, the old PNC shall send one or more Announce commands, as defined in 6.5.5.2, to the new PNC with the Group ID IEs, as defined in 6.4.18, that correspond to the GrpIDs that are currently in use.

The MSC for a DEV successfully joining a multicast group is illustrated in Figure 7-39.




Figure 7-39—MSC for a DEV successfully joining a multicast group

The MSC for leaving a multicast group is illustrated in Figure 7-40.




Figure 7-40—MSC for a DEV leaving a multicast group

7.6 Synchronization

All DEVs within a single piconet shall be synchronized to the PNC's clock. In addition, child or neighbor PNCs shall synchronize their piconet's time usage to the parent PNC's beacon and their CTA. The beacon sent at the beginning of every superframe contains the information necessary to time-synchronize the DEVs in the piconet. See 6.3.1 for the definition of the timing parameters sent in the beacon.

Each DEV in the piconet, including the PNC, shall reset its superframe clock to zero at the beginning of the beacon preamble, as shown in Figure 7-41. All times in the superframe shall be measured relative to the

beginning of the beacon preamble. If a DEV does not hear a beacon, it should reset its superframe clock to zero at the instant where it expected to hear the beginning of the beacon preamble.




Figure 7-41—Piconet timing relative to the beacon

7.6.1 Time accuracy

A compliant implementation shall maintain the accuracy of the timer to be at least as accurate as *pClockAccuracy*.

7.6.2 Beacon generation

The PNC shall send a beacon at the beginning of each superframe using the Beacon frame described in 6.3.1. The PNC may transmit quasi-omni beacons, as described in 7.6.5.1, between the PNC and a DEV and to enable DEVs located in different directional antenna coverage areas to join the same piconet if the PHY does not require support of omni-directional modes.

If the PNC determines that the Beacon frame is too large or if it is going to split the information in the Beacon frame, it may send one or more Announce commands with the SrcID set to the PNCID and the DestID set to the BcstID following the beacon. This is called an extended beacon. Unless it is specified otherwise, the term *beacon* applies to both the Beacon frame and the Announce commands that make up the extended beacon. The IFS between the Beacon frame, the first Announce command, and any additional Announce commands shall be less than a SIFS and greater than a MIFS. If the PNC sends some of the beacon information in the broadcast Announce commands, it shall set the More Data field to indicate more data in the Frame Control field of the Beacon frame and in all but the last Announce command frame used to communicate the IEs. The CAP or the CTAP, if the CAP is not present, begins after the last Announce command that is part of the extended beacon. The PNC shall send CTA IEs, BSID IE, and the Parent Piconet IE, if present, only in the Beacon frame and not in any of the broadcast Announce commands. The Announce commands are sent to the BcstID, and so the ACK Policy field shall be set to no-ACK in these frames.

If all of the CTA IEs contained in the beacon are identical to the CTA IEs transmitted in the previous beacon, the PNC may set the CTA IEs Unchanged field appropriately in the Piconet Mode field of the Piconet Synchronization Parameters field, 6.3.1.1.

If all of the IEs (other than CTA IEs) contained in the beacon are identical to the IEs (other than CTA IEs) transmitted in the previous beacon, the PNC may set the Other IEs Unchanged field appropriately in the Piconet Mode field of the Piconet Synchronization Parameters field, 6.3.1.1.

The PNC shall transmit the beacon such that the time between beacons is the superframe duration with an error of no more than *pClockAccuracy* times the superframe duration. The PNC changes the superframe position or duration using the procedures indicated in 7.11.1 and 7.11.2, respectively.

7.6.3 Beacon reception

All of the DEVs that are associated shall use the beacon start time, CAP end time, and the CTA IEs contained in the beacon to start their transmissions. The superframe duration and the CAP end time in the beacon, as described in 6.3.1, are used to accurately mark the beginning and the end of the CTAP. A lost beacon is defined as one for which the FCS is not valid or when a DEV has not received a beacon at the expected time.

If a DEV receives a Beacon frame with the CTA IEs Unchanged field set in the Piconet Mode field of the Piconet Synchronization Parameters field, and if the DEV has correctly received the previously transmitted beacon, the DEV may assume that the CTA IEs contained in the Beacon frame are identical to the CTA IEs received in the previous beacon. If the CTA IEs Unchanged field is not set, or if the previous Beacon frame was not correctly received, then the DEV shall consider that the CTA IEs in the current beacon contain new information.

If a DEV receives a Beacon frame with the Other IEs Unchanged field set in the Piconet Mode field of the Piconet Synchronization Parameters field, and if the DEV has correctly received the previously transmitted beacon, the DEV may assume that the IEs (other than CTA IEs) contained in the Beacon frame and any Announce commands that make up the beacon are identical to the IEs (other than CTA IEs) received in the previous beacon. If the Other IEs Unchanged field is not set, or if the previous beacon (Beacon frame and any Announce commands) was not correctly received, then the DEV shall consider that the IEs (other than CTA IEs) in the current beacon contain new information.

7.6.4 Beacon information announcement

The PNC sends several IEs in its beacons to inform the DEVs in the piconet about constant or temporary conditions. Some IEs are sent in every beacon, while others are only sent if certain operations are in use, such as power save or a dependent piconet. Some of these IEs are listed in Table 7-2.

Table 7-2—IEs included in beacons as needed

IE	Format	Usage
AS IE	6.4.7	7.15
Pending channel time map (PCTM) IE	6.4.8	7.14.2, 7.14.3
PS Status IE	6.4.13	7.14
CWB IE	6.4.14	7.14.2

Other IEs are only sent as an announcement of a changed condition in the piconet. These IEs could be for the benefit of all DEVs or for a particular DEV. IEs that are not sent in every beacon are called announcements and are listed in Table 7-3.

Table 7-3—Repeated beacon announcements

IE	Clause	Announced in	Intended for	Clause
DEV Association IE	6.4.4	<i>mMinBeaconInfoRepeat</i>	All DEVs	7.3.1, 7.3.4
PNC shutdown IE	6.4.5	<i>mMinBeaconInfoRepeat</i>	All DEVs	7.2.9.1

Table 7-3—Repeated beacon announcements (continued)

IE	Clause	Announced in	Intended for	Clause
Piconet Parameter Change IE	6.4.6	<i>mMinBeaconInfoRepeat</i>	All DEVs	7.11, 7.12.1, 7.12.2
PNC Handover IE	6.4.9	<i>mMinBeaconInfoRepeat</i>	All DEVs	7.2.3
CTA Status IE	6.4.10	<i>mMinBeaconInfoRepeat</i>	DestID	7.5.1.1, 7.5.1.2
Next PNC IE	6.4.20	<i>mMinBeaconInfoRepeat</i>	All DEVs	7.2.5

If the intended recipient of an IE is all DEVs, the following rules apply:

- The IE shall be sent in at least *mMinBeaconInfoRepeat* consecutive beacons.
- If any DEV is in PSPS, the IE announcement shall be made in a system wake beacon and in at least *mMinBeaconInfoRepeat*-1 consecutive beacons following the system wake beacon.


If the intended recipient of an IE is one individual DEV, the following rules apply:

- The IE shall be sent in at least *mMinBeaconInfoRepeat* consecutive beacons.
- If the DEV is in PSPS mode, the IE announcement shall be made in a system wake beacon and in at least *mMinBeaconInfoRepeat*-1 consecutive beacons following the system wake beacon.
- If the DEV is in DSPS mode, the IE announcement shall be made in one of the DEV's DSPS set wake beacons and in at least *mMinBeaconInfoRepeat*-1 consecutive beacons following the DEV's DSPS set wake beacon.

A CTA Status IE is considered to be intended for all DEVs if the DestID contained in that IE is the BcstID or McastID. Otherwise, the CTA Status IE is intended for the DEV defined by DestID.

7.6.5 Superframe support for directional PHYs

If a PHY does not support omni-directional modes of operation or if it allows DEVs that do not support omni-directional modes of operation, then the superframe, beacon and CAP need additional features to enable these DEVs to operate. In particular, the superframe structure for DEVs using these PHY modes will be different, as illustrated in Figure 7-42.

**Figure 7-42—Superframe structure for quasi-omni capable PHYs**

This subclause defines optional procedures to enable DEVs to associate and operate with PHYs that do not have omni-directional modes. A summary of the types of CTAs used in the standard is given in F.1.1.

7.6.5.1 Quasi-omni beacon

The quasi-omni beacon shall be structured as illustrated in Figure 7-43.

Quasi-omni beacon		
Quasi-omni beacon section ($I^{(1,t)}$ frames)	PNC quasi-omni tracking section ($I^{(1,t)}$ QTs)	Sector training section ($J^{(1,t)}$ STs)

Figure 7-43—Beacon structure for quasi-omni capable PHYs

The quasi-omni beacon section shall be formatted as illustrated in Figure 7-44.

Beacon frame TX direction # 0, $Q_0^{(1,t)}$	BSIFS	...	Beacon frame TX direction # $i^{(1,t)}$, $Q_{i^{(1,t)}}^{(1,t)}$	BSIFS	...	Beacon frame TX direction # $I^{(1,t)}-1$, $Q_{I^{(1,t)}-1}^{(1,t)}$	BSIFS

Figure 7-44—Quasi-omni beacon section structure

The PNC quasi-omni tracking section shall be formatted as illustrated in Figure 7-45.

QT # 0, $Q_0^{(1,t)}$	BSIFS	...	QT # $i^{(1,t)}$, $Q_{i^{(1,t)}}^{(1,t)}$	BSIFS	...	QT # $I^{(1,t)}-1$, $Q_{I^{(1,t)}-1}^{(1,t)}$	BSIFS

Figure 7-45—PNC quasi-omni tracking section structure

The sector training section shall be formatted as illustrated in Figure 7-46.

ST # 0, $S_0^{(1,t)}$	BSIFS	...	ST # $j^{(1,t)}$, $S_{j^{(1,t)}}^{(1,t)}$	BSIFS	...	ST # $J^{(1,t)}-1$, $S_{J^{(1,t)}-1}^{(1,t)}$	BSIFS

Figure 7-46—Sector training section structure

Quasi-omni coverage during beaconing shall be supported by repeating the Beacon frame $I^{(1,t)}$ times and sending each repetition with a different TX antenna, quasi-omni direction, or array pattern. Each repetition is followed by a BSIFS, as defined in 12.1. The $I^{(1,t)}$ antenna directions or array patterns shall be identified by indices zero through $I^{(1,t)}-1$. The number of Beacon frame repetitions $I^{(1,t)}$ and the index of the current beacon transmit direction are parameters in the Synchronization IE, as described in 6.4.22. In the following, the term *direction* shall be used to refer to an antenna direction or an array pattern.

When the PNC Quasi-omni Tracking Present field in the Synchronization IE is set to one, then a PNC quasi-omni tracking section shall follow the quasi-omni beacon section to allow devices in the piconet to track the PNC quasi-omni directions. The PNC quasi-omni tracking section consists of $I^{(1,t)}$ quasi-omni training (QT) sequences transmitted in the $I^{(1,t)}$ PNC quasi-omni directions. The $I^{(1,t)}$ QT sequences may be transmitted at once in a superframe or distributed over multiple superframes. The QT sequence shall be identical to the long preamble.

When the Sector Training field in the Synchronization IE is set to one, then a sector training section shall follow the PNC quasi-omni tracking section to enable proactive beam forming, as described in 12.7. The sector training section consists of $J^{(1,t)}$ sector training (ST) sequences transmitted in the $J^{(1,t)}$ PNC sector directions. The $J^{(1,t)}$ ST sequences may be transmitted once in a superframe or distributed over multiple superframes. The ST sequences shall be identical to the long preamble.

7.6.5.2 Directional CAP

The CAP period for directional communication may be divided into three sections—an association section, a regular sub-CAP (S-CAP) section, and a regular CAP section as shown in Figure 7-47. If PNC allows new association requests, an association section exists at the beginning of the CAP period. The association section may be further divided into a set of $I^{(1,r1)}$ equal size association S-CAPs corresponding to some of the $I^{(1,r)}$ PNC different quasi-omni receive directions within one superframe or distributed over multiple superframes ($r1 \leq r$). The regular S-CAP section is divided into the $I^{(1,r1)}$ equal size S-CAPs corresponding to some of the $I^{(1,r)}$ PNC different quasi-omni receive directions within one superframe or distributed over multiple superframes. Each S-CAP is received by the PNC using a different antenna receive direction.

CAP								
Association S-CAP				Regular S-CAP				Regular CAP
Association S-CAP #0	Association S-CAP #1	...	Association S-CAP # $I^{(1,r1)}-1$	S-CAP #0	S-CAP #1	...	S-CAP # $I^{(1,r1)}-1$	

Figure 7-47—Quasi-omni CAP structure

The association CAP shall be used solely for devices to send Association Request commands to the PNC and for the Imm-ACK from the PNC to the Association Request command. The regular CAP and regular S-CAP may be used for all other command and data exchanges. It is up to the implementer to determine the method used to achieve omni-directional communications for the SC and HSI PHY modes in the regular CAP.

If the AAS field in the piconet synchronization parameters field is set to zero, indicating that the PNC is SAS, the number of S-CAPs shall be one-to-one with the number of beacons ($I^{(1,t)} = I^{(1,r)}$). That is, the PNC transmit direction used for the i^{th} Beacon frame transmission shall be used for the i^{th} S-CAP, respectively.

If the AAS field in the piconet synchronization parameters field is set to one, indicating that the PNC is AAS, the number of S-CAPs shall be equal to $I^{(1,r)}$, which is specified in the Synchronization IE, as described in 6.4.22. The special case where $I^{(1,r)} = 1$ indicates that the PNC is omni capable on reception.

Before two peer DEVs communicate in the regular CAP and/or regular S-CAPs, the two DEVs may perform beam forming, as described in 12.5, if they do not know the antenna directions to point to each other.

To support directional peer communication in the regular CAP and/or regular S-CAPs, the source DEV may send the PNC an Announce command with Directional Peer IE with the Request/Release field set to one, as described in 6.4.36. If the PNC allows the directional peer communication, the PNC shall include the Directional Peer IE in the beacon to announce that the two devices will use regular CAP and/or regular S-CAPs for directional peer communication.

After the source and destination DEVs receive the Directional Peer IE in the beacon from the PNC, the two DEVs may switch their antenna directions and communicate with each other in the regular CAP and/or regular S-CAPs. After the two devices complete directional peer communication, the source device may

send the PNC an Announce command with Directional Peer IE with the Request/Release field set to zero, and the PNC shall remove Directional Peer IE from its beacon.

Directional peer-to-peer communication between two non-PNC DEVs is not bound by the S-CAP boundaries.

7.6.5.3 Directional association

A DEV that is not omni capable on reception and supports multiple receive directions shall implement directional association. Let $I^{(1,t)}$ and $I^{(1,r)}$ be the number of PNC quasi-omni transmit and receive directions, respectively, and let $I^{(2,t)}$ and $I^{(2,r)}$ be the number of DEV quasi-omni transmit and receive directions, respectively, of a DEV that wants to find a PNC and associate with that PNC.

The best and second best pair of PNC transmit and DEV receive or DEV transmit and PNC receive directions are referred to as the best antenna direction pair and second best antenna direction pair, respectively.

While searching for a PNC, a DEV shall listen to quasi-omni beacons at all $I^{(2,r)}$ quasi-omni receive directions to find the best and second best antenna direction pairs based on LQI measurement. The DEV shall use the best DEV quasi-omni receive direction to receive further transmissions from the PNC when quasi-omni transmission is used. The DEV shall include the information of the best PNC quasi-omni transmit direction in its Association Request commands, as defined in 6.5.1.1, to inform the PNC the best PNC quasi-omni transmit direction for further quasi-omni transmissions to the DEV.

The DEV shall track the best and second best antenna direction pairs during the quasi-omni beacon section and PNC quasi-omni tracking section of beacons based on LQI measurement. If the beacon quality in the second best antenna direction pair is better than the best antenna direction pair, the PNC and the DEV shall switch to the second best antenna direction pair, which then becomes the new best antenna direction pair.


The association procedure of a DEV depends on the antenna types at both the PNC and the DEV. The DEV shall transmit one or multiple Association Request commands, as described in 6.5.1.1. If both the DEV itself and the PNC utilize SAS antennas, this command shall be transmitted on DEV antenna quasi-omni direction $I^{(2,t)} = I^{(2,r)}$, during the S-CAP with index, $I^{(1,t)}$.

If either the DEV or the PNC is AAS, the DEV does not know which quasi-omni transmit direction to use, nor which S-CAP is the best. In this case, the DEV shall transmit the Association Request command at different antenna direction pairs during association S-CAPs in one or multiple superframes until it receives an Association Response command successfully or association timeout. The Association Request command shall be sent individually instead of back to back, with random backoff applied to each one. The backoff window used by the DEV to send the Association Request command shall be the same for the cycle of $I^{(2,t)}$ quasi-omni transmit directions. If a complete cycle of quasi-omni transmit direction fails, then the backoff window shall be increased, as defined in 7.4.2. The Association Response command shall include the information of the antenna direction pair from which it receives the Association Request command. The Association Response command may be sent in a CTA.


Once an Association Response command or an Imm-ACK to the Association Request command is received successfully, as defined in 7.3.1, the DEV shall cease the transmission of Association Request command in association S-CAPs and ignore other copies of the same Association Response command from PNC. The DEV shall use the regular S-CAP of the PNC quasi-omni receive direction carried in the Association Response command for all further CAP transactions with the PNC before the completion of the best antenna direction pair searching. In addition, the DEV shall use the DEV transmit quasi-omni direction carried in the Association Response command before the completion of the best antenna direction pair searching.

7.6.5.4 DEV quasi-omni transmit direction training and tracking


Since the antenna direction pair found at the association stage may not be the best pair due to possible collisions in CAP, the PNC shall reserve a CTA to search for the best and the second best antenna direction pairs when channel has free time. The CTA structure for DEV quasi-omni transmit direction training is illustrated in Figure 7-48.

**Figure 7-48—CTA structure for DEV quasi-omni transmit direction training**

The structure of the quasi-omni training from DEV to PNC is illustrated in Figure 7-49.

**Figure 7-49—Quasi-omni training DEV → PNC structure**

The structure of the quasi-omni training cycle is illustrated in Figure 7-50.

**Figure 7-50—Quasi-omni training cycle structure**

The quasi-omni training consists of $I^{(2,t)}$ cycles. During each cycle, the DEV shall send $I^{(1,r)}$ repetitions of a QT sequence, as defined in 7.6.5.1 in the same direction. Each cycle except the last one shall end with a BSIFS. The $I^{(2,t)}$ cycles shall be sent in $I^{(2,t)}$ different directions, namely $[Q_0^{(2,t)}, Q_1^{(2,t)}, \dots, Q_{I^{(2,t)}-1}^{(2,t)}]$.

During a cycle, the PNC shall attempt to receive each of the $I^{(1,r)}$ quasi-omni training sequences using a different direction. The $I^{(1,r)}$ different directions, $[Q_0^{(1,r)}, Q_1^{(1,r)}, \dots, Q_{I^{(1,r)}-1}^{(1,r)}]$, during a cycle shall correspond to the PNC's quasi-omni receive directions.

At the completion of the full $I^{(2,t)}$ cycles, the PNC will have had an opportunity to receive a QT sequence using each combination of DEV2 transmit quasi-omni direction (0 to $I^{(2,t)}-1$) and PNC receive quasi-omni direction (0 to $I^{(1,r)}-1$). Based on this information, the PNC selects the best quasi-omni pair, i.e., DEV2's optimal transmit quasi-omni direction, $I_{i(2,t)}^{(2,t)}$, and the PNC optimal transmit and receive sector, $I_{i(1,r)}^{(1,r)}$.

Following the quasi-omni training, the PNC shall transmit its quasi-omni feedback in a Feedback IE by sending an Announce command with Imm-ACK requested. The Announce command shall be sent in the optimal transmit quasi-omni direction, $Q_{i(2,t)}^{(1,t)}$, and DEV shall listen on its optimal receive direction, $Q_{i(2,t)}^{(1,r)}$. The Feedback IE informs DEV of its optimal transmit quasi-omni direction, $Q_{i(2,t)}^{(2,t)}$, second best quasi-omni direction, and the corresponding LQIs.

The PNC shall track the best and second best antenna direction pairs based on LQI measurement by allocating a tracking CTA periodically when channel has free time. The CTA structure for DEV quasi-omni

transmit direction tracking is illustrated in Figure 7-51. If the signal quality in the second best pair is better than the best pair, the PNC and the DEV shall switch to the second best pair, which becomes the new best antenna direction pair thereafter.

DEV → PNC Quasi-omni training	SIFS	PNC → DEV Announce command with Feedback IE	SIFS	DEV → PNC Imm-ACK	SIFS

Figure 7-51—CTA structure for DEV quasi-omni transmit direction tracking

The structure for quasi-omni tracking DEV → PNC is illustrated in Figure 7-52.

QT #i, $Q_{i,(2,t)}^{(2,t)}$	BSIFS	QT #j, $Q_{j,(2,t)}^{(2,t)}$

Figure 7-52—Quasi-omni direction training cycle structure

7.7 Fragmentation and defragmentation

Fragmentation may be performed at the transmitting DEV on each MSDU. In addition, certain commands, i.e., MAC command data units (MCDUs), may be fragmented, as indicated in 6.5. All fragments shall be of equal size, except the last fragment which may be shorter. Once the MSDU/MCDU is fragmented and a transmission attempted, it shall not be refragmented. The smallest size of a fragment, excluding the last fragment, shall be at least $pMinFragmentSize$. A DEV indicates its preferred fragment size for reception in the Preferred Fragment Size field in the DEV Capabilities field, as described in 6.4.11, that it sends to the PNC when the DEV associates with the piconet.

Each fragment shall be sent with the Last Fragment Number field set to the highest fragment number of the current MSDU/MCDU, which is one less than the total number of fragments of the current MSDU/MCDU.

The first fragment shall be sent with the Fragment Number field set to zero. Each subsequent fragment shall be sent with the Fragment Number field incremented. However, the Fragment Number field shall not be incremented when a fragment is retransmitted.

All fragments of the same MSDU/MCDU shall have the same MSDU/MCDU number.

Defragmentation of an MSDU/MCDU is the reassembly of the received fragments into the complete MSDU/MCDU. The MSDU/MCDU shall be completely reassembled in the correct order before delivering it to the frame convergence sublayer (FCSL). Any MSDU/MCDU with missing fragments shall be discarded.

The receiving DEV shall not deliver an MSDU/MCDU to the FCSL until all of the fragments have been obtained. The receiving DEV may discard the fragments of an MSDU/MCDU if it is not completely received within a timeout determined by the receiving DEV. The destination DEV may also discard the oldest incomplete MSDU/MCDU if otherwise a buffer overflow would occur. If the no-ACK policy is used, the destination DEV shall discard an MSDU/MCDU immediately if a fragment is missing.

A DEV shall support concurrent reception of fragments of at least three MSDU/MCDUs including isochronous streams, asynchronous data and commands.

For frames with Imm-ACK mechanism, a DEV shall not send another fragment or frame with the same stream index to the same DEV until the sending DEV has received an Imm-ACK frame response to that frame or it has timed out on sending the frame.

If Dly-ACK is used, unacknowledged fragments from multiple MSDUs belonging to the same stream may be retransmitted in the same burst. In this case it is the responsibility of the destination DEV to deliver the MSDUs in the correct order to the FCSL.

7.8 Aggregation

Aggregation may be performed for high-speed data/video transmission or low-latency bidirectional data transmission. Accordingly, there are two aggregation methods defined: standard aggregation and low-latency aggregation.

7.8.1 Standard aggregation

Figure 7-53 illustrates the aggregation process. The originating DEV, upon receiving an MSDU, maps it into a subframe payload. If the length of the MSDU exceeds the predetermined value indicated in the Preferred Fragment Size field in Capability IE, as defined in 6.4.11, the MSDU shall be fragmented and mapped into multiple subframe payloads. Each MSDU is assigned an unique MSDU number for identification. If fragmentation is adopted, each fragment is assigned a fragment number for identification within the MSDU. All the fragments of the same MSDU shall have the same MSDU number.




Figure 7-53—Aggregation at originating DEV

A subheader is created and configured, as defined in 6.2.8.1, for each subframe to contain the necessary information that helps the target DEV to retrieve the original data. The ACK Policy field in MAC header shall be set to Blk-ACK, as described in 6.2.1.4.

All the subheaders are combined together to form the MAC subheader. As specified in 6.2.8.1, up to 8 subframes aggregated into a single frame. All unused subframes have zero length in the Subframe Length field.

Figure 7-54 illustrates the deaggregation process. After receiving the aggregated frame, the target DEV divides it into subframes according to the information in MAC subheader, validates each subframe by FCS. To recreate the original MSDU, the target DEV uses the MSDU Number field and the Fragment Number field in the subheader, as defined in 6.2.8.1.




Figure 7-54—Deaggregation at target DEV

The standard aggregation supports unidirectional and bidirectional data transmission by attaching ACK information with data. For unidirectional transmission, the target DEV upon receiving an aggregated frame replies an empty Data frame with the Blk-ACK Bitmap field properly configured to support retransmission, as described in 7.9.5.1. For bidirectional data transmission, the target DEV sends ACK information with a Data frame that may contain a Data Payload field to the originating DEV.

To avoid buffer overflow, before starting real data transmission, the originating and target DEV may exchange empty data frame with the RX Buffer Size field properly configured, as defined in 6.2.8, to inform each other the available receiving buffer size. This information is used by the DEVs to adjust the subframe number and subframe length when sending real data.

The Fragmentation Control field in the MAC header of an aggregated frame shall always be used to carry the Receive Status field, as defined in 11.1.8.3.

7.8.2 Low-latency aggregation

Before starting low-latency aggregation, the originating and target DEVs shall exchange the desired CTA relinquish duration value by sending CTA Relinquish Duration IE, as described in 6.4.26, in an Announce command, as defined in 6.5.5.2. Low-latency aggregation shall not be allowed in CP because the usage of relinquish operation is not allowed in CP.

The DEV chooses the smaller value among its own CTA relinquish duration and that of the target DEV to start the bidirectional CTA relinquish duration (BCRD) timer.

In low-latency aggregation mode, the originating DEV receives MSDUs from the FCSL, and aggregates them into the Frame Payload field. Each MSDU is formatted, as described in 6.2.8.2. The originating DEV shall also aggregate zero-length MSDUs if there is no MSDU available from the FCSL until an MSDU is available, or until the BCRD timeout has expired. If the originating DEV has not received any MSDU from the FCSL until the beginning of the next BCRD period, it may transmit a zero-length frame with the ACK Policy field in MAC header set to No-ACK, the Aggregation field set to zero and the Low-latency Mode field set to one in the PHY header (for the purpose of maintaining the link alive). Each MSDU transmitted is assigned an MSDU number on its first transmission that is used to identify it uniquely in the MSDU sequence delivered to the destination FCSL. The sequence number of zero-length MSDU is assigned the

most recent sequence number transmitted by the source for which it has received an ACK. The MSDU sequence number in a zero-length MSDU shall be ignored by the destination. No retransmission is required for zero-length MSDU.

When the target DEV receives an aggregated MAC frame, it verifies the Low-latency Aggregation field, as defined in 11.2.3.2.1.

The target DEV reads the Frame Payload field, parsing each aggregated MSDU based on the received MSDU subheader. For each MSDU subheader, the destination validates its HCS. If the HCS is not valid, the destination shall continue searching for a match of a valid MSDU HCS, shifting the next expected MSDU on octet alignment until one is found. If the FCS is correct, the destination passes the MSDU in order to the FCSL, withholding any trailing MSDUs until all intermediate MSDU fragments are retransmitted, and received correctly.

The target DEV sets the corresponding bit in the next transmitted Blk-ACK Bitmap field according to its order offset taking the value of MSDU Response Number field of the transmitted MAC subheader as reference. The MSDU Response Number field is updated upon reception of a valid MAC subheader, from the MSDU Request Number field. In case the MSDU order offset is not sequential to the previous MSDU order-offset, the corresponding range of bits shall be cleared in the Blk-ACK Bitmap field.

The target DEV shall inform the originating DEV the available receiving buffer size using RX buffer size field, as defined in 6.2.8. The originating DEV adjusts the maximum number of MSDUs to send in the next frame to avoid buffer overflow.

If the MSDU is incorrectly received, the corresponding bit is set to zero in the next transmitted Blk-ACK Bitmap field according to its order offset in the current aggregated frame.

The source and destination interchange roles during the CTA, every BCRD or less, utilizing the CTA relinquish mechanism described in 7.4.3.8, to enable minimum latency retransmission of incorrectly received MSDUs.

7.9 Acknowledgment and retransmission

The acknowledgment types defined for this standard are as follows:

- No acknowledgment (no-ACK)
- Immediate acknowledgment (Imm-ACK)
- Delayed acknowledgment (Dly-ACK)
- Implied acknowledgment (Imp-ACK)
- Block acknowledgment (Blk-ACK)

7.9.1 No acknowledgment (no-ACK)

A transmitted frame with the ACK Policy field set to indicate no-ACK shall not be acknowledged by the intended recipient(s). The transmitting DEV assumes that the frame is successful for all its local management entities and proceeds to the next frame scheduled for transmission. The ACK Policy field in broadcast and multicast addressed frames shall be set to no-ACK upon transmission.

7.9.2 Immediate acknowledgment (Imm-ACK)

A directed frame that expects an Imm-ACK shall have the ACK Policy field in that directed frame set to indicate the same, as defined in 6.2.1.4. If the intended recipient of a directed frame correctly receives the frame, it shall send the Imm-ACK frame, as described in 7.4.1.

7.9.3 Delayed acknowledgment (Dly-ACK)

Delayed acknowledgment (Dly-ACK) shall be used only for directed stream data frames, i.e., isochronous connections, where the Dly-ACK mechanism has been set up with negotiation between the source and destination DEVs. The Dly-ACK mechanism is initiated by the source DEV sending a single data frame with the ACK Policy field set to Dly-ACK Request.

If the MAC header was correctly received, as described in 7.1, and the destination DEV accepts the use of Dly-ACK, it shall respond with a Dly-ACK frame, acknowledging the received data frame, if it was correctly received, and setting the Max Burst field to a value representing the maximum number of *pMaxFrameBodySize* MPDUs the source DEV may send in one burst. Because the receiver buffer requirement is equal to Max Burst field times *pMaxFrameBodySize*, the source may send as many smaller frames as will fit in the receive buffer window, up to a maximum of Max Frames, as provided in the Dly-ACK frame, as described in 6.3.2.2. The destination DEV may change the value of the Max Burst field and the Max Frames field in each Dly-ACK frame.

If the destination DEV wants to decline the use of the Dly-ACK mechanism, it shall reply with an Imm-ACK frame.

If the value of the Max Burst field is set to zero, the source DEV shall stop transmitting in the current CTA and reopen the Dly-ACK mechanism by sending a single frame with the ACK Policy field set to Dly-ACK Request in the next CTA for this stream. If the value is not zero, the source DEV may continue transmission in the current CTA, if the time is available.

A DEV shall send a Dly-ACK frame in response to a correctly received MAC header, as described in 7.1, that has the ACK Policy set to Dly-ACK request and the DestID set to the DEVID of this DEV or, if applicable, the PNCID. Because the Dly-ACK frame is sent even if the FCS failed on the data frame with the ACK Policy set to Dly-ACK request, the DEV shall check the contents of the Dly-ACK frame to determine which, if any, MSDUs were correctly received.

A DEV should not initiate a Dly-ACK procedure unless it has determined that the receiving DEV supports Dly-ACK by checking the DEV Capabilities field in the Capability IE, as defined in 6.4.11.

The source DEV may change the ACK policy in a stream from Dly-ACK to Imm-ACK or no ACK by sending a frame with the ACK Policy field set to one of those values. This has the effect of canceling the Dly-ACK policy and the source shall use the Dly-ACK negotiation procedure before restarting the Dly-ACK mechanism. The receiver shall no longer maintain the ACK status of any previous frames sent with the ACK Policy field set to Dly-ACK.

The source DEV may send any MPDUs including retransmissions and new MPDUs, as long as the total size of transmitted data is less than size indicated by the Max Burst value set by the destination DEV.

If the Dly-ACK frame is not received when requested, the last data frame of the burst is repeated until the Dly-ACK frame is received. The source DEV may send an empty data frame that was not in the original burst as an alternative to resending the last data frame, as long as the total number of frames, including the empty one, does not exceed Max Frames. The source DEV shall not start or resume burst transmissions until a Dly-ACK frame is received.

The destination MAC shall deliver MSDUs for each isochronous stream in ascending MSDU number order to its FCSL. If necessary to accomplish this, a destination MAC may discard correctly received, as described in 7.1, (and potentially acknowledged) frames.

Figure 7-55 illustrates the Dly-ACK negotiation. In the figure, the term *Alt* indicates that two alternate outcomes are illustrated in the MSC.




Figure 7-55—MSC for Dly-ACK negotiation

7.9.4 Implied acknowledgment (Imp-ACK)

Imp-ACK is one method that allows a CTA to be used for bidirectional data transfer. With Imp-ACK, the ACK is implied when the target DEV sends any frame, called the response frame, in response to a frame that has an ACK policy of Imp-ACK.

A DEV should not initiate an Imp-ACK procedure unless it has determined that the receiving DEV supports Imp-ACK by checking the DEV Capabilities field in the Capability IE, as defined in 6.4.11.

The originating DEV sets the ACK Policy to Imp-ACK in a frame to begin the Imp-ACK procedure. Only the DEV that is the source of a CTA shall initiate the Imp-ACK procedure. When the target DEV successfully receives a frame with the ACK policy set to Imp-ACK, the DEV shall respond with either an Imm-ACK frame or with a command or data frame. The target DEV shall only send a frame to the originating DEV, unless transmit control of the CTA has been relinquished, as described in 7.4.3.8.

If the target DEV sends a frame to the originating DEV in response to a frame with ACK policy of Imp-ACK, the target DEV shall set the ACK policy of the response frame to one of the following; no-ACK, Imm-ACK, or Imp-ACK.

If there is not sufficient time in the CTA for the response frame and any required acknowledgments before the end of the CTA, the target DEV shall respond with an Imm-ACK. If the target DEV does not know the end time of the current CTA, it shall only send an Imm-ACK frame as the response frame. The target DEV shall only send one frame in response to a frame with the ACK policy set to Imp-ACK. The response frame may be of any length as long as the frame and any required ACKs do not exceed the end of the current CTA.

If the target DEV successfully receives a MAC header for a frame with the ACK policy set to Imp-ACK but does not successfully receive the Frame Payload field, i.e., the FCS check fails, it may still send a data or command frame in response. In this case, the target DEV shall set the Imp-ACK NAK field, 6.2.1.7, to indicate that it did not successfully receive the frame from the originating DEV in the response frame.

As in any CTA, a DEV should only transmit to a DEV that it knows is listening to the CTA. Imp-ACK shall not be used for broadcast or multicast frames, i.e., frames with the DestID set to BcstID, McstID, or a GrpID. Imp-ACK shall not be used in the CAP or a contention CTA.

The use of Imp-ACK to enable bidirectional communication in a CTA does not imply that the stream is bidirectional as well. All streams in this standard are unidirectional. Only the originating DEV is allowed to send frames using the stream index assigned by the PNC. Frames transmitted in the opposite direction in the CTA shall use either the asynchronous stream index or a stream index that has been assigned to the DEV that is sending the frame.

If the destination DEV in a CTA has data to send to the source DEV of that CTA, it should set the More Data field in Imm-ACK frames that it sends to the source. This will indicate to the source that the destination is requesting that the source either use Imp-ACK or relinquish the CTA to the destination, as described in 7.4.3.8.

Figure 7-56 shows an example of the use of Imp-ACK where successive frames have an ACK policy of Imp-ACK.




Figure 7-56—Example of the use of Imp-ACK where the responding frame is sent with an ACK policy of Imp-ACK

Figure 7-57 shows an example of the use of Imp-ACK where one of the frames fails the FCS verification and the destination responds with the Imp-ACK NAK field set.




Figure 7-57—Example of the use of Imp-ACK where one of the frames fails FCS verification

7.9.5 Block acknowledgment (Blk-ACK)

Blk-ACK shall only be used with an aggregated frame. The destination, upon receiving an aggregated frame, checks each subframe. Based on the status of the subframe, either correctly or incorrectly received, the corresponding ACK bit or MSB/LSB ACK bit of the Blk-ACK Bitmap field of the MAC subheader in the Data frame will be set as described in 6.2.8.1. The originating DEV, after reading the Blk-ACK Bitmap field in MAC subheader, handles subframe retransmission.

For example, if the destination correctly receives subframes 0–4, 6, and 8–10, then the destination would set bits 0–4, 6, and 8–10 to one and the rest of the bits in the Blk-ACK Bitmap fields to zero.

The control of retransmission is different according to the different aggregation mode, standard or low-latency.

7.9.5.1 Blk-ACK for standard aggregation

For standard aggregation, the rules in this subclause apply.

If the MAC subheader is incorrectly received by the target DEV, all the subframes in the frame are considered invalid.

If the MAC subheader is correctly received, but some of the subframes are incorrectly received, then the target DEV responds with a Data frame with the Blk-ACK Bitmap field in the MAC subheader set to indicate the subframes that were incorrectly received, as described in 6.2.8. If the data transmission is unidirectional from originating DEV to target DEV, the Data frame sent back by the target DEV shall be an empty Data frame.

In any frame sent by the originating DEV, any retransmitted subframes should be put in the original order. The MSDU Number field in MAC header shall be set to the value of the first subframe, no matter it is a retransmission or a newly aggregated subframe.

If the originating DEV wants to drop certain retransmissions, it advances the MSDU number in MAC header to inform the target DEV that any subframe with the MSDU number older than this will not be retransmitted anymore.

In case that the Blk-ACK is lost, the originating DEV may retransmit the same frame after RIFS. The subframes contained will be the same as those of the previous transmitted frame. The Retry field in the MAC header shall be set to one to indicate to the target DEV that the same frame is retransmitted. For any frame that contains different subframes from previous transmission, the Retry field in MAC header shall be set to zero.

If the subframe contains only MSB information, the LSB ACK in the Blk-ACK Bitmap field in the MAC header shall be set to zero and shall be ignored upon reception. Likewise, if the a subframe contains only LSB information, the LSB ACK in the Blk-ACK Bitmap field shall be set to zero and shall be ignored upon reception.

7.9.5.2 Blk-ACK for low-latency aggregation

For low-latency aggregation, the rules in this subclause apply.

When the source receives the information in the Blk-ACK Bitmap field, it inserts the incorrectly received MSDUs for retransmission in the next frame to be sent.

The originating DEV of a Blk-ACK shall maintain the Blk-ACK Bitmap field according the rules in this subclause.

The originator maintains a list of recently transmitted MSDU numbers for retransmissions.

The list of recently transmitted MSDUs shall be updated per transmission frame, and with each Blk-ACK received, to reflect any outstanding MSDU not acknowledged.

The originating DEV may abort retransmissions of a specific MSDU by advancing the value of the MSDU Request Number field, as defined in 6.2.8.2, beyond the specific MSDU sequence number. When the MSDU Request Number is advanced beyond a missing MSDU sequence number the receiver updates its MSDU Response Number accordingly and assumes the skipped MSDU sequence number is discarded.

For each newly transmitted MSDU, its MSDU number is appended to the end of the list. The number of MSDUs in the list may grow up to the limit of the *mMaxSubframeNum*, as defined in 7.16.

If the MAC subheader is correctly received, the Blk-ACK Bitmap field is parsed. For each bit that is set, the entry of the MSDU number is removed from the list, updating the relative order of all remaining MSDUs (that were not acknowledged) and subsequently their offset order in the following retransmission attempt on the next transmission frame.

If the MAC subheader is not correctly received, the Blk-ACK Bitmap field is considered as corrupt. In this case, the source DEV may retransmit all the MSDUs from the previous frame.

If the subframe contains only MSB information, the LSB ACK in the Blk-ACK Bitmap field in the MAC header shall be set to zero and shall be ignored upon reception. Likewise, if the a subframe contains only LSB information, the LSB ACK in the Blk-ACK Bitmap field shall be set to zero and shall be ignored upon reception.

7.9.6 Retransmissions

During the CAP, retransmissions shall follow the backoff rules, as specified in 7.4.2.

During CTAs within the CTAP when an Imm-ACK or Dly-ACK is expected but is not received, the source DEV may start the retransmission of the frame (or the transmission of a new frame if the failed frame's retransmission limit has been met) a RIFS after the end of the last frame transmitted as long as there is enough channel time remaining in the CTA for the entire frame exchange.

A DEV determines the number of times a frame is retried before the DEV discards that frame. If the DEV gives up on a fragment of an MSDU/MCDU, the DEV shall discard all MPDUs of that MSDU/MCDU.

7.9.7 Duplicate detection

Because the DEV sending the data frame may not correctly receive an ACK, duplicate frames may be sent even though the intended recipient has already received and acknowledged the frame. Hence all DEVs shall detect such multiple receptions and indicate the data frames to the higher layers only once. The SrcID, Stream Index, Fragmentation Control field, Retry field, and PNID are used to detect multiple receptions of the same frame.

7.10 Peer discovery

Each DEV that is a member of the piconet may use the PNC Information Request command, as described in 6.5.4.1, to obtain information about other DEVs in the piconet. In addition the DEV may use the Probe Request command, as described in 6.5.4.5, to obtain other information required for peer-to-peer communication. The remote scan procedure is used by the PNC to determine channel conditions. All DEVs in the piconet are able to use the Channel Status Request, as described in 6.5.7.1, and Channel Status Response, as described in 6.5.7.2, commands to gather information about the quality of their link with another DEV.

7.10.1 PNC information request

A DEV may request information about either a single DEV in the piconet or about all of the DEVs in the piconet by sending a PNC Information Request command, as described in 6.5.4.2, to the PNC. If the DEV is requesting information about only a single DEV in the piconet, it shall set the DEVID in the command to the ID of that DEV. If the DEV is requesting information about all of the DEVs in the piconet, it shall set the DEVID in the command to the BcstID, as described in 6.2.3. If the originating DEV requests information about a single DEV that is not a member of the piconet, then the PNC shall send a PNC Information command with length zero. Otherwise, the PNC shall send the PNC Information command with the information requested by the originating DEV.

Figure 7-58 illustrates the sequence of messages involved in acquiring PNC information regarding a specific DEV or all of the DEVs that are members of the piconet.




Figure 7-58—MSC for acquiring information regarding a specific DEV or all of the DEVs from the PNC using the PNC Information Request and PNC Information commands

7.10.2 Probe request and response

The Probe Request command provides the ability to request IEs from a target DEV. If the target DEV receives the Probe Request command, it shall respond to the originator with Probe Response command(s) that shall have the IEs requested by the originator.

A DEV may request information about an isochronous stream by sending a Probe Request command requesting the CTA Status IE, as described in 6.4.10, with the Request Index field set to the stream index of the stream for which CTA information is requested. If the Request Index field is set to zero, the DEV is requesting information about all isochronous streams directed to the requesting DEV and to the BcstID and McstID. The PNC shall respond to a Probe Request command containing a request for the CTA Status IE by sending Probe Response command(s) containing the appropriate CTA Status IE(s).

If a DEV requests the CTA Status IE for a stream index that is not currently allocated by the PNC, the PNC shall respond with a Probe Response command containing a CTA Status IE with the following information:

- The SrcID and DestID fields both set to the UnassocID.
- The Stream Index field set to the requested stream index.
- All other fields set to zero.

