

INTERNATIONAL
STANDARD

ISO/IEC/
IEEE
8802-15-6

First edition
2017-10

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

**Part 15-6:
Wireless body area network**

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

Partie 15-6: Réseau corporel sans fil (BAN)



Reference number
ISO/IEC/IEEE 8802-15-6:2017(E)

© IEEE 2012

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-6:2017



COPYRIGHT PROTECTED DOCUMENT

© IEEE 2012

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 8802-15-6 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE Std. 802.15.6-2012). 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 8802-15-6:2017

IEEE Standard for Local and metropolitan area networks— Part 15.6: Wireless Body Area Networks

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 6 February 2012

IEEE-SA Standards Board

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 802-15-6:2017

Abstract: Short-range, wireless communications in the vicinity of, or inside, a human body (but not limited to humans) are specified in this standard. It uses existing industrial scientific medical (ISM) bands as well as frequency bands approved by national medical and/or regulatory authorities. Support for quality of service (QoS), extremely low power, and data rates up to 10 Mbps is required while simultaneously complying with strict non-interference guidelines where needed. This standard considers effects on portable antennas due to the presence of a person (varying with male, female, skinny, heavy, etc.), radiation pattern shaping to minimize the specific absorption rate (SAR) into the body, and changes in characteristics as a result of the user motions.

Keywords: BAN, body area network, IEEE 802.15.6

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

Copyright © 2012 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 29 February 2012. Printed in the United States of America.

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

PDF: ISBN 978-0-7381-7206-4 STD97213
Print: ISBN 978-0-7381-7227-9 STDPD97213

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.

Notice and Disclaimer of Liability Concerning the Use of IEEE Documents: IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. 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 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.

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

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

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

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 the official position of IEEE or any of its committees and shall not be considered to be, nor 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 to ensure 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. Any person who would like to participate in evaluating comments or revisions to an IEEE standard is welcome to join the relevant IEEE working group at <http://standards.ieee.org/develop/wg/>.

Comments on standards should be submitted to the following address:

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

Photocopies: Authorization to photocopy portions of any individual standard for internal or personal use is granted by The Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Notice to users

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

This document is copyrighted by the IEEE. It is made available for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making this document available for use and adoption by public authorities and private users, the IEEE does not waive any rights in copyright to this document.

Updating of IEEE documents

Users of IEEE Standards 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. 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://standards.ieee.org/index.html> or contact the IEEE at the address listed previously. For more information about the IEEE Standards Association or the IEEE standards development process, visit IEEE-SA Website at <http://standards.ieee.org/index.html>.

Errata

Errata, if any, for this and all other standards can be accessed 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 Patents Claims, or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this 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.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-6:2017

Participants

At the time this standard was submitted to the IEEE-SA Standards Board for approval, the 802.15 Working Group had the following membership:

Robert F. Heile, *Chair*
Rick Alfvén, *Vice Chair*
Patrick W. Kinney, *Vice Chair, Secretary*
James P. K. Gilb, *Technical Editor*
Clint Chaplin, *Treasurer*

Art Astrin, *Task Group 6 Chair*
Huan-Bang Li, *Task Group 6 Co-Vice Chair*
Ranjeet Patro, *Task Group 6 Co-Vice Chair*
Daniel Lewis, *Task Group 6 Editor*
Jin-Meng Ho, *Task Group 6 MAC and Security Editor*
Anuj Batra, *Task Group 6 Narrowband PHY Editor*
Marco Hernandez, *Task Group 6 UWB PHY Editor*
Seung-Hoon (Shannon) Park, *Task Group 6 HBC PHY Editor*
Igor Dotlic, *Task Group 6 Secretary*

Emad Afifi
 Gahng-Seop Ahn
 Roberto Aiello
 Takamasa Asano
 Taehan Bae
 Michael Bahr
 John Barr
 Tuncer Baykas
 Philip Beecher
 Ghulam Bhatti
 Gary Birk
 Mathew Boytim
 Peter Bradley
 Nancy Bravin
 David Britz
 Monique Brown
 Sverre Brubk
 Brian Buchanan
 John Buffington
 Kiran Bynam
 Brent Cain
 Edgar Callaway
 Chris Calvert
 Radhakrishna Canchi
 Ruben Cardozo
 Russell Chandler
 Kuor-Hsin Chang
 Soo-Young Chang
 Hind Chebbo
 In-Kyeong Choi
 Sangsung Choi
 David Cypher
 Matthew Dahl
 David Davenport
 Mark Dawkins
 Hendricus De Ruijter
 Gang Ding

Michael Dow
 Dietmar Eggert
 David Evans
 Charles S. Farlow
 John Farserotu
 Kory Fifield
 Stanislav Filin
 Will Filippo
 Michael Fischer
 George Flammer
 Ryosuke Fujiwara
 Noriyasu Fukatsu
 Kiyoshi Fukui
 John Geiger
 Gregory Gillooly
 Paul Gorday
 Elad Gottlieb
 Evan Green
 Robert Hall
 Shinsuke Hara
 Hiroshi Harada
 Shigekazu Harada
 Timothy Harrington
 Rodney Hemminger
 Ryoichi Higashi
 Garth Hillman
 Wei Hong
 Srinath Hosur
 David Howard
 Heqing Huang
 Sterling Hughes
 Jung-Hwan Hwang
 Ichirou Ida
 Tetsushi Ikegami
 Akio Iso
 Il Jang
 Adrian Jennings

Wuncheol Jeong
 Jorjeta Jetcheva
 Steven Jillings
 Seong-Soon Joo
 Tae-Gyu Kang
 Shuzo Kato
 Jeritt Kent
 Prithpal Khakuria
 Dae Kim
 Dong-Sun Kim
 Jinkyong Kim
 Youngsoo Kim
 Yunjoo Kim
 Jeffrey King
 Ryuji Kohno
 Fumihide Kojima
 Bruce Kraemer
 Raymond Krasinski
 Thomas Kuerner
 Masahiro Kuroda
 John Lampe
 Zhou Lan
 Khanh Tuan Le
 Cheolhyo Lee
 Hoosung Lee
 Hyungsoo Lee
 Myung Lee
 Yuro Lee
 Liang Li
 Sang-Kyu Lim
 Jeremy Link
 Lu Liru
 Michael Lynch
 Robert Mason
 Jeff McCullough
 Michael McInnis
 Michael McLaughlin

Charles Millet
Dino Miniutti
Siamak Mirnezami
Rishi Mohindra
Emmanuel Monnerie
Robert Moskowitz
Hamilton Moy
Peter Murray
Theodore Myers
Ken Naganuma
Chiu Ngo
Estelle Nguyen
Paul Nikolich
John Notor
Jong-Ee Oh
Okundu Omeni
Laurent Ouvry
Hyung-Il Park
Taejoon Park
Albert Petrick
Dalibor Pokrajac
Daniel Popa
Steve Pope
Clinton Powell
Richard Powell
Sridhar Rajagopal
Jayaram Ramasastry
Marc Reed
Ivan Reede

Emmanuel Riou
Richard Roberts
Craig Rodine
June Roh
Benjamin Rolfe
Didier Sagan
Kentaro Sakamoto
H. Sanderford
Naotaka Sato
Kamran Sayrafian-Pour
Timothy Schmidl
Michael Schmidt
Jean Schwoerer
Cristina Seibert
Kunal Shah
Steve Shearer
Stephen Shellhammer
Jie Shen
Shusaku Shimada
Chang Sub Shin
CheolHo Shin
Michael Sim
Jonathan Simon
Jaeseung Son
Ho-Jin Song
Paul Stadnik
Rene Struik
Chin-Sean Sum
Hui-Hsia Sung
Gu Sungi