Any DEV may send the Probe Request command with the Information Requested field set to zero and ACK Policy field set to Imm-ACK to any other DEV in the piconet to determine if the destination DEV is still present in the piconet and is within range of the sending DEV.

A DEV that is going to send a Probe Request command to a DEV operating in a power save mode should consider the operation of those modes, as described in 7.14, to determine the appropriate time to send the Probe Request command and the time to expect a response.

The types of IEs that are allowed to be requested or to be sent in the response depend on the status of the originator as either a DEV or PNC. The rules for requesting a specific IE are listed in Table 6-23 while the rules for responding to the request for a specific IE are listed in Table 6-24.

7.10.3 Announce

The Announce command provides the ability to send unrequested IEs to a target DEV. This command shall have one or more IEs that the originator is sending to the target.

A DEV that is going to send an Announce command to a DEV operating in a power save mode should consider the operation of those modes, as described in 7.14, to determine the appropriate time to send the Announce command.

The types of IEs that are allowed to be sent depend on the status of the originator as either a DEV or PNC. The rules for sending a specific IE are listed in Table 6-25.

7.10.4 Channel status request

The Channel Status Request command, as described in 6.5.7.1, may be used by any DEV in the piconet to get information from a target DEV about the link quality between the two DEVs. The originating DEV sends the Channel Status Request command to the target DEV to start the process. When the target DEV receives the request, it shall send a Channel Status Response command to the originating DEV with all of the information specified in 6.5.7.2.

The information conveyed in the channel status process is based on the results of an attempted data transfer between the two DEVs. Thus, the Channel Status Request command should only be used for DEVs that are actively participating in a data transfer as the information would not have much meaning otherwise. The PNC also uses this command to get the channel status information from the DEVs in the piconet, as described in 7.12.1, 6.5.7.1, and 6.5.7.2.

7.10.5 Remote scan

Remote scan is a procedure by which a PNC may request that a target DEV in the piconet initiate a channel scan on the PNC's behalf of the specified channels and to report the results of the scanned channels back to the PNC. The PNC may then use the results of the remote scan to initiate a change in maximum transmit power for the piconet, initiate a channel change to a channel with better channel characteristics, or other action. The algorithm for determining which procedure to execute is outside of the scope of this standard.

The PNC may optionally allocate channel time in the CTAP so that there is quiet time for the remote DEV to scan the channel for other piconets.

One of the reasons that the PNC requests a remote scan is because it determines that the current channel is unsatisfactory for continued operation of the piconet. This allows the PNC to get information about other wireless networks that may be out of range of the PNC but in range of some of the DEVs in the piconet. The PNC is able to get this information while it continues allocating channel time and generating beacons for its piconet.

The PNC initiates a remote scan by sending the Remote Scan Request command with a list of channels to a DEV in the piconet. The target DEV should accept the request to perform a channel scan on behalf of the PNC. If the DEV does not accept the request, it shall respond to the PNC by sending a Remote Scan Response command with a ReasonCode set to "request denied." The PNC upon receiving this response, may send a Remote Scan Request command to another piconet DEV, initiate its own channel scan or take other action.

In the case where the target DEV does accept the request, the target DEV shall initiate a series of OpenScan channel scans, as described in 7.2.1, based upon the channel list passed to it in the Remote Scan Request command. The target DEV shall scan each of the channels requested by the PNC. When the target DEV has finished scanning the channels, it shall respond via the Remote Scan Response command to the PNC with

the Remote Piconet Description Set parameters and the Channel Rating List field. The PNC upon receiving this information may then determine that there is a new channel with better characteristics than its current channel. The PNC may use this information to initiate the change channel procedure. The PNC may decide instead that a more appropriate solution to the current channel impairment is to increase/decrease transmission power level of the piconet. The PNC may also decide to do nothing or to take other unspecified action.

A DEV that is receiving beacons from more than one PNC may send an unsolicited Remote Scan Response command to its PNC with a Piconet Description Set representing the interfering PNCs. A DEV shall not report overlapping piconets if it determines that the beacons were received from a child or IEEE 802.15.3 neighbor piconet that is associated with the DEV's current piconet.

7.10.6 PNC channel scanning

PNC channel scanning is a procedure by which the PNC is able to determine the channel characteristics of not only its current operating channel but also the channel characteristics of other channels. The PNC may use the results of its channel scans to determine whether the current channel in which it is operating has acceptable characteristics or that there is another channel(s) with better channel characteristics than its current channel.

The PNC may allocate CTAs such that there is unallocated channel time in the CTAP. This provides quiet time for the PNC to scan channels for other IEEE 802.15.3 piconets, non-IEEE 802.15.3 wireless networks, or interference.

If the PNC initiates a scan of one or more alternate channels, the PNC shall not transmit a beacon for one or more beacon intervals. The PNC shall not suspend beacon transmissions for more than twice *mMinChannelScan*. The PNC, upon returning to its current channel and resuming the transmission of its beacons, shall increment the Time Token field by the number of beacons not sent during the time the PNC was scanning one or more alternate channels.

After scanning both its current channel and other channels, the PNC may initiate one of these options:

- a) Do nothing since the PNC has determined that none of the alternate channels were better than its current channel.
- b) Initiate the dynamic channel change procedure described in 7.12.1.
- c) Increase or decrease the max tx power level of the piconet, as described in 7.12.2.
- d) Initiate some other unspecified action.

The algorithm for determining when to initiate any of these actions is outside the scope of this standard.

7.10.7 Channel probing

Channel probing is an optional function to allow the source DEV choose the best MCS for the current channel. A DEV may use the Channel Status Request, as described in 6.5.7.1, and Channel Status Response, as described in 6.5.7.2, commands to gather information of the link quality with another DEV. In addition, when the DEVs are capable of beam forming, the quality of the link may be obtained by beam forming procedure, as described in 12.5. During data streaming, the destination DEV reports the channel status information, as defined in 11.1.8.3, to the source DEV. After learning the channel status, the source DEV may switch to an MCS that is best adapted to the channel condition. For the SC/HSI PHY, the channel status information may be written in the fragmentation field of the MAC header of a Imm-ACK or Dly-ACK, as defined in 6.3.4. For AV PHY, it may be written in the fragmentation field of the MAC header of an Imm-ACK or Dly-ACK frame.

7.11 Changing piconet parameters

This subclause describes the methods used to change certain key characteristics of the piconet.

A PNC shall not change either the pseudo-static CTAs or the pseudo-static CTA blocks during a piconet parameter change. If the parent needs to move a pseudo-static CTA because the superframe duration is being reduced, it shall do so prior to using the superframe duration change process, as described in 7.11.2. A child or IEEE 802.15.3 neighbor PNC shall use the value of the Change Beacon Number field in the Piconet Parameter Change IE from the parent's beacon to calculate the Change Beacon Number field in the Piconet Parameter Change IE in its own beacon. The exceptions to this are as follows:

- When the parent is changing its PNID or BSID.
- When a child or neighbor PNC decides not to change channels with the parent PNC and is shutting down, as described in 7.12.1.

When the Change Beacon Number field is used with changing parameters, the PNC shall ensure that the Piconet Parameter Change IE announcement complies with the rules for beacon announcements in 7.6.4. The PNC shall not place more than one Piconet Parameter Change IE, as described in 6.4.6, in any beacon.

The PNC shall not be required to wait until all of the DEVs in power save modes are in ACTIVE mode before changing a piconet parameter. A power save DEV is not required to switch to ACTIVE mode when there is a piconet parameter change in progress. However the power save DEV shall update its internal values for the new piconet parameter at the time indicated in the Piconet Parameter Change IE.

7.11.1 Moving beacon

The PNC may move the relative position of its beacon. Moving a beacon means that the superframe duration is unchanged while the position of the beacon is moved.

If the PNC wishes to move its beacon position, it shall insert the Piconet Parameter Change IE, as described in 6.4.6, into NbrOfChangeBeacons beacons with the change type set to MOVE and the superframe timing set to the delay of the first beacon after this sequence compared to previous beacon transmission time. See Figure 7-59.




Figure 7-59—Moving the beacon position

7.11.2 Changing superframe duration

The PNC may change the duration of its superframe. If the PNC wishes to change its superframe duration, it shall insert the Piconet Parameter Change IE, as described in 6.4.6, into NbrOfChangeBeacons beacons with the Change Type set to SIZE and the Superframe Timing set to the size of the superframe which follows the first beacon after this sequence. See Figure 7-60.




Figure 7-60—Changing superframe duration

7.11.3 Setting the PNID or BSID

The BSID is used to provide a way to identify the piconet. The PNC is able to change the BSID via the Piconet Parameter Change IE in the beacon, as described in 6.4.6, using the process described in this subclause. The BSID is preserved in the PNC handover process and it may be persistent when the PNC restarts a piconet that ended without handing over control to a PNC-capable DEV.

The PNC shall choose a PNID when it starts a piconet; the PNID should be selected randomly. An existing piconet's PNID shall be changed only if the PNC detects another piconet with the same PNID on any channel. The same PNID may be persistent when the PNC restarts a piconet that ended without handing over control to a PNC-capable DEV.

If the PNC detects that another piconet is using the same BSID in its operational area, it may change the BSID, but it is not required to change it. If the PNC detects that another piconet is using the same PNID within range of the PNC on any channel or if it is informed of this via the Overlapping PNID IE, as described in 6.4.15, the PNC shall choose another PNID and change it via the Piconet Parameter Change IE in the beacon. The PNC shall not simultaneously change both the PNID and the BSID.

If a DEV detects a piconet within its range on any channel with the same PNID, it shall send an Announce command, as described in 6.5.5.2, to the PNC including an Overlapping PNID IE, as described in 6.4.15, that contains the current PNID and channel index. Once this command has been sent successfully, the DEV shall not send this information again until after the current PNID has been changed by the PNC.

Before changing its PNID, a parent PNC shall scan for the PNIDs of other piconets, including all of its dependent piconets. The PNC shall not change its PNID to the same value as that of any other piconet that it detects.

If the PNC decides to change the PNID or BSID, the PNC shall send a beacon with the Piconet Parameter Change IE indicating the new PNID or BSID. The DEVs that received the beacon with Piconet Parameter Change IE shall change the PNID or BSID to the new value at the time of the beacon with a beacon number equal to the Change Beacon Number field in the previous Piconet Parameter change IEs.

The MSC in Figure 7-61 describes the BSID change process.




Figure 7-61—MSC for changing the BSID for a piconet

7.11.4 Maintaining synchronization in dependent piconets

A dependent PNC receiving a parent beacon with a Piconet Parameter Change IE, that has the Change Type field set to MOVE or SIZE, shall immediately insert the appropriate Piconet Parameter Change IE into its beacons. If the Change Type field is set to SIZE, a dependent PNC shall change the length of the private CTA used to reserve time for the operations of the parent piconet in the first beacon it broadcasts in the parent superframe of the new superframe size.

7.12 Interference mitigation

The PNC should periodically listen in the current channel to detect interference, the presence of other IEEE 802.15.3 piconets or the presence of other wireless networks. The PNC should also periodically listen to other overlapping channels for the presence of the same types of DEVs.

If the PNC detects the presence of another IEEE 802.15.3 piconet, four of the methods that are available to mitigate the interference between the two piconets are as follows:

- The PNC may join the other piconet to form a child piconet, as described in 7.2.7.
- The PNC may join the other piconet to form a neighbor network, as described in 7.2.8.
- The PNC may change channels to one that is unoccupied, as described in 7.12.1.
- The PNC may reduce the maximum transmit power in the piconet, as described in 7.12.2.

If the PNC determines that there is either an interferer or a non-IEEE 802.15.3 wireless network operating in the PNC's current channel or overlapping with the current channel, two of the methods that are available to mitigate the interference are as follows:

- The PNC may change channels to one that is unoccupied, as described in 7.12.1.
- The PNC may reduce the maximum transmit power in the piconet, as described in 7.12.2.

7.12.1 Dynamic channel selection

The PNC initiates dynamic channel selection if it determines that the current conditions of its channel are poor. The PNC may use one or more of the following methods to make this determination:

- The PNC may perform a PNC channel scanning procedure, defined in 7.10.6.
- The PNC may use the remote scan procedure, defined in 7.10.5.
- The PNC may collect the channel status from its member DEVs by sending a Channel Status Request command, defined in 6.5.7.1, to request that the DEVs provide their channel status via a Channel Status Response command, as described in 6.5.7.2.

The algorithm required to use the channel status information when deciding whether to change channels is outside the scope of this standard.

The PNC shall initiate a dynamic channel change procedure only after it has performed a PNC channel scan, as described in 7.10.6, and has determined that there is one or more channels with better characteristics than exist in its current operating channel. Note that in addition to the PNC channel scan, the PNC is able to use other methods, described above, to determine which channel to use as the new channel. If the PNC decides to initiate a dynamic channel change, the PNC shall broadcast the Piconet Parameter Change IE, as described in 6.4.6, with the change type set to CHANNEL in its current channel via its beacon for *NbrOfChangeBeacons* consecutive beacons. The Piconet Parameter Change IE shall contain the channel index of the new channel to which the PNC will be moving the piconet, and the Change Beacon Number field that contains the beacon number of the first beacon that will be sent on the new channel. The channel change shall take effect starting with the first beacon with a beacon number equal to the Change Beacon Number field in the previous Piconet Parameter Change IEs. The DEVs that received a beacon containing the Piconet Parameter Change IE shall change from their current channel to the new channel before the first expected beacon on the new channel.

No DEV shall transmit on the new channel until a beacon has been correctly received, as described in 7.1, on the new channel.

Dependent piconets shall either change to the new channel with the parent PNC or cease operations. If a dependent PNC ceases operation due to a channel change, it may attempt to restart its piconet in the original channel, as described in 7.2.2.

The PNC is not required to wait until all of the DEVs in power save modes are in ACTIVE mode before changing channels. A power save DEV is not required to switch to ACTIVE mode when there is a channel change in progress. However, the power save DEV should change channels at the time indicated in the Piconet Parameter Change IE.

In the last beacon transmitted prior to a channel change, the PNC should provide time at the end of the superframe for the DEVs in the piconet to change channels. The time required for a DEV to change channels is *pChannelSwitchTime*.

NOTE—One method to allow time for the channel change is for the PNC not to place CTAs just prior to the first beacon that will be sent on the new channel.

The PNC shall ensure that the Piconet Parameter Change IE announcement complies with the rules for beacon announcements in 7.6.4.

Figure 7-62 illustrates the message sequence involved in transitioning the current piconet to a new channel.




Figure 7-62—MSC for changing the piconet's channel

When a parent PNC changes channels, the dependent PNC may change channels as well. If the dependent PNC is going to change channels with the parent PNC, it puts the appropriate Piconet Parameter Change IE into its beacons when it receives a beacon from the parent indicating a pending channel change. The dependent PNC shall then switch channels to the channel indicated in the Piconet Parameter Change IE at the appropriate time.

A dependent PNC may change to a different channel even if the parent PNC does not change channels. In this case, the child or IEEE 802.15.3 PNC shall remove the Parent Piconet IE from its beacon after the channel change since it is no longer a child or neighbor of that piconet. The dependent PNC should also disassociate from the parent piconet when it changes channels without the parent PNC.

7.12.2 Transmit power control

Two independent forms of transmit power control (TPC) are available for IEEE 802.15.3 systems: a maximum power for the CAP, the beacon, and directed MCTAs and an adjustable power in a CTA. The goal of TPC in the CAP is to prevent one DEV from having better access to the medium in the CAP due to a higher transmit power level. Adjustable transmitter power in the CTA is intended to support reduced power usage as well as reducing the overall interference levels generated by the piconet. The transmit power used in a CTA is not limited by the transmit power used for the CAP, beacon, or directed MCTAs.

7.12.2.1 Maximum transmitter power for the CAP, beacon, and directed MCTAs

The PNC may choose a maximum transmit power level for the CAP, beacon, and directed MCTAs, excluding association MCTAs. The PNC shall convey this information to the DEVs via the Beacon frames

using the Max TX Power Level field in the beacon. The PNC shall not set the Max TX Power Level below the $pMinTpcLevel$ for the PHY, which is defined in 10.5.9 for the 2.4 GHz PHY. All DEVs within the piconet shall set their nominal transmit power level for frames in the CAP or directed MCTAs, excluding association MCTAs, to be no more than the value indicated in the Max TX Power Level field in the beacon, as described in 6.3.1. DEVs shall comply with the maximum transmit power within $10 \times mMaxLostBeacons$ superframes following the beacon in which the DEV detects the change. Likewise, the PNC shall set its nominal transmit power for the beacon to be no more than the value that is indicated in the beacon.

7.12.2.2 Adjustable transmitter power in a CTA

Each DEV participating in a CTA may request that the other DEV with which it is communicating in the CTA either increase or decrease its transmitter power level. The originating DEV shall use the Transmit Power Change command, as described in 6.5.7.5, to request a change in the power level setting of the target DEV for all CTAs assigned between the two DEVs. The target DEV shall increase or decrease its transmit power level as indicated in the Transmit Power Change command if the power level setting is supported by the target DEV. If the power level change is not supported by the target DEV, it shall use the closest implemented TX power level that is greater than the requested level provided that is within the allowable range. The target DEV shall apply the change in the power level for all CTAs assigned between the two DEVs.

A DEV may also change its transmit power based on its own estimation of the channel.

7.13 Multi-rate support

A compliant PHY may support more than one data rate. In each PHY there will be one mandatory base rate specified for the purposes described in this subclause. In addition to the base rate, the PHY may support rates that are both faster and slower than the base rate. A DEV shall send a frame with a particular data rate to a destination DEV only when the destination DEV is known to support that rate. The allowed data rates and the mandatory base rate are PHY dependent and defined as follows:

- In 10.3 for the 2.4 GHz PHY
- In 11.2.2.1 for the SC PHY mode
- In 11.3.2.1 for the HSI PHY mode
- In 11.4 for the AV PHY mode

In order to determine the rates that are supported by a target DEV in the piconet, the DEV shall use one of the following three methods:

- a) Check the capabilities of the target DEV broadcast by the PNC when the DEV becomes a member of the piconet.
- b) Send a Probe Request command, as described in 6.5.4.5, to the target DEV to request its Capability IE, as described in 6.4.11.
- c) Request the information from the PNC using the PNC Information Request command, as described in 6.5.4.1.

Similarly, each DEV may periodically use the Channel Status Request command, as described in 6.5.7.1, to obtain the channel status information from other DEVs and decide the PHY rate to be used in transmissions to those other DEVs. Additionally, the channel quality may be evaluated by the presence or absence of ACKs to transmitted frames. This information may be used to determine if the rate of transmission or the power level needs to change.

If the Dly-ACK mechanism is used, all frames in a burst shall be sent with the same rate. The Dly-ACK frame shall be sent with the same rate as the rate of the last frame in the burst that requested the Dly-ACK.

The allowed PHY rates for each of the different types of frames are listed in Table 7-4.

Table 7-4—Allowed PHY data rates for frames

Frame type	Allowed PHY data rates
All broadcast and multicast addressed frames (including the beacon and commands)	Base rate
Imm-ACK	Same rate as the frame that is being ACKed
Dly-ACK	Same rate as the last frame of the burst being ACKed
Association request	Base rate
Association response	Base rate
Disassociation request	Base rate
Directed command frame	Any rate supported by both the source and destination.
Directed data frame	Any rate supported by both the source and destination.

HRP frames in the AV PHY with the ACK policy set to Imm-ACK shall be ACKed with the directional ACK, as described in 6.2.9.5. The data rate for the payload in a directional ACK with payload is selected by the DEV sending the directional ACK. Blk-ACK policy shall not be used with HRP AV PHY frames.

7.14 Power management

There are four power management (PM) modes defined in this standard: ACTIVE, APS, PSPS, and DSPS. The latter three modes are collectively referred to as power save (PS) modes. A DEV that is in ACTIVE, APS, PSPS, or DSPS mode is said to be an ACTIVE DEV, an APS DEV, a PSPS DEV, or a DSPS DEV, respectively. In any given PM mode, a DEV may be in one of two power states, either AWAKE or SLEEP states. AWAKE state is defined as the state of the DEV where it is either transmitting or receiving. SLEEP state is defined as the state in which the DEV is neither transmitting nor receiving. A DEV, regardless of its PM mode, is allowed to enter the SLEEP state during a CTA for which it is neither the source nor the destination. A DEV is also allowed to enter the AWAKE state during any time when it is in a power save mode.

The wake beacon for a DEV is defined as the PNC-defined system wake beacon for DEVs in PSPS mode, as described in 7.14.1, and the wake beacon of the DSPS set for a DEV in DSPS mode, as described in 7.14.2. A DEV in DSPS mode may be in multiple DSPS sets and therefore may have multiple wake beacons because each of those DSPS sets may have its own wake beacon. The wake beacon for a DEV in APS mode occurs at times determined by the DEV and is unknown to the PNC and other DEVs in the piconet. Unlike the DSPS and PSPS wake beacons, the wake beacon of the DEV in APS mode is not periodic and is only guaranteed to happen once per ATP period for that DEV.

A DEV shall always establish membership with the piconet in ACTIVE mode.

Table 7-5 lists the rules for the four modes of operation defined in this standard. Each entry indicates the state required, either AWAKE or SLEEP, for the DEV.

The PNC shall support one PS set for APS and one PS set for PSPS. In addition the PNC shall support at least one DSPS set, i.e., a PS set with PS Set Index between 2 and 253.

Table 7-5—Power management rules for superframe elements

Superframe portion	ACTIVE	APS in wake superframe	PSPS or DSPS DEV in wake superframe
Beacon	AWAKE	AWAKE	AWAKE
CAP	AWAKE	May SLEEP	AWAKE
CTA with BcstID as DestID (including MCTAs)	AWAKE	May SLEEP	AWAKE
CTA with McstID as DestID (including MCTAs)	May SLEEP	May SLEEP	May SLEEP
CTA with DEV as SrcID or DestID (including MCTAs)	AWAKE	May SLEEP	AWAKE
All other CTAs and unallocated time (between CTAs)	May SLEEP	May SLEEP	May SLEEP

7.14.1 Piconet synchronized power save (PSPS) mode

A DEV in PSPS mode shall listen to all system wake beacons, as announced by PNC, and is required to be in the AWAKE state during system wake superframes, as indicated in Table 7-5. The wake beacon for PSPS DEVs is determined by the PNC. PSPS mode is identified by a PS Set Index equal to one. If a DEV in PSPS mode does not correctly receive the system wake beacon, it shall be in the AWAKE state during the expected beacon transmission times to receive the following beacons until a beacon is correctly received, as described in 7.1.

The PNC shall announce the system wake beacon in the Next Wake Beacon field in the PS Status IE with PS Set Index field equal to one in the beacon. If none of the DEVs are in PSPS mode, the PNC may omit the PS Status IE for PS Set Index one from the beacon. In that case every beacon is a system wake beacon for the purpose of beacon information announcements, as described in 7.6.4.

It is the responsibility of the DEV using PSPS mode to synchronize with the system wake beacon before entering the sleep state. Because the PSPS DEV at some point will need to send commands to the PNC, e.g., the PM Mode Change command, the PNC needs to take this into consideration when allocating MCTAs if the CAP is not available for sending commands.

Any DEV that is going to use the PSPS mode shall send the SPS Configuration Request command, as described in 6.5.8.3, to the PNC with the Operation Type field set to 'join', SPS Set Index set to one and the Wake Beacon Interval set to its desired system wake beacon interval. The valid range for a requested Wake Beacon Interval is defined in 6.5.8.3. Upon reception of this command, the PNC shall ACK the command and respond with an SPS Configuration Response command with the next wake beacon. The PNC uses the information in the Wake Beacon Interval field from all participating PSPS DEVs to determine the system wake beacon interval. The actual system wake beacon interval may not correspond to any of the PSPS DEVs requested wake beacon interval.

A DEV may send the SPS Configuration Request command more than once if its system wake beacon interval requirement changes. If the DEV no longer wishes to use PSPS mode, it shall send the SPS Configuration Request command with the Operation Type field set to 'leave' and the PS Set Index set to one to leave the PSPS set.

A DEV shall send a PM Mode Change command to the PNC with the PM Mode field set to SPS and receive the ACK from the PNC before entering the PSPS mode. When the PNC receives this command, it shall

terminate all super-rate streams for which the DEV is the destination, as described in 7.5.1.3, and set the DEVID Bitmap field in the PS Status IE appropriately, as described in 6.4.13.

The PS Status IE in the beacon for PS Set Index value of one with the bit for the DEV's DEVID set shall serve as indication to other DEVs in the piconet that its peer has switched into PSPS mode.

When the DEV is going to switch to ACTIVE mode, it shall send a PM Mode Change command to the PNC with the PM Mode field set to ACTIVE. Once this command is sent, the DEV shall regard itself as in the ACTIVE mode whether the command was acknowledged by the PNC or not. If the PNC does not set the DEVID Bitmap in the PS Status IE appropriately, the DEV should resend the PM Mode Change command to the PNC. The PNC is not required to align sub-rate allocations for a PSPS DEV with the system wake beacon.

7.14.2 Device synchronized power save (DPS) mode

DSPS mode allows a DEV that is sensitive to power utilization to synchronize its AWAKE state with other DEVs. The DSPS mode is based on grouping DEVs that have similar power save requirements into DSPS sets. These DSPS sets are managed by the PNC, but the parameters of the sets are determined by the DEVs.

7.14.2.1 Creation, use, and management of DSPS sets

In order to use DSPS mode, a DEV is required to first join a DSPS set. Each DSPS set has two associated parameters: the Wake Beacon Interval and the Next Wake Beacon. The DSPS set is identified by an index value called the DSPS Set Index that is between 2 and 253, inclusive. The Wake Beacon Interval is the number of superframes between two successive wake beacons of that DSPS set. This value is set by the DEV when it creates the DSPS set. The Next Wake Beacon parameter is the beacon number, as described in 6.3.1.1, corresponding to the immediate next wake beacon of that DSPS set. This parameter is set by PNC when it creates the DSPS set. Both of these parameters shall be maintained by the PNC.

Any DEV that is a member of the piconet may request the information about the existing DSPS sets by sending a PS Set Information Request command, as described in 6.5.8.1, to the PNC. The PNC shall respond by sending a PS Set Information Response command, as described in 6.5.8.2, that provides the parameters of all of the PS sets currently in use within the piconet.

The DEV may select a DSPS set to join. If there are not any DSPS sets currently in existence that match the DEV's requirements, the DEV may request the formation of a new DSPS set by setting the SPS Set Index field in the SPS Configuration Request command, as described in 6.5.8.3, to the "Unallocated DSPS set" value and the Operation Type field set to "join." The DEV shall set the Wake Beacon Interval field to its requested value. The valid range for the Wake Beacons Interval field is defined in 6.5.8.3. This value shall not be changed while the DSPS set has any members. The PNC shall respond to the request by sending the SPS Configuration Response command, as described in 6.5.8.4, to the DEV indicating success or the reason that the request failed. If the DSPS set is created, the PNC assigns a DSPS Set Index to it. The PNC shall assign a unique number between 2 and 253 for the DSPS set index. The PNC includes in the SPS Configuration Response command a value for the Next Wake Beacon field set to the beacon number, as described in 6.3.1.1, of the first wake beacon for the new DSPS set. Once a DSPS set is created, the PNC shall keep the Next Wake Beacon for that set updated at all times. The maximum number of DSPS sets supported by the PNC is implementation dependent up to a maximum of 252.

The PNC may require that all PS sets have a unique Wake Beacon Interval. For example, the PNC may reject a request to create a PS set with a Wake Beacon Interval of 4 if there is a PS set that already has this value. If the DEV requires this Wake Beacon Interval, it may join the existing PS set.

A DEV may join an existing DSPS set by sending the SPS Configuration Request command to the PNC with the SPS Set Index field set to the index of an existing DSPS set and the Operation Type field set to "join."

All other parameters shall be ignored. The PNC shall confirm or reject the request by sending the SPS Configuration Response command to the DEV. Since a DEV may support multiple applications with different requirements, a DEV may register in more than one DSPS set at a given time.

A DEV that no longer needs to be in a DSPS set shall send the SPS Configuration Request command to the PNC with Operation Type field set to "leave." The PNC shall not send the SPS Configuration Response command to the requesting DEV if the Operation Type field was set to "leave."

When the last member of a DSPS set has left, the DSPS set shall be terminated by the PNC. The PNC may also delete a DSPS set. The PNC deletes a DSPS set by placing the PCTM IE, 6.4.8, in the beacon and setting the bit for each of the DEVs that are members of that DSPS set and are in a PS mode. The PNC also removes from the beacon, if present, the PS Status IE of the DSPS set that the PNC is deleting. When the PNC determines that a DEV in that DSPS set is in AWAKE state, it shall send an SPS Configuration Response command to the DEV with SPS Set Index field set to the value of the DSPS set index that is being deleted and the Reason Code field set to "DSPS set deleted by PNC." If a DEV is in DSPS mode and it receives a beacon that does not have the PS Status IE for its DSPS set, the DEV should send a PS Set Configuration Request command to the PNC to determine if the DSPS set has been deleted.

7.14.2.2 Changing DSPS mode and operation

DSPS DEVs alternate between DSPS mode and ACTIVE mode depending on the amount and type of data traffic without leaving any of the DSPS sets that the DEV has joined. The PM Mode Change command, as described in 6.5.8.5, is used by a DEV to inform the PNC that it is changing its power management mode. A DEV shall have joined one or more DSPS sets before it will be allowed to switch into DSPS mode.

If the DEV is going to change its power management mode from ACTIVE to DSPS, the DEV shall send the PM Mode Change command, as described in 6.5.8.5, to the PNC with the PM Mode field set to SPS. The PNC shall then set the bit in the DEVID Bitmap field for the DEV in each PS Status IE that corresponds to an SPS set of which the DEV is a member. If the DEV is the source or destination of any streams, not including broadcast or multicast, which are not using a DSPS wake beacon interval, the PNC shall terminate those streams, as described in 7.5.1.3, when the DEV changes to DSPS mode. The PNC does not automatically terminate any streams when the DEV changes from DSPS to ACTIVE mode.

If the DEV has joined PS set one (PSPS) in addition to other DSPS sets before issuing the PM Mode Change command, the DEV shall transition into a combined PSPS and DSPS mode. In either case, the DEV shall not consider itself in either DSPS or PSPS mode until it has received an Imm-ACK to the PM Mode Change command.

The presence of a PS Status IE in the beacon with the bit for the DEV's DEVID set shall serve as indication to other DEVs in the piconet that its peer has switched into PSPS mode.

If the DEV is going to change its power save mode from DSPS to ACTIVE, the DEV shall send the PM Mode Change command, as described in 6.5.8.5, to the PNC with the PM Mode field set to ACTIVE. Once this command is sent the DEV shall regard itself as in the ACTIVE mode whether the command was acknowledged by the PNC or not. If the PNC does not set the DEVID Bitmap field in the PS Status IE appropriately, the DEV should resend the PM Mode Change command to the PNC. When the PNC correctly receives the PM Mode Change command with the PM Mode field set to ACTIVE, it shall no longer set the bit(s) for the DEV in the PS Status IE(s). If the PM Mode field indicates the current mode of the DEV, then no change is requested.

The PNC shall create one PS Status IE in the beacon for each DSPS set that has at least one member currently in DSPS mode. When a DSPS set is not in use, the PNC shall discontinue inserting its PS Status IE in the beacon. The PNC needs to ensure that the number and size of the PS Status IEs do not cause the beacon or extended beacon to exceed its maximum allowed size, as described in 6.2.

Other DEVs may use the information in the PS Status IE to learn when to transmit to a DSPS DEV. In addition, the PS Status IE informs DSPS capable DEVs which DSPS set is required in order to synchronize data transfers to the DEVs in DSPS mode.

It is possible that DEVs in DSPS mode will not receive broadcast messages. If a DEV requires the opportunity to receive all of the broadcast streams, it should not use DSPS mode. If a source DEV needs that the DEVs in DSPS mode have an opportunity to receive the broadcast frame, then the source DEV should send the frame in all of the superframes required to reach the DSPS DEVs in their wake beacons.

The PNC may grant an ACTIVE DEV's channel time request to create or modify a stream with PM Channel Time Request Type field set to ACTIVE and a DSPS mode DEV as the TargetID using the processes defined in 7.5.1.1 and 7.5.1.2.

If the CTA Rate Type field of a new allocation is set to sub-rate with the DSPS DEV as the destination but not aligned with one of the DEV's DSPS sets, the DSPS DEV may take any one of the following actions:

- Stay in DSPS mode while listening to the additional beacons required for the new allocation.
- Switch to ACTIVE mode using the PM Mode Change command.
- Terminate the stream, as described in 7.5.1.3.

7.14.2.3 CTA timing in DSPS mode

A DSPS Channel Time Request is a Channel Time Request command, as described in 6.5.6.1, with the PM Channel Time Request Type field set to DSPS. This requests that the PNC allocate channel time during the wake superframes of the specified DSPS set. An additional condition placed on the timing is that the value of the CTA Rate Factor field shall not be less than the number of superframes between wake beacons, i.e., the Wake Beacon Interval. Since the CTA Rate Factor field, like the Wake Beacon Interval, is a power of 2, the rate of DSPS CTAs also is a power of 2 sub-rate of the wake beacon rate, as illustrated in Figure 7-63. For example, in case 3 of Figure 7-63 the Wake Beacon Interval is $2=2^1$ while the CTA Rate Factor is $8=2^3$. Thus, the CTAs occur every fourth wake beacon, $4 = 8/2$. An example of values that are not allowed would be a Wake Beacon Interval of 4 and a CTA Rate Factor of 2. The reason that this is not allowed is that the CTAs would occur more often than the DEV was waking to listen to the beacon.

If the PNC grants a DSPS channel time request, the PNC shall allocate CTAs only in the wake beacons of that DSPS set unless there is insufficient channel time for the allocation. If the PNC determines that it is unable to provide the requested CTA in a wake beacon, the PNC shall not allocate the CTA and shall take one of the following three actions:

- a) Continue attempting to allocate the CTA at the appropriate times.
- b) Terminate the stream, as described in 7.5.1.3.
- c) Set the appropriate bits in the CWB IE, as described in 6.4.8, for the SrcID and DestID to indicate that those DEVs should wake up for the next beacon to see if there will be a CTA in the next beacon. The PNC may set the CWB bits in up to three consecutive beacons until it is able to allocate the CTA. Although this behavior is permitted, an undesirable reduction in battery life will result from unnecessary wake periods for DSPS DEVs.

A DSPS DEV shall listen to every wake beacon regardless of the frequency with which CTAs are allocated. If the DSPS DEV is the DestID of any CTA in the wake beacon, then the DSPS DEV shall listen during the assigned CTA in that wake superframe.






Case 1: Wake beacon interval = 2, CTA rate factor = 2**Case 2: Wake beacon interval = 2, CTA rate factor = 4****Case 3: Wake beacon interval = 2, CTA rate factor = 8****Case 4: Wake beacon interval = 4, CTA rate factor = 8****Figure 7-63—Illustration of DSPS set and sub-rate CTAs**

Figure 7-64 shows four equivalent PNC CTA arrangements for a Wake Beacon Interval of 2 and the CTA Rate Factor of 8. The PNC chooses the wake superframes to position the assigned CTAs at an overall rate of 1 CTA per 8 superframes.

**Figure 7-64—Equivalent CTAs for wake beacon interval = 2 and CTA rate factor = 8**

The example in Figure 7-65 shows that Wake Beacon Interval and CTA Rate Factor together provide the ability of the DEV to trade off power savings and superframe loading, i.e., the number of CTAs allocated to any one superframe. In this example, the PNC minimizes the superframe loading for the eight DSPS DEVs, if they are members of different DSPS sets, because there is one wake beacon every other superframe, rather than the minimum, one wake beacon every eight superframes if power saving were maximized.




**Figure 7-65—Minimum superframe loading of CTAs for 8 DEVs
requesting wake beacon interval = 2 and CTA rate factor = 8**

7.14.3 Asynchronous power save (APS) mode

APS mode allows a DEV to conserve power by remaining in SLEEP state for extended periods of time. The only responsibility of a DEV in APS mode is to communicate with the PNC before the end of its ATP in order to preserve its membership in the piconet.

In the APS mode the DEV is not required to listen to any beacons or other traffic until it changes to either ACTIVE or a different power save mode using the PM Mode Change command, as described in 6.5.8.5. APS mode shall not be used in combination with any other power save mode. A DEV shall not use the SPS Configuration Request command to set parameters for the APS Set Index.

All DEVs in APS mode need to send at least one acknowledged frame to the PNC during their ATP in order to avoid being disassociated from the piconet, as described in 7.3.4. Because the APS DEV will need to send a frame to the PNC at least once during its ATP, the PNC needs to take this into consideration when allocating MCTAs if the CAP is not available for sending commands.

A DEV shall send a PM Mode Change command to the PNC with the PM Mode field set to APS and receive the ACK before entering APS mode. When the PNC receives this command, it shall set the DEVID Bitmap field in the PS Status IE appropriately, as described in 6.4.14. The PNC shall terminate all streams and asynchronous data allocations that have the APS DEV as either the source or destination ID.

The PS Status IE in the beacon with the bit in the DEVID Bitmap field for the DEV's DEVID set shall serve as indication to other DEVs in the piconet that its peer has switched to APS mode. The PS Set Index of 0 shall only be used for APS DEVs. Although a PS Set Index is assigned to the DEVs in APS mode, the DEVs in this mode all act independently, unlike the DEVs that are members of other PS sets.

The DEV may leave APS mode by sending a PM Mode Change command to the PNC with the PM Mode field set to ACTIVE. Once this command is sent the DEV shall regard itself as in the ACTIVE mode whether the command was acknowledged by the PNC or not. If the PNC does not set the DEVID Bitmap in the PS Status IE appropriately, the DEV should resend the PM Mode Change command to the PNC.

The PNC may grant an ACTIVE mode DEV's channel time request, with PM Channel Time Request Type field set to ACTIVE and with an APS mode DEV as the DestID using the processes defined in 7.5.1.1 for isochronous allocations and in 7.5.2.1 for asynchronous allocations.

7.14.4 MSCs for power save modes

Figure 7-66 illustrates the message flow for a DEV inquiring about the PS sets that are currently defined within the PNC.




Figure 7-66—MSC for PS set information exchange

Figure 7-67 illustrates the message flow for a DEV requesting the creation of a new DSPS set.




Figure 7-67—MSC for DSPS set creation

Figure 7-68 illustrates the message flow for a DEV requesting to add itself to an existing SPS set.




Figure 7-68—MSC showing a DEV joining an existing SPS set

Figure 7-69 illustrates the message flow for a DEV requesting to leave an SPS set.




Figure 7-69—MSC showing a DEV leaving an SPS set

Figure 7-70 illustrates the message flow for a DEV DME requesting to change the current PM mode of operation from ACTIVE to an SPS mode when that DEV is a member of one or more SPS sets.




Figure 7-70—MSC showing DME initiated PM mode change from ACTIVE to an SPS mode

Figure 7-71 illustrates the message flow for a DEV DME requesting to change the current PM mode of operation from ACTIVE to APS mode.




Figure 7-71—MSC showing DME initiated PM mode change from ACTIVE to APS

Figure 7-72 illustrates the message flow for a DEV DME requesting to change the current PM mode of operation from one of the power save modes to ACTIVE mode.




Figure 7-72—MSC showing DME initiated PM mode change from any PS mode to ACTIVE

Figure 7-73 illustrates the message flow for a DEV MLME changing the PM mode of operation from DSPS to ACTIVE due to an ACTIVE mode channel time allocation.

7.15 AS IE operation

The AS IE is used to implement out-of-scope features that require additional functionality by both the PNC and one or more of its piconet member DEVs. The additional functionality is defined as an enhancement that does not violate the standard and allows DEVs that do not have the functionality to operate normally. The Application Specific Data field in this IE provides the messages that are only interpreted by the targeted DEV.

Multiple AS IEs may be placed in the beacon by the PNC. The designer should minimize the size of each AS IE used to support the custom application.

To place a new AS IE in the beacon, the DEV sends the AS IE Request command to the PNC with the Request Type field set to "add" and the AS IE Index field set to zero, as shown in Figure 7-74. The requesting DEV selects a unique value of the Request ID field in the AS IE Request command, 6.5.9.3, to be able to determine which AS IE Response command is associated with this request. The AS IE Index field shall be ignored by the PNC in an AS IE Request command when the value of the Request Type field is set to "add."

If the PNC grants the request to place the AS IE in the beacon, it shall respond with an AS IE Response command to the originating DEV with the Request ID set to the same value as in AS IE Request command, the AS IE Index set to an unused value, and the Reason Code set to "success."

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 7-73—MSC showing MLME initiated PM mode change from DSPS to ACTIVE in response to a new channel time allocation




Figure 7-74—MSC for a DEV adding an AS IE to the beacon

If the PNC refuses the request to place the AS IE in the beacon, it shall send an AS IE Response command to the originating DEV with the Request ID set to the same value as in the AS IE Request command and the Reason Code set to the appropriate value. The AS IE Index shall be set to zero by the PNC and shall be ignored by the originating DEV upon reception when the Reason Code is other than “success.”

A DEV may modify an existing AS IE by sending the AS IE Request command to the PNC with the Request Type field set to “modify,” the AS IE Index field set to the value assigned by the PNC for that AS IE, and the Request ID field set to a value unique for that DEV. This process is illustrated in Figure 7-75.




Figure 7-75—MSC for a DEV modifying an AS IE in the beacon

If the PNC allows the request to modify an existing AS IE, it shall respond with an AS IE Response command to the originating DEV with the Request ID and AS IE Index set to the same values as in AS IE Request command and the Reason Code set to “success.”

If the PNC refuses the request to modify the AS IE in the beacon, it shall send an AS IE Response command to the originating DEV with the Request ID and the AS IE Index set to the same values as in the AS IE Request command and the Reason Code set appropriately. If the PNC refuses a request to modify an AS IE, it does not delete the AS IE, but it shall continue placing the existing AS IE in the beacon.

A DEV may also request that the PNC remove the AS IE from the beacon by sending the AS IE Request command to the PNC with the Request Type field set to “remove” and the AS IE Index field set to value for the AS IE that is to be removed. If the PNC successfully receives this command with the AS IE Index value that corresponds to an AS IE that is being currently managed by the PNC, it shall no longer place that AS IE in the beacon and shall send an AS IE Response command to the DEV with the Reason Code set to “success” and the Request ID and AS IE Index set to the same values as in AS IE Request command. This procedure is illustrated in Figure 7-76.

If the value of the AS IE Index in an AS IE Request command with Request Type set to “modify” or “remove” does not correspond to an AS IE that is currently being managed by the PNC, the PNC shall refuse the request with the Reason Code set to “Unassigned AS IE Index.”

Only the DEV that requested an AS IE be added to the beacon shall be allowed to modify or remove that AS IE.




Figure 7-76—MSC for a DEV removing an AS IE from a beacon

Once the PNC has allowed a request to add or modify an AS IE, it shall place the AS IE in every beacon until the originating DEV either has been disassociated or has requested that the PNC remove the AS IE from the beacon. The PNC shall remove all AS IEs that were sent by a DEV when that DEV disassociates from the piconet. AS IEs are not transferred during handover, and so DEVs that need AS IEs in the beacon will need to send AS IE Request commands to the new PNC once it has taken over the piconet.

7.16 MAC sublayer parameters

The parameters that define some of the MAC characteristics are given in Table 7-6.

Table 7-6—MAC sublayer parameters

Parameter	Value
<i>mMinChannelScan</i>	<i>mMaxSuperframeDuration</i>
<i>mBroadcastDevInfoDuration</i>	$64 \times mMaxSuperframeDuration$
<i>mAssocRespConfirmTime</i>	$4 \times mMaxSuperframeDuration$
<i>mMinSuperframeDuration</i>	1000 μ s
<i>mMaxSuperframeDuration</i>	65535 μ s
<i>mMaxLostBeacons</i>	4
<i>mMctaAssocPeriod</i>	150 ms
<i>mFirstCtaGap</i>	100 μ s
<i>mMinBeaconInfoRepeat</i>	4
<i>mAsyncRequestLifetime</i>	1 s
<i>mMaxKeyChangeDuration</i>	65535 ms
<i>mMaxTimeTokenChange</i>	65535

Table 7-6—MAC sublayer parameters (continued)

Parameter	Value
<i>mMaxNumValidDEVs</i>	243
<i>mCtaRelinquishTimeout</i>	$3 \times \text{SIFS} + 2 \times (\text{duration of Imm-ACK frame})$
<i>mMaxSubframeNum</i>	For standard aggregation mode = 8, for low-latency mode = 256
<i>mMaxSubframeSize</i>	For standard aggregation mode = 10478575, for low-latency mode = 256

Additional characteristics that are PHY dependent are indicated in Table 7-7 for the 2.4 GHz PHY.

Table 7-7—MAC sublayer parameters—2.4 GHz PHY dependent

Parameter	Subclause
SIFS	10.2.7.1
MIFS	10.2.7.1
BIFS	10.2.7.1
RIFS	10.2.7.1
<i>pBackoffSlot</i>	10.2.7.1
<i>pChannelSwitchTime</i>	10.2.7.5
<i>pMinTpcLevel</i>	10.5.9
<i>pMaxTransferUnitSize</i>	10.2.8.2
<i>pMaxFrameBodySize</i>	10.2.8.1
<i>pClockAccuracy</i>	10.5.5
<i>pLengthHcs</i>	10.2.9
<i>pMinFragmentSize</i>	10.2.8.3

Additional characteristics that are PHY dependent are listed in Table 7-8 for the SC PHY.

Table 7-8—MAC sublayer parameters—SC PHY dependent

Parameter	Subclause
SIFS	11.2.6.1
MIFS	11.2.6.1
BIFS	11.2.6.1
RIFS	11.2.6.1
<i>pBackoffSlot</i>	11.2.6.1
<i>pMaxTransferUnitSize</i>	11.2.7.2

Table 7-8—MAC sublayer parameters—SC PHY dependent (continued)

Parameter	Subclause
<i>pMaxFrameBodySize</i>	11.2.7.1
<i>pClockAccuracy</i>	11.2.4.3
<i>pLengthHcs</i>	11.2.3.2.2
<i>pMinFragmentSize</i>	11.2.7.3

Additional characteristics that are PHY dependent are listed in Table 7-9 for the HSI PHY.

Table 7-9—MAC sublayer parameters—HSI PHY dependent

Parameter	Subclause
SIFS	11.3.5.5
MIFS	11.3.5.8
BIFS	11.3.5.6
RIFS	11.3.5.6
<i>pBackoffSlot</i>	11.3.5.6
<i>pMaxTransferUnitSize</i>	11.3.6.4
<i>pMaxFrameBodySize</i>	11.3.6.3
<i>pClockAccuracy</i>	11.3.4.2
<i>pLengthHcs</i>	11.3.3.4.1
<i>pMinFragmentSize</i>	11.3.6.5

Additional characteristics that are PHY dependent are listed in Table 7-10 for the AV PHY.

Table 7-10—MAC sublayer parameters—AV PHY dependent

Parameter	Subclause
SIFS	11.4.1.2.2
MIFS	11.4.1.2.4
BIFS	11.4.1.2
RIFS	11.4.1.2
<i>pBackoffSlot</i>	11.4.1.2
<i>pMaxTransferUnitSize</i>	11.4.1.3.2
<i>pMaxFrameBodySize</i>	11.4.1.3.1
<i>pClockAccuracy</i>	11.4.4.3

Table 7-10—MAC sublayer parameters—AV PHY dependent (continued)

Parameter	Subclause
<i>pLengthHcs</i>	11.4.1.4
<i>pMinFragmentSize</i>	11.4.1.3.3

7.17 Unequal error protection (UEP)

The following three types of UEP may be used for flexible error protection of a frame consisting of MSB and LSB subframes:

- UEP Type 1 protects an aggregated frame that consists of MSB subframes, LSB subframes, or both MSB and LSB subframes by using the different FECs.
- UEP Type 2 protects an aggregated frame that consists of MSB subframes, LSB subframes, or both MSB and LSB subframes by using different MCSs. The UEP Specific IE, as defined in 6.4.24, is used in UEP Type 2 only for the originating DEV to inform the target DEV of the separating position of the MSBs and LSBs. If necessary, a DEV using UEP Type 2 may send an Announce command with the UEP Specific IE to dynamically change the separating position.
- UEP Type 3 protects MSB and LSB in a subframe unequally by either applying different FEC coding rates to MSBs and LSBs or mapping MSBs and LSBs to a skewed constellation. The capability of supporting UEP by FEC coding and/or by the skewed constellation is indicated in the Capability IE, as defined in 6.4.11.

All UEP types, UEP Type 1, UEP Type 2, and UEP Type 3, may be used in SC PHY, as described in 11.2.2.9. UEP Type 3 may be used in HSI PHY, as described in 11.3.2.1, and AV PHY, as described in 11.4.

It is required that originating and target DEVs understand the UEP type and MCS information of each other before sending frames that use UEP. This is accomplished by the DEVs reporting their UEP capability in the DEV Capability IE in the association process. If the source and destination do not support the same UEP type, then only EEP frames shall be exchanged between the DEVs.

7.18 Sync frame transmission and virtually dependent piconet

Sync frame transmission is a function that mitigates co-channel interference due to a hidden PNC node. It also provides a method to obtain synchronization among independent piconets. Sync frame transmission is mandatory for PNC-capable DEVs when operating as a mmWave PNC, as defined in 11.1.9, and is optional for non-PNC-capable mmWave DEVs. A mmWave DEV indicates if it is capable of sending Sync frames by setting the Sync Frame Capable field in the Capabilities IE, 6.4.11.

The PNC controls the transmission of Sync frames by DEVs in the piconet by sending a Sync Frame Frequency IE, as defined in 6.4.35, in an Announce command to a DEV that is capable of sending Sync frames. The PNC sets the frequency with which the frames will be sent by setting the Sync Frame TX Frequency field to the appropriate value.

When a DEV that is capable of sending Sync frames receives the Sync Frame Frequency IE from the PNC, it sends a Sync frame in the first transmission opportunity in the next superframe. The DEV continues sending the Sync frame every Sync Frame TX Frequency superframes. DEVs may use the response frame in an Imp-ACK exchange to send the Sync frame. DEVs may use either data or Imp-ACK transmit opportunities for sync frame transmission.

A DEV that receives a Sync frame from another piconet may utilize the information in the frame to obtain network synchronization and mitigate interference. A PNC may use Sync frames that it receives from DEVs in another piconet to synchronize to PNC timing and forming a virtually dependent piconet. The PNC may then adjust the CTAs in its piconet to avoid interference with the other piconet.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017

8. Security

8.1 Security mechanisms

Security mechanisms provided by this standard allow security services to be implemented to control the admission of DEVs into a security relationship between the PNC and a DEV or between two ordinary DEVs and protect the information and integrity of communications between DEVs in a security relationship. This standard also provides a symmetric cryptography mechanism to assist in providing security services. Additional security services need to be provided by the higher layers to ensure proper management and establishment of the symmetric keys used in this standard.

The background assumptions used in designing this security solution are outlined in C.1.

8.1.1 Security membership and key establishment

The method by which a DEV becomes a member of a security relationship and obtains the appropriate key is outside of the scope of this standard. The Security Message command has been included as a special command to assist in the implementation of vendor defined protocols for establishing security relationships and any related data. It can be achieved with higher layer protocols that are not specified in this standard. The MAC/MLME is informed of changes to the membership of a security relationship and the key for that relationship with the MLME-MEMBERSHIP-UPDATE primitive, as described in 5.3.7.1.

Because the association protocol is unprotected, to maintain security, an authentication algorithm needs to include the DEVID among the items that are validated.

8.1.2 Key transport

All keys that are transmitted from one DEV to another shall be encrypted as specified in the key request, as described in 8.4.3, and distribute key protocols, as described in 8.4.2. For example, key transport is used to provide a copy of the piconet group data key to a DEV.

8.1.3 Data encryption

Data encryption uses a symmetric cipher to protect data from being read by parties without the cryptographic key. Data may be encrypted either by using a key shared by all associated DEVs or by using a key shared between only two DEVs.

8.1.4 Data integrity

Data integrity uses an integrity code to protect data from being modified by parties without the cryptographic key. It further provides assurance that data came from a party with the cryptographic key. Integrity may be provided using a key shared by all piconet DEVs or using a key shared between only two DEVs. All secure data frames that fail integrity checks are passed to the DME using MLME-SECURITY-ERROR.indication, and no other action is taken on the frame by the MLME.

8.1.5 Beacon integrity protection

The beacon may be integrity-protected. This integrity protection provides evidence to all the DEVs in the piconet that the PNC of the secure piconet transmitted the beacon. Under normal operations, the integrity check on the beacon provides evidence that the piconet is operating properly and that no security changes have occurred. When not scanning for other piconets, if the integrity check on the beacon fails, the DEV is alerted to the fact that the DEV does not have its security state synchronized with the PNC.

8.1.6 Command integrity protection

The integrity of commands may be protected just like any other data. Integrity protected commands sent between the PNC and a DEV shall be protected using the PNC-DEV management key. All secure commands that fail integrity checks are passed to the DME using MLME-SECURITY-ERROR.indication, and no other action is taken on the frame by the MLME.

8.1.7 Freshness protection

To prevent replay of old messages, a strictly-increasing time token is included in the beacon. A DEV shall reject as invalid a received beacon with a time token less than or equal to the current time token. In addition, the time token is included in the nonce, as described in 9.2.4, for each secure frame, as described in 6.2, so the integrity check will fail if a frame is replayed in a different superframe. A DEV maintains two values for freshness. The CurrentTimeToken is the time token value found in the beacon for the current superframe and is used to protect all messages sent and check all messages received during that superframe. The LastValidTimeToken is used by the DEV to ensure that the security of the beacons have not been compromised.

8.2 Security modes

The security mode indicates whether a DEV is currently implementing frame protection in the piconet. The security mode in use is determined by the *macSecurityOptionImplemented* entry in the MAC PIB.

8.2.1 Security mode 0

A DEV operating in security mode 0 shall not perform any cryptographic operations on MAC frames. While in this mode, if the MAC receives a frame with the SEC field set to one, the MAC shall discard the frame, and the MLME shall return an MLME-SECURITY-ERROR.indication to the higher layer with the ReasonCode set to UNAVAILABLE-KEY.

8.2.2 Security mode 1

Security mode 1 provides a mechanism for a device to perform cryptographic security on frames transmitted in the piconet. DEVs operating in security mode 1 use symmetric-key cryptography to protect frames using encryption and integrity.

While in mode 1, the cryptographic operations used for secure frames exchanged with the PNC and with other members of the piconet security group shall be performed as specified by the symmetric key security operations. While in this mode, if the MAC receives a frame with the SEC field in the Frame Control field set to a value different than expected, as defined in Table 6-22, the MLME shall generate an MLME-SECURITY-ERROR.indication with the ReasonCode set to INVALID-SEC-VALUE.

8.3 Security support

The security policies determine the actions taken to preserve the security of the piconet. Subclauses 8.3.1 through 8.3.9 specify the methods that are provided in this standard to support specific security policies.

8.3.1 PNC handover

When a PNC chooses to handover the PNC role to another DEV in the piconet, the security relationships with the old PNC no longer apply to the new PNC. When the old PNC hands over the piconet information using a PNC Information command, as described in 6.5.4.2, the list of associated DEVs is passed to the new PNC.

PNC handover does not affect the group membership, so it does not require a rekey of the group key. However, in a piconet with payload protection, the command functions of the PNC that relate to specific DEVs are not supported until the new PNC has established secure membership with each DEV in the piconet. When the PNC role has been handed over, the new PNC should create CTAs for each of the associated DEVs to establish secure membership with the new PNC.

The old PNC may send security information about the new PNC to the other DEVs in the piconet and send security information about all of the DEVs that are secure members in the piconet to the new PNC when it hands over the role of the PNC. This is accomplished by sending a directed Security Information command, as described in 6.5.4.4, to the new PNC with the security information of the piconet in it and by sending a broadcast Security Information command or a directed Security Information command to each member of the piconet with the security information of the new PNC.

8.3.2 Changes in the piconet group data key

When the PNC changes the piconet group data key, the PNC shall transmit the new key to all of the members of the piconet that are in ACTIVE mode using the Distribute Key Request command, as described in 6.5.2.3. Once the Distribute Key Request command has been issued for all of the members of the piconet that are in ACTIVE mode, the PNC may change the SECID in the beacon. When a DEV receives a valid Distribute Key Request command, as described in 6.5.2.3, from the PNC, the DEV shall use the new key for all outgoing secure frames that require the use of the piconet group data key once it sees the corresponding SECID in the beacon. The DEV may continue to accept frames protected by the old piconet group data key for up to *mMaxKeyChangeDuration* since the DEV last received a valid beacon protected by the old piconet-wide group data key.

If a DEV receives a beacon with a time token greater than the last known time token, but with a SECID that does not match the SECID of the known key, the DEV shall send a Key Request command to the PNC to obtain the new key.

8.3.3 Joining a secure piconet

If a DEV wishes to join a secure piconet, it should associate with the PNC in order to be assigned a local DEVID. Once the DEV is associated, the PNC shall allocate an MCTA if commands are not allowed in the CAP. The DEV or PNC may choose to send Probe Request and/or Announce commands to each other to either request or transmit IEs, including Vendor Defined IEs. The DEV and PNC may also exchange additional data frames or Security Message commands. After the DEV has associated and exchanged the desired information with the PNC, the DEV shall establish secure membership. The process by which secure membership is established is outside of the scope of this standard.

8.3.4 Membership update

When the DME determines that there has been a change of membership status with a particular DEV or when a management key is changed, the DME shall issue an MLME-MEMBERSHIP-UPDATE.request to its MLME. This membership status change or key change may be the result of a successful establishment of a security relationship, key update process, termination of a security relationship, or some other event. The process by which this change occurs is outside the scope of this standard.

When the MLME receives the MLME-MEMBERSHIP-UPDATE.request, it shall first examine the TrgtId to determine the membership relationship to modify. If the TrgtId is the PNCID, the management key corresponds to the management key for the relationship with the PNC, and the MembershipStatus indicates whether the DEV is a secure member of the piconet. Otherwise, the management key is for a peer-to-peer relationship with the DEV indicated by the TrgtId, and the MembershipStatus indicates whether the DEV shares a secure relationship with that peer DEV. If the OrigId is the PNCID, then the management key

corresponds to the management key for a secure relationship with a DEV in the piconet and will be used for PNC-related frames.

If the TrgtId is the PNCID and the MembershipStatus is set to MEMBER, the DEV is a secure member of the piconet. If the TrgtId is the PNCID and the MembershipStatus is set to NON-MEMBER, the DEV is no longer a secure member of the piconet.

The MembershipStatus field indicates to the MLME whether the DEV is currently maintaining secure relationship information with the target DEV. If the MembershipStatus is set to NON-MEMBER, the MLME shall securely delete the management key, the data key, and the related SECID, key type, and key originator values corresponding to that TrgtId. When a DEV is not a member of a security relationship with a peer DEV, the DEV shall select keys for secure frame processing as if the DEV does not have an individual relationship with that peer DEV, as described in 8.3.9. When a DEV is not a member with the PNC, the DEV is not a secure member of the piconet and shall select keys for secure processing as if the DEV does not have a piconet group data key or PNC-DEV management key, as described in 8.3.9.

If the MembershipStatus is set to MEMBER, the MLME shall examine the KeyInfoLength field to determine if a new key is being added or a key is being deleted. If the KeyInfoLength field is set to zero, the MLME shall securely delete the key and SECID corresponding to the management key for that relationship. When the key is deleted, the DEV is unable to transmit or successfully receive frames to any DEV that require protection with the management key, as described in 8.3.9, but the DEV may continue to use the data key corresponding to that relationship and the piconet group data key. This may occur, for instance, during PNC handover, in which the management key with the PNC is no longer valid (since the PNC has changed), but the piconet group data key is still valid.

If the MembershipStatus is set to MEMBER and the KeyInfoLength field is not 0, the MLME shall set the SECID, key originator field, and key for the management key of this relationship to the values in the SECID, KeyOriginator, and KeyInfo fields, respectively, from the MLME-MEMBERSHIP-UPDATE.request.

8.3.5 Secure frame generation

When a DEV wishes to send a secure frame, it shall use the keying material required for the type of frame and by the relationship between the sending DEV and the receiving DEV. For each security relationship, there are two keys used to protect secure frames: a management key and a data key. Table 8-1 provides a listing of which of the keys shall be used to protect secure frames and which frames shall be sent without security. A DEV shall not send a secure frame if the only key selection in Table 8-1 is "none." A DEV shall not send an unprotected frame or a frame with an incorrect SECID when security is required for that frame. If the DEV is unable to find the corresponding key that is to be used, the MLME shall return an MLME-SECURITY-ERROR.indication to the DME with the ReasonCode set to UNAVAILABLE-KEY and shall not transmit the requested frame.

A PNC in a piconet using security shall send secure beacons protected with the piconet group data key stored in the MAC/MLME. For each superframe, the PNC shall increment the time token and transmit a secure beacon with the SEC field in the Frame Control field set to one.

Key selection for secure frames is described in 8.3.9.

If the DEV is able to obtain the appropriate keying material, the DEV shall use the CurrentTimeToken and secure frame counter for the corresponding SECID to construct the CCM nonce, Figure 9-1, used to protect the secure frame. The SECID included in the frame shall be the value corresponding to the keying material being used. The integrity code shall be computed as specified in 9.3.2. The result of the integrity code computation shall be encrypted as specified in 9.2.2 and placed in the Integrity Code field in the secure frame. The encryption operation shall be applied only to the integrity code, the key that is transmitted in a Distribute Key command or Request Key Response command and the payload of data frames. The result of

the encryption operation shall be inserted into the frame in the place of the data that was encrypted. The DEV shall then compute the FCS over the modified frame.

A DEV shall send only frames that have increasing SFCs in a superframe, except for frames that are retransmitted with the same SFC without any intervening frames having been sent.

8.3.6 Updating CurrentTimeToken

If the DEV is able to determine that it missed a beacon or that the beacon was corrupted and if CurrentTimeToken is less than LastValidTimeToken + $mMaxTimeTokenChange - 1$, the DEV should increment the CurrentTimeToken to maintain synchronization with other DEVs in the piconet.

8.3.7 Secure frame reception

Before any security operations have been performed on a received frame, the DEV shall check the FCS. Table 8-1 provides a listing of the keys that shall be used to protect secure frames and the frames that shall be sent without security. A DEV may ignore any secure frame if the only key selection in Table 8-1 is "none." A DEV shall ignore any non-secure frame or a secure frame with an incorrect SECID when security is required.

An associated device that has not yet received the piconet group data key shall accept all secure beacons and ignore the integrity code, SECID, and secure frame counter. When the DEV has received the piconet group data key, it shall set the LastValidTimeToken and CurrentTimeToken to be the time token in that beacon.

When a DEV receives a secure Beacon frame, as defined in 6.3.1.2, the DEV shall determine if the received time token is greater than the CurrentTimeToken and less than the LastValidTimeToken + $mMaxTimeTokenChange$. If not, the MLME shall return an MLME-SECURITY-ERROR.indication to the DME with the ReasonCode set to BAD-TIME-TOKEN and shall not perform any additional operations on the received beacon. The DEV shall also determine if the SECID matches the SECID of the piconet group data key stored in the MAC/MLME, or the SECID of a valid old piconet group data key, as described in 8.3.5. If the SECID does not match, the DEV may request a new piconet group data key, as described in 8.3.2. If both of these checks succeed, the DEV shall check the integrity code on the beacon using the piconet group data key. If this succeeds, the DEV shall accept the beacon and set the LastValidTimeToken and CurrentTimeToken to be the time token in the beacon.

When a DEV receives a secure non-Beacon frame, it shall use the appropriate keying material depending on the type of frame, SECID, and SrcID found in the frame. If the SECID in the frame does not correspond to known keying material in the receiving DEV, the MLME shall return an MLME-SECURITY-ERROR.indication to the DME with the ReasonCode set to UNAVAILABLE-KEY and shall not perform any additional operations on the received frame. A DEV shall reject all frames that do not have an SFC that is strictly greater than the last SFC received from that DEV in that superframe.

If there are no previous security errors in the processing of the frame, the DEV shall apply the operations defined by the symmetric key security operations to the frame, as defined in 9.3.2. If any of the security operations fail, the MLME shall return an MLME-SECURITY-ERROR.indication to the DME with the ReasonCode set to FAILED-SECURITY-CHECK and shall not perform any additional operations on the received frame. If the security operations have been successfully performed and the frame has been modified appropriately, the DEV may then continue to process the frame.

While operating in mode 1, if the MAC receives a command frame with the SEC field in the Frame Control field set to a value different than expected, as defined in Table 6-22, the MLME shall generate an MLME-SECURITY-ERROR.indication with the ReasonCode set to INVALID-SEC-VALUE.

8.3.8 Selecting the SECID for a new key

For each management and data key used in the piconet, the key originator in the relationship shall select the 2-octet SECID, as described in 6.2.7.2, that identifies the key.

When a PNC-capable DEV starts a secure piconet, as described in 7.2.2, it shall select a SECID and a symmetric key to be used for beacon protection. Because there are no other DEVs in the piconet when the PNC-capable DEV starts a piconet, this key is not distributed to any other DEVs. Once another DEV joins the piconet, the PNC will update the key and SECID, as indicated in 8.3.3.

8.3.9 Key selection

The key used to protect a particular frame depends on the purpose of the frame and the membership states of the DEV. If the DEV is a member of a secure piconet (i.e., the DEV is the PNC or the DEV is a secure member with the PNC), the DEV will have entries for the piconet group data key and for the PNC-DEV management key. If the DEV has a secure relationship with a peer-DEV (i.e., the DEV is a secure member with a peer DEV), the DEV will have entries for a peer-to-peer data key and a peer-to-peer management key that it shares with that DEV. For any given frame, the DEV shall either send the frame without security or with the single key that is required for that frame, as indicated in Table 8-1. All secure commands between the PNC and other DEVs shall be protected with the PNC management key. All secure data frames with the PNC as either the DestID or SrcID, all secure broadcast frames and all secure beacons shall be protected with the piconet group data key. If two DEVs in a secure piconet do not have a peer-to-peer security relationship, they shall use the piconet group data key for commands that are required to be sent securely and they shall use the piconet group data key for secure data frames transmitted between them. Table 8-1 summarizes the keys that shall be used for each type of frame.

Table 8-1—Key selection for secure frames

Frame type or command	None	PNC-DEV mgmt key	Piconet group data key	Peer-to-peer mgmt key	Peer-to-peer data key	Comment
Beacon frame			X			All secure Beacon frames shall be protected by the piconet group data key.
Imm-ACK frame	X					Immediate acknowledgment frames shall not be secured with any key.
Dly-ACK frame	X					Delayed acknowledgment frames shall not be secured with any key.
Data frame			X		X	Only secure data frames shall be exchanged between DEVs that have a secure relationship. Secure data frames between DEVs that share a peer-to-peer relationship shall use the peer-to-peer data key; otherwise, they shall use the piconet group data key.

Table 8-1—Key selection for secure frames (*continued*)

Frame type or command	None	PNC-DEV mgmt key	Piconet group data key	Peer-to-peer mgmt key	Peer-to-peer data key	Comment
Multi-protocol Data frame			X		X	Only secure Multi-protocol Data frames shall be exchanged between DEVs that have a secure relationship. Secure Multi-protocol Data frames between DEVs that share a peer-to-peer relationship shall use the peer-to-peer data key; otherwise, they shall use the piconet group data key.
Announce Response command	X	X	X	X		If the Announce Response command is sent to or from the PNC before the DEV becomes a secure member of the piconet, the command shall not be secured by any key. If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key (peer-to-peer or PNC-DEV) for the relationship shall be used.
PM Mode Change Response command		X				
AS IE Request command		X				
AS IE Response command		X				
Multicast Configuration Request command		X				
Multicast Configuration Response command		X				
Association Request command	X					Association Request commands shall not be secured with any key.
Association Response command	X					Association Response commands shall not be secured with any key.
Disassociation Request command	X	X				Disassociation Request commands shall not be secured with any key before the DEV establishes secure membership in the piconet and shall be protected by the PNC-DEV management key otherwise.
Request Key command		X		X		The management key for the relationship (peer-to-peer or PNC-DEV) shall be used for this command.

Table 8-1—Key selection for secure frames (*continued*)

Frame type or command	None	PNC-DEV mgmt key	Piconet group data key	Peer-to-peer mgmt key	Peer-to-peer data key	Comment
Request Key Response command		X		X		The management key for the relationship (peer-to-peer or PNC-DEV) shall be used for this command.
Distribute Key Request command		X		X		The management key for the relationship (peer-to-peer or PNC-DEV) shall be used for this command.
Distribute Key Response command		X		X		The management key for the relationship (peer-to-peer or PNC-DEV) shall be used for this command.
PNC Handover command		X				
PNC Handover Response command		X				
PNC Handover Information command		X				
PNC Information Request command		X				
PNC Information command		X	X			If the PNC Information command is sent as a directed frame from the PNC to a DEV, the PNC-DEV management key shall be used. If the PNC Information command is sent as a broadcast frame, the piconet group data key shall be used.
Probe Request command	X	X	X	X		If the Probe Request command is sent to or from the PNC before the DEV becomes a secure member of the piconet, the command shall not be secured by any key. If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key (peer-to-peer or PNC-DEV) for the relationship shall be used.

Table 8-1—Key selection for secure frames (*continued*)

Frame type or command	None	PNC-DEV mgmt key	Piconet group data key	Peer-to-peer mgmt key	Peer-to-peer data key	Comment
Probe Response command	X	X	X	X		If the Probe Response command is sent to or from the PNC before the DEV becomes a secure member of the piconet, the command shall not be secured by any key. If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key (peer-to-peer or PNC-DEV) for the relationship shall be used.
Piconet Services command	X					
Announce	X	X	X	X		If the Announce command is sent to or from the PNC before the DEV becomes a secure member of the piconet, the command shall not be secured by any key. If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key (peer-to-peer or PNC-DEV) for the relationship shall be used.
Channel Time Request command	X	X				If the communicating parties are the PNC and a neighbor PNC, the Channel Time Request command shall not be protected with any key. Otherwise, the PNC-DEV management key shall be used.
Channel Time Response command	X	X				If the communicating parties are the PNC and a neighbor PNC, the Channel Time Response command shall not be protected with any key. Otherwise, the PNC-DEV management key shall be used.
Channel Status Request command		X	X	X		If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key for the relationship (peer-to-peer or PNC-DEV) shall be used.
Channel Status Response command		X	X	X		If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key for the relationship (peer-to-peer or PNC-DEV) shall be used.
Remote Scan Request command		X				

Table 8-1—Key selection for secure frames (*continued*)

Frame type or command	None	PNC-DEV mgmt key	Piconet group data key	Peer-to-peer mgmt key	Peer-to-peer data key	Comment
Remote Scan Response command		X				
Transmit Power Change command		X	X	X		If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key (peer-to-peer or PNC-DEV) for the relationship shall be used.
PM Mode Change command		X				
SPS Configuration Request command		X				
SPS Configuration Response command		X				
PS Set Information Request command		X				
PS Set Information Response command		X				
Security Message	X					
Vendor Defined		X	X	X		If the DEVs do not share an individual relationship, the piconet group data key shall be used. Otherwise, the management key (peer-to-peer or PNC-DEV) for the relationship shall be used.

8.4 Protocol details

The following protocol details include all cryptographic components and headers for the frames. The headers should be interpreted as being headers in the MAC frames. In addition, each element should be interpreted as specified in Clause 6. Note that all frames transmitted in this subclause are sent with the ACK Policy field set to Imm-ACK unless specified otherwise. The ACK frames do not affect the security of the protocols and are omitted from all diagrams.

8.4.1 Security information request and distribution

A DEV establishing membership in a security relationship may request or send security information to another DEV. This most often is done directly before or during the PNC handover process, but may be done at any time.

Figure 8-1 illustrates the message flows for Security Information Request and Security Information commands between two peer DEVs.




Figure 8-1—MSC for DEV-DEV Security Information request

Figure 8-2 illustrates the message flows for an Security Information Request from the new PNC to the old PNC.




Figure 8-2—MSC for new PNC-old PNC security information transfer

Figure 8-3 illustrates the message flows for an Security Information Request from the new PNC to the old PNC.




Figure 8-3—MSC for sending security information with the Security Message command

8.4.2 Key distribution protocol

In a secure piconet or in a secure peer-to-peer relationship, the key originator may wish to update the current data protection key by initiating the key distribution protocol described here. For a change in the piconet group data key, the PNC sends the new piconet group data key to each member of the piconet using the Distribute Key Request command. DEVs do not respond to a Distribute Key Request command sent by the PNC, other than with an Imm-ACK if the frame FCS is valid. Note that the Imm-ACK does not indicate that the Distribute Key Request command has passed cryptographic verification, only that the FCS was valid. For a change in a peer data key, the key originator in the relationship initiates the key distribution protocol. The key originator sends the Distribute Key Request command to the DEV with which it is updating the key. A DEV that successfully receives a Distribute Key Request command that also passes the data authentication, 10.4.2, shall respond to the key originator with the Distribute Key Response command. The key originator should initiate this protocol with each DEV with their respective shared key whenever the key is updated.

8.4.3 Key request protocol

In a secure piconet, if a DEV receives a frame or beacon with an unknown SECID, it may initiate the request key protocol in order to obtain the unknown key from the key originator of the relationship. The DEV initiates the protocol by sending the Request Key command to the key originator. When the key originator receives a Request Key command that has a valid Integrity Code, it checks to see whether it has a secure relationship with the requesting DEV. If there is a secure relationship, the key originator sends the Request Key Response command to the requesting DEV using the management key for that secure relationship.

9. Security specifications

9.1 Modes for security

When symmetric key security operations are selected, DEVs perform secure operations in mode 1. This mode is defined in 8.2. Symmetric key security operations are not defined for mode 0.

9.2 Symmetric cryptography building blocks

9.2.1 Notational conventions

When transmitting and interpreting security material in this standard, the first byte transmitted shall be the first byte of the security material and represented on the left of the other bytes. The bit ordering within the byte for security operations shall be MSB first and LSB last. The octet ordering shall be most significant octet first and least significant octet last. This ordering shall be irrespective of the transmission order of the bits. See Figure 6-1 for the mapping of bit transmission order to most significant or least significant.

A bit string is a sequence of symbols over the set $\{0,1\}$. The notation $x \parallel y$ is the right-concatenation of two bit strings x and y of length m and n whose result is the string z of length $m + n$ that is constructed with x on its leftmost (most significant) m symbols and with y on its rightmost (least significant) n symbols.

9.2.2 CCM combined encryption and data authentication

The combined symmetric encryption and data authentication mechanisms used in the symmetric key security operations consists of the generation of an integrity code followed by the encryption of plaintext data and the integrity code. The output consists of the encrypted data and the encrypted integrity code.

The symmetric authentication operation consists of the generation of an integrity code using a block cipher in CBC mode computed on a nonce followed by (optional) padded authentication data followed by (optional) padded plaintext data. The verification operation consists of the computation of this integrity code and comparison to the received integrity code.

The symmetric encryption operation consists of the generation of a key stream using a block cipher in counter mode with a given key and nonce and performing an XOR of the key stream with the plaintext and integrity code. The decryption operation consists of the generation of the key stream and the XOR of the key stream with the ciphertext to obtain the plaintext and integrity code.

All of the above operations shall be performed as specified in 9.4. The parameters for these operations shall be as specified in 9.2.3.

9.2.3 CCM parameters

The CCM operations shall be parameterized by the following selections: the AES encryption algorithm as specified in 9.2.5, the length in octets of the length field L shall be 2 octets, the length of the authentication field M shall be 8 octets, the nonce shall be formatted as specified in 9.2.4.

9.2.4 Nonce value

In order to preserve the security of the symmetric algorithms, the nonce used for CCM encryption and authentication shall be unique. As a result, the DEV shall not reuse any Secure Frame Counter field value within a single superframe that is intended for a particular DEVID (as this would cause a repeated nonce). This uniqueness is guaranteed by the use of the SrcID, which guarantees that different DEVs sharing the

same key will use a different nonce, by the time token, which is different for every superframe with a given key and by the DestID and secure frame counter, which guarantee uniqueness within a superframe as long as a DEV does not send more than 65536 frames to a particular DestID within that superframe. If a frame is retransmitted and a single bit in the header or Frame Payload field has changed, a new nonce shall be used. To ensure this, each time a frame is retransmitted, the value of the Secure Frame Counter field shall be incremented.

The nonce that is input to the CCM algorithm shall be formatted as illustrated in Figure 9-1.

Octets: 1	1	6	2	3
SrcID	DestID	Time Token	Secure Frame Counter	Fragmentation Control Field

Figure 9-1—CCM nonce format

Figure 9-2 specifies the format of the nonce for aggregated frame.

Octets: 1	1	6	2	1	2
SrcID	DestID	Time Token	Secure Frame Counter	Secure Subframe Counter	Padding

Figure 9-2—CCM nonce format for aggregated frame

The SrcID field, DestID field, Secure Frame Counter field, and Fragmentation Control field shall be set to the values that are included in the frame that is being protected.

The Time Token field shall be set to the value of CurrentTimeToken.

The Secure Subframe Counter field shall be included to generate an unique nonce for each subframe that is aggregated. The counter shall be incremented for each subframe in a frame.

The Padding field shall be set to 0xCFCF.

9.2.5 AES encryption

The advanced encryption standard (AES) encryption algorithm used for symmetric key security operations shall be performed as specified in NIST FIPS Pub 197. This encryption algorithm is parameterized by the use of 128-bit keys and 128-bit block size.

9.3 Symmetric cryptography implementation

9.3.1 Symmetric cryptography data formats

Table 9-1 specifies the length and meaning of the symmetric cryptography-related specific data elements from Clause 6. The operations performed to obtain the variable data values are specified in 9.3.2.

9.3.2 Symmetric cryptographic operations

Figure 9-3 specifies the length information and data input to the CCM operation for secure beacons. The Auth Data Length $l(a)$ shall be set to the length of all of the Piconet Sync Parameters field plus the sum of

Table 9-1—Symmetric cryptography frame object formats

Notation	Length	Value	Description
Encrypted key	16	Variable	The encrypted key consists of the result of the encryption of a 16-octet key (not including the integrity code) using CCM encryption, as specified in 9.2.2.
Integrity code	8	Variable	The integrity code consists of the encrypted integrity code that is the result of a CCM computation, as specified in 9.2.2, that is computed along with the encrypted seed.
Encrypted data	Variable	Variable	The encrypted data consists of the result of the encryption of the specified data (not including the integrity code) using CCM encryption, as specified in 9.2.2.

the lengths of the IEs that are included in the beacon. The Enc Data Length, $l(m)$, shall be set to zero. The data input to CCM shall be taken in the order it is received in the frame, omitting the HCS, FCS and integrity code.

Octets: 10	2	2	13	L_1	...	L_n	2	2
Frame Header	SECID	SFC	Piconet Sync Parameters	IE-1	...	IE-n	Auth Data Length	Enc Data Length

Figure 9-3—CCM input for secure beacons

Figure 9-4 specifies the length information and data input to the CCM operation for secure commands. For all commands except for the Request Key Response command and Distribute Key Request command, the Auth Data Length, $l(a)$, shall be set to the length of all of the protected data and the Enc Data Length, $l(m)$, shall be set to zero. For the Request Key Response command and Distribute Key Request command, the Auth Data Length, $l(a)$, shall be set to the length of all of the protected data minus 16 (the length of the key) and the Enc Data Length, $l(m)$, shall be set to 16 (the length of the key). The data input to CCM shall be taken in the order it is received in the frame, omitting the HCS, FCS, and integrity code.

Octets: 10	2	2	2	2	L_1	L_2	2	2
Frame Header	SECID	SFC	Command Type	Length	Auth Data	Enc Data	Auth Data Length	Enc Data Length

Figure 9-4—CCM input for secure commands

Figure 9-5 specifies the length information and data input to the CCM operation for Secure Data frames. The Auth Data Length, $l(a)$, shall be set to 14 and the Enc Data Length, $l(m)$, shall be set to the length of the Data Payload field. The data input to CCM shall be taken in the order it is received in the frame, omitting the HCS, FCS, and integrity code.

Octets: 10	2	2	L_1	2	2
Frame header	SECID	SFC	Data Payload	Auth Data Length	Enc Data Length

Figure 9-5—CCM input for Secure Data frames

Figure 9-6 specifies the length information and data input to the CCM operation for Secure Multi-Protocol Data frames. The Auth Data Length, $l(a)$, shall be set to 14 and the Enc Data Length, $l(m)$, shall be set to the sum of the lengths of the Data ID field, the Data Header field, and the Data Payload field. The data input to CCM shall be taken in the order it is received in the frame, omitting the HCS, FCS, and integrity code.

Octets: 10	2	2	2	L_1	L_2	2	2
Frame header	SECID	SFC	Data ID	Data Header	Data Payload	Auth Data Length	Enc Data Length

Figure 9-6—CCM input for Secure Multi-Protocol Data frames

9.4 CCM mode

CCM is a generic authenticate-and-encrypt block cipher mode. CCM is defined for use with block ciphers with a 128-bit block size, such as AES.

For CCM mode there are two parameter choices to be made (see Table 9-2). The first choice is M , the size of the authentication field. The choice of the value for M involves a trade-off between message expansion and the probability that an attacker will be able to undetectably modify a message. Valid values for M are 4, 6, 8, 10, 12, 14, and 16 octets. The second choice is L , the size of the length field. This value requires a trade-off between the maximum message size and the size of the nonce. Different applications require different trade-offs, so L is a parameter. Valid values for L are 2 to 8 octets (the value $L = 1$ is reserved).

Table 9-2—Parameters of CCM mode

Name	Description	Field size	Encoding of field
M	Number of octets in authentication field	3 bits	$(M-2)/2$
L	Number of octets in length field	3 bits	$L-1$

9.4.1 Inputs

To send a message, the sender shall provide the following information (see Table 9-3):

- An encryption key K suitable for the block cipher.
- A nonce N of $15-L$ octets. Within the scope of any encryption key K , the nonce value shall be unique. That is, the set of nonce values used with any given key shall not contain any duplicate values. Using the same nonce for two different messages encrypted with the same key destroys the security properties of this mode.
- The message m , consisting of a string of $l(m)$ octets where $0 \leq l(m) < 2^{8L}$. The length restriction ensures that $l(m)$ will be able to be encoded in a field of L octets.
- Additional authenticated data a , consisting of a string of $l(a)$ octets where $0 \leq l(a) < 2^{64}$. This additional data is authenticated but not encrypted, and is not included in the output of this mode. It may be used to authenticate plaintext headers, or contextual information that affects the interpretation of the message. If there is no additional data to authenticate, the string shall be zero length.

Table 9-3—Inputs for CCM

Name	Description	Field size	Encoding of field
K	Block cipher key	Depends on block cipher	String of octets
N	Nonce	$15-L$ octets	Not specified
m	Message to be encrypted and sent	$l(m)$ octets	String of octets
a	Additional authenticated data	$l(a)$ octets	String of octets

9.4.2 Data authentication

The first step is to compute the authentication field T . This is done using cipher block chaining-message authentication code (CBC-MAC). First a sequence of blocks $B_0, B_1, \dots, B_n$ is defined and then CBC-MAC is applied to these blocks.

The blocks shall be ordered as shown in Figure 9-7.

Octets: 16	16	$16 \times (n-2)$	16
B_0	B_1	B_2 to B_{n-1}	B_n

Figure 9-7—Authentication block ordering

The first block B_0 shall be constructed as follows:

$$B_0 = \text{Flags} \parallel \text{Nonce} \parallel l(m)$$

The value $l(m)$ is 16 bits and is encoded in most-significant-octet first order.

The Flags field is 8 bits and shall be constructed as follows:

$$\text{Flags} = \text{Reserved} \parallel \text{Adata} \parallel M \parallel L$$

The L field is 3 bits and encodes the size of the length field used to store $l(m)$. The parameter L may take on the values from 2 to 8 (the value $L=1$ is reserved). This value is encoded in the field using the values from 1 to 7 by choosing the field value as $L-1$ (the zero value is reserved).

The M field is 3 bits encodes the value of M as $(M-2)/2$. As M may take on the even values from 4 to 16, the field may take on the values from 1 to 7.

The Adata field is 1 bit and is set to zero if $l(a)=0$, and set to one if $l(a)>0$.

The Reserved field is 1 bit and shall be set to zero.

If $l(a)>0$ (as indicated by the Adata field) then one or more blocks of authentication data are added. These blocks contain $l(a)$ and a encoded in a reversible manner. The string that encodes $l(a)$ shall be constructed as follows:

- If $0 < l(a) < 2^{16}-2^8$ then the length field shall be encoded as two octets which contain the value $l(a)$ in most-significant-octet first order.

- If $2^{16}-2^8 \leq l(a) < 2^{32}$ then the length field shall be encoded as six octets consisting of the octets 0xff, 0xfe, and four octets encoding $l(a)$ in most-significant-octet-first order.
- If $2^{32} \leq l(a) < 2^{64}$ then the length field shall be encoded as ten octets consisting of the octets 0xff, 0xff, and eight octets encoding $l(a)$ in most-significant-octet-first order.

This is summarized in Table 9-4. Note that all fields are interpreted in most-significant-octet first order.

Table 9-4—Length encoding for additional authentication data

First two octets	Followed by	Comment
0x0000		Reserved
0x0001 ... 0xFEFF		For $0 < l(a) < 2^{16} - 2^8$
0xFF00 ... 0xFFFD		Reserved
0xFFFFE	four octets $l(a)$	For $2^{16} - 2^8 \leq l(a) < 2^{32}$
0xFFFFF	eight octets $l(a)$	For $2^{32} \leq l(a) < 2^{64}$

The blocks encoding a shall be formed by the string that encodes $l(a)$ followed by a itself, and splitting the result into 16-octet blocks, padding the last block with zeros if necessary. These blocks shall be appended as the octets following the first block B_0 . These blocks, if created shall be formatted as shown in Figure 9-8, where the length of a is $16 \times (k-1) - L_1 + L_2$.

Octets: $L_2 = 2, 6$ or 10	$16 - L_1$	$16 \times (k-2)$	$L_2 = 1$ to 16	$16 - L_2$
$l(a)$	First octets of a	Next octets of a	Final octets of a	0
B_1		B_2 to B_{k-1}	B_k	

Figure 9-8—Authentication block ordering for additional authentication blocks

After the (optional) additional authentication blocks have been added, the next step is to form the message blocks. The message blocks are formed by splitting the message m into 16-octet blocks, padding the last block with zeros if necessary. If the message m consists of the empty string, then no blocks shall be added in this step.