Ronald Tabroff
Kenichi Takizawa
Hirokazu Tanaka
Larry Taylor
James Tomcik
Ichihiko Toyoda
Jana van Greunen
Hartman Van Wyk
Billy Verso
Bhupender Virk
Khurram Waheed
Joachim Walewski
Junyi Wang
Xiang Wang
Andy Ward
Scott Weikel
Nicholas West
Mark Wilbur
Ludwig Winkel
Eun Tae Won
Alan Wong
Tao Xing
Wen-Bin Yang
Yang Yang
Kazuyuki Yasukawa
Kamya Yazdandoost
Kaoru Yokoo
Mu Zhao
Bin Zhen

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

Jon Adams	Noriyuki Ikeuchi	Jayaram Ramasastry
Rajni Agarwal	Paul Isaacs	Ivan Reede
Rick Alfvin	Akio Iso	Maximilian Riegel
Nobumitsu Amachi	Atsushi Ito	Robert Robinson
Butch Anton	Raj Jain	Benjamin Rolfe
Art Astrin	Oyvind Janbu	Randall Safier
Taehan Bae	Junghoon Jee	Didier Sagan
Jay Bain	Tal Kaitz	Kazuyuki Sakoda
Anuj Batra	Shinkyo Kaku	Shigenobu Sasaki
Nancy Bravin	Piotr Karocki	Naotaka Sato
Vern Brethour	Ruediger Kays	Bartien Sayogo
Sverre Brubk	Stuart J. Kerry	Timothy Schmidt
John Buffington	Brian Kiernan	Cristina Seibert
Kiran Bynam	Yongbum Kim	Peretz Shekalm
William Byrd	Patrick W. Kinney	Jie Shen
Edgar Callaway	Bruce Kraemer	Shusaku Shimada
Dave Cavalcanti	Yasushi Kudoh	Gil Shultz
Clint Chaplin	Thomas Kurihara	Jaeseung Son
Keith Chow	Jeremy Landt	Kapil Sood
Charles Cook	Khanh Tuan Le	Robert Soranno
David Davenport	David G. K. Li	Thomas Starai
Mark Dawkins	Jan-Ray Liao	Rene Struik
Joseph Decuir	Arthur Light	Walter Struppler
Wael Diab	Lu Liru	Mark Sturza
Patrick Diamond	HaiTao Liu	Chin-Sean Sum
Thomas Dineen	William Lumpkins	Steven Thoen
Sourav Dutta	Greg Luri	Ichihiko Toyoda
Richard Edgar	Michael Lynch	Anna Urra
John Egan	Elvis Maculuba	Sunil Vadgama
Charles S. Farlow	Michael McInnis	Dmitri Varsanofiev
Avraham Freedman	Michael McLaughlin	Prabodh Varshney
Monisha Ghosh	Apurva Mody	John Vergis
James P. K. Gilb	Kenichi Mori	Billy Verso
Randall C. Groves	Peter Murray	Bhupender Virk
Michael Gundlach	Michael Newman	George Vlantis
Jose Gutierrez	Paul Nikolich	Khurram Waheed
C. Guy	John Notor	Stanley Wang
Rainer Hach	Satoshi Obara	Xiang Wang
Hiroshi Hamano	Okundu Omeni	Andy Ward
Timothy Harrington	Satoshi Oyama	Hung-Yu Wei
Rodney Hemminger	Ranjeet Patro	Andreas Wolf
Marco Hernandez	Clinton Powell	Eun Tae Won
Jin-Meng Ho	Richard Powell	Ariton Xhafa
Srinath Hostur	Venkatesha Prasad	Tao Xing
Heqing Huang	Mohammad Azizur Rahman	Yang Yang
Ichirou Ida	Sridhar Rajagopal	Oren Yuen
Tetsushi Ikegami		Daidi Zhong

When the IEEE-SA Standards Board approved this standard on 6 February 2012, it had the following membership:

Richard H. Hulett, *Chair*
John Kulick, *Vice Chair*
Robert M. Grow, *Past Chair*
Judith Gorman, *Secretary*

Masayuki Ariyoshi
 William Bartley
 Ted Burse
 Clint Chaplin
 Wael Diab
 Jean-Philippe Faure
 Alexander Gelman
 Paul Houzé