These blocks, if created shall be formatted as shown in Figure 9-9, where the length of a is $16 \times (k-1) - L_1 + L_2$.

Octets: 16	$16 \times (n-k-2)$	L_3	$16 - L_3$
First octets of m	Next octets of m	Final octets of m	0
B_{k+1}	B_{k+2} to B_{n-1}	B_n	

Figure 9-9—Authentication block ordering for message blocks

The result is a sequence of blocks $B_0, B_1, \dots, B_n$. The CBC-MAC shall be computed by the following:

$$X_1 = E(K, B_0)$$

$$X_{i+1} = E(K, X_i \oplus B_i) \text{ for } i=1, \dots, n$$

$$T = \text{first-}M\text{-octets}(X_{n+1})$$

where $E()$ is the block cipher encryption function and T is the integrity code value. Note that the last block B_n is XORed with X_n and encrypted with the block cipher to give T .

9.4.3 Encryption

The message data shall be encrypted with CTR mode. The key stream blocks are defined by the following:

$$S_i = E(K, A_i)$$

for $i=0, 1, 2, \dots$.

The values A_i are constructed as follows:

$$A_i = \text{Flags} \parallel \text{Nonce} \parallel \text{counter } i$$

The counter i field is 16 bits and is encoded as most-significant-octet first order.

The Flags field is 8 bits and shall be constructed as follows:

$$\text{Flags} = \text{Reserved} \parallel \text{Null} \parallel L$$

The L field is 3 bits and is encoded using the same encoding as in B_0 .

The Null field is 3 bits and shall be set to zero. This ensures that all the A blocks are distinct from B_0 , which has the non-zero encoding of M in this position.

The Reserved field is 2 bits and shall be set to zero.

Bit 6 corresponds to the Adata field in the B_0 block, but this field is not used here.

The message is encrypted by XORing the octets of message m with the first $l(m)$ octets of $S_1, S_2, S_3, \dots, S_{n-k}$ ordered as shown in Figure 9-10. Note that S_0 is not used to encrypt the message.

Octets: 16	16	$16 \times (n - k - 2)$	L_3
S_1	S_2	S_3 to S_{n-k-1}	First L_3 octets of S_{n-k}

Figure 9-10—Block ordering for encryption

The authentication value U shall be computed by encrypting T with the key stream block S_0 and truncating it to the desired length.

$$U = T \oplus \text{first-}M\text{-octets}(S_0)$$

9.4.4 Output

The final result c consists of the encrypted message, followed by the encrypted authentication value U .

9.4.5 Decryption

To decrypt a message, the following information is required:

- The encryption key K
- The nonce N
- The additional authenticated data a
- The encrypted and authenticated message c

Decryption starts by recomputing the key stream to recover the message m and the integrity code value T . The message and additional authentication data is then used to recompute the CBC-MAC value and check T .

If the T value is not correct, the receiver shall not reveal any information except for the fact that T is incorrect. In particular, the receiver shall not reveal the decrypted message, the value T , or any other information.

9.4.6 Restrictions

All implementations shall limit the total amount of data that is encrypted with a single key. The sender shall ensure that the total number of block cipher encryption operations in the CBC-MAC and encryption together shall not exceed 2^{61} . (This allows close to 2^{64} octets to be encrypted and authenticated using CCM.) Receivers that do not expect to decrypt the same message twice may also implement this limit.

The recipient shall verify the CBC-MAC before releasing any information such as the plaintext. If the CBC-MAC verification fails, the receiver shall destroy all information, except for the fact that the CBC-MAC verification failed.

9.4.7 List of symbols

Table 9-5 provides a list of the symbols used for the above specification of CCM.

Table 9-5—List of symbols

Name	Description	Size	Comment
a	Additional authenticated data	$l(a)$ octets	Use empty string if not desired.
A_i	Counter block to generate key stream	16 octets	Contains block counter, nonce, and flags.
B_i	Input block for CBC-MAC	16 octets	Together encode N , L , M , m , and a uniquely.
c	Ciphertext	$l(m) + M$ octets	Includes the encrypted integrity code.
K	Block cipher key	N/A	At least 128 bits, preferably 256 bits.
L	Number of octets in length field	3 bits	Values 1 ... 8, encoded in 3 bits as $L-1$.
m	Message to be encrypted and sent	$l(m)$ octets	Subject to $0 \leq l(m) < 2^{8L}$.
M	Number of octets in authentication field	3 bits	Values 4, 6, 8, ..., 16. Encoded value is $(M-2)/2$.
N	Nonce	$15-L$ octets	Nonce should never be repeated for same key.

Table 9-5—List of symbols (*continued*)

Name	Description	Size	Comment
S_i	Block of the encryption key stream	16 octets	Use $S_0, S_1, S_2, \dots$ to encrypt m and T .
T	Unencrypted authentication tag	M octets	
U	Encrypted authentication tag	M octets	Appended as the higher order octets to the message after encryption.
X_i	Intermediate value of CBC-MAC	16 octets	

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017

10. PHY specification for the 2.4 GHz band

10.1 Overview of the 2.4 GHz PHY

This clause specifies the PHY for a single carrier system that supports up to five modulation formats with coding at 11 Mbaud to achieve scalable data rates. The formats, coding, and data rates are given in Table 10-1.

Table 10-1—Modulation, coding, and data rates for 2.4 GHz PHY

Modulation type	Coding	Data rate
QPSK	8-state TCM	11 Mb/s
DQPSK	none	22 Mb/s
16-QAM	8-state TCM	33 Mb/s
32-QAM	8-state TCM	44 Mb/s
64-QAM	8-state TCM	55 Mb/s

This standard is based on the established regulations for Europe, Japan, Canada and the United States. The regulatory documents listed below are for information only and are subject to change or revision at any time. The regulatory domains are enumerated in a vector called *phyRegDomainsSupported*, specified in 10.7, and are indicated by the parameter *phyCurrentRegDomain*.

Europe (except France and Spain):

Approval standards: European Telecommunications Standards Institute (ETSI)
Documents: ETS 300-328 [B1], ETS 300-826
Approval authority: National type approval authorities

Japan:

Approval standards: Association of Radio Industries and Businesses (ARIB)
Document: ARIB STD-T66
Approval authority: Ministry of Post and Telecommunications (MPT)

United States:

Approval standards: Federal Communications Commission, USA
Documents: 47 CFR, Part 15, Sections 15.205, 15.209, 15.249

Canada:

Approval standards: Industry Canada, Canada
Document: GL36

10.2 General requirements

10.2.1 Operating frequency range

This PHY operates in the 2.4–2.4835 GHz frequency range as allocated by the regulatory agencies in Europe, Japan, Canada, and the United States as well as any other areas where the regulatory bodies have allocated this band.

10.2.2 Radio frequency (RF) power measurements

Unless otherwise stated, all RF power measurements for the purposes of this standard, either transmit or receive, shall be made at the appropriate transceiver to antenna connector. The measurements shall be made with equipment that is either matched to the impedance of the antenna connector or is corrected for any mismatch. For DEVs without an antenna connector, the measurements shall be interpreted as EIRP (i.e., a 0 dBi gain antenna) and any radiated measurements shall be corrected to compensate for the antenna gain in the implementation.

10.2.3 Channel assignments

A total of 5 channels in two sets are assigned for operation. The first set is the high-density mode, which allocates 4 channels while the second is an IEEE Std 802.11b-1999 [B3] co-existence mode, which allocates 3 channels. Since the two outer channels of the sets overlap, there are a total of 5 channels allowed for operation. The assigned channels are shown in Table 10-2. A compliant IEEE 802.15.3 implementation shall support all 5 channels.

Table 10-2—2.4 GHz PHY channel plan

CHNL_ID	Center frequency	High-density	IEEE 802.11b coexistence
1	2.412 GHz	X	X
2	2.428 GHz	X	
3	2.437 GHz		X
4	2.445 GHz	X	
5	2.462 GHz	X	X

The *phyCurrentChannel* is the CHNL_ID of the current channel. For the purpose of the Remote Scan Request and Remote Scan Response commands, as described in 6.5.7.3 and 6.5.7.4, respectively, the Channel Index field is the CHNL_ID in Table 10-2.

10.2.4 Scanning channels

A DEV may, in the course of a scan, change to an IEEE 802.11b channel for the purpose of detecting the presence of IEEE 802.11b networks.

When a DEV is scanning to start a piconet, it should scan all five channels to decrease the probability of choosing an occupied channel.

If a DEV is capable of identifying an IEEE 802.11b network and it does identify an IEEE 802.11b network while scanning, it should use the IEEE 802.11b coexistence channel set. It should also rate the channels

where IEEE 802.11b networks were identified as the worst channels. If multiple IEEE 802.11b networks are detected, the DEV should order them based on an estimate of the amount of traffic and the power level in the channel.

10.2.5 Unwanted emissions

Conformant implementations shall comply with the in-band and out-of-band emissions for all operational modes as set by the applicable regulatory bodies.

10.2.6 Operating temperature range

A conformant implementation shall meet all of the specifications in this standard for ambient temperatures from 0 °C to 40 °C.

10.2.7 PHY layer timing

The values for the PHY layer timing parameters are defined Table 10-3.

Table 10-3—PHY layer timing parameters

PHY parameter	Value	Subclause
<i>pMifsTime</i>	2 μ s	10.2.7.4
<i>pSifsTime</i>	10 μ s $\pm$ 0.5 μ s	10.2.7.2
<i>pCcaDetectTime</i>	5 $\times$ 16/11 μ s	10.6.5
<i>pChannelSwitchTime</i>	500 μ s	10.2.7.5

10.2.7.1 Interframe space

A conformant implementation shall support the IFS parameters, as described in 7.4.1, given in Table 10-4.

Table 10-4—IFS parameters

MAC parameter	Corresponding PHY parameter	Definition
MIFS	<i>pMifsTime</i>	10.2.7.4
SIFS	<i>pSifsTime</i>	10.2.7.2
pBackoffSlot	<i>pSifsTime</i> + <i>pCcaDetectTime</i>	10.6.5
BIFS	<i>pSifsTime</i> + <i>pCcaDetectTime</i>	10.2.7.2, 10.6.5
RIFS	2 $\times$ <i>pSifsTime</i> + <i>pCcaDetectTime</i>	10.2.7.2, 10.6.5

10.2.7.2 Receive-to-transmit turnaround time

The RX-to-TX turnaround time shall be *pSifsTime*, including the power-up ramp specified in 10.5.7.

The RX-to-TX turnaround time shall be measured at the air interface from the trailing edge of the last symbol received until the first symbol of the PHY preamble is present at the air interface.

10.2.7.3 Transmit-to-receive turnaround time

The TX-to-RX turnaround time shall be less than $pSifsTime$, including the power-down ramp specified in 10.5.7.

The TX-to-RX turnaround time shall be measured at the air interface from the trailing edge of the last transmitted symbol until the receiver is ready to begin the reception of the next PHY frame.

10.2.7.4 Time between successive transmissions

The minimum time between successive transmissions shall be $pMifsTime$, including the power-up ramp specified in 10.5.7.

The $pMifsTime$ shall be measured at the air interface from the trailing edge of the last symbol transmitted until the first symbol of the PHY preamble is present at the air interface.

10.2.7.5 Channel switch time

The channel switch time is defined as the time from when the last valid bit is received at the antenna on one channel until the DEV is ready to transmit or receive on a new channel. The channel switch time shall be less than $pChannelSwitchTime$.

10.2.8 Data size restrictions

The PHY definitions creates restrictions on the maximum frame size, maximum transfer unit size, and minimum fragmentation size that will be supported. These parameters are defined in this subclause.

10.2.8.1 Maximum frame length

The maximum frame length allowed, $pMaxFrameBodySize$, shall be 2048 octets. This total includes the frame payload and FCS but not the PHY preamble, PHY header, MAC header, or HCS. The maximum frame length also does not include the tail symbols, as described in 10.4.7, or the stuff bits, as described in 10.4.6.

10.2.8.2 Maximum transfer unit size

The maximum size data frame passed from the upper layers, $pMaxTransferUnitSize$, shall be 2044 octets. If security is enabled for the data connection, the upper layers should limit data frames to 2044 octets minus the security overhead, as defined in 6.3.4.2.

10.2.8.3 Minimum fragment size

The minimum fragment size, $pMinFragmentSize$, that is allowed with the 2.4 GHz PHY shall be 64 octets.

10.2.9 Header check sequence

The combined PHY and MAC headers shall be protected with a CCITT CRC-16 header check sequence (HCS). The value $pLengthHcs$ shall be 2 for this PHY. The CCITT CRC-16 HCS shall be the one's complement of the remainder generated by the modulo-2 division of the protected combined PHY and MAC headers by the following polynomial:

$$x^{16}+x^{12}+x^5+1$$

The protected bits shall be processed in transmit order. All HCS calculations shall be made prior to data scrambling. A schematic of the processing is shown in Figure 10-1.




Figure 10-1—CCITT CRC-16 Implementation

As an example, consider the following 32-bit length sequence to be protected by the CRC-16:

0101 0000 0000 0000 0011 0000 0000
b0.....b31

The leftmost bit (b0) is transmitted first in time.

The one's complement HCS for this sequence would be the following:

0101 1011 0101 0111
b0.....b15

The leftmost bit (b0) is transmitted first in time. Bit b0 corresponds to X^{15} in the Figure 10-1.

An illustrative example of the CCITT CRC-16 HCS using the information from Figure 10-1 is shown in Figure 10-2.

Data	CRC Registers	
	MSB	LSB
	1111111111111111	; Initialize preset to ones
0	1110111110111111	
1	1101111110111110	
0	10101111101011101	
1	0101111010111010	
0	1011110101110100	
0	0110101011001001	
0	1101010110010010	
0	1011101100000101	
0	0110011000101011	
0	1100110001010110	
0	1000100010001101	
0	0000000100111011	
0	0000001001110110	
0	0000010011101100	
0	0000100111011000	
0	0001001110110000	
0	0010011101100000	
0	0100111011000000	
0	1001110110000000	
0	0010101100100001	
0	0101011001000010	
0	1010110010000100	
1	0101100100001000	
1	101001000110001	
0	0101010001000011	
0	1010100010000110	
0	0100000100101101	
0	1000001001011010	
0	0001010010010101	
0	0010100100101010	
0	0101001001010100	
0	1010010010101000	

Figure 10-2—Example of CRC calculation

The CRC-16 described in this subclause is the same one used in IEEE Std 802.11b-1999 [B3].

10.2.10 Channel access methods

A PNC-capable DEV compliant to this standard shall allow the use of the CAP for contention based access for association, data, and commands, as described in 6.3.1, when using the 2.4 GHz PHY. A DEV that is compliant to this standard shall support the use of the CAP when using the 2.4 GHz PHY.

10.3 Modulation and coding

The 2.4 GHz PHY uses uncoded DQPSK, uncoded QPSK, and 16/32/64-QAM with trellis coding (see Ungerboeck [B11]). A 2.4 GHz DEV shall, at a minimum, support DQPSK modulation. In addition, if a 2.4 GHz DEV supports a given modulation format other than DQPSK, it shall also support all of the lower

modulation formats. For example, if an IEEE 802.15.3 implementation supports 32-QAM, it shall also support 16-QAM and QPSK-TCM as well as the DQPSK modulation formats.

The symbol rate for all modulations shall be 11 Mbaud. Based on this symbol rate and the coding, the raw PHY data rates supported are 11 Mb/s, 22 Mb/s, 33 Mb/s, 44 Mb/s, 55 Mb/s (QPSK-TCM, DQPSK, 16/32/64-QAM –TCM, respectively) as shown in Table 10-1. The data rates are respectively the entries to the `phyDataRateVector`, as described in 10.7.

10.3.1 Base data rate

The base data rate of the 2.4 GHz PHY shall be 22 Mb/s operating in the uncoded DQPSK mode.

The DQPSK mode is used as a base rate instead of the 11 Mb/s QPSK-TCM mode to reduce the overhead due to the duration of the PHY and MAC headers. Also, DQPSK capability is necessary to implement the PHY preamble, 10.4.2.

The QPSK-TCM mode is implemented in assigned CTAs to help maintain connections of DEVs that are in range of the PNC, but which may be more distant from each other.

10.3.2 Signal constellations

Figure 10-3 illustrates the signal constellations used in encoding bit streams into discrete signal levels sent through the common air interface.




Figure 10-3—DQPSK, QPSK, 16/32/64-QAM signal constellations

The average power within a frame including the PHY preamble and header is required to be a constant, regardless of the modulation format. Thus, a conformant implementation shall scale the constellation such that the PHY header and the MPDU have the same average power. One method to calculate the normalization is as follows: The output values, d , are formed by multiplying the resulting $(I + jQ)$ value by a normalization factor, K_{MOD} , i.e.,

$$d = (I + jQ) \times K_{MOD}$$

The normalization factor, K_{MOD} , depends on the base modulation mode and is given in for each of the modulation formats in Table 10-5. The purpose of the normalization factor is to achieve the same average power for all mappings. In practical implementations, an approximate value of the normalization factor may be used, as long as the DEV conforms with the modulation accuracy requirements described in 10.5.2.

Table 10-5—Normalization factor for PHY modulation formats

Modulation	K_{MOD}
DQPSK	1
QPSK-TCM	1
16-QAM-TCM	$1/(\sqrt{5})$
32-QAM-TCM	$1/(\sqrt{10})$
64-QAM-TCM	$1/(\sqrt{21})$

10.3.3 DQPSK modulation

No coding shall be applied to the DQPSK modulation. The mapping of the bit pairs to DQPSK symbols shall be implemented as specified in Table 10-6. In Table 10-6, a “+j ω ” phase change shall be defined as a counterclockwise rotation. The differential encoding shall apply only to the DQPSK mode. In this mode the entire frame, with the exception of the PHY preamble, shall be encoded differentially. The phase change of the first symbol is determined relative to the phase of the last symbol in the CAZAC sequence, as described in 10.4.2.

The differential encoding is provided to allow for non-coherent receiver implementations.

Table 10-6—DQPSK encoding table

Bit pattern (d0,d1) d0 is first in time	Phase change (+j ω)
0,0	0
0,1	$\pi/2$
1,1	π
1,0	$3\pi/2$ ($-\pi/2$)

10.3.4 QPSK and 16/32/64-QAM with trellis coding

The QPSK and 16/32/64-QAM formats shall use an 8-state two-dimensional (2-D) trellis code. The encoder shown in Figure 10-4 shall be used in implementing the 8-state 2-D trellis code.




Figure 10-4—QPSK, 16/32/64-QAM 8-state trellis encoder

The pair $([x_n^5, x_n^4, x_n^3, x_n^2, x_n^1], [y_n^5, y_n^4, y_n^3, y_n^2, y_n^1, y_n^0])$ represents the input and output relationship of the trellis encoder, where 5 input bits per symbol interval are encoded into 6 output bits for 64-QAM symbol mapping. Likewise, the input/output pair $([x_n^4, x_n^3, x_n^2, x_n^1], [y_n^4, y_n^3, y_n^2, y_n^1, y_n^0])$ represents the 32-QAM symbol mapping case, where 4 input bits per symbol interval are encoded into 5 output bits. Similarly, for the case of 16-QAM symbol mapping, 3 input bits per symbol interval are encoded into 4 output bits as denoted by the input/output pair $([x_n^3, x_n^2, x_n^1], [y_n^3, y_n^2, y_n^1, y_n^0])$. Finally, for the QPSK symbol mapping, 1 input bit per symbol interval is encoded into 2 output bits as denoted by the input/output pair $([x_n^1], [y_n^1, y_n^0])$. The 8 states generated by the trellis encoder shown in Figure 10-4 are denoted as $S_0, S_1, S_2, \dots, S_7$, which correspond to the binary representations of $s_n^2 s_n^1 s_n^0$. The input-output relations and state transitions of the trellis encoder shown above is given in Table 10-7 for the QPSK modulation format and in Table 10-8 for the 16/32/64-QAM formats. Trellis symbol subsets C0, C1, C2, C3 for the QPSK constellation and D0, D1, D2, ..., D7 for the 16/32/64-QAM constellations will be described next in the context of “mapping-by-set-partitioning” concept.

Table 10-7—Input-output relations and state transitions
of QPSK trellis encoder

Current state $s_n^2 s_n^1 s_n^0$	Input bit x_n^1	Output bits		Next state $s_{n+1}^2 s_{n+1}^1 s_{n+1}^0$
		$y_n^1 y_n^0$	Subset number	
0 0 0 (S0)	0	0 0	C0	0 0 0 (S0)
	1	1 0	C2	0 0 1 (S1)
0 0 1 (S1)	0	0 1	C1	1 1 0 (S6)
	1	1 1	C3	1 1 1 (S7)
0 1 0 (S2)	0	0 0	C0	0 0 1 (S1)
	1	1 0	C2	0 0 0 (S0)

**Table 10-7—Input-output relations and state transitions
of QPSK trellis encoder (continued)**

Current state $s_n^2 s_n^1 s_n^0$	Input bit x_n^1	Output bits		Next state $s_{n+1}^2 s_{n+1}^1 s_{n+1}^0$
		$y_n^1 y_n^0$	Subset number	
0 1 1 (S3)	0	0 1	C1	1 1 1 (S7)
	1	1 1	C3	1 1 0 (S6)
1 0 0 (S4)	0	0 0	C0	0 1 0 (S2)
	1	1 0	C2	0 1 1 (S3)
1 0 1 (S5)	0	0 1	C1	1 0 0 (S4)
	1	1 1	C3	1 0 1 (S5)
1 1 0 (S6)	0	0 0	C0	0 1 1 (S3)
	1	1 0	C2	0 1 0 (S2)
1 1 1 (S7)	0	0 1	C1	1 0 1 (S5)
	1	1 1	C3	1 0 0 (S4)

**Table 10-8—Input-output relations and state transitions of
16/32/64-QAM trellis encoder**

Current state $s_n^2 s_n^1 s_n^0$	Input bits $x_n^2 x_n^1$	Output bits		Next state $s_{n+1}^2 s_{n+1}^1 s_{n+1}^0$
		$y_n^2 y_n^1 y_n^0$	Subset number	
0 0 0 (S0)	0 0	0 0 0	D0	0 0 0 (S0)
	0 1	0 1 0	D2	0 0 1 (S1)
	1 0	1 0 0	D4	0 1 0 (S2)
	1 1	1 1 0	D6	0 1 1 (S3)
0 0 1 (S1)	0 0	0 0 1	D1	1 1 0 (S6)
	0 1	0 1 1	D3	1 1 1 (S7)
	1 0	1 0 1	D5	1 0 0 (S4)
	1 1	1 1 1	D7	1 0 1 (S5)
0 1 0 (S2)	0 0	0 0 0	D0	0 0 1 (S1)
	0 1	0 1 0	D2	0 0 0 (S0)
	1 0	1 0 0	D4	0 1 1 (S3)
	1 1	1 1 0	D6	0 1 0 (S2)
0 1 1 (S3)	0 0	0 0 1	D1	1 1 1 (S7)
	0 1	0 1 1	D3	1 1 0 (S6)
	1 0	1 0 1	D5	1 0 1 (S5)
	1 1	1 1 1	D7	1 0 0 (S4)

**Table 10-8—Input-output relations and state transitions of
16/32/64-QAM trellis encoder (continued)**

Current state $s_n^2 s_n^1 s_n^0$	Input bits $x_n^2 x_n^1$	Output bits		Next state $s_{n+1}^2 s_{n+1}^1 s_{n+1}^0$
		$y_n^2 y_n^1 y_n^0$	Subset number	
1 0 0 (S4)	0 0	0 0 0	D0	0 1 0 (S2)
	0 1	0 1 0	D2	0 1 1 (S3)
	1 0	1 0 0	D4	0 0 0 (S0)
	1 1	1 1 0	D6	0 0 1 (S1)
1 0 1 (S5)	0 0	0 0 1	D1	1 0 0 (S4)
	0 1	0 1 1	D3	1 0 1 (S5)
	1 0	1 0 1	D5	1 1 0 (S6)
	1 1	1 1 1	D7	1 1 1 (S7)
1 1 0 (S6)	0 0	0 0 0	D0	0 1 1 (S3)
	0 1	0 1 0	D2	0 1 0 (S2)
	1 0	1 0 0	D4	0 0 1 (S1)
	1 1	1 1 0	D6	0 0 0 (S0)
1 1 1 (S7)	0 0	0 0 1	D1	1 0 1 (S5)
	0 1	0 1 1	D3	1 0 0 (S4)
	1 0	1 0 1	D5	1 1 1 (S7)
	1 1	1 1 1	D7	1 1 0 (S6)

The symbol mapper shown in Figure 10-4 provides a one-to-one mapping between an output bit vector of the trellis encoder and a two-dimensional signal point of the signal constellation given in Figure 10-3. For a given output bit vector of the trellis encoder, a QPSK or 16/32/64-QAM constellation point is selected based on the set partitioning rule illustrated in Figure 10-5 and Figure 10-6, respectively. The lower order output bits y_n^1, y_n^0 and y_n^2, y_n^1, y_n^0 are used in determining the symbol subsets C0, C1, C2, C3 and D0, D1, D2, ..., D7 for QPSK and 16/32/64-QAM constellations, respectively. The subsets C0, C1, C2, C3 each contain 1 symbol for the QPSK modulation, whereas the subsets D0, D1, D2, ..., D7 each contain 2, 4, and 8 symbols for 16-QAM, 32-QAM, and 64-QAM cases, respectively. Therefore, as shown in Figure 10-4, the remaining output bits y_n^5, y_n^4, y_n^3 select one of the 8 symbols from the subsets D0, ..., D7 in the 64-QAM case, and the output bits y_n^4, y_n^3 select one of the 4 symbols from the subsets D0, ..., D7 in the 32-QAM case, and finally the output bit y_n^3 selects one of the 2 symbols from the subsets D0, ..., D7 in the 16-QAM case. Figure 10-7 and Figure 10-8 show the assignment of signal subsets to the QPSK and 16/32/64-QAM constellations, respectively. Furthermore, specific bit mappings to constellation points are given in Figure 10-9 and Figure 10-10 for the respective constellations. Binary representations below the subset numbers correspond to $y_n^1 y_n^0$ in the QPSK case and to $y_n^5 y_n^4 y_n^3 y_n^2 y_n^1 y_n^0$ in the 16/32/64-QAM cases. The lower order 2 bits correspond to the subset numbers for the QPSK modulation. Likewise, the lower order 3 bits correspond to the subset numbers for the 16/32/64-QAM. The higher order 3 bits for the 16/32/64-QAM cases are assigned within each signal subset (D0, ..., D7) such that decimal representation of the bit mapping goes from low to high as the constellation points are traced from center outward. This rule ensures that the decimal representations of the bit mappings from 0 to 15 belong to 16-QAM constellation, and 0 to 31 belong to 32-QAM constellation, and 0 to 63 belong to 64-QAM constellation.




Figure 10-5—QPSK set partitioning




Figure 10-6—16/32/64-QAM set partitioning




Figure 10-7—Assignment of subsets to QPSK constellation symbols



Figure 10-8—Assignment of subsets to 16/32/64-QAM constellation symbols

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 10-9—QPSK constellation bit mappings




Figure 10-10—16/32/64-QAM constellation bit mappings

Finally, based on Table 10-7 and Table 10-8, the state-transition diagram of the 8-state trellis code is shown for QPSK modulation in Figure 10-11 and for 16/32/64-QAM in Figure 10-12.




Figure 10-11—State transition diagram of 8-state QPSK trellis code




Figure 10-12—State transition diagram of 8-state 16/32/64-QAM trellis code

10.4 PHY frame format

10.4.1 Frame format

The PHY frame format for the 22 Mb/s, 33 Mb/s, 44 Mb/s, and 55 Mb/s modes is illustrated in Figure 10-13. The PHY layer prepends the PHY header, as described in 10.4.5, to the MAC header, as described in 6.2, calculates the HCS, as described in 10.2.9, over the combined PHY and MAC headers, and appends the HCS to the end of the MAC header. If the size of the MAC Frame Body field, as defined in 6.2, is not an integer multiple of the bits/symbol, then stuff bits are added following the MAC Frame Body field, as described in 10.4.6. The PHY preamble, as described in 10.4.2, is sent first, followed by the PHY header, MAC header and HCS, followed by the frame payload, the FCS, the SBs, if necessary, and finally the tail symbols, as described in 10.4.7. As shown in Figure 10-13, for the 22 Mb/s, 33 Mb/s, 44 Mb/s, and 55 Mb/s

transmission modes, the PHY preamble, as described in 10.4.2, is modulated with the 22 Mb/s QPSK mode. The PHY header, MAC header, and the HCS is modulated in the 22 Mb/s DQPSK mode. Finally, the frame payload, FCS, stuff bits (if necessary) and the tail symbols are modulated at the desired rates of either 22 Mb/s, 33 Mb/s, 44 Mb/s, or 55 Mb/s.




Figure 10-13—PHY frame formatting for 22 Mb/s, 33 Mb/s, 44 Mb/s, and 55 Mb/s modes

The PHY frame format for the 11 Mb/s mode is slightly different from the other modes and is illustrated in Figure 10-14. As in the other modes, the PHY layer prepends the PHY header, as described in 10.4.5, to the MAC header, as described in 6.2, calculates the HCS, as described in 10.2.9, over the combined PHY and MAC headers, and appends the HCS to the end of the MAC Frame Body field in this mode. Since the 11 Mb/s mode is modulated at 1 bit per symbol, stuff bits are not needed for the MAC Frame Body field in this mode. The concatenation of the PHY header, MAC header and HCS is repeated in the 11 Mb/s transmission mode. That is, the PHY preamble, as described in 10.4.2, is sent first in the frame, followed by two repetitions of the combined PHY header, MAC header and HCS, followed by the frame payload, the FCS and finally the tail symbols, as described in 10.4.7.

As illustrated in Figure 10-14, for the 11 Mb/s transmission mode, the PHY preamble, as described in 10.4.2, is modulated in the 22 Mb/s QPSK mode since the CAZAC sequence used, as described in 10.4.2, is based on 4-phase symbols. The first repetition of the combined PHY header, MAC header and the HCS, is modulated in the 22 Mb/s DQPSK mode. The second repetition of the combined PHY header, MAC header and HCS, is modulated in the 11 Mb/s QPSK-TCM mode. The repetition process ensures that the header error rate is significantly lower than the frame error rate. Finally, the Frame Payload field, FCS, and the tail symbols are all modulated with the 11 Mb/s QPSK-TCM mode. Note that although the concatenation of the PHY header, MAC header and HCS is repeated twice in the PHY frame format, the PHY layer shall only provide one copy of the correctly decoded MAC header to the MAC sublayer at the receiver side.

10.4.2 PHY preamble

A PHY preamble shall be added prior to the PHY header to aid receiver algorithms related to synchronization, carrier-offset recovery, and signal equalization. The preamble shall consist of multiple periods of a special sequence of 16 QPSK symbols called a CAZAC sequence (see Milewski [B8]), which demonstrates a constant amplitude zero auto-correlation property. The CAZAC sequence shall be denoted as




Figure 10-14—PHY frame formatting for 11 Mb/s mode

$\{C_0, C_1, C_2, \dots, C_{15}\}$. Each element, C_i , of the CAZAC sequence shall have a complex value representing the in-phase and quadrature components of a QPSK-type sequence, as shown in Table 10-9.


Table 10-9—CAZAC sequence

CAZAC sequence element	Value
C_0	$1 + j$
C_1	$1 + j$
C_2	$1 + j$
C_3	$1 + j$
C_4	$-1 + j$
C_5	$-1 - j$
C_6	$1 - j$
C_7	$1 + j$
C_8	$-1 - j$
C_9	$1 + j$
C_{10}	$-1 - j$
C_{11}	$1 + j$
C_{12}	$1 - j$
C_{13}	$-1 - j$

Table 10-9—CAZAC sequence (continued)

CAZAC sequence element	Value
C_{14}	$-1 + j$
C_{15}	$1 + j$

The PHY preamble shall be constructed by successively appending 12 periods, denoted as $\{P_0, P_1, P_2, \dots, P_{10}, E\}$, of the CAZAC sequence defined in Table 10-9, except for the 12th period where each element of the CAZAC sequence shall be negated, or equivalently, rotated by 180 degrees. $P_0, P_1, \dots, P_{10}$ are all identical vectors containing the CAZAC symbols defined in Table 10-9, and the vector E denotes the end-of-preamble delimiter that is the 180 degrees rotated version of the CAZAC sequence given in Table 10-9. The complete 2.4 GHz PHY layer preamble is shown in Figure 10-15.

**Figure 10-15—2.4 GHz PHY preamble format**

10.4.3 Header modulation

The 2.4 GHz PHY header and MAC header shall be modulated using DQPSK modulation for all modulation formats except the 11-Mb/s QPSK-TCM mode.

The header is usually much shorter than the MAC Frame Body field. Consequently, it is more probable to correctly receive the header than the MAC Frame Body field even without a FEC for the header that is modulated in the DQPSK format. This is true for all modulation formats other than the 11-Mb/s QPSK-TCM.

For the QPSK-TCM mode, the combined PHY header, MAC header and the HCS shall be repeated twice as shown in Figure 10-14 to ensure that the header error rate is lower than the frame payload error rate. The first occurrence of the combined PHY header, MAC header and HCS shall be encoded using the DQPSK modulation. The second occurrence of the combined PHY header, MAC header, and HCS shall be encoded using the 11-Mb/s QPSK-TCM format.

10.4.4 Scrambling

A side-stream scrambler (see Figure 10-16) shall be used for the MAC Header field, HCS field, MAC Frame Body field and, if present, the stuff bits. The PHY Preamble field, PHY Header field and tail symbols shall

not be scrambled. The polynomial generator, $g(D)$, for the pseudo-random binary sequence (PRBS) generator shall be as follows:

$$g(D) = 1 + D^{14} + D^{15}$$

where D is a single bit delay element. The polynomial forms not only a maximal length sequence, but also is a primitive polynomial (see Peterson, et al. [B9]). By the given generator polynomial, the corresponding PRBS, x_n , is generated as follows:

$$x_n = x_{n-14} \oplus x_{n-15}$$

where " $\oplus$ " denotes modulo-2 addition.




Figure 10-16—Realization of side-stream scrambler by linear feedback shift registers

The following sequence defines the initialization sequence, x_{init} , which is specified by the parameter "seed value" in Table 10-10.

$$x_{init} = [x_{n-1}^i, x_{n-2}^i, x_{n-3}^i, x_{n-4}^i, x_{n-5}^i, x_{n-6}^i, x_{n-7}^i, x_{n-8}^i, x_{n-9}^i, x_{n-10}^i, x_{n-11}^i, x_{n-12}^i, x_{n-13}^i, x_{n-14}^i, x_{n-15}^i]$$

where x_{n-k}^i represents the binary initial value at the output of the k^{th} delay element.

The scrambled data bits, s_n , are obtained as follows:

$$s_n = b_n \oplus x_n$$

where b_n represents the unscrambled data bits. The side-stream de-scrambler at the receiver shall be initialized with the same initialization vector, x_{init} , used in the transmitter scrambler. The initialization vector is determined from the seed identifier contained in the PHY header of the received frame.

The 15-bit seed value chosen shall correspond to the seed identifier, as described in 10.4.5, and as shown in Table 10-10. The seed identifier value is set to 00 when the PHY is initialized and is incremented in a 2-bit rollover counter for each frame that is sent by the PHY. The value of the seed identifier that is used for the frame is sent in the PHY header, as described in 10.4.5.

Table 10-10—Scrambler seed selection

Seed identifier b1, b0	Seed value $x_{14} \dots x_0$
0, 0	0011 1111 1111 111
0, 1	0111 1111 1111 111

Table 10-10—Scrambler seed selection

Seed identifier b1, b0	Seed value $x_{14} \dots x_0$
1, 0	1011 1111 1111 111
1, 1	1111 1111 1111 111

The 15-bit seed value is configured as follows. At the beginning of each PHY frame, the register is cleared, the seed value is loaded in using x_{init} , where n is set to 15, and the first scrambler bit is calculated. The PHY preamble and PHY header shall not be scrambled. The first bit of data of the MAC header, is modulo-2 added with the first scrambler bit, followed by the rest of the bits in the MAC Header field and the MAC Frame Body field. In the 11 Mb/s mode, the PHY header of the second repetition of the PHY + MAC header + HCS shall not be scrambled. In this mode, the scrambler shall be re-initialized with the same seed used for the first header when it begins scrambling the second header. The scrambler shall continue its operation as normal for the MAC Frame Body field following the second header structure.

10.4.5 PHY header

The PHY header consists of two octets that identify the number of octets in the frame payload (which does not include the FCS, as described in 6.2), the data rate of the MAC Frame Body field, and the seed identifier for the data scrambler. The frame payload length does not include the tail symbols, as described in 10.4.7, or the stuff bits, as described in 10.4.6. The fields for the PHY header are shown in Table 10-11. Bit b0 is sent over the air first and the other bits follow sequentially.

Table 10-11—PHY header

Bits	Content	Description
b1–b0	Seed identifier	A field that selects the seed for the data scrambler, defined in Table 10-10.
b4–b2	MAC Frame Body field data rate	A field that indicates the data rate at which the MAC Frame Body field is sent. The data rate encodings are defined in Table 10-12.
b15–b5	Payload length	A field that contains the length of the frame payload, in octets, MSB is b15, LSB is b5, e.g., 5 octets of data, is encoded as 0b00000000101 (MSB on left, LSB on right). A zero-length frame payload is encoded as 0b00000000000, and there is no FCS for this frame.

The encoding for the MAC Frame Body field data rate is defined in Table 10-12.

Table 10-12—MAC Frame Body field data rate encoding

Modulation	Data rate	b4	b3	b2
QPSK-TCM	11 Mb/s	0	0	0
DQPSK	22 Mb/s	0	0	1

Table 10-12—MAC Frame Body field data rate encoding (continued)

Modulation	Data rate	b4	b3	b2
16-QAM-TCM	33 Mb/s	0	1	0
32-QAM-TCM	44 Mb/s	0	1	1
64-QAM-TCM	55 Mb/s	1	0	0

10.4.6 Stuff bits

If the total length of the MPDU is not an integer multiple of the number of bits/symbol that will be used to modulate the MPDU, then stuff bits shall be added to the end of the MPDU prior to modulation. The stuff bits may be set to either 0 or 1 and shall be ignored when the frame is received. Note that the stuff bits are not a part of either the HCS or FCS calculation. A compliant PHY shall add enough stuff bits so that the MPDU plus stuff bits is an integer multiple of the bits/symbol that is to be used to modulate the MPDU. A compliant PHY shall add less than the number of bits than are contained in a single symbol, i.e., less than 3 for the 33-Mb/s mode and less than 5 for the 55-Mb/s mode.

Since the MPDU is an integer number of octets, this requirement applies only to the 33- and 55-Mb/s modes. For the 11-, 22-, and 44-Mb/s modes, the bits/symbol (1, 2, and 4, respectively) allow the MPDU to be directly mapped into an integer number of symbols.

10.4.7 Tail symbols

Tail symbols shall be added to the end of the MAC Frame Body field, i.e., after either the FCS or the stuff bits, if they are present, for all modulation formats. The tail symbols are used for trellis coded modulation formats in order to terminate the encoded trellis sequence in a known state to aid the decoding process.

For 11-Mb/s QPSK-TCM format, three tail symbols, each containing 1 bit, shall be appended to the end of the MAC Frame Body field. As shown in Figure 10-11, the lowest order input bit x_n^1 solely determines the state transitions. Encoded trellis sequences shall be terminated in state 0, i.e., S0 at the end of each transmission frame as illustrated in Figure 10-17. Table 10-13 shows the assignment of trellis bits based on the last state visited at the end of the MAC Frame Body field.

Table 10-13—Assignment of trellis tail symbol for QPSK-TCM

Last state	1st tail bit x_n^1	2nd tail bit x_n^1	3rd tail bit x_n^1
S0	0	0	0
S1	0	1	1
S2	1	0	0
S3	1	1	1
S4	0	1	0
S5	0	0	1
S6	1	1	0
S7	1	0	1




Figure 10-17—Assignment of tail symbols for QPSK-TCM

For the 16/32/64-QAM formats, the lower order two input bits x_n^2, x_n^1 to the trellis encoder determine the trellis code state transitions as illustrated in Figure 10-12. In this case, only two trellis symbols shall be appended, as shown in Figure 10-18, to the end of the MAC Frame Body field in order to terminate the trellis sequence in state 0, i.e., S0. Consequently, the one higher order bit, x_n^3 , for 16-QAM does not affect the outcome. Likewise, the higher order bits, x_n^4, x_n^3 and x_n^5, x_n^4, x_n^3 , for 32-QAM and 64-QAM, respectively, are irrelevant in determining the final state. The resultant bit assignments to trellis tail symbols shall take on the form given in Table 10-14, for which x_n^3, x_n^2, x_n^1 refer to 16-QAM, $x_n^4, x_n^3, x_n^2, x_n^1$ refer to 32-QAM, and $x_n^5, x_n^4, x_n^3, x_n^2, x_n^1$ refer to 64-QAM.

For the DQPSK modulation format, the two tail symbols shall be constructed from the binary sequence 0101 where 0 is the first bit of the bit pair.




Figure 10-18—Assignment of tail symbols for 16/32/64-QAM-TCM

Table 10-14—Assignment of trellis tail symbols for 16/32/64-QAM

Last State	1st Tail symbol $x_n^5, x_n^4, x_n^3, x_n^2, x_n^1$	2nd Tail symbol $x_n^5, x_n^4, x_n^3, x_n^2, x_n^1$
S0	XXX00	XXX00
S1	XXX00	XXX11
S2	XXX01	XXX00
S3	XXX01	XXX11
S4	XXX10	XXX00
S5	XXX10	XXX11

Table 10-14—Assignment of trellis tail symbols for 16/32/64-QAM (continued)

Last State	1st Tail symbol $x_n^5, x_n^4, x_n^3, x_n^2, x_n^1$	2nd Tail symbol $x_n^5, x_n^4, x_n^3, x_n^2, x_n^1$
S6	XXX11	XXX00
S7	XXX11	XXX11

10.5 Transmitter specifications

10.5.1 Error vector magnitude definition

The modulation accuracy of a compliant IEEE 802.15.3 transmitter is determined with an error vector magnitude (EVM) measurement (see Figure 10-19). In order to calculate this measurement, a time record of N received signal co-ordinate pairs (I_j, Q_j) is captured. For each received symbol, a decision is made as to which symbol was transmitted. The ideal position of the chosen symbol (the center of the decision box) is represented by the vector (I_j, Q_j) . The error vector $(\delta I_j, \delta Q_j)$ is defined as the distance from this ideal position to the actual position of the received symbol.




Figure 10-19—Error vector calculation

Thus, the received vector is the sum of the ideal vector and the error vector.

$$(\tilde{I}_j, \tilde{Q}_j) = (I_j, Q_j) + (\delta I_j, \delta Q_j)$$

The EVM for this standard is defined as follows:

$$EVM \equiv \sqrt{\frac{\frac{1}{N} \sum_{j=1}^N (\delta I_j^2 + \delta Q_j^2)}{S_{max}^2}}$$

where S_{max} is the magnitude of the vector to the outermost constellation point and $(\delta I_j, \delta Q_j)$ is the error vector.

10.5.2 EVM calculated values

A compliant transmitter shall have EVM values of less than those given in Table 10-15 for all of the modulation levels supported by the PHY when measured for 1000 symbols. The error vector measurement shall be made on baseband *I* and *Q* data after recovery through a ideal reference receiver system. The ideal reference receiver shall perform carrier lock, symbol timing recovery and amplitude adjustment while making the measurements. The ideal reference receiver shall have a data filter impulse response that approximates that of an ideal root raised cosine filter with 30% excess bandwidth.

Table 10-15—EVM values for various modulations

Modulation	EVM (%)
64 QAM	3.3
32 QAM	4.8
16 QAM	7.5
DQPSK	9.2
QPSK-TCM	20.0

10.5.3 Transmit power spectral density (PSD) mask

The transmitted spectral products shall be less than the limits specified in Table 10-16. The power shall be measured in a 100 kHz bandwidth relative to the highest average power in a 100 kHz bandwidth measured within ± 6 MHz of the center frequency.

Table 10-16—Transmit PSD limits

Frequency	Relative limit
$7.5 \text{ MHz} < f - f_c < 15 \text{ MHz}$	-30 dBr
$15 \text{ MHz} < f - f_c < 22 \text{ MHz}$	$-10/7[f - f_c \text{ (MHz)} + 13] \text{ dBr}$
$22 \text{ MHz} < f - f_c $	-50 dBr

The transmitter may also have one in-band image, i.e., within 2.4–2.4835 GHz, with a relaxed transmit PSD requirement of -40 dBr over a 15 MHz bandwidth.

A graphical (informative) representation of the transmit PSD is shown in Figure 10-20.

10.5.4 Transmit center frequency tolerance

The transmitted center frequency tolerance shall be $\pm 25 \times 10^{-6}$ maximum.

10.5.5 Symbol rate

The PHY shall be capable of transmitting at a symbol rate of 11-Mbaud $\pm 25 \times 10^{-6}$.

The MAC parameter *pClockAccuracy* shall be $\pm 25 \times 10^{-6}$.




Figure 10-20—Transmit power spectral density mask

10.5.6 Clock synchronization

The transmit center frequency and the symbol rate shall be derived from the same reference oscillator.

10.5.7 Transmit power-on and power-down ramp

The transmit power-on ramp is defined as the time it takes for the RF power emitted by the compliant DEV to rise from less than 10% to greater than 90% of the maximum power to be transmitted in the frame.

The transmit power-on ramp shall be less than 2 μ s.

The transmit power-down ramp is defined as the time it takes for the RF power emitted by the compliant DEV to fall from greater than 90% to less than 10% of the maximum power to be transmitted in the frame.

The transmit power-down ramp shall be less than 2 μ s.

The transmit power ramps shall be constructed such that the emissions conform with the unwanted emissions specification defined in 10.2.5.

10.5.8 RF carrier suppression

The RF carrier suppression, measured at the channel center frequency, shall be at least 15 dB below the peak $\sin(x)/x$ power spectrum. The RF carrier suppression shall be measured while transmitting a repetitive 01 data sequence with the scrambler disabled using DQPSK modulation. A 100 kHz resolution bandwidth shall be used to perform this measurement.

10.5.9 Transmit power

The maximum allowable output power, as measured in accordance with practices specified by the appropriate regulatory bodies, is shown in Table 10-17. A compliant DEV may use any transmit power level up to the applicable limits in the geographical region. In the USA, the radiated emissions should also conform to the ANSI uncontrolled radiation emission standards (see IEEE Std C95.1™, 1999 Edition [B7]). In other geographical regions, compliant DEVs shall also conform to any applicable radiation emission standards.

Table 10-17—Maximum transmit power levels

Geographical region	Power limit	Regulatory document
Japan	10 mW	ARIB STD-T66
Europe (except Spain and France)	100 mW EIRP 10 mW/MHz peak power density	ETS 300-328 [B1]
USA	50 mV/m at 3 m in at least a 1 MHz resolution bandwidth ^a	47 CFR 15.249

^aElectric field strength measurement rather than conducted power measurement.

A compliant transmitter that is capable of transmitting more than 0 dBm shall be capable of reducing its power to less than 0 dBm in monotonic steps no smaller than 3 dB and no larger than 5 dB. The steps shall form a monotonically decreasing sequence of transmit power levels. A compliant DEV shall have its supported power levels indicated in its PHY PIB based on its maximum transmit power and power level step size.

The minimum TX power level required to support TPC, $pMinTpcLevel$, shall be 0 dBm.

10.6 Receiver specifications

10.6.1 Error rate criterion

The error rate criterion shall be a frame error rate (FER) of less than 8% with a frame payload length of 1024 octets of pseudo-random data generated with a PN23 sequence as defined by $x_{n+1} = x_n^{23} + x_n^5 + 1$.

Note that the frames used for measuring the error rate criterion include not only the frame payload of 1024 octets, but also the PHY preamble, PHY header, MAC header, HCS, and the FCS.

10.6.2 Receiver sensitivity

The receiver sensitivity is the minimum power level of the incoming signal, in dBm, present at the input of the receiver for which the error rate criterion in 10.6.1 is met. The error ratio shall be determined after any error correction has been applied. Compliant systems may have a lower actual sensitivity than the reference sensitivity. A compliant DEV shall achieve at least the reference sensitivity listed in Table 10-18 for each of the modulation formats that the DEV supports.

Table 10-18—Reference sensitivity levels for modulation formats

Modulation	Reference sensitivity
QPSK-TCM	−82 dBm
DQPSK	−75 dBm
16-QAM-TCM	−74 dBm
32-QAM-TCM	−71 dBm
64-QAM-TCM	−68 dBm

10.6.3 Receiver maximum input level

The receiver maximum input level is the maximum power level of the incoming signal present at the input of the receiver for which the error rate criterion in 10.6.1 is met. A compliant receiver shall have a receiver maximum input level of at least -10 dBm for each of the modulation formats that the DEV supports.

10.6.4 Receiver jamming resistance

The jamming resistance levels are given in Table 10-19. The high-density channel plan shall be used for this test. The adjacent channel is the one on either side of the desired channel that is closest in frequency to the desired channel for the high-density channel mode defined in 10.2.3. The alternate channel is one removed from the adjacent channel in that channel mode. For example, when channel 1 is the desired channel, channel 2 is the adjacent channel and channels 4 and 5 are the alternate channels. When channel 2 is the desired channel, channels 1 and 4 are the adjacent channels and channel 5 is the alternate channel. Channel 3 is not used for the purposes of this test.

Table 10-19—Receiver jamming resistance requirements

Modulation format	Adjacent channel rejection	Alternate channel rejection
QPSK-TCM	33 dB	48 dB
DQPSK	26 dB	41 dB
16-QAM-TCM	25 dB	40 dB
32-QAM-TCM	22 dB	37 dB
64-QAM-TCM	19 dB	34 dB

The list of adjacent and alternate channels is given in Table 10-20.

Table 10-20—Adjacent and alternate channels for receiver jamming resistance test

Desired channel number	Adjacent channel number	Alternate channel number
1	2	4, 5
2	1, 4	5
4	2, 5	1
5	4	1, 2

The adjacent channel rejection shall be measured as follows. The desired signal shall be a conformant IEEE Std 802.15.3 2.4 GHz PHY signal of pseudo-random data modulated with one of the five modulation types. The desired signal is input to the receiver at a level 6 dB above the reference sensitivity for that modulation as given in Table 10-18. In either the adjacent or alternate channel a conformant IEEE 802.15.3 signal is input at the level specified in Table 10-19 relative to the reference sensitivity for that modulation as given in Table 10-18. For example, for QPSK-TCM the desired signal is input at -76 dBm while the adjacent channel interferer would be input at a level of -49 dBm. The interfering signal shall be DQPSK modulated with pseudo-random data uncorrelated in time with the desired signal. The receiver shall meet the error rate criteria defined in 10.6.1 under these conditions.

A compliant implementation shall satisfy the receiver jamming test for all of the modulations types supported by the DEV.

The desired and interfering signals shall conform to the transmit PSD mask specified in 10.5.3. In addition, the test shall be performed for only one interfering signal at a time.

10.6.5 Receiver CCA performance

A compliant receiver provides CCA capability by performing energy detection in the received signal bandwidth. The start of a valid preamble sequence at a receive level equal to or greater than the minimum sensitivity for the DQPSK base rate, as described in 10.6.2, shall cause CCA to indicate medium busy with a probability of >90% within five CAZAC periods as described in 10.4.2. The receiver CCA function shall in all circumstances report medium busy with any signal 20 dB above the minimum sensitivity for the DQPSK base rate.

The CCA detection time shall be equal $pCcaDetectTime$, which is five CAZAC periods. The CCA shall be maintained as busy until the end of the frame for which the inverted CAZAC sequence was detected.

10.6.6 Receiver RSSI

RSSI is defined as the power relative to the maximum receiver input power level, as described in 10.6.3, in 8 steps of 8 dB with ± 4 dB step size accuracy. The range covered shall be a minimum of 40 dB. The steps shall be monotonic. The RSSI power shall be the average power measured in the last CAZAC sequence of the PHY preamble, as described in 10.4.2.

10.6.7 Link quality indication

The link quality indication (LQI) shall be reported for the TCM coded QAM modes using a signal-to-noise ratio (SNR) estimation. The SNR shall be measured at the decision point in the receiver. The SNR includes the thermal noise, distortion, uncorrected interference, and other signal impairments at the decision point in the receiver. The receiver shall report the SNR as a 5-bit number that covers a range from 6 dB to 21.5 dB of SNR. The value 0x00000 shall correspond to less than or equal to 6 dB SNR and 0x11111 shall correspond to more than or equal to 21.5 dB SNR with equal steps in between. The LQI SNR shall be measured during the TCM MAC Frame Body field and shall be reported after the last FCS symbol.

10.7 PHY management

The PHY PIB comprises the managed objects, attributes, actions, and notifications required to manage the PHY layer of a DEV. The encoding of the PHY data rates used in the Supported Data Rates field in the Capability IE, as described in 6.4.11, is given in Table 10-21.

Table 10-21—2.4 GHz PHY supported data rate encoding

Rates supported (Mb/s)	b0	b1	b2	b3	b4
22	0	0	0	0	0
11, 22	1	0	0	0	0
11, 22, 33	0	1	0	0	0
11, 22, 33, 44	1	1	0	0	0
11, 22, 33, 44, 55	0	0	1	0	0

The encoding of the supported PHY data rates into an octet number is accomplished by adding bits b5–b7, all set to zero, to the encoding given in Table 10-21. Bit b0 is the LSB while bit 7 is the MSB. Thus a DEV that supports 11 Mb/s, 22 Mb/s, and 33 Mb/s would have a Supported Data Rates field 01000 (LSB to MSB) and an octet encoding of 0x2.

The encoding of the preferred fragment size used in the Capability IE, as described in 6.4.11, is given in Table 10-22.

Table 10-22—2.4 GHz PHY preferred fragment size encoding

Field value	Preferred fragment size (octets)
0	<i>pMaxFrameBodySize</i>
1	1792
2	1536
3	1280
4	1024
5	512
6	256
7	<i>pMinFragmentSize</i>

The Fragmentation Control field, 6.2.4, in an Imm-ACK or Dly-ACK frame may be used by the DEV to report the status of the frame that is being ACKed. In the case of a Dly-ACK frame, the status is valid only for the frame that prompted the Dly-ACK frame and not for other frames indicated in the Dly-ACK frame, 6.3.2.2. DEVs are neither required to send receive status information in Imm-ACK and Dly-ACK frames nor required to decode the field in Imm-ACK or Dly-ACK frames. If the Fragmentation Control field is used to send receive status, for the 2.4 GHz PHY it shall be formatted as illustrated in Figure 10-21.

Bits: b0–b2	b3–b7	b8–b23
RSSI	LQI	Reserved

Figure 10-21—Format of fragmentation control field when used for receive status information

The RSSI field and LQI field are set to the values determined for the frame that is being ACKed. The values for RSSI and LQI are defined in 10.6.6 and 10.6.7, respectively.

The PHY-dependent PIB values for the 2.4 GHz PHY are given in Table 10-23 and Table 10-24. The PHY PIB characteristics group, Table 10-23, contains information that is common to most 2.4 GHz implementations.

The PHY PIB implementation group, Table 10-24, contains information that is more characteristic of a particular PHY implementation than of the PHY as a whole.

Table 10-23—PHY PIB characteristics group parameters

Managed object	Octets	Definition	Access
<i>phyType</i>	1	0x00=2.4 GHz	Read/write
<i>phyRegDomainsSupported</i>	Variable	One octet for each regulatory domain supported, as defined for <i>phyCurrentRegDomain</i>	Read/write
<i>phyCurrentRegDomain</i>	1	0x00 = European Telecommunications Standards Institute 0x01 = Federal Communications Commission 0x02 = Industry Canada 0x03 = Association of Radio Industries and Businesses 0x04–0xFF = Reserved	Read/write
<i>phyDataRateVector</i>	1	Encodes the data rates, defined in Table 10-21 and 10.7.	Read/write
<i>phyNumChannelsSupported</i>	1	Value = 0x05, see 10.2.3.	Read/write
<i>phyCurrentChannel</i>	1	Indicates the channel that is currently being used, see 10.2.3.	Read/write
<i>phyCcaThreshold</i>	1	The CCA threshold in dBm, encoded in twos complement format. The value is implementation dependent but no larger than the value listed in 10.6.5.	Read/write
<i>phyMaxFrameLength</i>	2	<i>pMaxFrameBodySize</i> , see 10.2.8.1.	Read/write
<i>phyPreferredFragmentSize</i>	1	The preferred fragment size of a DEV. The encoding is defined in Table 10-22.	Read/write

Table 10-24—PHY PIB implementation group parameters

Managed object	Octets	Definition	Access
<i>phyDiversitySupported</i>	1	Numeric entry that indicates the number of antennas that are available.	Read/write
<i>phyMaxTXPower</i>	1	The maximum TX power that the DEV is capable of using, 6.4.11, implementation dependent.	Read/write
<i>phyTXPowerStepSize</i>	1	The step size for power control supported by the DEV, 6.4.12, implementation dependent.	Read/write
<i>phyNumPMLevels</i>	1	Number of power management levels supported. The range is 1 to 8 and the value is implementation dependent.	Read/write
<i>phyPMLLevelReturn</i>	Variable	Table of vectors with number of entries given by <i>phyNumPMLLevels</i> . Each vector is the time required to change between power saving states of the PHY. Vector number 0 is the time required to change the PHY from the off state to a state where it is ready to receive commands. Other values are implementation dependent.	Read/write

11. PHY specification for millimeter wave

11.1 General requirements

A compliant millimeter wave (mmWave) PHY shall implement at least one of the following PHY modes:

- a) Single Carrier mode in mmWave PHY (SC PHY), as defined in 11.2.
- b) High Speed Interface mode in mmWave PHY (HSI PHY), as defined in 11.3.
- c) Audio/Visual mode in mmWave PHY (AV PHY), as defined in 11.4.

Unless otherwise specified, all reserved fields shall be set to zero on transmission and shall be ignored upon reception.

Unless otherwise stated, in all figures in this clause the ordering of the octets and bits as they are presented to the PHY for modulation is the same, as defined in 6.1.

11.1.1 Regulatory information

The mmWave PHY operating frequency is within the 57.0–66.0 GHz range as allocated by the regulatory agencies in Europe, Japan, Canada, and the United States. This band will also be available in other areas where allocated by the regulatory bodies.

The documents listed in Table 11-1 are provided as a reference for various geographic regulatory regions. The list is neither exhaustive nor complete. It is the responsibility of the implementer to verify that the DEV complies with all regulatory requirements in the geographic region where the device is deployed or sold.

Table 11-1—Documents for selected geographic regulatory regions

Region	Regulatory document
Australia	Radio communications class license 2000
Canada	RSS-210, Issue 6, September 2005
Japan	Regulations for the enforcement of radio law, 6-4.2 specified low power radio station (17) 59–66 GHz band
USA	47 CFR 15.255

The maximum allowable output power, as measured in accordance with practices specified by the appropriate regulatory bodies, is shown in Table 11-2. A compliant DEV may use any transmit power level up to the applicable limits in the geographical region.

Table 11-2—Maximum transmit power levels in selected geographical regions

Region	Power limit	EIRP limit	Regulatory document
USA	—	Maximum indoor EIRP: 27 dBi Maximum outdoor EIRP: 40 dBi	47 CFR 15.255
Japan	Maximum output power: 10dBm Maximum bandwidth: 2.5 GHz	Maximum EIRP: 57 dBi	ARIB STD-T69, ARIB STD-T74
Australia	Maximum output power: 10dBm	Maximum EIRP: 51.8 dBi	Radiocommunications Class License 2000

11.1.2 RF power measurements

Unless otherwise stated, all RF power measurements for the purpose of this standard, either transmit or receive, shall be made based on EIRP and any radiated measurements shall be corrected to compensate for the antenna gain in the implementation. The gain of the antenna is the maximum estimated gain by the manufacturer.

11.1.3 Unwanted emissions

Conformant implementations shall comply with the in-band and out-of-band emissions for all operational modes as set by the applicable regulatory bodies.

11.1.4 Operating temperature range

A conformant implementation shall meet all of the specifications in this standard for ambient temperatures from 0 °C to 40 °C.

11.1.5 RF channelization

The mmWave PHYs use the channels defined in Table 11-3.

Table 11-3—mmWave PHY channelization

CHNL_ID	Start frequency ^a	Center frequency	Stop frequency ^a
1	57.240 GHz	58.320 GHz	59.400 GHz
2	59.400 GHz	60.480 GHz	61.560 GHz
3	61.560 GHz	62.640 GHz	63.720 GHz
4	63.720 GHz	64.800 GHz	65.880 GHz

^aThe start and stop frequencies are nominal values. The frequency spectrum of the transmitted signal needs to conform to the transmit power spectral density (PSD) mask for the PHY mode as well as any regulatory requirement

11.1.6 Transmit PSD mask

The transmitted spectrum shall adhere to the transmit spectrum mask shown in Figure 11-1. For the transmit mask measurements, the resolution bandwidth is set to 3 MHz and video bandwidth to 300 kHz. During OOK modulation, transmitters shall meet the same PSD mask, except for the single line spectra of 40 dB above the 0 dB line in Figure 11-1 within the frequency band of [−6 MHz,+6 MHz] from the carrier frequency.

11.1.7 Error vector magnitude calculation**11.1.7.1 SC PHY**

The error vector measurement shall be made on baseband *I* and *Q* data after recovery through an ideal reference receiver system. An ideal receiver is a receiver that is capable of converting the transmitted signal into a stream of complex samples at sufficient rate or more, with sufficient accuracy in terms of *I/Q* arm amplitude and phase balance, DC offsets, and phase noise. It shall perform carrier lock, symbol timing




Figure 11-1—Transmit spectral mask

recovery and amplitude adjustment while making the measurements. For SC PHY EVM, measuring 1000 chips at the chip rate is recommended. The equation for EVM is given by the following:

$$EVM \equiv \sqrt{\frac{1}{1000P_{avg}} \sum_{i=1}^{1000} [(I_i - I_i^*)^2 + (Q_i - Q_i^*)^2]}$$

where

- P_{avg} is the average power of the constellation
- (I_i^*, Q_i^*) is the complex coordinates of the i^{th} measured chip
- (I_i, Q_i) is the complex coordinates of the nearest constellation point for the i^{th} measured chip

The measuring device should have the accuracy of at least 20 dB better than the EVM value to be measured. The test equipment should use a root-raised cosine filter with roll-off factor of 0.25 for the pulse shaping filter when conducting EVM measurement.

11.1.7.2 AV and HSI PHYs

The EVM of an OFDM mmWave transmitter is measured by calculating the Euclidean distance between the received symbol and the closest constellation point. The EVM measurement has the following requirements:

- a) The measurement shall be made over the OFDM data symbols starting with the PHY header.
- b) The phase and frequency for the data subcarriers shall be corrected for using the pilot subcarriers.
- c) Channel estimation shall have been performed for each of the data subcarriers.
- d) The error vector magnitude calculation for the HSI and AV PHY modes shall be averaged over a minimum of 20 frames.

The EVM is then the average the RMS errors in the frame over a number of frames. In equation form, it is given by the following:

$$EVM_{rms} \equiv \frac{1}{N_p} \sqrt{\sum_{i=1}^{N_p} \left[\frac{1}{P_{avg} N_s N_{dsc}} \sum_{j=1}^{N_s} \left\{ \sum_{k=1}^{N_{dsc}} [(I_{ijk} - I_{ijk}^*)^2 + (Q_{ijk} - Q_{ijk}^*)^2] \right\} \right]}$$

where

- N_p is the number of frames in the measurement
- N_s is the number of symbols per frame in the measurement
- N_{dsc} is the number of data subcarriers in the OFDM modulation
- P_{avg} is the average power of the constellation
- (I_{ijk}^*, Q_{ijk}^*) is the complex coordinates of the j^{th} measured symbol in the k^{th} subcarrier of the i^{th} frame
- (I_{ijk}, Q_{ijk}) is the complex coordinates of the nearest constellation point for the j^{th} measured symbol of the k^{th} subcarrier of the i^{th} frame

11.1.8 Common PHY management for mmWave PHY modes

11.1.8.1 Supported MCSs

The Supported Data Rates field in the Capability IE, as described in 6.4.11, shall be formatted as illustrated in Figure 11-2.

Bits: 0	b1	b2	b3	b4	b5	b6	b7
SC spreading	SC QPSK	SC 8-PSK	SC 16-QAM	SC LDPC 1	SC LDPC 2	AV 16-QAM	HSI QAM

Figure 11-2—Supported data rates field format for mmWave PHY modes

The SC spreading field shall be set to one if spreading factors 2, 4, and 6 are supported by the SC PHY DEV and shall be set to zero otherwise.

The SC QPSK field shall be set to one if QPSK modulation is supported by the SC PHY DEV and shall be set to zero otherwise.

The SC 8-PSK field shall be set to one if 8-PSK modulation is supported by the SC PHY DEV and shall be set to zero otherwise.

The SC 16-QAM field shall be set to one if 16-QAM modulation is supported by the SC PHY DEV and shall be set to zero otherwise.

The SC LDPC 1 field shall be set to one if the optional FECs LDPC(672, 336), LDPC(672, 504), and LDPC(672, 588) are supported by the SC PHY DEV and shall be set to zero otherwise.

The SC LDPC 2 field shall be set to one if the optional FEC LDPC(1440, 1344) is supported by the SC PHY DEV and shall be set to zero otherwise.

The AV 16-QAM field shall be set to one if 16-QAM modulation is supported by the AV PHY DEV and shall be set to zero otherwise.

The HSI QAM field shall be set to one if the 16-QAM and 64-QAM modulations are supported by the HSI PHY DEV and shall be set to zero otherwise.

11.1.8.2 Preferred fragment size

The encoding of the Preferred Fragment Size field, as described in 6.4.11, is given in Table 11-4.

Table 11-4—Preferred fragment size field encoding

Field value	Preferred fragment size (octets)
0b000	1048576
0b001	262144
0b010	65536
0b011	16384
0b100	4096
0b101	2048
0b110	512
0b111	Reserved

11.1.8.3 Receive Status field

The Receive Status field is used to send information about the received frame to the transmitter, as described in 6.2.4. The Receive Status field shall be formatted as illustrated in Figure 11-3.

Bits: 0	b1–b4	b5–b8	b9–b10	b11–b14	b15–b16	b17	b18–b19	b20–b23
Valid	RSSIR	SINR	Suggested Preamble Type	FER	Suggested Pilot Word Type	Suggested PCES	Suggested MCS Status	Reserved

Figure 11-3—Receive Status field format for mmWave PHY

The Valid field shall be set to one if the information in the Receive Status field is to be accessed by the receiver and shall be set to zero otherwise.

The RSSIR field contains the difference of the received signal power in dBm and the sensitivity in dBm of the selected MCS. The range of the RSSIR field is from 0 dB to 28 dB in 2 dB steps with 0b0000 corresponding to less than or equal to 0 dB, and 0b1111 corresponding to greater than 28 dB. For example, an RSSIR value that is greater than or equal to 8 dB but less than 10 dB would be encoded as 0b0101.

The SINR field contains the estimated signal-to-interference-plus-noise ratio of the received frame, and is defined to be $E_b/(I + N)$ in dB where E_b is the energy per bit, I is the interference power and N is the noise power over the bandwidth. The range of the SINR field is from 2 dB to 28 dB in 2 dB steps with 0b0001 corresponding to less than or equal to 2 dB, 0b1111 corresponding to greater than 28 dB, and 0b0000 reserved. For example, an SINR value that is greater than or equal to 18 dB but less than 20 dB would be encoded as 0b1010.

Valid values of the Suggested Preamble Type field are as follows:

0b00 → Suggest CMS preamble
0b01 → Suggest SC preamble
0b10–0b11 → Reserved

Valid values of the Suggested Pilot Word Type field are as follows:

0b00 → Suggest SC pilot word length of 64
 0b01 → Suggest SC pilot word length of 8
 0b10 → Suggest SC pilot word length of 0
 0b11 → Reserved

The Suggested Pilot Channel Estimation Sequence (PCES) field shall be set to zero to suggest that no PCES should be sent and shall be set to one to suggest the use of PCES.

The FER field contains the exponent of the estimate of the FER ranging from 10^{-1} to 10^{-10} in steps of 10^{-1} with 0b0000 corresponding to an FER exponent of less than or equal to -1 (i.e., an FER greater than or equal to 10^{-1}), 0b1010 corresponding to an FER exponent of greater than -10 (i.e., an FER of less than 10^{-10}), and 0b1011–0b1111 reserved. For example, an FER field value of 0b0110 indicates an exponent of -6 and that the FER was less than or equal to 10^{-6} but was greater than 10^{-7} .

Valid values of the Suggested MCS Status field are as follows:

0b00 → Use an MCS with lower required SNR
 0b01 → Use current MCS
 0b10 → Use an MCS with higher required SNR
 0b11 → Reserved

11.1.9 Requirements for mmWave PNCs

In order to promote coexistence among DEVs using different PHYs, the following MAC rules have defined:

- An AV PNC-capable DEV, when operating as a PNC, shall transmit an AV beacon and a Sync frame using CMS in every superframe.
- An HSI PNC-capable DEV, when operating as a PNC, shall transmit an HSI beacon and a Sync frame using CMS in every superframe.
- An AV PNC-capable DEV shall be able to receive Sync frames and MAC Command frames that use CMS.
- An HSI PNC-capable DEV shall be able to receive Sync frames and MAC Command frames that use CMS.
- A DEV capable of transmitting a Sync frame may do so in the first granted CTA in a superframe, and in every pre-defined number of superframes for the directions indicated in the Sync Frame Frequency IE, as indicated in 6.4.35.

A DEV is defined as being able to receive a CMS frame if it can at least successfully perform detection of the CMS preamble when the signal power is greater than the receiver sensitivity.

11.1.10 CP operation

A PNC-capable DEV compliant to this standard shall allow the use of the CAP for contention based access for association, data and commands when using the mmWave PHY, as described in 6.3.1. A DEV compliant to this standard shall support the use of the CAP. CSMA/CA shall be used for all CPs in a mmWave PHY piconet.

The CPs in a piconet shall all be conducted using the same mmWave PHY mode, one of SC, HSI or AV, as the beacon. Any MCS that is supported by both the source and destination and that is in the same PHY mode as the beacon may be used in a CP, with the exception of the AV HRP modes, which shall not be used in a CP. The usage of directional CAPs is described in 7.6.5.2.

For medium access in the CP, the following rules apply:

- A DEV attempting to access the medium in an SC CAP for data transmission, shall always use frames consisting of the CMS preamble, as described in 11.1.12.4, the MR header, as described in 11.2.2.2, and any SC MCS, as described in 11.2.2.1, using the best TX direction over regular CAP and S-CAP.
- A DEV attempting to access the medium in an HSI CAP for data transmission, shall always do so by using frames consisting of the HSI MCS0 long preamble, as described in 11.3.3.1, and any HSI MCS, as described in 11.3.2.1, using the best TX direction over regular CAP and S-CAP.
- A DEV attempting to access the medium in an AV CP, shall always do so by using omni LRPDU, as described in 11.4.3.

11.1.11 mmWave PHY mode usage in CTA

In a CTA, any PHY mode may be used that is supported by both the source and destination DEVs.

Examples of how this is used in a mmWave PHY piconet are given in F.3.

11.1.12 Common mode signaling (CMS)

Common mode signaling (CMS) is a low data rate SC PHY mode specified to enable interoperability among different PHY modes. The CMS is used for transmission of the Beacon frame defined in 7.6.2, and, if supported, the Sync frame defined in 6.3.6 and 11.1.9. The CMS is also used for transmission of command frame and training sequence in the beam forming procedure, as defined in Clause 12, for the SC and HSI PHYs.

The CMS frame shall be formatted as illustrated in Figure 11-4.



Figure 11-4—CMS frame format

The CMS PHY preamble is defined in 11.1.12.4.

The CMS Frame Header shall be formatted as illustrated in Figure 11-5. The CMS frame header shall be constructed as described in 11.1.12.5 and 11.1.12.6.



Figure 11-5—CMS Frame Header format

The CMS PHY Payload field of a shall be constructed as described in 11.1.12.7.

The chip rate of CMS is 1760 Mchip/s. The entire CMS frame shall be modulated with $\pi/2$ BPSK, as specified in 11.2.2.5.1. The FEC for the CMS frame header and MAC Frame Body field shall be as specified in 11.1.12.1. The CMS Frame Header field and MAC Frame Body field shall be spread as specified in 11.1.12.2. The CMS preamble shall be excluded from the spreading process. The scrambling process shall be specified in 11.1.12.3.

The header rate shall be as defined in Table 11-12 while the PHY payload rate shall be as defined in Table 11-10.

11.1.12.1 Forward error correction for CMS


The FEC scheme for CMS shall be RS coding. The RS(255,239), which is the mother code, shall be used for encoding the MAC Frame Body field of CMS. The RS($n+16, n$), a shortened version of RS(255,239) where n is the number of octets in the frame header, shall be used for encoding the frame header of CMS. Details of the coding are provided in 11.2.2.6.1.

11.1.12.2 Code spreading for CMS

To increase robustness in the frame header and MAC Frame Body field of the CMS, code spreading shall be applied using Golay sequences. The code spreading factor shall be 64, and the Golay sequence specified in Table 11-5 shall be used. The frame header and the MAC Frame Body field shall be spread according to Figure 11-6. Note that in each hexadecimal-equivalent 4-binary-digit group, the leftmost bit shall be the MSB, and the rightmost bit, the LSB. For example, 3 is denoted as 0011. The order of the octets and bits over the air is the same as defined in 6.1.

Table 11-5—Golay sequences with length 64

Sequence name	Sequence value
a_{64}	0x63AF05C963500536
b_{64}	0x6CA00AC66C5F0A39

**Figure 11-6—Realization of the CMS code spreading**


As shown in Figure 11-6, a pseudo-random bit sequence (PRBS) generated using an linear-feedback shift register (LFSR) is used to vary the spreading code from one bit to another, i.e., each bit is spread with Golay code a_{64} or Golay code b_{64} depending on the LFSR output. The LFSR shall be run at a rate equals to 1/64 the bit rate and shall be initialized with the following 15-bit seed value: $[x_{-1}, x_{-2}, \dots, x_{-15}] = [0101\ 0000\ 0011\ 111]$. Each input bit shall be held for 64 chips as it is XORed with Golay code a_{64} or Golay code b_{64} operating at the chip rate.

11.1.12.3 Scrambling for CMS

Scrambling shall be employed to whiten the CMS frames. Scrambling shall be used for the MAC header, HCS, and the entire PHY Payload field as detailed in 11.2.2.10.

11.1.12.4 PHY preamble for CMS

A PHY preamble is used to aid receiver algorithms related to AGC setting, timing acquisition, frame synchronization, and channel estimation. The CMS preamble shall be formatted as illustrated in Figure 11-7.

**Figure 11-7—PHY preamble structure for CMS**

The Sync field is used primarily for frame detection and shall consist of 48 repetitions of $\mathbf{b}_{128}$.

The SFD field is used to establish frame timing and shall consist of [+1 −1 +1 +1 −1 −1 −1] spread by $\mathbf{b}_{128}$.

The CES field is used for channel estimation and shall consist of $\mathbf{b}_{128} \mathbf{b}_{256} \mathbf{a}_{256} \mathbf{b}_{256} \mathbf{a}_{256} \mathbf{a}_{128}$ with $\mathbf{a}_{128}$ first in time.

The Golay complimentary sequences of length 128, denoted by $\mathbf{a}_{128}$ and $\mathbf{b}_{128}$, are shown in Table 11-6. The Golay code complimentary sequences of length 256, denoted by $\mathbf{a}_{256}$ and $\mathbf{b}_{256}$, are defined as follows:

$$\mathbf{a}_{256} = [\mathbf{a}_{128} \mathbf{b}_{128}]$$

$$\mathbf{b}_{256} = [\mathbf{a}_{128} \overline{\mathbf{b}_{128}}]$$

where the number on the right ($\mathbf{b}_{128}$ and $\overline{\mathbf{b}_{128}}$) are the first in time and the binary complement of a sequence x is denoted by an overline on x (i.e., $\overline{x}$).

Alternatively, the CMS preamble may be used by the SC PHY for data transmission in the CAP, as specified in 11.2.3.1.

Table 11-6—Golay sequences with length 128

Sequence name	Sequence value
$\mathbf{a}_{128}$	0x0536635005C963AFFAC99CAF05C963AF
$\mathbf{b}_{128}$	0x0A396C5F0AC66CA0F5C693A00AC66CA0

11.1.12.5 Frame header for CMS

The frame header conveys information in the PHY and MAC headers necessary for successfully decoding the frame. The construction of the CMS header is shown in Figure 11-8.

The detailed process of the construction is as follows:

- Construct the PHY header, as described in 11.1.12.6.
- Compute the HCS as described in 11.2.3.2.2 over the combined PHY and MAC headers.
- Append the HCS to the MAC header.
- Scramble the combined MAC header and HCS, as described in 11.1.12.3.
- Compute the RS Parity field by encoding the concatenation of the PHY Header field, scrambled MAC Header field, and scrambled HCS field into a shortened RS block code, as described in 11.1.12.1.




Figure 11-8—Frame header construction process for CMS

- f) Form the base frame header by concatenating the PHY Header field, scrambled MAC Header field, scrambled HCS field and RS Parity field.
- g) Spread the frame header, as described in 11.1.12.2.

11.1.12.6 PHY header for CMS

The CMS PHY header shall be formatted as illustrated in Figure 11-22. The description of each field is provided in 11.2.3.2.1. In this subclause, the field values for the CMS PHY header are specified.

The Scrambler Seed ID field contains the scrambler seed identifier value, as defined in 11.1.12.3.

The Aggregation field shall be set to zero.

The UEP field shall be set to zero.

The MCS field shall be set to 0b000000.

The Frame Length field shall be an unsigned integer that indicates the number of octets in the MAC Frame Body field, excluding the FCS.

The Preamble Type field shall be set to 0b00.

The Beam Tracking field shall be set to one if the training sequence for beam tracking is following the current frame, and shall be set to zero otherwise.

The Low-latency mode field shall be set to zero.

The Pilot Word Length field shall be set to 0b10.

The PCES field shall be set to zero.

11.1.12.7 PHY Payload field for CMS

The PHY Payload field is the last component of the CMS frame, and is constructed as shown in Figure 11-9.




Figure 11-9—PHY Payload field construction process for CMS

The PHY Payload field of the CMS shall be constructed as follows:

- Scramble the MAC Frame Body field according to 11.1.12.3.
- Encode the scrambled MAC Frame Body field, as specified in 11.1.12.1.
- Spread the encoded and scrambled MAC Frame Body field using the spreading code, as detailed in 11.1.12.2.

11.1.12.8 Receiver clear channel assessment performance for CMS

The start of a valid preamble sequence at a receive level equal to or greater than the minimum sensitivity for the CMS, as described in 11.2.5.2, shall cause CCA to indicate medium busy with a probability of > 90% within *pCcaDetectTime* for the SC PHY mode. The receiver CCA function shall in all circumstances report medium busy with any signal 20 dB above the minimum sensitivity for the CMS. The CCA shall be maintained as busy until the end of the frame.

11.1.13 mmWave PHY PIB

The PHY-dependent PIB values for the mmWave PHY are given in Table 11-7 and Table 11-8. The PHY PIB characteristics group, Table 11-7, contains information that is common to most implementations.

Table 11-7—PHY PIB characteristics group parameters

Managed Object	Octets	Definition	Access
<i>phyType</i>	1	0x01 = mmWave PHY	Read/write
<i>phyMode</i>	1	bit 1 = SC PHY bit 2 = HSI PHY bit 3 = AV PHY bit 4 = OOK mode bit 5 = DAMI mode bit 6–8 = Reserved A bit is set to one if the associated PHY is supported, and is set to zero otherwise.	Read/write
<i>phyRegDomainsSupported</i>	Variable	One octet for each regulatory domain supported, as defined for <i>phyCurrentRegDomain</i> .	Read/write
<i>phyCurrentRegDomain</i>	1	0x00 = European Telecommunications Standards Institute 0x01 = Federal Communications Commission 0x02 = Industry Canada 0x03 = Association of Radio Industries and Businesses	Read/write

Table 11-7—PHY PIB characteristics group parameters (continued)

Managed Object	Octets	Definition	Access
<i>phyDataRateVector</i>	Variable	One octet for each supported MCS. The three MSBs indicate the mmWave PHY mode, as in <i>phyMode</i> , and the last five LSBs contain the MCS supported for that mode using the encoding for that PHY mode.	Read/write
<i>phyNumChannelsSupported</i>	Variable	Value = 0x04, as defined in 11.1.5.	Read/write
<i>phyCurrentChannel</i>	1	Indicates the channel that is currently being used, as defined in 11.1.5.	Read/write
<i>phyCcaThreshold</i>		The CCA threshold in dBm, encoded in twos complement format. The value is implementation dependent but no larger than the value listed in 11.1.12.8.	Read/write
<i>phyFrameLengthMax</i>	2	The value of <i>pMaxFrameBodySize</i>	Read/write

The PHY PIB implementation group, Table 11-8, contains information that is more characteristic of a particular PHY implementation than of the PHY as a whole.

Table 11-8—PHY PIB Implementation group parameters

Managed Object	Octets	Definition	Access
<i>phyDiversitySupported</i>	1	Numeric entry that indicates the number of antennas that are available.	Read/write
<i>phyMaxTXPower</i>	1	The maximum TX power that the DEV is capable of using, 6.4.11, implementation dependent.	Read/write
<i>phyTXPowerStepSize</i>	1	The step size for power control supported by the DEV, 6.4.12, implementation dependent.	Read/write
<i>phyNumPMLlevels</i>	1	Number of power management levels supported. The range is 1 to 8 and the value is implementation dependent.	Read/write
<i>phyPMLlevelReturn</i>	Variable	Table of vectors with number of entries given by <i>phyNumPMLlevels</i> . Each vector is the time required to change between power saving states of the PHY. Vector number 0 is the time required to change the PHY from the off state to a state where it is ready to receive commands. Other values are implementation dependent.	Read/write

11.2 Single Carrier Mode of mmWave PHY

The Single Carrier Mode in mmWave PHY (SC PHY) provides three classes of modulation and coding schemes (MCSs) targeting different wireless connectivity applications. Class 1 is specified to address the low-power low-cost mobile market while maintaining a relatively high data rate of up to 1.5 Gb/s. Class 2 is specified to achieve data rates up to 3 Gb/s. Class 3 is specified to support high performance applications with data rates in excess of 5 Gb/s. Table 11-9 summarizes the MCS classes. There are two mandatory

MCSs for all SC DEVs except for the optional OOK/DAMI mode DEVs, the common mode signaling (CMS) and the mandatory PHY rate (MPR). The optional OOK/DAMI modes, as described in 11.2.8, are allowed for low complexity SC DEVs. The CMS shall be the base rate for the SC PHY. The CMS is used for transmission of beacon and command frame in the association procedure as well as for transmission of command frame and training sequence in the beam forming procedure.

Table 11-9—MCS categorization for the SC PHY

Class	Categorization
Class 1	Data rates < 1.5 Gb/s ^a
Class 2	1.5 Gb/s < Data rates < 3 Gb/s
Class 3	Data rates > 3 Gb/s

^aThe data rates are based on subblock length of 512 and pilot word length of 64.

The SC PHY is specified with a high degree of flexibility in order to allow implementers the ability to optimize for different applications. The SC PHY supports NLOS, as well as line-of-sight (LOS), operation, with or without equalization. The data is divided into blocks, each block is divided into subblocks, as described in 11.2.3.4.1, and each subblock may be equalized in time, frequency or in hybrid time-frequency domain. A subblock consists of pilot word and data.

The SC PHY supports $\pi/2$ BPSK, $\pi/2$ QPSK, $\pi/2$ 8-PSK, and $\pi/2$ 16-QAM modulations. Two main FEC schemes are specified: RS block codes and LDPC block codes. The RS(255,239) and the shortened RS(33,17) block codes are mandatory, whereas all the other FECs are optional.

11.2.1 PHY operating specifications of SC PHY

11.2.1.1 Channelization

The RF channels are defined in Table 11-3. A compliant implementation shall support at least 1 channel from the channels allocated for operation by its corresponding regulatory body.

The *phyCurrentChannel* is the CHNL_ID of the current channel. For the purpose of the Remote Scan Request and Remote Scan Response commands, as described in 6.5.7.3 and 6.5.7.4, respectively, the Channel Index field is the CHNL_ID in Table 11-3.

11.2.1.2 Scanning channels

When a DEV is scanning to start a piconet, it should scan all channels it supports in respective regions to decrease the probability of choosing an occupied channel.

11.2.2 Modulation, forward error correction, and spreading

11.2.2.1 MCS-dependent parameters

The chip rate for all SC PHY MCS is given in Table 11-14. The MCS-dependent parameters shall be set according to Table 11-10. The data rates in the table are approximate and are calculated to three significant figures. The CMS and MPR are part of Class 1, and are listed as MCS identifier 0 and 3, respectively.

Table 11-10—MCS-dependent parameters

MCS class	MCS identifier	Data rate (Mb/s) with pilot word length = 0	Data rate (Mb/s) with pilot word length = 64	Modulation	Spreading factor, L_{SF}	FEC type
Class1	0	25.8 (CMS)	—	$\pi/2$ BPSK	64	RS(255,239)
	1	412	361		4	
	2	825	722		2	
	3	1650 (MPR)	1440		1	
	4	1320	1160	$\pi/2$ BPSK	1	LDPC(672,504)
	5	440	385	$\pi/2$ BPSK	2	LDPC(672,336)
	6	880	770		1	
Class2	7	1760	1540	$\pi/2$ QPSK	1	LDPC(672,336)
	8	2640	2310	$\pi/2$ QPSK	1	LDPC(672,504)
	9	3080	2700	$\pi/2$ QPSK	1	LDPC(672,588)
	10	3290	2870	$\pi/2$ QPSK	1	LDPC(1440,1344)
	11	3300	2890	$\pi/2$ QPSK	1	RS(255,239)
Class3	12	3960	3470	$\pi/2$ 8-PSK	1	LDPC(672,504)
	13	5280	4620	$\pi/2$ 16-QAM	1	LDPC(672,504)

The FEC rate, R_{FEC} , is for an RS(n, m) code or an LDPC(n, m) code is m/n . The subblock length $L_{subblock}$ shall be 512 chips. The pilot word length L_{PW} shall be 0, 8, or 64 chips. For subblocks with $L_{PW} = 8$, the effective length of the pilot word and data is equivalent to that with $L_{PW} = 64$, as described in 11.2.3.4.1.

The MCS-dependent parameters for the optional OOK/DAMI modes are given in Table 11-11.

Table 11-11—MCS-dependent parameters for optional OOK/DAMI modes

Device Type	MCS identifier	Data rate ^a (Mb/s)	Modulation	Spreading factor	FEC type	Support for CMS
PNC-capable DEVs	OOK	25.8 (CMS)	$\pi/2$ BPSK	64	RS(255,239)	Mandatory
		818	OOK	2		
		1640		1		
	DAMI	25.8 (CMS)	$\pi/2$ BPSK	64		
		3270	DAMI	1		
Non PNC-capable DEVs	OOK	818	OOK	2	RS(255,239)	Not mandatory
		1640		1		
	DAMI	3270	DAMI	1		

^aThe data rate of CMS is calculated in a manner similar to Table 11-10. The data rates of two OOK modes are calculated based on pilot word design as in G.5. The data rate of DAMI mode is calculated based on subblock length of 512 and pilot word length of 4. The data rates are approximate and are rounded to three significant figures.

11.2.2.2 Header rate-dependent parameters

The base header rate-dependent parameters shall be set according to Table 11-12. The base headers use a shortened RS code, as defined in 11.2.3.2.3.

Table 11-12—Base header rate-dependent parameters

Header class	Header rate (Mb/s)	Modulation scheme	Spreading factor, L_{SF}	Pilot word length (chips), L_{PW}	Coded bits per subblock, L_{CBPS}	Number of occupied subblocks, $N_{subblock\_hdr}$	Number of stuff bits, L_{STUFF}
CMS rate	12.5	$\pi/2$ BPSK	64	0	8	33	0
Medium rate (MR)	82.5	$\pi/2$ BPSK	6	176 (effective)	56	5	16
High rate (HR)	206	$\pi/2$ BPSK	2		168	2	72

For the MR and HR headers, the header subblock is divided into eight sub-subblocks. The first sub-subblock is a pilot word with $L_{PW} = 64$. This is followed by six sub-subblocks, each with a prepended pilot word $L_{PW} = 8$. The last sub-subblock consists of eight pilot words of $L_{PW} = 8$. This gives the effective pilot word length of $L_{PW} = 176$.

There are three types of MAC subheaders: the standard aggregation subheader, the low-latency EEP subheader, and the low-latency UEP subheader, as defined in 6.2.8. When the MAC subheader is present, the subheader shall be modulated with $\pi/2$ BPSK and the header rate-dependent parameters shall be set according to Table 11-13. The subheader subblock is the same as for the base header, giving the same effective pilot word length of 176 chips. The MAC subheaders for standard aggregation use a shortened RS code for the FEC, as defined in the 11.2.3.2.5. If the base header uses either the CMS or MR, the MAC subheader shall be sent using the subheader MR. If the base header uses HR, the MAC subheader shall be sent using subheader HR. For the standard aggregation frames, each subframe may use a different MCS. For the low-latency frames, the MAC subheaders and the MSDU subheaders shall use the same MCS for the MAC Frame Body field, thus the MCS remains the same within a low-latency aggregation frame.

Table 11-13—MAC subheader rate-dependent parameters for standard aggregation

Frame type	Subheader class	Header rate ^a (Mb/s)	Spreading factor, L_{SF}	Coded bits per subblock, L_{CBPS}	Number of occupied subblocks	Number of stuff bits, L_{STUFF}
Non-secure frame	Subheader MR	131	6	56	9	16
	Subheader HR	394	2	168	3	16
Secure frame	Subheader MR	132	6	56	10	32
	Subheader HR	330	2	168	4	144

^aOnly the parameters for the MAC subheader for standard aggregation is presented in the table. The MAC subheader for low-latency aggregation uses the same MCS as for the MAC Frame Body field and therefore is not listed in this table.

11.2.2.3 Timing-related parameters

Table 11-14 lists the general timing parameters associated with the SC PHY.

Table 11-14—Timing-related parameters

Parameter	Description	Value			Unit	Formula
R_c	Chip rate	1760			Mchip/s	
T_c	Chip duration	~0.568			ns	$1/R_c$
L_{subblock}	Subblock length	512			chips	
L_{PW}	Pilot word length	0	8 ^a	64	chips	
T_{PW}	Pilot word duration	0	4.5	~37.0	ns	
L_{DC}	Length of data chips per subblock	512	56	448	chips	
T_{subblock}	Subblock duration	~290.9			ns	$L_{\text{subblock}} \times T_c$
R_{subblock}	Subblock rate	~3.44			MHz	$1/T_{\text{subblock}}$

^aDetails of the subblock with pilot word length 8 are given in 11.2.2.1 and 11.2.2.2.

11.2.2.4 Frame-related parameters

The frame parameters associated with the PHY are listed in Table 11-15 where CEIL is the ceiling function, which returns the smallest integer value greater than or equal to its argument. The maximum frame duration occurs when the number of octets in the PHY Payload field is 1048576.

Table 11-15—Frame-related parameters

Parameter	Description	Value
N_{SYNC}	Number of code repetitions ^a in the Sync sequence	14
T_{SYNC}	Duration of the Sync sequence	~1 μ s
N_{SFD}	Number of code repetitions in the SFD	4
T_{SFD}	Duration of the SFD	~0.29 μ s
N_{CES}	Number of code repetitions in the CES	9
T_{CES}	Duration of the CES	~0.65 μ s
N_{pre}	Number of code repetitions in the PHY preamble	27
T_{pre}	Duration of the PHY preamble	~1.96 μ s
L_{hdr}	Length of the base header in octets	33
$N_{\text{subblock\_hdr}}$	Number of subblocks in the base frame header	$\text{CEIL}[L_{\text{hdr}} \times 8 \times L_{\text{SF}} / (L_{\text{subblock}} - L_{\text{PW}}^b)]$
T_{hdr}	Duration of the base frame header	$N_{\text{subblock\_hdr}} \times T_{\text{subblock}}$
L_{payload}	Length of frame payload in octets	variable
L_{FCS}	Length of FCS in octets	4

Table 11-15—Frame-related parameters (continued)

Parameter	Description	Value
L_{MFB}	Length of the MAC Frame Body field in octets	$L_{\text{payload}} + L_{\text{FCS}}$
N_{PCES}	Number of code repetitions in the PCES	9
T_{PCES}	PCES duration	0.65 μs
$N_{\text{PCES\_frame}}$	Number of PCESs per frame	$\text{CEIL}[(L_{\text{MFB}} \times 8) / (L_{\text{subblock}} \times 64)] - 1$
$T_{\text{PCES\_interval}}$	Interval of PCES insertion	$T_{\text{subblock}} \times 64 + T_{\text{PW}}$
L_{CBPS}	Number of coded bits per subblock in the MAC Frame Body field	$(L_{\text{subblock}} - L_{\text{PW}}^{\text{C}}) / L_{\text{SF}}$
$N_{\text{subblock\_MFB}}$	Number of subblocks in the MAC Frame Body field	$\text{CEIL}[(L_{\text{MFB}} \times 8) / (R_{\text{FEC}} \times L_{\text{CBPS}})]$
T_{MFB}	Duration of the MAC Frame Body field	$N_{\text{subblock\_MFB}} \times T_{\text{subblock}}$
$T_{\text{datafield}}$	Duration of the PHY Payload field	$T_{\text{MFB}} + (N_{\text{PCES\_frame}} + 1) \times T_{\text{PW}} + N_{\text{PCES\_frame}} \times T_{\text{PCES}}$
T_{frame}	Duration of the frame	$T_{\text{pre}} + T_{\text{hdr}} + T_{\text{datafield}}$

^aA code repetition is defined as a code with length of 128 chips.

^bThe value for L_{PW} given here is the effective pilot word length, as defined in 11.2.2.2.

^cThe value for L_{PW} given here is the effective pilot word length, as defined in 11.2.2.1.

NOTE—Regardless of the maximum frame length, no frame is allowed to exceed the timing boundaries, e.g., CAP end time or CTA end time.

11.2.2.5 Modulation

After channel encoding and spreading, the bits shall be inserted into the constellation mapper. The constellations used for the SC PHY are illustrated in Figure 11-10. The mapping rules and diagrams for skewed constellation are given in 11.3.2.6.

11.2.2.5.1 $\pi/2$ BPSK

The $\pi/2$ BPSK modulation is a binary phase modulation with $\pi/2$ phase shift counter-clockwise. Figure 11-10(a) shows the signal constellation of $\pi/2$ BPSK signals, and Figure 11-11 shows the $\pi/2$ BPSK modulator. The data chips, g_n , at the output of PCES inserter shown in Figure 11-23 are mapped to $d_1, d_2, \dots, d_N$, where $d_N = 2 \times g_N - 1$. The d_n values are mapped onto constellation points z_n as follows:

$$z_n = j^n \times d_n, n = 1, 2, \dots, N$$

where j denotes $\pi/2$ phase rotation. c_n in Figure 11-11(a) denotes a complex envelope of filtered $\pi/2$ BPSK signal.

The filtered waveform of the signal shall satisfy transmit PSD mask as in 11.1.6.

11.2.2.5.2 $\pi/2$ QPSK

The $\pi/2$ QPSK constellation diagram is shown in Figure 11-10(b), with four points equally spaced on a circle of radius one, representing four phases. QPSK shall encode 2 bits per symbol, with input bit d_1 being




Figure 11-10—Constellation maps for modulations: (a) $\pi/2$ BPSK, (b) $\pi/2$ QPSK, (c) $\pi/2$ 8-PSK, (d) $\pi/2$ 16-QAM




Figure 11-11—Possible $\pi/2$ BPSK modulator realizations

the earliest in the stream. The $\pi/2$ shift is employed to obtain a simple implementation aligning with the $\pi/2$ BPSK. The $\pi/2$ rotation is performed in the same manner as in 11.2.2.5.1. As illustrated in Figure 11-10(c), Gray encoding shall be employed. The normalization factor, K_{MOD} , is 1.

11.2.2.5.3 $\pi/2$ 8-PSK

The $\pi/2$ 8-PSK constellation diagram is shown in Figure 11-10(c), equally spaced on a circle of radius one, representing eight phases. The $\pi/2$ 8-PSK shall encode 3 bits per symbol, with input bit d_1 being the earliest in the stream. The $\pi/2$ rotation is performed in the same manner as in 11.2.2.5.1. Gray encoding shall be employed in the mapping of $\pi/2$ 8-PSK. The normalization factor, K_{MOD} , is 1.

11.2.2.5.4 $\pi/2$ 16-QAM

The $\pi/2$ 16-QAM constellation diagram is shown in Figure 11-10(d). The serial bit stream shall be divided into groups of four bits with input bit d_1 being the earliest in the stream. The $\pi/2$ rotation is performed in the same manner as in 11.2.2.5.1. The normalization factor for $\pi/2$ 16-QAM constellation is $1/\sqrt{10}$. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements.

11.2.2.6 FEC

The FEC schemes are specified in this subclause. Support for RS block codes is mandatory, whereas support for LDPC block codes is optional.

11.2.2.6.1 Reed-Solomon block codes in $GF(2^8)$

The RS(255,239), which is the mother code, is used in payloads of CMS, MPR and MCS identifier 1 MCSs in Table 11-10. A shortened version of RS(255,239) is used for the base frame header and MAC subheader, as defined in 11.2.2.2.

The systematic RS code shall use the following generator polynomial:

$$g(x) = \prod_{k=1}^{16} (x + \alpha^k)$$

where $\alpha = 0x02$ is a root of the binary primitive polynomial $p(x) = 1 + x^2 + x^3 + x^4 + x^8$. As notation, the element $M = b_7x^7 + b_6x^6 + b_5x^5 + b_4x^4 + b_3x^3 + b_2x^2 + b_1x^1 + b_0$, has the binary representation $b_7b_6b_5b_4b_3b_2b_1b_0$, where b_7 is the MSB and b_0 is the LSB.

The mapping of the information octets $\mathbf{m} = (m_{238}, m_{237}, \dots, m_0)$ to codeword octets $\mathbf{c} = (m_{238}, m_{237}, \dots, m_0, r_{15}, r_{14}, \dots, r_0)$ is achieved by computing the remainder polynomial $r(x)$:

$$r(x) = \sum_{k=0}^{15} r_k x^k = x^{16} m(x) \bmod g(x)$$

where $m(x)$ is the information polynomial:

$$m(x) = \sum_{k=0}^{238} m_k x^k$$

and $r_k, k = 0, \dots, 15$, and $m_k, k = 0, \dots, 238$, are elements of $GF(2^8)$. The message order is as follows: m_{238} is the first octet of the message and m_0 is the last octet of the message.

For a shortened RS($L_{inf} + 16, L_{inf}$), $239 - L_{inf}$ zero elements are appended to the incoming L_{inf} octet message as follows:

$$m_k = 0, k = L_{inf}, \dots, 238$$

These inserted zero elements are not transmitted. A shift-register implementation of the RS encoder RS($L_{inf} + 16, L_{inf}$) is shown in Figure 11-12, with additions and multiplications over $GF(2^8)$. After m_0 has

been inserted into the shift register, the switch shall be moved from the message polynomial input connection to the shift register output connection (right-to-left).




Figure 11-12—Reed Solomon encoder GF(2⁸)

11.2.2.6.2 Irregular LDPC codes

The irregular LDPC codes are used as a high-performance error correction coding technique. The supported FEC rates, information block lengths, and codeword block lengths are described in Table 11-16.

Table 11-16—Irregular LDPC parameters

FEC rate, R_{FEC}	LDPC information block length (bits), L_{inf}	LDPC codeword block length (bits), L_{FEC}
1/2	336	672
3/4	504	672
7/8	588	672

The LDPC encoder is systematic, i.e., it encodes an information block of size k , $i = (i_0, i_1, \dots, i_{(k-1)})$, into a codeword c of size n , $c = (i_0, i_1, \dots, i_{(k-1)}, p_0, p_1, \dots, p_{(n-k-1)})$, by adding $n-k$ parity bits obtained so that $Hc^T = 0$, where H is an $(n-k) \times n$ parity check matrix.

Each of the parity-check matrices can be partitioned into square subblocks (submatrices) of size $z \times z$ ($z = 21$). These submatrices are either cyclic-permutations of the identity matrix or null (all-zero) submatrices.

The cyclic-permutation matrix p^i is obtained from the $z \times z$ identity matrix by cyclically shifting the columns to the left by i elements. The matrix p^0 is the $z \times z$ identity matrix.

In the following, an example of cyclic-permutation matrices with $z = 21$ is shown. The matrix $\mathbf{p}^1$ and $\mathbf{p}^2$ are produced by cyclically shifting the columns of the identity matrix $\mathbf{I}_{21 \times 21}$ to the left by 1 and 2 places, respectively.

$$\mathbf{p}^0 = \begin{bmatrix} 1 & 0 & \dots & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \dots & 0 & \dots & 0 & \dots \\ 0 & \dots & 0 & 1 & 0 \\ 0 & \dots & \dots & 0 & 1 \end{bmatrix}, \mathbf{p}^1 = \begin{bmatrix} 0 & \dots & \dots & 0 & 1 \\ 1 & 0 & \dots & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \dots & 0 & 1 & 0 & 0 \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix}, \mathbf{p}^2 = \begin{bmatrix} 0 & \dots & 0 & 1 & 0 \\ 0 & \dots & \dots & 0 & 1 \\ 1 & 0 & \dots & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$$

where all the above matrices have dimension 21×21 .

Due to the cyclic permutation, $\mathbf{p}^{21} = \mathbf{p}^0 = \mathbf{I}_{21 \times 21}$.

Figure 11-13 displays the matrix permutation indices of parity-check matrices for all three FEC rates at block length $n = 672$ bits. The integer i denotes the cyclic-permutation matrix $\mathbf{p}^i$, as explained in above example. The ‘-’ entries in the table denote null (all zero) submatrices.

For shortened LDPC operation, the $k-l$ zero elements are appended to the incoming l message bits as follows: $r_i = 0$ for $i = l, l+1, \dots, k-1$. These inserted zero elements are not transmitted.

11.2.2.6.3 Rate 14/15 LDPC code

The rate 14/15 LDPC(1440,1344) code is also systematic, i.e., the LDPC encoder encodes an information block of size k , $\mathbf{i} = (i_0, i_1, \dots, i_{(k-1)})$, into a codeword $\mathbf{c}$ of size n , $\mathbf{c} = (i_0, i_1, \dots, i_{(k-1)}, p_0, p_1, \dots, p_{(n-k-1)})$, by adding $n - k$ parity bits obtained so that $\mathbf{H}\mathbf{c}^T = 0$, where $\mathbf{H}$ is an $(n - k) \times n$ parity check matrix. Denote the 96×1440 parity check matrix as $\mathbf{H} = (h_{i,j})$, where $h_{i,j}$ consists of $\{0,1\}$, $0 \leq i < 96$ and $0 \leq j < 1440$. Table 11-17 shows the matrix elements whose values are ‘1’ in the first 15 columns of parity check matrix.

For $15 \leq j$, the matrix element can be obtained by using the following:

$$h_{i,j} = h_{\text{mod}(i + \lfloor j/15 \rfloor, 96), \text{mod}(j, 15)}$$

where $\text{mod}(x, y)$ is the modulo function and is defined as $x - n \times y$ where n is the nearest integer less than or equal to x/y .

The LDPC(1440, 1344) code is a quasi-cyclic code such that every cyclic shift of a codeword by 15 symbols yields another codeword.

For shortened LDPC operation, the $k-l$ zero elements are appended to the incoming l message bits as follows: $r_i = 0$ for $i = l, l+1, \dots, k-1$. The message order is r_{k-1} as the first bit of the message with r_0 as the last bit of the message. These inserted zero elements are not transmitted.

11.2.2.7 Stuff bits

Stuff bits shall be added to the end of the encoded MAC Frame Body field if the number of the encoded data bits is not an integer multiple of the length of the data portion in the subblock. The number of stuff bits is computed for each subframe if standard aggregation is employed. The calculation of stuff bits is as follows.

In the encoded MAC Frame Body field, the number of FEC codewords, N_{FEC} is given by the following:

$$N_{FEC} = \text{CEIL}[(L_{MFB} \times 8)/(L_{FEC} \times R_{FEC})]$$

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks

(672,336), Code rate: 1/2																																	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
1	-	-	-	5	-	18	-	-	-	-	3	-	10	-	-	-	-	-	-	5	-	-	-	-	-	-	-	5	-	7	-	-	
2	0	-	-	-	-	-	16	-	-	-	-	6	-	-	-	0	-	7	-	-	-	-	-	-	-	10	-	-	-	-	-	19	
3	-	-	6	-	7	-	-	-	-	2	-	-	-	-	9	-	20	-	-	-	-	-	-	-	-	-	19	-	10	-	-	-	
4	-	18	-	-	-	-	-	0	10	-	-	-	-	16	-	-	-	-	9	-	-	-	-	-	4	-	-	-	-	-	17	-	
5	5	-	-	-	-	-	18	-	-	-	-	3	-	10	-	-	5	-	-	-	-	-	-	-	-	-	-	-	-	-	7	-	
6	-	0	-	-	-	-	-	16	6	-	-	-	0	-	-	-	-	-	7	-	-	-	-	-	-	-	-	-	19	-	-	-	
7	-	-	-	6	-	7	-	-	-	-	2	-	-	-	-	9	-	20	-	-	-	-	-	-	-	-	-	-	-	10	-	-	
8	-	-	18	-	0	-	-	-	-	10	-	-	-	-	16	-	-	-	-	9	-	-	-	-	-	-	-	-	-	-	-	17	-
9	-	5	-	-	-	-	-	18	3	-	-	-	-	-	10	-	-	5	-	-	4	-	-	-	-	5	-	-	-	-	-	-	
10	-	-	0	-	16	-	-	-	-	6	-	-	-	0	-	-	-	-	-	7	-	4	-	-	-	-	-	10	-	19	-	-	
11	6	-	-	-	-	-	7	-	-	-	-	2	9	-	-	-	-	-	20	-	-	-	4	-	19	-	-	-	-	-	10	-	
12	-	-	-	18	-	0	-	-	-	-	10	-	-	-	-	16	9	-	-	-	-	-	-	12	-	-	4	-	17	-	-	-	
13	-	-	5	-	18	-	-	-	-	3	-	-	-	-	10	-	-	5	-	-	-	-	-	-	-	-	5	-	-	-	-	-	
14	-	-	-	0	-	16	-	-	-	-	6	-	-	-	0	-	7	-	-	-	-	-	-	-	10	-	-	-	-	-	-	-	
15	-	6	-	-	-	-	-	7	2	-	-	-	-	9	-	-	-	-	-	20	-	-	-	-	-	19	-	-	-	-	-	-	
16	18	-	-	-	-	-	0	-	-	-	-	10	16	-	-	-	-	9	-	-	-	-	-	-	-	-	-	4	-	-	-	-	

(672,504), Code rate: 3/4

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
1	0	-	-	5	-	18	16	-	-	-	3	6	10	-	-	0	-	7	-	5	-	-	4	4	-	10	-	5	-	-	-	
2	-	18	6	-	7	-	-	0	10	2	-	-	-	16	9	-	20	-	9	-	4	12	-	-	4	-	19	-	-	-	-	
3	5	0	-	-	-	-	18	16	6	-	-	3	0	10	-	-	5	-	7	-	4	-	-	4	5	-	10	-	19	-	-	
4	-	-	18	6	0	7	-	-	-	10	2	-	-	-	16	9	-	20	-	9	-	4	12	-	-	4	-	19	-	10	-	-
5	-	5	0	-	16	-	-	18	3	6	-	-	-	0	10	-	-	5	-	7	4	4	-	-	-	5	-	-	-	-	-	
6	6	-	-	18	-	0	7	-	-	-	10	2	9	-	-	16	9	-	20	-	-	-	4	12	19	-	-	-	-	-	-	
7	-	-	5	0	18	16	-	-	-	3	6	-	-	-	0	10	7	-	5	-	-	4	4	-	10	-	5	-	7	-	19	-
8	18	6	-	-	-	-	0	7	2	-	-	10	16	9	-	-	9	-	20	12	-	-	4	-	19	-	4	-	17	-	10	-

(672,588), Code rate: 7/8

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
1	0	18	6	5	7	18	16	0	10	2	3	6	10	16	9	0	20	7	9	5	4	12	4	4	4	10	19	5	10	-	-	-
2	5	0	18	6	0	7	18	16	6	10	2	3	0	10	16	9	5	20	7	9	4	4	12	4	5	4	10	19	19	10	-	-
3	6	5	0	18	16	0	7	18	8	6	10	2	9	0	10	16	9	5	20	7	4	4	4	12	19	5	4	10	17	19	10	-
4	18	6	5	0	18	16	0	7	2	3	6	10	16	9	0	10	7	9	5	20	12	4	4	4	10	19	5	4	7	17	19	10

Figure 11-13—Matrix permutation indices of structured parity check matrices

where:

L_{FEC} is the FEC codeword length

L_{MFB} is the length of the MAC Frame Body field in octets

R_{FEC} is the FEC rate

The FEC codeword length, L_{FEC} , is 2040 for the RS FEC specified in 11.2.2.6.1, 672 for the irregular LDPC specified in 11.2.2.6.2, and 1440 for the rate 14/15 LDPC specified in 11.2.2.6.3.

The encoded MAC Frame Body field shall be concatenated with stuff bits of length L_{STUFF} so that the resulting MAC Frame Body field is aligned on the subblock symbol boundary. The stuff bits shall be set to zero and then scrambled using the continuation of the scrambler sequence that scrambled the MAC Frame Body field in 11.2.2.10. The length of bits in the encoded MAC Frame Body field, L_{EBITS} , is given by the following:

**Table 11-17—Positions of ‘1’s in the first 15 columns of parity check matrix H
(codeword block length $L_{\text{FEC}} = 1440$)**

$h_{0,0}$ $h_{1,0}$ $h_{4,0}$
$h_{32,1}$ $h_{34,1}$ $h_{39,1}$
$h_{64,2}$ $h_{70,2}$ $h_{78,2}$
$h_{8,3}$ $h_{18,3}$ $h_{95,3}$
$h_{31,4}$ $h_{42,4}$ $h_{54,4}$
$h_{63,5}$ $h_{76,5}$ $h_{91,5}$
$h_{14,6}$ $h_{45,6}$ $h_{94,6}$
$h_{30,7}$ $h_{47,7}$ $h_{83,7}$
$h_{17,8}$ $h_{62,8}$ $h_{80,8}$
$h_{28,9}$ $h_{48,9}$ $h_{82,9}$
$h_{22,10}$ $h_{60,10}$ $h_{81,10}$
$h_{27,11}$ $h_{49,11}$ $h_{84,11}$
$h_{7,12}$ $h_{53,12}$ $h_{77,12}$
$h_{19,13}$ $h_{44,13}$ $h_{85,13}$
$h_{6,14}$ $h_{46,14}$ $h_{75,14}$

$$L_{\text{EBITS}} = 8 \times L_{\text{MFB}} + N_{\text{FEC}} \times (1 - R_{\text{FEC}}) \times L_{\text{FEC}}$$

The number of subblocks in the encoded MAC Frame Body field, $N_{\text{subblock-encMFB}}$, and the length of stuff bits, L_{STUFF} , are given by the following:

$$N_{\text{subblock-encMFB}} = \text{CEIL}(L_{\text{EBITS}}/L_{\text{CBPS}})$$

$$L_{\text{STUFF}} = N_{\text{subblock-encMFB}} \times L_{\text{CBPS}} - L_{\text{ebits}}$$

where L_{CBPS} is the number of coded bits per subblock as given in Table 11-18 for each MCS.

Table 11-18—Rate-dependent bits per symbol

MCS identifier	Coded bits per subblock, L_{CBPS} (pilot word length = 0)	Coded bits per subblock, L_{CBPS} (pilot word length = 64)
0	8	—
1	128	112
2	256	224
3	512	448
4	512	448
5	256	224

Table 11-18—Rate-dependent bits per symbol (continued)

MCS identifier	Coded bits per subblock, L_{CBPS} (pilot word length = 0)	Coded bits per subblock, L_{CBPS} (pilot word length = 64)
6	512	448
7	1024	896
8	1024	896
9	1024	896
10	1024	896
11	1024	896
12	1536	1344
13	2048	1792

For the stuff bits in the frame headers, the values are given in Table 11-12 and Table 11-13.

11.2.2.8 Code spreading

To increase robustness in header and MAC Frame Body field, Golay and pseudo random binary sequence (PRBS) codes by linear feedback shift register (LFSR) are applied for code spreading. The following two categories of spreading are defined:


- For spreading factor of 64, Golay sequences shall be used.
- For Class 1 MCSs with spreading factor of 2 and 4 and for headers with spreading factor of 2 and 6, the LFSR shall be used.

11.2.2.8.1 Golay sequences

In the base rate mode, the frame header and MAC Frame Body field shall be spread as shown in Figure 11-6. The Golay sequences for spreading factor 64 are given in Table 11-5.

11.2.2.8.2 PRBS generation with LFSR

For a spreading factor of length 2, 4, or 6, the data bits shall be spread with a PRBS generated using an LFSR, as shown in Figure 11-14. Since the output of the spreader is a factor of L_{SF} larger than the input, the input shall hold while the feedback and output clock.

**Figure 11-14—PRBS generation by LFSR**

The 15-bit seed value of the LFSR shall be $[x_{-1}, x_{-2}, \dots, x_{-15}] = [0101\ 0000\ 0011\ 111]$.


11.2.2.9 Unequal error protection

The UEP MCSs for SC PHY are shown in Table 11-19. UEP is an optional function. The data rate listed in the table is approximate and is rounded to three significant figures.

Table 11-19—UEP MCS for SC PHY

UEP MCS	Modulation	FEC		Data rate (Mb/s)	Supported UEP Type
		MSB	LSB		
0b00001	$\pi/2$ BPSK	RS (255,239)		1420	UEP Type 1 and UEP Type 2
0b00010		LDPC (672,336)		756	UEP Type 1 and UEP Type 2
0b00011		LDPC (672,504)		1130	UEP Type 1 and UEP Type 2
0b00100	$\pi/2$ QPSK	LDPC (672,336)		1510	UEP Type 1 and UEP Type 2
0b00101		LDPC (672,504)		2270	UEP Type 1 and UEP Type 2
0b00110		LDPC (672,588)		2650	UEP Type 1 and UEP Type 2
0b00111	$\pi/2$ QPSK	LDPC(672,336)	LDPC(672,504)	2040	UEP Type 3
0b01000		LDPC(672,504)	LDPC(672,588)	2650	UEP Type 3

In UEP Type 3, when the transmitter chooses one of the UEP MCS types, the transmitted data is divided into two groups in octet, MSB and LSB, as shown in Figure 11-15. The MSB and LSB groups are processed in parallel to adopt different LDPC coding rates. The outputs of the two LDPC encoders are multiplexed to generate a single bit stream for the symbol mapping.

**Figure 11-15—Parallel encoders for UEP**

The bit multiplexing/interleaving method depends on the MSB and LSB coding rates as illustrated in Figure 11-16. When LDPC(672,336) is used for MSB group and LDPC(672,504) is used for LSB group, the encoded bits shall be multiplexed every 10 bits with 6 bits as the encoded MSBs and 4 bits as the encoded LSBs. A1, A2, A3, A4, A5, A6 are used in an increasing order in time to label the 6 encoded MSBs from encoder LDPC(672,336), while B1, B2, B3, B4 are used in an increasing order in time to label the four encoded LSBs from encoder LDPC(672,504). The bit multiplexing/interleaving shall be performed such that the output pattern is A1, B1, A2, B2, A3, A4, B3, A5, B4, A6 with A1 being the earliest bit at the input of the symbol mapper while A6 being the latest, as illustrated in the upper part of Figure 11-16.

When LDPC(672,504) is used for MSB group and LDPC(672,588) is used for LSB group, the encoded bits shall be multiplexed every 13 bits with 7 bits as the encoded MSBs and 6 bits as the encoded LSBs. A1, A2,

A3, A4, A5, A6, A7 are used in an increasing order in time to label the seven encoded MSBs from encoder LDPC(672,504), while B1, B2, B3, B4, B5, B6 are used in an increasing order in time to label the six encoded LSBs from encoder LDPC(672,588). The bit multiplexing/interleaving shall be performed such that the output pattern is A1, B1, A2, B2, A3, B3, A4, B4, A5, B5, A6, B6, A7 with A1 being the earliest bit at the input of the symbol mapper while A7 being the latest, as illustrated in the lower part of Figure 11-16.




Figure 11-16—UEP bit multiplexing/interleaving

The effect of unequal error protection can also be obtained by using a skewed constellation, which is shown in Figure 11-29. A longer distance in x-axis than in y-axis is given between the two symbols in the skewed constellation, so that more energy is given to the MSB group. The mapping rules and the constellation are given in 11.3.2.6.

To apply the skewed constellation in the symbol mapping, both LDPC encoders in Figure 11-15 shall have the same coding rates among LDPC (672,336), LDPC (672,504), and LDPC (672,588), and the bit multiplexing/interleaving is 1:1 alternating, as illustrated in Figure 11-17.




Figure 11-17—UEP bit multiplexing/interleaving for skewed constellations

11.2.2.10 Scrambling

The frames shall be scrambled by modulo-2 addition of the data with the output of a PRBS generator, as illustrated in Figure 11-14 with $L_{SF} = 1$.

The scrambler shall be used for the MAC header, HCS, MAC subheader, HCS (for subheader), and MAC Frame Body field. The PHY preamble, PHY Header field, and RS field shall not be scrambled. The polynomial for the PRBS generator used by the scrambler shall be as follows:

$$g(D) = 1 + D^{14} + D^{15}$$

where D is a single bit delay element. The polynomial forms not only a maximal length sequence, but also is a primitive polynomial. By the given generator polynomial, the corresponding PRBS, is generated as follows:

$$x_n = x_{n-14} \oplus x_{n-15}, n = 0, 1, 2, \dots$$

The initialization sequence is defined by the following:

$$x_{init} = [x_{-1}x_{-2}x_{-3}x_{-4}x_{-5}x_{-6}x_{-7}x_{-8}x_{-9}x_{-10}x_{-11}x_{-12}x_{-13}x_{-14}x_{-15}]$$

The scrambled data bits, s_n , are obtained as follows:

$$s_n = b_n \oplus x_n$$

where b_n represents the unscrambled data bits. The side-stream de-scrambler at the receiver shall be initialized with the same initialization vector, x_{init} , used in the transmitter scrambler. The initialization vector is determined from the Scrambler Seed ID field contained in the PHY header of the received frame.

The 15-bit seed value chosen shall be computed from the Scrambler Seed ID field as follows:

$$[x_{-1}x_{-2}\dots x_{-15}] = [11010000101\ S1\ S2\ S3\ S4]$$

The seed identifier value is set to 0000 when the PHY is initialized and is incremented in a 4-bit rollover counter for each frame that is sent by the PHY. The value of the seed identifier that is used for the frame is sent in the PHY header.

For a Scrambler Seed ID field set to all zero, the first 16 bits should be as follows:

$$[x_0x_1\dots x_{15}] = [0001111000111010]$$

The 15-bit seed value is configured as follows. At the beginning of each PHY frame, the register is cleared, the seed value is loaded, and the first scrambler bit is calculated. The first bit of the data of the MAC header is modulo-2 added with the first scrambler bit, followed by the rest of the bits in the MAC header, MAC subheader, and MAC Frame Body field. The pilot word and pilot channel estimation sequences shall be excluded from the scrambling process.

11.2.3 SC PHY frame format

The SC PHY frame shall be formatted as illustrated in Figure 11-18.



Figure 11-18—SC PHY frame format

The Frame Header field for the PHY frame shall be formatted as illustrated in Figure 11-19.




Figure 11-19—Frame header format

The PHY preamble is described in 11.1.12.4. The MAC header is defined in 6.2 and the MAC subheader is defined in 6.2.8. The PHY header is defined in 11.2.3.2.1, and the HCS is defined in 11.2.3.2.2. The PHY Payload field consisting of the MAC Frame Body field, the PCES and stuff bits, is described in 11.2.3.3. The PCES is described in 11.2.3.4.2. The stuff bits are described in 11.2.2.7.

When transmitting a frame, the PHY preamble is sent first, followed by the base frame header, and then followed by the optional frame header, and finally the PHY Payload field.

11.2.3.1 PHY preamble

A PHY preamble shall be added prior to the frame header to aid receiver algorithms related to AGC setting, antenna diversity selection, timing acquisition, frequency offset estimation, frame synchronization, and channel estimation.

The PHY preamble shall be transmitted at the chip rate defined in Table 11-14.

In the CTAP, the CMS may be used. Transmission using the CMS in a CTAP shall be done using the frame consisting of the CMS preamble described in 11.1.12.4, the CMS header described in 11.2.2.2, and the CMS payload described in 11.2.2.1.

In the CAP, the CMS may be used. For this purpose, the SC data frame shall be specified as in 11.1.12, except that a different pattern, $[+1 +1 +1 -1 -1 +1 -1]$ spread by $\mathbf{b}_{128}$, shall be used in the SFD field.

Figure 11-20 shows the structure of the PHY preamble.




Figure 11-20—SC PHY preamble structure

11.2.3.1.1 Sync field

The Sync field is used for frame detection and uses a repetition of codes for a higher of robustness. The Sync field shall consist of 14 code repetitions of $\mathbf{a}_{128}$. Table 11-6 shows the sequence for $\mathbf{a}_{128}$ used for the Sync field.

11.2.3.1.2 SFD field

The SFD is used to establish frame timing as well as the header rate, either MR or HR. The SFD for the two header rates are as follows:

- The MR header shall use an SFD with $[+1 -1 +1 -1]$ spread by $\mathbf{a}_{128}$
- The HR header shall use an SFD with $[+1 +1 -1 -1]$ spread by $\mathbf{a}_{128}$

11.2.3.1.3 Channel estimation sequence (CES)

The CES field, used for channel estimation, shall consist of $[\mathbf{b}_{128} \mathbf{b}_{256} \mathbf{a}_{256} \mathbf{b}_{256} \mathbf{a}_{256}]$ where the right most sequence, $\mathbf{a}_{256}$, is first in time. The sequences $\mathbf{a}_{128}$, $\mathbf{a}_{256}$, and $\mathbf{b}_{256}$ are specified in 11.1.12.4.

11.2.3.2 Frame header

A frame header shall be added after the PHY preamble. The frame header conveys information in the PHY and MAC headers necessary for successfully decoding the frame. The frame header consists of a base frame header followed by an optional frame header. The construction of the frame header is shown in Figure 11-21.

The detailed process of the construction is as follows:




Figure 11-21—Frame header construction process

- a) Form the base frame header as follows:
 - 1) Construct the PHY header based on information provided by the MAC.
 - 2) Compute the HCS over the combined PHY and MAC headers.
 - 3) Append the HCS to the MAC header.
 - 4) Scramble the combined MAC header and HCS, as described in 11.2.2.10.
 - 5) Compute the RS Parity field by encoding the concatenation of the Header field, scrambled MAC Header and scrambled HCS into a shortened RS block code, as described in 11.2.2.6.1.
 - 6) Form the base frame header by concatenating the PHY Header field, scrambled MAC Header, field, scrambled HCS, RS Parity field, and scrambled stuff bits.
- b) Form the optional frame header as follows:
 - 1) Compute the HCS over the MAC subheader.
 - 2) Append the HCS to the MAC subheader.
 - 3) Scramble the combined MAC subheader and HCS, as described in 11.2.2.10.
 - 4) Compute the RS Parity field by encoding the concatenation of scrambled MAC Subheader field and scrambled HCS into a shortened RS block code, as defined in 11.2.2.6.1.
 - 5) Form the optional frame header by concatenating the scrambled MAC subheader, scrambled HCS, RS Parity field, and scrambled stuff bits.
- c) Form the frame header by concatenating the base frame header and optional frame header. The resulting frame header shall be modulated as shown in Figure 11-21.
- d) Spread the frame header, as described in 11.1.12.2.
- e) Map the frame header onto $\pi/2$ BPSK, as described in 11.2.2.5.1.
- f) Build subblocks from the resulting frame header, as described in 11.2.3.4.1.

The LFSR for the spreader is reset between the header and payload.

11.2.3.2.1 SC PHY header

The SC PHY header shall be formatted as illustrated in Figure 11-22.

Bits: b0–b3	b4	b5	b6–b10	b11–b30	b31–b32	b33	b34	b35–b36	b37	b38–b39
Scrambler seed ID	Aggregation	UEP	MCS	Frame length	Preamble type	Beam tracking	Low-latency mode	Pilot word length	PCES	Reserved

Figure 11-22—PHY header format for CMS and SC PHY

The Scrambler Seed ID field contains the scrambler seed identifier value, as defined in 11.2.2.10.

The Aggregation field shall be set to one if aggregation is used; it shall be set to zero otherwise.

The UEP field shall be set to one if UEP is used; it shall be set to zero if otherwise.

The MCS field shall be set according to the values in Table 11-20.

Table 11-20—Modulation and coding scheme

MCS	MCS identifier
0b00000	0
0b00001	1
0b00010	2
0b00011	3
0b00100	4
0b00101	5
0b00110	6
0b00111	7
0b01000	8
0b01001	9
0b01010	10
0b01011	11
0b01100	12
0b01101	13
0b01110–0b11111	Reserved

The Frame Length field shall be an unsigned integer equal to the number of octets in the MAC Frame Body field of a regular frame, excluding the FCS; it shall be set to zero for an aggregated frame.

The Preamble Type field indicates the type of the PHY preamble used in the next frame, as defined in Table 11-21.

Table 11-21—Preamble type field definition

Preamble type	Type of preamble used for next frame
0b00	CMS preamble
0b01	SC preamble
0b10–0b11	Reserved

The Beam Tracking field shall be set to one if training sequences for beam tracking are present following the current frame; it shall be set to zero otherwise.

The Low-latency Mode field shall be set to one if the frame is using the low-latency aggregation mode, and it shall be set to zero otherwise. If the Low-latency Mode field is set to one, then the Aggregation field shall also be set to one.

The Pilot Word Length field indicates the length of the pilot word used in the current frame and shall be encoded as defined Table 11-22.

Table 11-22—Pilot word length field definition

Pilot word length	Pilot word length in subblock
0b00	64 (mandatory)
0b01	8
0b10	0 (mandatory)
0b11	Reserved

The PCES field shall be set to one if the frame includes PCES; it shall be set to zero if otherwise.

11.2.3.2.2 Base header HCS

The combination of the PHY header and MAC header shall be protected with an ITU-T CRC-16 base HCS. The ITU-T CRC-16 is described in 10.2.9.

11.2.3.2.3 Base header FEC

The concatenation of the PHY header, scrambled MAC header and scrambled HCS shall use shortened systematic $RS(n+16,n)$, for the FEC, where n is the number of octets in the combined PHY header, MAC header and HCS. The 128 RS parity bits are appended after the scrambled HCS as shown in Figure 11-21.

11.2.3.2.4 MAC subheader HCS

The MAC subheader shall be protected with an ITU-T CRC-16 HCS. This shall be computed using the method specified in 11.2.3.2.2.

11.2.3.2.5 MAC subheader FEC

The scrambled MAC subheader and scrambled HCS shall be encoded using a shortened RS($n+16,n$), for the FEC, where n is the number of octets in the MAC subheader and HCS. The 128 RS parity bits are appended after the scrambled HCS, as shown in Figure 11-21.

NOTE—The length of the MAC subheader is different for secure and non-secure frames and hence the length of the shortened RS code will be different as well.

11.2.3.3 SC PHY Payload field

The SC PHY Payload field is the last component of the frame, and is constructed as shown in Figure 11-23.




Figure 11-23—SC PHY Payload field construction process

The PHY Payload field shall be constructed as follows:

- Scramble the MAC Frame Body field according to 11.2.2.10.
- Encode the scrambled MAC Frame Body field, as specified in 11.2.2.6.
- Add stuff bits to the encoded and scrambled MAC Frame Body field according to 11.2.2.7.
- Spread the encoded and scrambled MAC Frame Body field using the spreading code, as detailed in 11.1.12.2.
- Map the resulting MAC Frame Body field onto the appropriate constellation, as described in 11.2.2.5.
- Build subblocks from the resulting MAC Frame Body field according to 11.2.3.4.1.
- Insert PCES periodically, as described in 11.2.3.4.2.
- Apply a chip-level $\pi/2$ continuous rotation to the resulting MAC Frame Body field, as described in 11.2.2.5.

11.2.3.3.1 SC PHY Payload scrambling

The SC PHY payload shall use the scrambling process defined in 11.2.2.10.

11.2.3.3.2 Modulation

Modulation for the MAC Frame Body field is defined in 11.2.2.5.

11.2.3.3.3 FEC


FEC for the MAC Frame Body field is defined in 11.2.2.6.

11.2.3.3.4 Code spreading

Code spreading for the MAC Frame Body field is defined in 11.1.12.2.

11.2.3.4 Pilot word and PCES**11.2.3.4.1 Subblocks and pilot word**

Pilot words are used in SC PHY for timing tracking, compensation for clock drift and compensation for frequency offset error. Furthermore, pilot words act as a known cyclic prefix and enables frequency domain equalization if desired. In frequency domain equalization, the data is handled in the unit of subblocks. The building of the data blocks and subblocks is illustrated by Figure 11-24.

**Figure 11-24—Frame format with pilot word**

A block shall contain 64 subblocks with the exception of the last block (i.e., the N_b -th block). A subblock is formed by appending a pilot word to the data. The possible pilot word lengths are 0, 8, and 64. For pilot word lengths 0 and 64, the length of the data is $L_{DC} = L_{subblock} - L_{PW}$ symbols. For $L_{PW} = 8$, every subblock consists of 8 sub-subblocks each with $L_{PW} = 8$ and $L_{DC} = L_{subblock}/8 - L_{PW}$, thus giving the effective length of pilot word and data that is equivalent to that with $L_{PW} = 64$.

The Golay sequence for pilot word length 8 are given in Table 11-23 while for pilot word of length 64 they are shown in Table 11-5.

Table 11-23—Golay sequences of length 8

Sequence name	Sequence value
a₈	0xEB
b₈	0xD8

Even number blocks shall use pilot word sequences of type **a**. Odd number blocks shall use pilot word sequences of type **b**. Furthermore, an LFSR shall be used to change the polarity of pilot word from one block to another. The LFSR used shall be the same as the LFSR described in 11.2.2.8.2 with the same initial state, but shall be run at the appropriate rate, i.e., one LFSR output per block. The last subblock of the block shall be followed by a pilot word as well. The pilot word is modulated with $\pi/2$ BPSK.

11.2.3.4.2 Pilot channel estimation sequence (PCES)

Pilot channel estimation sequence (PCES) insertion is an optional feature that allows a DEV to periodically re-acquire the channel. To add the PCES, the scrambled, encoded, spread and modulated MAC Frame Body field is divided into data blocks. Each data block, as shown in Figure 11-24, shall be preceded by a PCES, with the exception of the first block. The duration and insertion interval of PCES are specified in Table 11-

15. The PCES field shall be the CES field in the preamble defined in 11.2.3.1.3. Similarly, the PCES is modulated with $\pi/2$ BPSK.

11.2.4 Transmitter specifications

11.2.4.1 Error vector magnitude

A compliant transmitter shall have EVM values of less than those given in Table 11-24 for the MCS classes listed.

Table 11-24—EVM for SC PHY MCS classes

MCS	EVM (dB)
Class 1	−7
Class 2	−14
Class 3	−21

11.2.4.2 Transmit center frequency tolerance

The transmitted center frequency tolerance shall be $\pm 25 \times 10^{-6}$ maximum.

11.2.4.3 Symbol rate

The SC PHY shall be capable of transmitting at the chip rate defined in Table 11-14 to within $\pm 25 \times 10^{-6}$. The MAC parameter, *pClockAccuracy*, shall be $\pm 25 \times 10^{-6}$. The transmit center frequency and symbol clock frequency shall be derived from the same reference oscillator (locked).

11.2.4.4 Transmit power-on and power-down ramp

The transmit power-on ramp is defined as the time it takes for the RF power emitted by the compliant DEV to rise from less than 10% to greater than 90% of the maximum power to be transmitted in the frame.

The transmit power-on ramp shall be less than 9.3 ns.

The transmit power-down ramp is defined as the time it takes for the RF power emitted by the compliant DEV to fall from greater than 90% to less than 10% of the maximum power to be transmitted in the frame.

The transmit power-down ramp shall be less than 9.3 ns.

The transmit power ramps shall be constructed such that the emissions conform to the unwanted emissions specification defined in 11.1.3.

11.2.5 Receiver specifications

11.2.5.1 Error rate criterion

The error rate criterion shall be a frame error rate (FER) of less than 8% with a frame payload length of 2048 octets. The error rate should be determined at the PHY SAP interface after any error correction methods (excluding retransmission) required in the proposed device has been applied. The measurement shall be performed in AWGN channel.

11.2.5.2 Receiver sensitivity

The receiver sensitivity is the minimum power level of the incoming signal, in dBm, present at the input of the receiver for which the error rate criterion in 11.2.5.1 is met. The error ratio shall be determined after any error correction has been applied. A compliant DEV that implements the SC PHY shall achieve at least the reference sensitivity listed in Table 11-25.

Table 11-25—Reference sensitivity levels for MCS

MCS Identifier	Receiver sensitivity
0	−70 dBm
1	−61 dBm
2	−58 dBm
3	−55 dBm
4	−59 dBm
5	−65 dBm
6	−62 dBm
7	−58 dBm
8	−56 dBm
9	−54 dBm
10	−53 dBm
11	−52 dBm
12	−50 dBm
13	−46 dBm

11.2.5.3 Receiver maximum input level

The receiver maximum input level is the maximum power level of the incoming signal, in dBm, present at the input of the receiver for which the error rate criterion in 11.2.5.1 is met. A compliant receiver shall have a receiver maximum input level of at least −10 dBm for each of the modulation formats that the DEV supports.

11.2.5.4 Receiver clear channel assessment performance

A compliant receiver provides CCA capability by performing energy detection in the received signal bandwidth. The start of a valid preamble sequence at a receive level equal to or greater than the minimum sensitivity for the CMS, as described in 11.2.5.2, shall cause CCA to indicate medium busy with a probability of >90% within $pCcaDetectTime$. The receiver CCA function shall in all circumstances report medium busy with any signal 20 dB above the minimum sensitivity for the CMS.

The CCA detection time shall be equal to $pCcaDetectTime$. The CCA shall be maintained as busy until the end of the frame.

11.2.6 PHY layer timing

The values for the PHY layer timing parameters are defined Table 11-26.

Table 11-26—PHY layer timing parameters

PHY parameter	Value	Subclause
<i>pMifsTime</i>	0.2 μ s, 0.5 μ s (default), 2.0 μ s	11.2.6.4
<i>pSifsTime</i>	0.2 μ s, 2.0 μ s, 2.5 μ s (default)	11.2.6.3
<i>pCcaDetectTime</i>	4 μ s	11.2.5.4
<i>pChannelSwitchTime</i>	100 μ s	11.2.6.5

11.2.6.1 Interframe space

A conformant implementation shall support the IFS parameters, as described in 7.4.1, given in Table 11-27.

Table 11-27—IFS parameters

MAC parameter	Corresponding PHY parameter	Definition
MIFS	<i>pMifsTime</i>	11.2.6.4
SIFS	<i>pSifsTime</i>	11.2.6.3
<i>pBackoffSlot</i>	<i>pSifsTime</i> + <i>pCcaDetectTime</i>	11.2.6.3
BIFS	<i>pSifsTime</i> + <i>pCcaDetectTime</i>	11.2.6.3, 11.2.5.4
RIFS	$2 \times pSifsTime + pCcaDetectTime$	11.2.6.3, 11.2.5.4

11.2.6.2 Receive-to-transmit turnaround time

The receive to transmit turnaround time shall be *pSifsTime*, including the power-up ramp specified in 11.2.4.4. The receive to transmit turnaround time shall be measured at the air interface from the trailing edge of the last symbol received until the first symbol of the PHY preamble is present at the air interface.

11.2.6.3 Transmit-to-receive turnaround-time

The transmit to receive turnaround time shall be less than *pSifsTime*, including the power-down ramp specified in 11.2.4.4.

11.2.6.4 Time between successive transmissions

The minimum time between successive transmissions shall be *pMifsTime*, including the power-up ramp specified in 11.2.4.4. The *pMifsTime* shall be measured at the air interface from the trailing edge of the last symbol transmitted until the first symbol of the PHY preamble is present at the air interface.

11.2.6.5 Channel switch

The channel switch time is defined as the time from the last valid bit is received at the antenna on one channel until the DEV is ready to transmit or receive on a new channel. The channel switch time shall be less than $pChannelSwitchTime$.

11.2.7 PHY management for SC PHY

The PHY PIB comprises the managed objects, attributes, actions, and notifications required to manage the SC PHY layer of a DEV.

11.2.7.1 Maximum frame size

The maximum frame length allowed, $pMaxFrameBodySize$, shall be 8388608 octets. This total includes the MAC subheader and the MAC Frame Body field, but not the PHY preamble, base header, (PHY header, MAC header, and HCS). The maximum frame length also does not include the stuff bits.

11.2.7.2 Maximum transfer unit size

The maximum size data frame passed from the upper layers, $pMaxTransferUnitSize$, shall be 8388576 octets. If security is enabled for the data connection, the upper layers should limit data frames to 8388576 octets minus the security overhead, as defined in 6.3.4.2, 6.2.8.1.2, or 6.2.8.2.2.

11.2.7.3 Minimum fragment size

The minimum fragment size, $pMinFragmentSize$, allowed with the SC PHY shall be 512 octets.

11.2.8 Optional OOK/DAMI modes

Besides the MCS classes in 11.2.2.1, optional low complexity and low power consumption MCSs, which are important for SC applications, may be employed within child piconets. As optional modes, OOK and DAMI may be employed for these applications. This subclause describes the OOK and DAMI MCSs should an implementer chose to implement these optional modes.

All PNC-capable OOK/DAMI DEVs shall be able to transmit and receive CMS signals.

If a PNC-capable OOK/DAMI DEV starts an independent piconet, it shall start the piconet in SC mode. OOK/DAMI PNC-capable DEVs may use OOK/DAMI signals in CTAs allocated for child piconet to support communication with non-PNC-capable OOK/DAMI DEVs. Details on child piconet creation and usage are described in G.2.

The summary of the MCS for OOK and DAMI is given in Table 11-11. All PNC-capable OOK/DAMI DEVs shall transmit CMS beacons and conduct CP in CMS, with $\pi/2$ BPSK and RS(255,239). These PNC-capable OOK/DAMI DEVs may create child piconet for respective non-PNC-capable DEVs by using respective MCS-formatted signals in Table 11-11. For OOK non-PNC-capable DEVs, OOK modulation and RS(255,239) shall be used. For DAMI non-PNC-capable devices, DAMI modulation and RS(255,239) shall be used.

11.2.8.1 OOK

The OOK modulation shall use variable amplitudes to represent the data. As shown in Figure 11-25(a), OOK shall be represented by two points in the constellation map. The simplest form of OOK represents a binary '1' with the presence of the signal, and a binary '0' with the absence of it. The normalization factor, K_{MOD} , shall be $\sqrt{2}$.

IEEE Std 802.15.3-2016
IEEE Standard for High Data Rate Wireless Multi-Media Networks




Figure 11-25—Constellation diagram for (a) OOK and (b) DAMI

11.2.8.2 DAMI

DAMI modulation is shown in Figure 11-25(b). The transmitted RF signal for a DAMI system shall be a single-sideband (SSB) modulated signal accompanied by two low-power pilot tones, as described in G.3.

11.2.8.3 FEC

The forward error correction scheme for OOK and DAMI shall be RS(255,239), as described in 11.2.2.6.

11.2.8.4 OOK/DAMI spreading

The spreading scheme for OOK with spreading factor of 2 shall use a simple bit repetition in which each bit shall be repeated twice.

11.3 High Speed Interface mode of mmWave PHY

The High Speed Interface mode of mmWave PHY (HSI PHY) is designed for devices with low-latency, bidirectional high-speed data and uses orthogonal frequency domain multiplexing (OFDM). HSI PHY supports a variety of modulation and coding schemes (MCSs) using different frequency-domain spreading factors, modulations, and LDPC block codes.

11.3.1 General operating specifications

11.3.1.1 Operating frequency bands

The set of operating channels is defined in Table 11-5. A compliant IEEE 802.15.3c implementation that implements the HSI PHY shall support at least CHNL_ID 2 or CHNL_ID 3.

11.3.2 HSI PHY modulation, forward error correction, and spreading

11.3.2.1 MCS-dependent parameters

The HSI PHY MCS-dependent parameters are listed in Table 11-28. For the FEC rates of 1/2, 5/8, 3/4, and 7/8, LDPC(672,336), LDPC(672,504), LDPC(672,420), and LDPC(672,588) codes are used, respectively. For the UEP MCSs (MCS index 8-11), different coding schemes are applied to the MSB octets and LSB octets. The data rates specified in Table 11-28 assume a cyclic prefix length of 64 chips.

Table 11-28—HSI PHY MCS-dependent parameters

MCS index	Data rate (Mb/s)	Modulation scheme	Spreading factor (L_f)	Coding mode	FEC rate (R_{FEC})	
					MSB 8b	LSB 8b
0	32.1	QPSK	48	EEP	1/2	
1	1540	QPSK	1		1/2	
2	2310	QPSK	1		3/4	
3	2695	QPSK	1		7/8	
4	3080	16-QAM	1		1/2	
5	4620	16-QAM	1		3/4	
6	5390	16-QAM	1		7/8	
7	5775	64-QAM	1		5/8	
8	1925	QPSK	1	UEP	1/2	3/4
9	2503	QPSK	1		3/4	7/8
10	3850	16-QAM	1		1/2	3/4
11	5005	16-QAM	1		3/4	7/8

The number of spread, coded, and data information bits per OFDM symbol are listed in Table 11-29 for the various MCSs.

Table 11-29—HSI PHY MCS-dependent bits per OFDM symbol

MCS index	Spread and coded bits/OFDM symbol (N_{SCBPOS})	Coded bits/OFDM symbol (N_{CBPOS})	Information bits/OFDM symbol (N_{IBPOS})	
			MSB 8b	LSB 8b
0	672	14	7	
1	672	672	336	
2	672	672	504	
3	672	672	588	
4	1344	1344	672	
5	1344	1344	1008	
6	1344	1344	1176	
7	2016	2016	1260	
8	672	672	336	504
9	672	672	504	588
10	1344	1344	672	1008
11	1344	1344	1008	1176

The HSI PHY frame header rate-dependent modulation parameters are listed in Table 11-30. The frame header is QPSK modulated and encoded with shortened LDPC(672,336) code.

Table 11-30—HSI PHY frame header rate-dependent parameters

Header rate type	Header rate (Mb/s)		Spreading factor	Coded bits /OFDM symbol	Number of occupied OFDM symbols		Number of stuff bits ^a	
	Main	Optional			Main	Optional	Main	Optional
MCS 0	16.8	29.6	48	14	35	47	10	2
MCS 1–11	587	1363	1	672	1	1	192	16

^aStuff bits are inserted after the LDPC encoding and before constellation mapping.

The base rate shall be MCS index 0 for HSI PHY and the corresponding header rates shall be as indicated in Table 11-30.

A DEV that supports the HSI PHY shall support HSI MCS index 1 and either CMS or HSI MCS index 0, as defined in 11.1.12.

11.3.2.2 HSI PHY timing-related parameters

Table 11-31 lists the timing-related parameters.

Table 11-31—Timing-related parameters

Parameters	Description	Value	Formula
f_s	Reference sampling rate/chip rate	2640 MHz	
T_C	Sample/chip duration	~0.38 ns	$1/f_s$
N_{sc}	Number of subcarriers/FFT size	512	
N_{dsc}	Number of data subcarriers	336	
N_P	Number of pilot subcarriers	16	
N_G	Number of guard subcarriers	141	
N_{DC}	Number of DC subcarriers	3	
N_R	Number of reserved subcarriers	16	
N_U	Number of used subcarriers	352	$N_{dsc} + N_P$
N_{GI}	Guard interval length in samples	64	
Δf_{sc}	Subcarrier frequency spacing	5.15625 MHz	f_s/N_{sc}
BW	Nominal used bandwidth	1815 MHz	$N_U \times \Delta f_{sc}$
T_{FFT}	IFFT and FFT period	~193.94 ns	$1/\Delta f_{sc}$
T_{GI}	Guard interval duration	~24.24 ns	$N_{GI} \times T_C$
T_S	OFDM Symbol duration	~218.18 ns	$T_{FFT} + T_{GI}$

Table 11-31—Timing-related parameters (continued)

Parameters	Description	Value	Formula
F_S	OFDM Symbol rate	~4.583 MHz	$1/T_S$
N_{CPS}	Number of samples per OFDM symbol	576	$N_{sc} + N_{GI}$

11.3.2.3 HSI PHY frame-related parameters


The frame-related parameters are listed in Table 11-32.

Table 11-32—OFDM frame-related parameters

Parameter	Description	Value	
N_{pre}	Number of symbols in the PHY preamble (A preamble symbol is 512 chips long)	Long Preamble	16
		Short Preamble	6.75
T_{pre}	Duration of the PHY preamble	Long Preamble	~3.15 μ s
		Short Preamble	~1.31 μ s
T_{HDR}	Duration of the header	Main header only for MCS 0	~7.64 μ s
		Main header only for MCS 1–11	~0.22 μ s
		Main header and optional header for MCS 0	~17.89 μ s
		Main header and optional header for MCS 1–11	~0.44 μ s
N_{OSMF}	Number of OFDM symbols in the MAC Frame Body field	variable	
T_{OSMF}	Duration of the MAC Frame Body field	$N_{OSMF} \times T_S$	
N_{frame}	Number of OFDM symbols in the frame	$N_{pre} + N_{HDR} + N_{OSMF}$	
T_{frame}	Duration of the frame	$T_{pre} + T_{HDR} + T_{OSMF}$	

11.3.2.4 HSI PHY FEC

The HSI PHY FEC process is illustrated in Figure 11-26. In the figure, the bit interleaver block is drawn with a dashed line because it is an optional part of the HSI PHY FEC process.

**Figure 11-26—FEC process for the HSI PHY**

11.3.2.4.1 LDPC block code

The supported LDPC block FEC rates, information block lengths, L_{INF} , and codeword block lengths, L_{FEC} , are described in Table 11-33.

Table 11-33—LDPC parameters

R_{FEC}	L_{INF} (bits)	L_{FEC} (bits)
1/2	336	672
5/8	420	
3/4	504	
7/8	588	

The LDPC encoder with rates 1/2, 3/4, and 7/8 is described in 11.2.2.6.2. For rate 5/8, the matrix permutation indices of parity-check matrix is given in Figure 11-27.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
1	0	-	-	5	-	18	16	-	-	-	3	6	10	-	-	0	-	7	-	5	-	-	4	4	-	10	-	5	-	-	-	-	
2	-	-	6	-	7	-	-	-	-	2	-	-	-	-	9	-	20	-	-	-	4	-	-	-	-	-	19	-	-	-	-	-	
3	-	18	-	-	-	-	-	0	10	-	-	-	-	16	-	-	-	-	9	-	-	12	-	-	4	-	-	-	-	-	-	17	-
4	5	0	-	-	-	-	18	16	6	-	-	3	0	10	-	-	5	-	7	-	4	-	-	-	-	-	-	-	-	-	-	-	
5	-	-	-	6	-	7	-	-	-	-	2	-	-	-	-	9	-	20	-	-	-	4	-	-	-	-	-	-	-	-	-	-	
6	-	-	18	-	0	-	-	-	-	10	-	-	-	-	16	-	-	-	-	9	-	-	12	-	-	-	-	-	-	-	-	-	
7	-	5	0	-	16	-	-	18	3	6	-	-	-	0	10	-	-	5	-	7	4	4	-	-	-	5	-	-	-	-	-	-	
8	6	-	-	-	-	-	7	-	-	-	-	2	9	-	-	-	-	-	20	-	-	-	-	4	-	19	-	-	-	-	-	-	
9	-	-	-	18	-	0	-	-	-	-	10	-	-	-	-	16	9	-	-	-	-	-	-	12	-	-	-	-	-	-	-	-	
10	-	-	5	0	18	16	-	-	-	3	6	-	-	-	0	10	7	-	5	-	-	4	4	-	10	-	5	-	7	-	-	-	
11	-	6	-	-	-	-	-	7	2	-	-	-	-	-	9	-	-	-	-	20	-	-	-	4	-	19	-	-	-	-	-	10	
12	18	-	-	-	-	-	0	-	-	-	-	10	16	-	-	-	-	9	-	-	12	-	-	-	-	-	-	4	-	17	-	-	

Figure 11-27—Matrix permutation indices of parity check matrix of rate-5/8 LDPC code**11.3.2.4.2 EEP data multiplexer**

For the EEP MCSs, as defined in Table 11-28, the two LDPC encoders use the same rate, and the outputs of the LDPC encoders shall be multiplexed to form a single data stream. Let a_n and b_n be the outputs bits of the MSB and LSB encoders, respectively. The multiplexer output shall be $a_0, b_0, a_1, b_1, \dots$

11.3.2.4.3 UEP data multiplexer

For UEP Type 3 MCSs, as defined in Table 11-28, the two LDPC encoders have different rates. The method used to multiplex the encoded bits is dependent on the LDPC MSB and LSB rates.

For UEP MCSs with MSB encoder rate 1/2 and LSB encoder rate 3/4, the encoded bits shall be multiplexed every 5 bits. During the length 5 multiplexing cycle, a group multiplexer shall be used with group size three for the MSB encoder and group size two for the LSB encoder. a_0, a_1, a_2 are used to label the three encoded bits (in increasing order in time) from the MSB encoder, and b_0, b_1 are used to label the two encoded bits from the LSB encoder. The multiplexer output shall be a_0, b_0, a_1, b_1, a_2 .

For UEP MCSs with MSB encoder rate 3/4 and LSB encoder rate 7/8, the encoded bits shall be multiplexed every 13 bits. During the length 13 multiplexing cycle, a group multiplexer shall be used with group size 7 for the MSB encoder and group size 6 for the LSB encoder. $a_0, a_1, a_2, \dots, a_6$ are used to label the 7 encoded bits (in increasing order in time) from the MSB encoder, and $b_0, b_1, b_2, \dots, b_5$ are used to label the 6 encoded bits from the LSB encoder. The multiplexer output shall be $a_0, b_0, a_1, b_1, \dots, a_5, b_5, a_6$.

11.3.2.4.4 Bit interleaver

After the data multiplexer, the bits shall be interleaved by a block interleaver if the Bit Interleaver field in the PHY header is set to one. The interleaving is performed upon encoded bits included within an interleaving depth covering 4 LDPC codewords, i.e., more than 2688 bits.

The block interleaving process is performed using a permutation rule $L(k)$. That is, the k^{th} output, written to location k in the output vector, is read from location $L(k)$ in the input vector.

The block interleaving algorithm $L(k) = I_{p,q}^j(k)$ is described by four parameters: the block size $K_B = 2688$, an integer parameter p setting the partition size, an integer parameter q and the iteration j governing the interleaving spreading. The relationship between the block of K_B coded bits, $a_0, a_1, \dots, a_{K-1}$, and the block of K_B interleaved bits, $b_0, b_1, \dots, b_{K-1}$, is given by

$$b(k) = a[I_{p,q}^j(k)]$$

To realize the interleaver stage, it is recommended to implement a lookup table that contains the interleaving rule.

The interleaving rule is based on an iterative structure, as illustrated in Figure 11-28, in order to increase the scalability of the interleaver.




Figure 11-28—Turbo-based bit interleaver structure

The interleaving rules for the 1st and the j^{th} iteration are defined as

$$I^1_{p,q}(k) = \text{mod}[K_B - p + k + q \times p \times \text{mod}(-k - p \times k, K_B), K_B]$$

$$I^j_{p,q}(k) = \text{mod}[K_B - p + k + q \times p \times \text{mod}(-k - p \times I^{j-1}_{p,q}(k), K_B), K_B]$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

The interleaver parameters are selected in order to optimize the interleaving spreading between successive samples. The interleaving spreading $\Delta L(s)$ is defined as the minimum distance between interleaved bits separated by a distance $s - 1$, and is expressed as

$$\Delta L(s) = \min_k [|I^j_{p,q}(k + s) - I^j_{p,q}(k)|]$$

The interleaving spreading is calculated in an algebraic way and allows the selection of interleaving parameters $\{p, q, j\}$ for each interleaving block size K_B . The binary interleaving parameters shall be $p = 24$, $q = 2$, $j = 1$, and $\Delta L(s = 1, 2, 4, 6) = \{1199, 290, 580, 870\}$.

11.3.2.5 Stuff bits

Stuff bits shall be appended in the MAC Frame Body field after scrambling, encoding and interleaving so that the resulting MAC Frame Body field is aligned with the boundaries of an OFDM symbol. The stuff bits shall be set to zero and then scrambled using the continuation of the scrambler sequence that scrambled the MAC Frame Body field. To calculate the number of stuff bits, L_{STUFF} , that shall be inserted, first the number of codewords should be calculated as follows:

$$N_{\text{FEC}} = \text{CEIL}[(L_{\text{MFB}} \times 8)/(L_{\text{FEC}} \times R_{\text{FEC}})]$$

where L_{MFB} is the length of the uncoded MAC Frame Body field in octets. FEC codeword length is denoted by L_{FEC} and R_{FEC} is the coding rate. Number of the encoded bits L_{EBITS} in the MAC Frame Body field is given by the following

$$L_{\text{EBITS}} = L_{\text{MFB}} \times 8 + N_{\text{FEC}} \times (1 - R_{\text{FEC}}) \times L_{\text{FEC}}$$

Number of OFDM symbols N_{OSMF} in MAC Frame Body field is equal to the following:

$$N_{\text{OSMF}} = \text{CEIL}(N_{\text{EBITS}}/N_{\text{SCBPOS}})$$

where N_{SCBPOS} is the number of spread and coded bits per OFDM symbol. N_{SCBPOS} for different MCSs is given in Table 11-29. The number of stuff bits for a MAC Frame Body field is equal to the following:

$$L_{\text{PAD}} = N_{\text{OSMF}} \times N_{\text{SCBPOS}} - N_{\text{EBITS}}$$

Stuff bits shall be added to main and optional headers as well. The number of stuff bits for each header is given in Table 11-30.

11.3.2.6 Constellation mapping

The coded and interleaved binary serial input data, b_i , where $i = 0, 1, 2, \dots$, shall be modulated using QPSK, 16-QAM or 64-QAM modulation, depending on the MCS requested. The binary serial stream shall be divided into groups of N_{BPSC} (2, 4, or 6) bits and converted into complex numbers representing QPSK, 16-QAM or 64-QAM constellation points. The conversion shall be performed according to Gray-coded constellation mappings, illustrated in Figure 11-29 with the input bit, b_0 , being the earliest in the stream. The output values, a_k , where $k = 0, 1, 2, \dots$, are formed by multiplying the resulting value $(I_k + jQ_k)$ by a normalization factor, K_{MOD} .

$$a_k = (I_k + jQ_k) \times K_{\text{MOD}}$$

The normalization factor, K_{MOD} , depends on the modulation, as prescribed in Table 11-34. The purpose of the normalization factor is to achieve the same average power for all mappings. In practical implementations, an approximate value of the normalization factor can be used, as long as the device conforms to the modulation accuracy requirements described in 11.3.4.1.

Table 11-34—Modulation-dependent normalization factor

Modulation	K_{MOD}
QPSK	$1/\sqrt{1+d^2}$
16-QAM	$1/\sqrt{5(1+d^2)}$
64-QAM	$1/\sqrt{21(1+d^2)}$

An optional skewed constellation is also specified in Figure 11-29. A parameter d is introduced to distinguish between normal and skewed constellation. Its value is given by the following:

$$d = \begin{cases} 1 & \text{normal constellation} \\ 1.25 & \text{skewed constellation} \end{cases}$$

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-3:2017




Figure 11-29—QPSK, 16-QAM and 64-QAM normal and skewed constellation bit encoding

11.3.2.7 HSI spreader

The spreading factors in relationship to the data rate in the PHY Payload field and Frame Header field are specified in Table 11-28 and Table 11-30, respectively. The spreading rules shall be the same for all three fields.

11.3.2.7.1 Spreader for spreading factor of 1

For a spreading factor of 1, the modulated QPSK and QAM complex values a_k , where $k = 0, 1, 2, \dots$ at the output of the constellation mapper shall be grouped into sets of 336 complex numbers. Each group shall be assigned to an OFDM symbol. This is denoted by writing the complex number $b_{k,n}$.

$$b_{k,n} = a_{k+n \times 336}, \text{ for } k = 0, 1, \dots, 335, n = 0, 1, 2, \dots$$

where n is the OFDM symbol number. Each group shall be passed to the tone interleaver before being modulated by the OFDM modulator into an OFDM symbol.

11.3.2.7.2 Spreader for spreading factor of 48

For a spreading factor of 48, the modulated QPSK complex values $a(k)$, where $k = 0, 1, 2, \dots$ at the output of the constellation mapper shall be grouped into sets of seven complex numbers. This is denoted by writing the complex number $a_{k,n}$.

$$a_{k,n} = a_{k+n \times 28}, \text{ for } k = 0:6, n = 0, 1, 2, \dots$$

where n is the group number. Each group shall be spread by a factor of 48 to generate a block of 336 complex numbers given by the following:

$$b_{k,n} = q_{\text{floor}(k/28)} a_{k,n}, \text{ for } k = 0:167, \text{ and}$$

$$b_{k,n} = b_{335-k,n}^*, \text{ for } k = 168:335$$

where q is a length 12 sequence given by the following:

$$q = [+1 +j -1 +j +j +1 -1 +j -j +j -1 -j -1 +1 +1 +1 +j -j -1 -1 -1 +j -j +j]$$

and the function floor() rounds its argument to the nearest integer toward minus infinity. Each spread group shall be passed to the tone interleaver before being modulated into an OFDM symbol.

11.3.2.8 Tone interleaver

All bits shall be interleaved by a block interleaver with a block size corresponding to the size of FFT in a single OFDM symbol, N_{sc} . The interleaver is used so that the adjacent data symbols are mapped onto separate subcarriers.

At the transmitter side, the interleaver permutation shall be defined as follows: Let k be the index of the tones (including data tones, pilot tones, DC tones, and null tones) before permutation ranging between 0 and $N_{sc}-1$. Let i be the index of the interleaved tones over the same range (including data tones, pilot tones, DC tones and null tones) after permutation. Let

$$k = \sum_{j=0}^{L-1} a_j 2^j$$

where $L = \log_2(N_{sc}) - 1$, with $[a_L, \dots, a_0]$ being the binary representation of integer k . Then the binary representation of integer i can be written as $[a_0, \dots, a_L]$.

$$i = \sum_{j=0}^L a_j 2^{L-j}$$

DC, null, and pilot tones shall be inserted in the bit-reversal position before the tone interleaver. This ensures that after permutation, the DC, null, and pilot tones appear in the pre-specified positions.

11.3.2.9 HSI PHY OFDM modulator

The stream of interleaved complex numbers is divided into groups of N_{dsc} data complex numbers. This is denoted by writing the complex number $d_{k,n}$, which corresponds to data subcarrier k of OFDM symbol n .

$$d_{k,n} = d_{k+n \times N_D}, \quad k = 0, 1, 2, \dots, N_D - 1, \quad n = 0, 1, \dots, N_{hdr} + N_{OSMF} - 1$$

where $N_{hdr} + N_{OSMF}$ is the number of OFDM symbols occupied by the header and PHY Payload field.

The discrete-time signal during the n^{th} OFDM symbol is given by the following:

$$s_{k,n} = \frac{1}{\sqrt{N_{sc}}} \left[\sum_{m=0}^{N_D-1} d_{m,n} e^{j2\pi \frac{k \times M_D(m)}{N_{sc}}} + x_n \sum_{m=0}^{N_P-1} p_{m,n} e^{j2\pi \frac{k \times M_P(m)}{N_{sc}}} + \sum_{m=0}^{N_G-1} g_{m,n} e^{j2\pi \frac{k \times M_G(m)}{N_{sc}}} \right]$$

where

$$k \in [0:N_{FFT} - 1]$$

N_{dsc} is the number of data subcarriers

N_P is the number of pilot subcarriers

N_R is the number of reserved subcarriers

N_G is the number of guard subcarriers

N_{sc} is the number of total subcarriers

$d_{m,n}$, $p_{m,n}$, and $g_{m,n}$ are the complex numbers placed on the m^{th} data, pilot, and guard subcarriers of the n^{th} OFDM symbol, respectively

The functions $M_{dsc}(m)$, $M_P(m)$, and $M_G(m)$ define a mapping between the indices $[0:N_{dsc} - 1]$, $[0:N_P - 1]$, $[0:N_R - 1]$, and $[0:N_G - 1]$ into the logical frequency offset index $[-N_{sc}/2:N_{sc}/2 - 1]$. The definition for the mapping functions are given by the following:

$$M_{dsc}(m) = \begin{cases} m - 177 + \text{round}(m/21) & 0 \leq m \leq 167 \\ m - 174 + \text{round}[(m+1)/21] & 168 \leq m \leq 335 \end{cases}$$

$$M_P(m) = \begin{cases} -166 + m \times 22 & 0 \leq m \leq 7 \\ 12 + (m - 8 \times 22) & 8 \leq m \leq 15 \end{cases}$$

$$M_G(m) = \begin{cases} -185 + m & 0 \leq m \leq 7 \\ 170 + m & 8 \leq m \leq 15 \end{cases}$$

where the function $\text{round}()$, rounds the input argument to the nearest integer.

The mapping of data and pilot subcarriers within an OFDM symbol is illustrated in Figure 11-30. The mapping is further summarized in Table 11-35. As shown in Figure 11-30, there are 16 groups of subcarriers where each group is constituted of 21 data subcarriers and one pilot subcarrier.




Figure 11-30—Subcarrier frequency allocation

Table 11-35—Subcarrier frequency allocation

Subcarriers type	Number of subcarriers	Logical subcarriers indexes
Null subcarriers	141	$[-256: -186] \cup [186:255]$
DC subcarriers	3	$-1, 0, 1$
Pilot subcarriers	16	$[-166:22: -12] \cup [12:22:166]$
Guard subcarriers	16	$[-185: -178] \cup [178:185]$
Data subcarriers	336	All others

The common way to implement the inverse Fourier transform is by an inverse Fast Fourier Transform (IFFT) algorithm. If, for example, a 512-point IFFT is used, the logical subcarriers 2 to 185 are mapped to the same numbered IFFT inputs, while the logical frequency subcarriers -185 to -2 are copied into IFFT inputs 327 to 510. The rest of the inputs, 186 to 326 and the 0 (DC) input, 1, and 511 are set to zero. The subcarriers 0, 1, and 511 are set to zero to avoid difficulties in D/A and A/D converter offsets and carrier feed through in the RF system. This mapping is illustrated in Figure 11-31. After performing an IFFT, the output is cyclically extended to the desired length.

11.3.2.9.1 Pilot subcarriers

In all OFDM symbols following the frame preamble, sixteen of the subcarriers shall be dedicated to pilot signals in order to allow for coherent detection and to provide robustness against frequency offsets and phase noise. These pilot signals shall be placed into logical frequency subcarriers $-166, -144, -122, -100, -78, -56, -34, -12, 12, 34, 56, 78, 100, 122, 144$, and 166 . The information for the m^{th} pilot subcarrier of the n^{th} OFDM symbol shall be defined as follows:

$$p_m = \begin{cases} (1+j)/\sqrt{2} & \text{for } m = 0, 3, 5, 7, 9, 13, 15 \\ (1-j)/\sqrt{2} & \text{for } m = 1, 2, 4, 6, 8, 10, 11, 12, 14 \end{cases}$$

The polarity of the pilot subcarriers is controlled by the sequence, x_n , generated by the linear feedback shift register described in 11.3.2.11.

The 15-bit seed value chosen is $x_{-1:-15} = [1101\ 0000\ 1010\ 000]$. The first 16 output values should be $x_{0:15} = [0001\ 1110\ 0011\ 1010]$.




Figure 11-31—Input and output relationship of the IFFT

11.3.2.9.2 Guard subcarriers

In all OFDM symbols following the frame preamble, there shall be 16 guard subcarriers, 8 on each edge of the occupied frequency band, at logical frequency subcarriers $-185, -184, \dots, -178$ and $178, 179, \dots, 185$. The data on these subcarriers shall be left to the implementer. Individual implementations may exploit these guard subcarriers for various purposes, including relaxing the specifications on analog transmit and analog receive filters, and possibly peak to average power ratio reduction.

11.3.2.10 PCES insertion

The pilot channel estimation sequence (PCES) symbols are used for channel re-acquisition or tracking. This is an optional field which content is identical to the CES of the short preamble prepended by c_{128} .

If PCES is used then the PCES symbols shall be inserted periodically in the PHY payload field as shown in Figure 11-32. The value of the exact period N_{PCES} shall be 96 OFDM symbols.



Figure 11-32—PCES positions in the PHY Payload field

11.3.2.11 HSI PHY scrambling

HSI PHY uses same scrambling method of SC PHY, which is explained in 11.2.2.10.

11.3.3 HSI PHY frame format

The HSI PHY frame shall be formatted as illustrated in Figure 11-33.



Figure 11-33—HSI PHY frame format

The PHY preamble is defined in 11.3.3.1.

The Frame Header field for the HSI PHY frame shall be formatted as illustrated in Figure 11-34. The goal of the frame header is to convey necessary information about both the PHY and the MAC to aid decoding of the PHY Payload at the receiver.




Figure 11-34—HSI frame header format

The construction of main frame header and optional frame header is explained in 11.3.3.2. PHY header is defined in 11.3.3.3. MAC header is defined in 6.2, and the MAC subheader is defined in 6.2.8.

The PHY Payload field is formed by adding any necessary stuff bits to the MAC Frame Body field, as defined in 6.2, to align the data stream on the boundary of an OFDM symbol. The optional PCES may be added to the PHY Payload field.

The frame header is sent at the appropriate header rate, as defined in Table 11-30. The PHY Payload field is transmitted at the desired data rate, as defined in Table 11-28.

11.3.3.1 PHY preamble

A PHY preamble shall be added prior to the frame header to aid receiver algorithms related to frame detection, AGC setting, timing acquisition, frequency recovery, frame synchronization, and channel estimation.

The PHY preamble shall be transmitted at the rate equal to the subcarrier frequency spacing, as defined in Table 11-31, i.e., $R_S = \Delta f_{sc}$. A preamble symbol is defined as a sequence of length 512 chips that corresponds to the FFT length.

Two preambles are defined for the HSI PHY mode: the long preamble and optional short preamble. The Preamble Type field in the PHY header, as described in 11.3.3.3, indicates the type of preamble that shall be used in the next frame. The long preamble is used with MCS index 0 and is the same structure as defined for the CMS in Figure 11-7 using the chip rate of the HSI PHY mode. The short preamble is the same structure as defined for the SC PHY mode in Figure 11-20 using the chip rate of the HSI PHY mode and the Sync defined in 11.2.3.1.1, the SFD defined in 11.2.3.1.2, and the CES defined in 11.2.3.1.3. The durations of the parts of the preambles are provided in Table 11-31.

11.3.3.2 Frame header

A frame header shall be added after the PHY preamble. It conveys information in the PHY and the MAC headers necessary for a successful decoding of the frame. The frame header is constructed from a main frame header followed by an optional frame header.

The frame header shall be constructed as follows:

- a) Form the main frame header as follows:
 - 1) Construct the PHY header based on information provided by the MAC.
 - 2) Compute the HCS over the combined PHY and MAC headers, as described in 11.3.3.4.1.
 - 3) Append the HCS to the MAC header.
 - 4) Scramble the combined MAC header and HCS, as described in 11.3.2.5.
 - 5) Form the main frame header by concatenating the PHY header, scrambled MAC header, scrambled HCS.
 - 6) The resulting main frame header shall be modulated as follows:
 - i) The main frame header is further encoded with shortened LDPC encoder of rate 1/2, as described in 11.3.2.4.1.
 - ii) The resulting main frame header has stuff bits added, as indicated in Table 11-30.
 - iii) The resulting main frame header is then mapped to a QPSK constellation, as described in 11.3.2.6.
 - iv) The resulting complex values are passed to the OFDM modulator, as described in 11.3.2.9. The cyclic prefix during the header is fixed to 128 chips.
- b) Form the optional frame header as follows:
 - 1) Compute the HCS over the MAC subheader, as described in 11.3.3.4.1.
 - 2) Append the HCS to the MAC subheader.
 - 3) Scramble the combined MAC subheader and HCS, as described in 11.3.2.5.
 - 4) Form the optional frame header by concatenating the scrambled MAC subheader, scrambled HCS.
 - 5) The resulting optional frame header shall be modulated as follows:
 - i) The optional frame header is further encoded with shortened LDPC encoder of rate 1/2, as described in 11.3.2.4.1.
 - ii) The resulting optional frame header has stuff bits added, as indicated in Table 11-30.
 - iii) The resulting optional frame header is then mapped to a QPSK constellation, as described in 11.3.2.6.
 - iv) The resulting complex values are passed to the OFDM modulator, as described in 11.3.2.9. The cyclic prefix during the header is fixed to 128 chips.
- c) Form the frame header by concatenating the main frame header and optional frame header.

The construction and modulation of the main frame header for HSI PHY is illustrated in Figure 11-35.




Figure 11-35—HSI main frame header encoding process

The construction and modulation of the optional header for HSI is illustrated in Figure 11-36.




Figure 11-36—HSI optional frame header encoding process

11.3.3.3 HSI PHY header

The HSI PHY header field shall be formatted as illustrated in Figure 11-37.

Bits: b0–b3	b4	b5	b6–b10	b11–b30	b31–b32
Scrambler ID	Aggregation	UEP	MCS	Frame Length	Preamble Type

b33	b34	b35	b36	b37	b38–b47
Beam Tracking	Low-latency Mode	Bit Interleaver	PCES Mode	Skewed Constellation	Reserved

Figure 11-37—HSI PHY header format

The scrambling process is described in 11.3.2.5. The MAC shall set the bits defined in 11.3.2.11 according to the scrambler seed identifier value.

The Aggregation field shall be set to one if aggregation is used; it shall be set to zero otherwise.

The UEP field shall be set to one if any type of UEP is used; it shall be set to zero otherwise.

The MCS field shall be set to the MCS used in the MAC Frame Body field, according to the values in Table 11-36. MCS field also identifies which header rate should be used, as defined in Table 11-30. The MCS identifiers in relation to the data rate, modulation, spreading code length, FEC, pilot word, and burst length are provided in Table 11-28.

Table 11-36—MCS field

MCS identifier	MCS index
0b00000	0
0b00001	1
0b00010	2
0b00011	3
0b00100	4
0b00101	5
0b00110	6
0b00111	7
0b01000	8
0b01001	9
0b01010	10
0b01011	11
0b01100–0b11111	Reserved

The Frame Length field shall be an unsigned integer number that indicates the number of octets in the Frame Payload field (which does not include the FCS).

The Preamble Type field indicates the type of the PHY preamble (long or short) used in the next frame, as defined in Table 11-37.

Table 11-37—Preamble type field

Preamble type	Type of preamble used for next frame
0b00	Long preamble
0b01	Short preamble
0b10–0b11	Reserved

The Beam Tracking field shall be set to one if training sequences for beam tracking are present following the current frame. It shall be set to zero otherwise.

The Low-latency Mode field shall be set to one if the frame is using the low-latency aggregation mode, and shall be set to zero otherwise.

If the Low-latency Mode field is set to one, then the Aggregation field shall also be set to one.

The Bit Interleaver field shall be set to one to indicate that bit interleaving is used in the payload of the frame and shall be set to zero otherwise.

The PCES Mode field shall be set to one to indicate that PCES symbols are added periodically to the PHY payload field, and it shall be set to zero otherwise.

The Skewed Constellation field is optional. When used, this field shall be set to one if UEP uses constellation mapping and shall be set to zero to indicate that UEP uses coding. If the field is not used, it shall be set to zero.

11.3.3.4 Header check sequences

11.3.3.4.1 Main Header HCS

The combination of the PHY header and the MAC header shall be protected with an ITU-T CRC-16 header check sequence (HCS). The ITU-T CRC-16 is described in 10.2.9.

11.3.3.4.2 MAC Subheader HCS

The MAC subheader shall be protected with an ITU-T CRC-16 header check sequence (HCS). This shall be computed in the same way, as specified in 10.2.9.

11.3.3.5 PHY Payload field

The PHY Payload field is the last component of the frame, and is constructed as shown in Figure 11-38.




Figure 11-38—PHY Payload field construction process

The PHY Payload field shall be encoded as follows:

- Scramble the MAC Frame Body field according to 11.3.2.11.
- Encode and interleave the scrambled MAC Frame Body field, as specified in 11.3.2.4.
- If necessary, add stuff bits, as defined in 11.3.2.5.
- Map the MAC Frame Body field onto the appropriate constellation, as described in 11.3.2.6.
- Spread the modulated signal, as described in 11.3.2.7.
- Interleave the MAC Frame Body field with the tone interleaver, as described in 11.3.2.8.
- Modulate the resulting complex values with the OFDM modulator described in 11.3.2.9.
- Insert PCES periodically, as described in 11.3.2.10.

Note that when the frame payload length is zero, the length of the PHY Payload field shall also be zero.

11.3.4 Transmitter specifications**11.3.4.1 EVM requirement**

The EVM of a compliant transmitter shall be measured and calculated as defined in 11.1.7 and shall not exceed the values given in Table 11-38.

Table 11-38—Allowed relative constellation error versus data rate

Data rate	Relative constellation RMS error
Up to 1.5 Gb/s	−7 dB
2.1 Gb/s to 2.7 Gb/s	−14 dB
2.8 Gb/s to 5.3 Gb/s	−21 dB
Above 5.4 Gb/s	−23 dB

The relative constellation RMS error calculation shall be performed using a device capable of converting the transmitted signal into a stream of complex samples at the sampling rate, f_s , as defined in Table 11-31, or higher, with sufficient accuracy in the I/Q imbalance, DC offset, phase noise, etc. The sampled signal shall then be processed in a manner similar to that of an ideal receiver including adding the 64 samples of the cyclic prefix to the received OFDM symbol.

11.3.4.2 Chip rate and clock alignment

The transmitted center frequency and chip clock frequency tolerances shall be $\pm 20 \cdot 10^{-6}$ maximum. The transmit center frequency and symbol clock frequency shall be derived from the same reference oscillator (locked).

11.3.5 Receiver specifications**11.3.5.1 Receiver sensitivity**

For a BER of 10^{-6} , the minimum receiver sensitivity shall be -50 dBm for MCS index 1 and shall be -70 dBm for MCS index 0. The BER shall be measured at the MAC/PHY interface after PHY level FEC. A PN sequence defined by $x^{55} + x^{24} + 1$ shall be used for the BER measurement.

11.3.5.2 Receiver CCA performance

The start of a valid transmission at a receiver equal to or greater than the minimum sensitivity for MCS index 0, if MCS index 0 is implemented, shall cause CCA to indicate busy with a probability $> 90\%$ within *pCcaDetectTime*, as defined in Table 11-26. The receiver CCA function shall in all circumstances report the medium busy with any signal 20 dB above the minimum sensitivity for MCS index 0.

If CMS is supported, then the receiver CCA performance is described in 11.1.12.8.

11.3.5.3 Receiver maximum input level

The receiver maximum input level is the maximum power level of the incoming signal present at the input of the receiver for which the error rate criterion is met. A compliant receiver shall have a receiver maximum input level of at least -25 dBm for each of the modulation formats that the device supports.

11.3.5.4 PHY layer timing

The values for the PHY layer timing parameters are given in Table 11-39.

Table 11-39—PHY layer timing parameters

PHY parameter	Value	Subclause
<i>pMifsTime</i>	0.2 μ s, 0.5 μ s (default), 2.0 μ s	11.3.5.8
<i>pSifsTime</i>	0.2 μ s, 2.0 μ s, 2.5 μ s (default)	11.3.5.6
<i>pCcaDetectTime</i>	2.5 μ s	11.3.5.2
<i>pChannelSwitchTime</i>	100 μ s	11.3.5.9

11.3.5.5 Interframe spacing

The interframe spacing parameters are given in Table 11-40.

Table 11-40—HSI interframe spacing parameters

MAC parameter	PHY parameter	Definition
MIFS	$pMifsTime$	10.2.7.4
SIFS	$pSifsTime$	10.2.7.2
$pBackoffSlot$	$pSifsTime + pCcaDetectTime$	10.2.7.1
BIFS	$pSifsTime + pCcaDetectTime$	10.2.7.1
RIFS	$2 \times pSifsTime + pCcaDetectTime$	10.2.7.1

11.3.5.6 Receive-to-transmit turnaround time

The RX-to-TX turnaround time shall be less than $pSifsTime$. The RX-to-TX turnaround time shall be measured at the air interface from the trailing edge of the last symbol received until the first symbol of the PHY preamble is present at the air interface.

11.3.5.7 Transmit-to-receive turnaround time

The TX-to-RX turnaround time shall be less than $pSifsTime$. The TX-to-RX turnaround time shall be measured at the air interface from the trailing edge of the last transmitted symbol until the receiver is ready to begin the reception of the next PHY frame.

11.3.5.8 Time between successive transmissions

The time between successive transmissions shall be $pMifsTime$, including the power-up ramp specified in 10.5.7. The $pMifsTime$ shall be measured at the air interface from the trailing edge of the last symbol transmitted until the first symbol of the PHY preamble is present at the air interface.

11.3.5.9 Channel switch time

The channel switch time is defined as the time from when the last valid bit is received at the antenna on one channel until the DEV is ready to transmit or receive on a new channel. The channel switch time shall be less than $pChannelSwitchTime$.

11.3.6 HSI PHY management

The PHY PIB comprises the managed objects, attributes and notifications required to manage the PHY layer of a DEV.

11.3.6.1 PHY supported data rate encoding

The encoding of the PHY data rates used in the Supported Data Rates field in the Capability IE, as described in 6.4.11, is given in 11.1.8.1.

11.3.6.2 HSI PHY fragment size encoding

The encoding of the preferred fragment size used in the Capability IE, as described in 6.4.11, is given in Table 11-4.

The PHY definitions create restrictions on the maximum frame size, maximum transfer unit size, and minimum fragmentation size that will be supported.

11.3.6.3 Maximum frame length

The maximum frame length allowed, $pMaxFrameBodySize$, shall be $2^{20} - 1$ octets. This total includes the Frame Body field and FCS but not the PHY preamble, PHY header, or MAC header. The maximum frame length also does not include the tail symbols or the stuff bits.

11.3.6.4 Maximum transfer unit size

The maximum size data frame passed from the upper layers, $pMaxTransferUnitSize$, shall be the same as $pMaxFrameBodySize$, as defined in 11.3.6.3. If security is enabled for the data connection, the upper layers should limit data frames to $pMaxFrameBodySize$ minus the security overhead, as defined in 6.3.4.2.

11.3.6.5 Minimum fragment size

The minimum fragment size, $pMinFragmentSize$, shall be 512 octets.

11.4 Audio/Visual mode of mmWave PHY

The Audio/Visual (AV) PHY is implemented with two PHY modes: the high-rate PHY (HRP) and low-rate PHY (LRP), both of which use orthogonal frequency domain multiplexing (OFDM). The data rates supported by the HRP are defined in Table 11-41.

Table 11-41—HRP data rates and coding

HRP mode index	Coding mode	Modulation	Inner code rate		Data rate (Gb/s)
			MSB	LSB	
0	EEP	QPSK	1/3		0.952
1		QPSK	2/3		1.904
2		16-QAM	2/3		3.807
3	UEP	QPSK	4/7	4/5	1.904
4		16-QAM	4/7	4/5	3.807
5	MSB-only retransmission	QPSK	1/3	N/A	0.952
6		QPSK	2/3	N/A	1.904

If a DEV supports the use of the HRP, it shall support the use of HRP mode index 0 and HRP mode index 1.

HRP mode index 5 and HRP mode index 6 transmit only the four MSBs of each octet. In this case the data rate refers to MSBs only. The four LSBs in each octet are discarded by the transmitter and only the MSB portions of the baseband, as shown in Figure 11-39, shall be used.

Typical video and audio consumer electronics are configured either as a source of data, e.g., a video disc player, or as a sink, e.g., a display. For these applications, the data flow is highly asymmetric in the same direction. Thus, AV PHY DEVs shall implement one of the following configurations:

- HR0: The DEV implements LRP transmit and receive functions.
- HRRX: The DEV implements HRP receive, LRP transmit, and LRP receive functions.

- HRTX: The DEV implements HRP transmit, LRP transmit, and LRP receive functions.
- HRTR: The DEV implements HRP transmit, HRP receive, LRP transmit, and LRP receive functions.

The data rates supported by the LRP are defined in Table 11-42.

Table 11-42—LRP data rates and coding

LRP mode index	Modulation	FEC	Data rate (Mb/s)	Repetition
0	BPSK	1/3	2.5	8×
1		1/2	3.8	8×
2		2/3	5.1	8×
3		2/3	10.2	4×

A DEV that supports the AV PHY mode shall support LRP modes 0, 1, and 2. All broadcast and multi-cast frames using the AV PHY shall be sent with an LRP mode.

11.4.1 General requirements

11.4.1.1 AV PHY channelization

The HRP mode uses the channels defined in Table 11-3. A compliant IEEE 802.15.3c implementation that implements the AV PHY shall support at least channel number 2.

In each of the HRP channels, three LRP channels are defined. In a piconet, only one HRP channel and one LRP channel is used at a time. Each of the LRP channels is defined relative to the center frequency of the current HRP channel, $f_{c(HRP)}$, as defined in Table 11-3. The LRP channels are defined in Table 11-43.

Table 11-43—LRP channelization

LRP channel index	Start frequency ^a	Center frequency	Stop frequency ^a
1	$f_{c(HRP)} - 207.625 \text{ MHz}$	$f_{c(HRP)} - 158.625 \text{ MHz}$	$f_{c(HRP)} - 109.625 \text{ MHz}$
2	$f_{c(HRP)} - 49 \text{ MHz}$	$f_{c(HRP)}$	$f_{c(HRP)} + 49 \text{ MHz}$
3	$f_{c(HRP)} + 109.625 \text{ MHz}$	$f_{c(HRP)} + 158.625 \text{ MHz}$	$f_{c(HRP)} + 207.625 \text{ MHz}$

^aThe start and stop frequencies are nominal values. The TX mask requirements for the LRP is defined in 11.4.4.1.

The *phyCurrentChannel* is the CHNL_ID of the current channel. For the purpose of the Remote Scan Request and Remote Scan Response commands, 6.5.7.3 and 6.5.7.4, respectively, the Channel Index field is the CHNL_ID in Table 11-44.

Table 11-44—Mapping HRP/LRP channel index to CHNL_ID

CHNL_ID	HRP channel index	LRP channel index
1	1	1
2	2	1
3	3	1
4	4	1
5	1	2
6	2	2
7	3	2
8	4	2
9	1	3
10	2	3
11	3	3
12	4	3

11.4.1.2 PHY layer timing

The values for the AV PHY layer timing are defined in Table 11-45.

Table 11-45—AV PHY layer timing parameters

PHY parameter	Value	Subclause
<i>pMifsTime</i>	<i>pSifsTime</i>	11.4.1.2.4
<i>pSifsTime</i>	2 μ s $\pm$ 24 samples at LRP reference sampling rate	11.4.1.2.2
<i>pCcaDetectTime</i>	9 μ s	11.4.1.2.6
<i>pChannelSwitchTime</i>	1000 μ s	11.4.1.2.5

11.4.1.2.1 Interframe space

A conformant AV PHY implementation shall support the IFS parameters, as described in 7.4.1, given in Table 11-46.

Table 11-46—AV PHY IFS parameters

IEEE 802.15.3 MAC parameter	Corresponding AV PHY parameter	Definition
MIFS	$pMifsTime$	11.4.1.2.4
SIFS	$pSifsTime$	11.4.1.2.2
$pBackoffSlot$	$pSifsTime + pCcaDetectTime$	11.4.1.2.2, 11.4.1.2.6
BIFS	$pSifsTime + pCcaDetectTime$	11.4.1.2.2, 11.4.1.2.6
RIFS	$2 \times pSifsTime + pCcaDetectTime$	11.4.1.2.2, 11.4.1.2.6

11.4.1.2.2 Receive-to-transmit turnaround time

The RX-to-TX turnaround time shall be $pSifsTime$. The RX-to-TX turnaround time shall be measured at the air interface from the trailing edge of the last symbol received until the first symbol of the PHY preamble is present at the air interface.

11.4.1.2.3 Transmit-to-receive turnaround time

The TX-to-RX turnaround time shall be less than $pSifsTime$. The TX-to-RX turnaround time shall be measured at the air interface from the trailing edge of the last transmitted symbol until the receiver is ready to begin the reception of the next PHY frame.

11.4.1.2.4 Time between successive transmissions

The time between successive transmissions shall be $pMifsTime$. The $pMifsTime$ shall be measured at the air interface from the trailing edge of the last symbol transmitted until the first symbol of the PHY preamble is present at the air interface.

11.4.1.2.5 Channel switch time

The channel switch time is defined as the time from when the last valid bit is received at the antenna on one channel until the DEV is ready to transmit or receive on a new channel. The channel switch time shall be less than $pChannelSwitchTime$.

11.4.1.2.6 CCA detect time

The LRP shall be able to detect the presence of an LRP preamble at an input power level equivalent to the sensitivity of the LRP mode 1 within $pCcaDetectTime$. CCA detection is not used for the HRP mode because only the LRP is used in CPs.

11.4.1.3 Data size restrictions

The PHY definitions creates restrictions on the maximum frame size, maximum transfer unit size, and minimum fragmentation size that will be supported. These parameters are defined in this subclause.

11.4.1.3.1 Maximum frame length

The maximum frame length allowed, $pMaxFrameBodySize$, shall be $2^{20} - 1$ octets for HRP subframes and $2^{12} - 1$ octets for LRP frames. This total includes the frame payload and FCS but not the PHY preamble,

PHY header, MAC header, or HCS. The maximum frame length also does not include the tail symbols or the stuff bits.

11.4.1.3.2 Maximum transfer unit size

The maximum size data frame passed from the upper layers, $pMaxTransferUnitSize$, shall be the same as $pMaxFrameBodySize$, as defined in 11.4.1.3.1. If security is enabled for the data connection, the upper layers should limit data frames to $pMaxFrameBodySize$ minus the security overhead, as defined in 6.3.4.2.

11.4.1.3.3 Minimum fragment size

The minimum fragment size, $pMinFragmentSize$, that is allowed with the AV PHY shall be 512 octets.

11.4.1.4 Header check sequence

The HCS shall be a 32-bit CRC that is equivalent to the one used for the FCS, as defined 6.2.7.6. For the AV PHY, the MAC parameter $pLengthHcs$ shall be four. The HCS shall be transmitted in the order specified in 6.1. At both the receiver and transmitter, the initial state of the remainder shall be set to all ones.

11.4.2 AV PHY modulation and forward error correction

The modulation parameters for the HRP are given in Table 11-47.

Table 11-47—HRP modulation parameters

Parameter	Value	Symbol
Occupied bandwidth	1.76 GHz	N/A
Reference sampling rate	2.538 GHz	$f_{s(HR)}$
Number of subcarriers	512	$N_{sc(HR)}$
FFT period	$N_{sc(HR)}/f_{s(HR)} \sim 202$ ns	$T_{FFT(HR)}$
Subcarrier spacing	$1/T_{FFT(HR)} \sim 4.96$ MHz	$\Delta f_{SC(HR)}$
Guard interval	$64/f_{s(HR)} \sim 25.2$ ns	$T_{GI(HR)}$
Symbol duration	$T_{FFT(HR)} + T_{GI(HR)} \sim 227$ ns	$T_{S(HR)}$
Number of data subcarriers	336	$N_{dsc(HR)}$

The modulation parameters for the LRP are given in Table 11-48.

Table 11-48—LRP modulation parameters


Parameter	Value	Symbol
Occupied bandwidth	92 MHz	N/A
Reference sampling rate	317.25 MHz	$f_{s(LR)}$
Number of subcarriers	128	$N_{sc(LR)}$

Table 11-48—LRP modulation parameters (continued)


Parameter	Value	Symbol
FFT period	$N_{sc(LR)}/f_{s(LR)} \sim 403$ ns	$T_{FFT(LR)}$
Subcarrier spacing	$1/T_{FFT(LR)} \sim 2.48$ MHz	$\Delta f_{SC(LR)}$
Guard interval	$28/f_{s(LR)} \sim 88.3$ ns	$T_{GI(LR)}$
Symbol duration	$T_{FFT(LR)} + T_{GI(LR)} \sim 492$ ns	$T_{S(LR)}$
Number of data subcarriers	30	$N_{dsc(LR)}$

Unless otherwise specified, the terms *sampling rate* or *sample* are in terms of the reference sampling rate for that mode.

A reference implementation of the HRP baseband is illustrated in Figure 11-39.

**Figure 11-39—HRP reference implementation block diagram**

A reference implementation of the LRP baseband is illustrated in Figure 11-40.

**Figure 11-40—LRP reference implementation block diagram**

The purpose of the reference implementation is to provide a reference for defining the encoding of the incoming bit stream into the appropriate RF signal. Implementations may use different architectures as long as the resulting RF signal matches the one that would be generated by a reference implementation.

11.4.2.1 AV PHY base rate

The base data rate of the AV PHY shall be LRP mode 0, 1, or 2.

11.4.2.2 Repetition coding and spatial diversity

The LRP transmitter utilizes spatial diversity and repetition coding for the omni-directional and directional modes. Up to eight antenna directions or phased array patterns may be used that shall be identified by indices zero through seven.

For omni LRPDUs, the repetition coding is implemented by repeating the OFDM symbol and its associated cyclic prefix four or eight times as indicated by the LRP mode index and sending each repetition with a different TX antenna direction or pattern. For the case of eight times repetition coding, each OFDM symbol in the Channel Estimation field and the OFDM symbols that follow in the frame are repeated using patterns zero through seven. For the four times repetition coding each OFDM symbol in the MAC header, HCS and MAC Frame Body field is repeated using patterns zero through three, where patterns zero through three are the first four patterns used in the Channel Estimation field of the omni LRP preamble. The assignment of indices to the patterns is implementation dependent and arbitrary. However, for all omni LRPDUs that use the short preamble, with the exception of LRP mode 3, the assignment of indices shall remain unchanged after the system is powered on. For omni LRPDUs that use the long omni LRP preamble, or omni LRPDUs that use the short preamble with LRP mode 3, the assignment of indices may be different from frame to frame. However, in these cases the assignment of indices shall remain unchanged for the duration of a frame. These TX antenna directions or patterns are selected such that the transmission covers the region of space that is of interest. The patterns that are used do not need to be unique as long as the same set of patterns is used for all omni LRPDUs. For example, for implementations with less than eight independently controllable elements, some of the patterns are repeated during the eight times repetition. In the case of a single TX antenna, a single direction or pattern is used for all of the repetitions.

In the Directional mode, the preamble is used to perform channel estimation and receiver training for only one optimum TX antenna direction, or phase array pattern. The same TX optimum pattern is also used for the header and payload. This optimum direction is selected while an omni LRPDU with short preamble is received. The Directional LRPDU Header field shall use the eight times directional repetition coding while the directional repetition coding of the data symbols is either eight times or four times as determined by the LRP mode index. The directional repetition coding for Directional LRPDUs is performed by adding one additional OFDM symbol as cyclic prefix to the eight or four times repeated 128-sample OFDM symbols depending on the mode index. Therefore, for eight times directional repetition, the 128-sample OFDM symbol is repeated nine times, while it is repeated five times for four times directional repetition. In this case, no other cyclic prefix is used.

11.4.2.3 Stuff bits

In order that an integer number of OFDM symbols are created and, for the HRP, complete the outer interleaver units, the PHY adds additional bits to the bit stream, called stuff bits, prior to performing any operations on the incoming data. Stuff bits for the HRP shall be set to zero prior to adding them to the end of the bit stream. Stuff bits for the LRP shall be set equal to a sequence generated by the scrambler, defined in 11.4.2.5, initialized to an arbitrary state. In LRPDUs, stuff bits only occur at the end of a frame, i.e., after the Frame Body field or the Directional LRP Payload field.

The PHY shall add the minimum number of stuff bits necessary to create an integer number of OFDM symbols and, for HRP modes, complete the outer interleaver unit for the combination of the PHY Header field, MAC Header field and HCS field. In addition, the HRP PHY shall add the minimum number of stuff bits necessary to create an integer number of OFDM symbols and complete the outer interleaver unit for each combination of subframes that have the same HRP mode index and for the last subframe in the frame. These additional bits shall be discarded by the receiver upon reception. The last outer interleaver unit for each of the subframes that end on a HRP mode change and for the last subframe may be shortened, as specified in 11.4.2.7.

11.4.2.4 HRP splitter and scrambler

For the HRP header and MSB only mode, no re-ordering is applied to the input octet stream before it is sent to the scrambler. For all other HRP modes, prior to scrambling, the input octets are re-ordered into upper and lower branches. If the input to the splitter is an array of octets, $in(i, b)$, where i is the index of the input octets, $i = 0, 1, \dots, N - 1$, where N is the number of octets in the input stream and b is the index of the bit in the octet where $b = 0$ is the LSB and $b = 7$ is the MSB. The output of the scrambler is two arrays of octets, $upper(n, b)$ and $lower(n, b)$, where $n = 0, 1, \dots, N - 1$. Note that adding the stuff bits guarantees that the input octet stream is an even number of octets.

In EEP mode, the output arrays are constructed according to the following equations:

$$lower(n, b) = in(\{\text{floor}[b/4] + \text{floor}[n/2] \times 4\}, \{\text{mod}[b, 4] + \text{mod}[n, 2] \times 4\})$$

$$upper(n, b) = in(\{\text{floor}[b/4] + \text{floor}[n/2] \times 4 + 2\}, \{\text{mod}[b, 4] + \text{mod}[n, 2] \times 4\})$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

In UEP mode, the output arrays are constructed according to the following equations:

$$lower(n, b) = in(\{\text{floor}[b/4] + n \times 2\}, \{\text{mod}[b, 4]\})$$

$$upper(n, b) = in(\{\text{floor}[b/4] + n \times 2\}, \{\text{mod}[b, 4] + 4\})$$

When the octets are re-ordered, the first 8 bits of the scrambler are applied to the lower branch and the next 8 bits to the upper branch, switching every 8 bits between the two branches.

The HRP scrambler shall use the generator polynomial $P(x) = x^{15} + x^{14} + 1$. A reference implementation is illustrated in Figure 11-14.

The initial value of scrambler shall be set by the four variable seeds, S_0, S_1, S_2 and S_3 , and the 11 fixed seeds as $[x_{-1}, x_{-2}, \dots, x_{-15}] = [1101\ 0000\ 101\ S_3\ S_2\ S_1\ S_0]$. All fields in the HRP header, MAC header and HCS, as illustrated in 11.4.3, shall be scrambled as with the variable seeds set to $S_0 = 0, S_1 = 1, S_2 = 0, S_3 = 1$. The entire data stream following the HCS, including the stuff bits, is scrambled using the seed bits specified in the PHY Control field, as illustrated in Figure 11-55.

11.4.2.5 LRP scrambler

Scrambling applies to the MAC header, HCS, and Frame Body field only. A scrambler with six state bits, s_0 through s_5 , and wired according to the polynomial $x^6 + x + 1$ is initialized with a four-bit random value combined with 0b01 as illustrated in Figure 11-41. Bit s_5 from successive states of the scrambler are exclusive-ORed with each data bit in sequence to create the scrambled data. The four bit random value that is used to initialize the scrambler is placed in the fields S_0, S_1, S_2 , and S_3 in Figure 11-60.




Figure 11-41—LRP scrambler reference implementation

NOTE—The directional LRPDU, which does not contain the MAC Header, HCS, or MAC Frame Body field, is not scrambled.

11.4.2.6 HRP outer code

For HRP modes, the output from the scrambler is split into two branches, MSB and LSB. The data in each branch shall have an outer code using a Reed-Solomon code with parameters (224, 216, $t = 4$). The last portion of the data shall be coded by a further shortened Reed-Solomon code. The tail bits for the inner convolutional code are inserted in the outer interleaver by further shortening of the Reed-Solomon code. The parameters of these shortened codes are defined in 11.4.2.7.

The field generator polynomial is as follows:

$$g_{field}(x) = x^8 + x^4 + x^3 + x^2 + 1$$

and the code generator polynomial is as follows:

$$g_{RS}(x) = (x+\lambda)(x+\lambda^2)(x+\lambda^3)(x+\lambda^4)(x+\lambda^5)(x+\lambda^6)(x+\lambda^7)(x+\lambda^8)$$

where:

$$\lambda = 0x02$$

The combination of PHY Header field, MAC Header field, and HCS shall use only the MSB branch.

The input and output order for the Reed-Solomon encoder is specified in 11.4.2.7 in combination with the outer interleaver.

11.4.2.7 HRP outer interleaver

The outer interleaver shall output the octets from $i = 0, k = 0$ first to $i = depth-1, k = N-1$ last, where $depth$ is the depth of the outer interleaver and N is the length of RS code. With M parallel convolutional inner encoders for each RS codeword and $b(n, m)$ is the output of the RS encoder, the outer interleaver shall give the RS octets of $b(0,0), \dots, b(depth-1,0)$ to the first convolutional encoder with LSB first. All octets of $b(i, k \times M + m), i = 0, \dots, depth-1, k = 0, 1, \dots, N/M-1$, shall be output to the m^{th} convolutional encoder. The number of parallel convolutional encoders is specified in 11.4.2.8.

The outer block interleaver shall be operated with a $depth = 4$ for HRP data and a $depth = 2$ for the combination of the HRP Header, MAC Header and HCS fields. LRP modes do not use an outer block interleaver.

The combination of HRP header, MAC header and HCS field has 92 octets that are encoded into 112 octets by adding 16 parity octets for error protection and 4 tail octets to terminate the convolutional code.

For the combination of the HRP Header, MAC Header and HCS fields, the first 48 octets are encoded using RS(56, 48, $t = 4$) while the next 44 octets are encoded using RS(52, 44, $t = 4$). The second codeword is followed by 4 tail octets set to zero. The transmitted order of those octets, the method to insert the tail octets, and the method of interleaving for the outer interleaver are the same as those used for data.

The outer interleaver inserts the tail bits for the convolutional encoder. For the outer interleaver with tail bits, to improve the efficiency, the number of rows of the outer interleaver may be reduced to a minimum number that is an integer multiple of 28. At the columns of $i = 0$ to $i = depth - 2$, a shortened RS($28 \times n, 28 \times n - 8, t = 4$) code may be used, where $n = 1$ to 8.

11.4.2.8 Convolutional encoder

Scrambled data, tail bits and stuff bits shall be encoded with a convolutional encoder of code rate 1/3. The convolutional encoder shall use constraint length $K = 7$, delay memory 6, generator polynomial $g_0 = 133_0$, $g_1 = 171_0$, $g_2 = 165_0$, mother code rate 1/3. A detailed schematic diagram of the convolutional encoder is shown in Figure 11-42.




Figure 11-42—Convolutional encoder reference implementation

The initial value of the delay register shall be zero at the beginning of every HRP header and at the beginning of the Frame Body field.

The HRP transmitter uses eight parallel convolutional encoders, labeled A through H. The first four encoders, labeled A through D, are for the first outer Reed-Solomon coding branch and the last four encoders, labeled E through H are for the second outer Reed-Solomon coding branch.

The HRP header shall use the same convolutional encoder with a code rate of 1/3. However, only four parallel convolutional encoders, labeled A through D, are used.

For all LRP Frame Body fields other than with the directional LRP payload, six tail bits with value zero are added. For LRP PHY header, the 1/2 and 1/3 rate coding with tail-biting are used. In the tail biting method, the initial encoder state is set equal to the last six information bits, and no tail bits are appended. Therefore, the initial and final encoder states are equal. The directional LRPDU payload shall use either tail-biting or tail bits, as specified in 11.4.3.8.

11.4.2.9 Puncturing

Convolutional encoded data is punctured to make the desired code rate using the puncturing pattern indicated in Table 11-49. In Table 11-49, puncturing pattern '1' means to send the bit while a '0' means to omit (or do not transmit) the corresponding bit.

The output of puncturing block shall be serialized with sequential order as shown in Table 11-49.

11.4.2.10 HRP data multiplexer and bit interleaver

The output of the 8 encoders, labeled A through H, shall be multiplexed to form a single data stream prior to the bit interleaver, defined in 11.4.2.11, as illustrated in Figure 11-39. The method used to multiplex the encoded bits is dependent on the type of HRP mode, either EEP or UEP.

Table 11-49—Puncturing table

Code rate	Puncturing pattern	Transmitted sequence
1/3	X: 1 Y: 1 Z: 1	X1 Y1 Z1
1/2	X: 1 Y: 1 Z: 0	X1 Y1
4/7	X: 1 1 1 1 Y: 1 0 1 1 Z: 0 0 0 0	X1 Y1 X2 X3 Y3 X4 Y4
2/3	X: 1 1 Y: 1 0 Z: 0 0	X1 Y1 X2
4/5	X: 1 1 1 1 Y: 1 0 0 0 Z: 0 0 0 0	X1 Y1 X2 X3 X4

11.4.2.10.1 EEP data multiplexer

In the EEP mode, all the 8 encoders shall use the same inner code rate. The encoded bits shall be multiplexed and bit-interleaved every 48 bits.

During the length 48 multiplexing/interleaving cycle, a group multiplexer shall be used first with fixed group size 6 for all eight encoders. A1, A2, A3, A4, A5, A6 are used to label the 6 encoded bits (in increasing order in time) from encoder A, and similarly for B1 through B6, C1 through C6, D1 through D6, E1 through E6, F1 through F6, G1 through G6, and H1 through H6 from encoders B, C, D, E, F, G, and H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 11-50.

Table 11-50—Multiplexing scheme for EEP mode

numbering	0	1	2	3	4	5	6	7	8	9	10	11
labeling	A1	A2	A3	A4	A5	A6	B1	B2	B3	B4	B5	B6
numbering	12	13	14	15	16	17	18	19	20	21	22	23
labeling	C1	C2	C3	C4	C5	C6	D1	D2	D3	D4	D5	D6
numbering	24	25	26	27	28	29	30	31	32	33	34	35
labeling	E1	E2	E3	E4	E5	E6	F1	F2	F3	F4	F5	F6
numbering	36	37	38	39	40	41	42	43	44	45	46	47
labeling	G1	G2	G3	G4	G5	G6	H1	H2	H3	H4	H5	H6

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \dots, 47$ and $y = 0, \dots, 47$ be the index at the input and output of the bit interleaver, respectively. The bit interleaver in the EEP mode shall implement the relation given by the following:

$$y = \text{mod}((6 \times \text{floor}(x/6) - 5 \times \text{mod}(x,6)), 48)$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

The overall read out order of the multiplexer and interleaver for EEP is illustrated in Figure 11-43.




Figure 11-43—EEP multiplexing and bit interleaving pattern

At the receiver side, after demodulation, the received bits shall be deinterleaved. Let $y = 0, \dots, 47$ and $z = 0, \dots, 47$ be the index at the input and output of the bit deinterleaver, respectively. The bit deinterleaver in the EEP mode shall implement the following relation:

$$z = \text{mod}((6 \times \text{floor}(y/6) - 7 \times \text{mod}(y,6)), 48)$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

11.4.2.10.2 UEP coding data multiplexer

The UEP coding data multiplexer shall be implemented by DEVs that support UEP HRP modes. It shall not be implemented by DEVs that do not support any UEP HRP modes.

In the UEP coding mode, top 4 encoders (or encoders A, B, C, D) shall use rate 4/7 convolutional codes, and bottom 4 encoders (or encoders E, F, G, H) shall use rate 4/5 convolutional codes. The encoded bits shall be multiplexed and bit interleaved every 96 bits. The total length of 96 bits is divided into two half cycles, with each half cycle multiplexing and interleaving 48 bits, in a slightly different manner.

In the first half cycle, a group multiplexer with group size 7, 7, 7, 5, 5, 5, 5 for all eight encoders, A through H, respectively, shall be used. In this half cycle, A1, A2, A3, A4, A5, A6, A7 are used to label the 7 encoded bits (in increasing order in time) from encoder A, and similarly B1 through B7, C1 through C7, D1 through D7, E1 through E5, F1 through F5, G1 through G5, and H1 through H5 from encoders B, C, D, E, F, G, and H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 11-51.

Table 11-51—Multiplexing scheme in the first half cycle for UEP coding mode

numbering	0	1	2	3	4	5	6	7	8	9	10	11
labeling	A1	A2	A3	A4	A5	A6	A7	B1	B2	B3	B4	B5
numbering	12	13	14	15	16	17	18	19	20	21	22	23
labeling	B6	B7	C1	C2	C3	C4	C5	C6	C7	D1	D2	D3
numbering	24	25	26	27	28	29	30	31	32	33	34	35
labeling	D4	D5	D6	D7	E1	E2	E3	E4	E5	F1	F2	F3
numbering	36	37	38	39	40	41	42	43	44	45	46	47
labeling	F4	F5	G1	G2	G3	G4	G5	H1	H2	H3	H4	H5

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \dots, 47$ be the index at the input of the bit interleaver, and $y = 0, \dots, 47$ be the index at the output of the bit interleaver. The bit interleaver in the first half cycle of the UEP coding mode shall implement the following relation:

$$y = \text{mod}((6 \times \text{floor}(x/6) - 5 \times \text{mod}(x,6)), 48)$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

The overall read out order of the multiplexer and interleaver in the first half cycle for UEP is illustrated in Figure 11-44.




Figure 11-44—UEP coding mode first half cycle multiplexing and bit interleaving pattern

At the receiver side, after demodulation, the received bits shall be deinterleaved. Let $y = 0, \dots, 47$ and $z = 0, \dots, 47$ be the index at the input and output of the bit deinterleaver, respectively. The bit deinterleaver in the first half cycle of the UEP coding mode shall implement the following relation:

$$z = \text{mod}((6 \times \text{floor}(y/6) - 7 \times \text{mod}(y,6)), 48)$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

In the second half cycle, a group multiplexer with group size 7, 7, 7, 5, 5, 5, 5 for all eight encoders shall be used. In this half cycle, A8, A9, A10, A11, A12, A13, A14 are used to label the 7 encoded bits (in increasing order in time) from encoder A, and similarly B8 through B14, C8 through C14, D8 through D14, E6 through E10, F6 through F10, G6 through G10, and H6 through H10 from encoders B, C, D, E, F, G, and H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 11-52.

Table 11-52—Multiplexing scheme in the second half cycle for UEP coding mode

numbering	0	1	2	3	4	5	6	7	8	9	10	11
labeling	B8	B9	B10	B11	B12	B13	B14	C8	C9	C10	C11	C12
numbering	12	13	14	15	16	17	18	19	20	21	22	23
labeling	C13	C14	D8	D9	D10	D11	D12	D13	D14	A8	A9	A10
numbering	24	25	26	27	28	29	30	31	32	33	34	35
labeling	A11	A12	A13	A14	F6	F7	F8	F9	F10	G6	G7	G8
numbering	36	37	38	39	40	41	42	43	44	45	46	47
labeling	G9	G10	H6	H7	H8	H9	H10	E6	E7	E8	E9	E10

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \dots, 47$ be the index at the input of the bit interleaver, and $y = 0, \dots, 47$ be the index at the output of the bit interleaver. The bit interleaver in the first half cycle of the UEP coding mode shall implement the following relation:

$$y = \text{mod}((6 \times \text{floor}(x/6) - 5 \times \text{mod}(x,6)), 48)$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

The overall read out order of the multiplexer and interleaver in the second half cycle for UEP is illustrated in Figure 11-45.

column-wise readout (from top to bottom) in this direction →

B8	F6	D14	C13	H8	F10	A11	B12	E7	G9	D10	C9
C8	G8	A14	D13	E10	H7	F9	C12	B11	E6	A10	D9
D8	B14	G7	A13	D12	E9	H6	F8	C11	B10	H10	A9
A8	C14	H9	G6	A12	B13	E8	G10	D11	C10	B9	F7

Figure 11-45—UEP coding mode second half cycle multiplexing and bit interleaving pattern

At the receiver side, after demodulation, the received bits shall be deinterleaved. Let $y = 0, \dots, 47$ and $z = 0, \dots, 47$ be the index at the input and output of the bit deinterleaver, respectively. The bit deinterleaver in the second half cycle of the UEP coding mode shall implement the following relation:

$$z = \text{mod}((6 \times \text{floor}(y/6) - 7 \times \text{mod}(y,6)), 48)$$

where $\text{mod}(x, y)$ is the modulo function and is defined in 11.2.2.6.3.

11.4.2.10.3 UEP mapping data multiplexer

The UEP mapping data multiplexer shall be implemented by DEVs that support UEP HRP modes. It shall not be implemented by DEVs that do not support any UEP HRP modes.

In the UEP mapping mode, all eight encoders shall use the same coding rate and the encoded bits shall be multiplexed and bit interleaved every 48 bits.

During the length 48 multiplexing/interleaving cycle, a group multiplexer shall be used with fixed group size 6 for all eight encoders. A1, A2, A3, A4, A5, A6 are used to label the 6 encoded bits (in increasing order in time) from encoder A, and similarly B1 through B6, C1 through C6, D1 through D6, E1 through E6, F1 through F6, G1 through G6, and H1 through H6 from encoders B, C, D, E, F, G, H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 11-53.

Table 11-53—Multiplexing scheme for UEP mapping mode

numbering	0	1	2	3	4	5	6	7	8	9	10	11
labeling	A1	A2	A3	A4	A5	A6	B1	B2	B3	B4	B5	B6
numbering	12	13	14	15	16	17	18	19	20	21	22	23
labeling	C1	C2	C3	C4	C5	C6	D1	D2	D3	D4	D5	D6
numbering	24	25	26	27	28	29	30	31	32	33	34	35
labeling	E1	E2	E3	E4	E5	E6	F1	F2	F3	F4	F5	F6
numbering	36	37	38	39	40	41	42	43	44	45	46	47
labeling	G1	G2	G3	G4	G5	G6	H1	H2	H3	H4	H5	H6

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \dots, 47$ be the index at the input of the bit interleaver. At the output of the bit interleaver, 24 bits are mapped to the I branch of the constellation, and the other 24 bits are mapped to the Q branch of the constellation. Let $yI = 0, \dots, 23$ be the index of those bits mapped to the I branch, and let $yQ = 0, \dots, 23$ be the index of those bits mapped to the Q branch. The bit interleaver in the UEP mapping mode shall implement the following relations:

$$yI = \text{mod}((6 \times \text{floor}(x/6) - 5 \times \text{mod}(x,6), 24), \text{ if } 0 \leq x \leq 23$$

$$yQ = \text{mod}((6 \times \text{floor}(x/6) - 5 \times \text{mod}(x,6), 24), \text{ if } 24 \leq x \leq 47$$

The overall readout order of the multiplexer and bit interleaver for UEP mapping is illustrated in Figure 11-46. Columns are read out from left to right, while the read out order inside each column depends on the constellation used.




Figure 11-46—UEP mapping mode overall multiplexing and bit interleaving pattern

In the QPSK case, the read out order inside each column is (row 1, row 3, row 2, row 4). Take the first column for example, and the readout is ordered as A1, E1, B2, F2, where A1, B2 are mapped to the Q branch, and E1, F2 are mapped to the I branch.

In the 16-QAM case, the read out order inside each column is (row 1, row 2, row 3, row 4). Take the first column for example, and the readout is ordered as A1, B2, E1, F2, where A1, B2 are mapped to the I branch, and E1, F2 are mapped to the Q branch.

At the receiver side, after demodulation, the received bits shall be deinterleaved. Let $yI = 0, \dots, 23$, $yQ = 0, \dots, 23$ be the index at the inputs of the deinterleaver I-branch and Q-branch, respectively, and $z = 0, \dots, 47$ be the index at the output of the bit deinterleaver. The bit deinterleaver in the UEP mapping mode shall implement the following relations:

$$z = \text{mod}((6 \times \text{floor}((yI)/6) - 5 \times \text{mod}(yI,6), 24), \text{ if } 0 \leq z \leq 23$$

$$z = \text{mod}((6 \times \text{floor}(yI/6) - 7 \times \text{mod}(yQ,6), 24), \text{ if } 24 \leq z \leq 47$$

11.4.2.10.4 HRP header data multiplexer

The data multiplexer and bit interleaver for the HRP header is similar to that for EEP mode described in 11.4.2.10.1. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 11-54.

After the data multiplexer, the bit interleaver, tone interleaver, and other operations for the PHY followed that for HRP mode index 0.

Table 11-54—HRP header multiplexing scheme

numbering	0	1	2	3	4	5	6	7	8	9	10	11
labeling	A1	A3	A5	A7	A9	A11	B1	B3	B5	B7	B9	B11
numbering	12	13	14	15	16	17	18	19	20	21	22	23
labeling	C1	C3	C5	C7	C9	C11	D1	D3	D5	D7	D9	D11
numbering	24	25	26	27	28	29	30	31	32	33	34	35
labeling	A2	A4	A6	A8	A10	A12	B2	B4	B6	B8	B10	B12
numbering	36	37	38	39	40	41	42	43	44	45	46	47
labeling	C2	C4	C6	C8	C10	C12	D2	D4	D6	D8	D10	D12

11.4.2.11 Bit reversal tone interleaver

All bits shall be interleaved by a block interleaver with a block size corresponding to the size of FFT in a single OFDM symbol, $N_{sc(HR)}$ or $N_{sc(LR)}$. The interleaver is used so that the adjacent data symbols are mapped onto separate subcarriers.

At the transmitter side, the interleaver permutation shall be defined as follows: Let k be the index of the tones (including data tones, pilot tones, DC tones and null tones) before permutation ranging between 0 and $N_{sc(HR)} - 1$ for HRP modes and between 0 and $N_{sc(LR)} - 1$ for LRP modes. Let i be the index of the interleaved tones over the same range (including data tones, pilot tones, DC tones and null tones) after permutation. Let

$$k = \sum_{j=0}^{L-1} a_j 2^j$$

where $L = \log_2(N_{sc(HR)}) - 1$ for HRP modes and $L = \log_2(N_{sc(LR)}) - 1$ for LRP modes, with $[a_L, \dots, a_0]$ being the binary representation of integer k . Then the binary representation of integer i can be written as $[a_0, \dots, a_L]$,

$$i = \sum_{j=0}^{L-1} a_j 2^{L-j}$$

DC, null, and pilot tones shall be inserted in the bit-reversal position before the tone interleaver. This makes sure that after permutation, the DC, null, and pilot tones appear in the pre-specified positions.

11.4.2.12 Signal constellations

The bits for the LRP mode for each subcarrier shall be mapped to -1 when the bit is zero and shall be mapped to $+1$ when the bit is one.

For QPSK modulation, serial bits shall be divided into groups of $N_{BPSK} = 2$ bits and converted into complex numbers representing QPSK constellation points. The subcarrier mapping for the HRP mode for QPSK modulation is illustrated in Figure 11-47 where the input bit b_0 is the earliest in time.