Jim Hughes
 Joseph L. Koepfinger*
 David J. Law
 Thomas Lee
 Hung Ling
 Oleg Logvinov
 Ted Olsen

Gary Robinson
 Jon Walter Rosdahl
 Sam Sciacca
 Mike Seavey
 Curtis Siller
 Phil Winston
 Howard L. Wolfman
 Don Wright

*Member Emeritus

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

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

Michelle D. Turner
IEEE Standards Program Manager, Document Development

Patricia A. Gerdon
IEEE Standards Program Manager, Technical Program Development

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-6:2017

Introduction

This introduction is not part of IEEE Std 802.15.6-2012, IEEE Standard for Local and metropolitan area networks—Part 15.6: Wireless Body Area Networks.

With the decreasing size and increasing capability of electronic devices, thanks to the Moore's Law, it was inevitable that small and portable devices would be developed for communications around human bodies. Some devices are wearable and some are implementable for medical purposes. These devices need to communicate with their remote controllers. IEEE Std 802.15.6-2012 is a standard for short-range, wireless communications in the vicinity of, or inside, a human body (but not limited to humans). It uses ISM and other bands as well as frequency bands in compliance with applicable medical and communication regulatory authorities. It allows devices to operate on very low transmit power for safety to minimize the specific absorption rate (SAR) into the body and increase the battery life. It supports quality of service (QoS), for example, to provide for emergency messaging. Since some communications can carry sensitive information, it also provides for strong security.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-15-6:2017

Contents

1. Overview	1
1.1 Scope	1
1.2 Purpose	1
2. Normative references.....	2
3. Definitions, acronyms, and abbreviations	2
3.1 Definitions	2
3.2 Special terms	7
3.3 Acronyms and abbreviations	7
4. General framework elements	10
4.1 General	10
4.2 Network topology	10
4.3 Reference model	11
4.4 Time base	12
4.5 MAC and security state diagrams	12
4.6 Security paradigm	15
5. MAC frame formats	16
5.1 Conventions	16
5.2 General format	17
5.3 Management type frames	27
5.4 Control type frames	55
5.5 Data type frames	60
5.6 MAC/PHY Capability fields	61
5.7 Information elements	65
6. MAC functions	77
6.1 General	77
6.2 Frame processing	78
6.3 Access classification and division	88
6.4 BAN creation/operation and node connection/disconnection	90
6.5 Random access	92
6.6 Improvised access and unscheduled access	98
6.7 Scheduled access and scheduled-polling access	107
6.8 Access continuation, termination, and timeout	110
6.9 MICS band communication	116
6.10 Two-hop star topology extension	121
6.11 Clock synchronization and guard time provisioning	132
6.12 Power management	140
6.13 Coexistence and interference mitigation	142
6.14 MAC/PHY capability handling/interaction and Application Specific IE usage	147
6.15 MAC sublayer parameters	148
7. Security services	151
7.1 Security association and disassociation	151
7.2 PTK creation and GTK distribution	163
7.3 Message security	165
7.4 Optional cipher functions	172

8. Narrowband PHY specification	172
8.1 Data-rate-dependent parameters	173
8.2 PLCP preamble	176
8.3 PLCP header	178
8.4 PSDU	181
8.5 Constellation mapping	185
8.6 General requirements	187
8.7 PHY layer timing	188
8.8 Transmitter specifications	190
8.9 Receiver specifications	194
9. Ultra wideband PHY specification	196
9.1 Definition of hubs and devices	196
9.2 Modes of operation	197
9.3 Rules for use of modes and options	197
9.4 Pulse shape option	198
9.5 UWB PHY frame format	198
9.6 PSDU construction	198
9.7 PHR construction	201
9.8 Synchronization header	204
9.9 IR-UWB symbol structure	206
9.10 UWB modulations	209
9.11 IR-UWB PSDU timing parameters	215
9.12 Operating frequency bands	217
9.13 Transmit spectral mask	218
9.14 IR-UWB pulse shapes	219
9.15 Type II hybrid ARQ mechanism	224
9.16 FM-UWB	228
9.17 General UWB PHY requirements	231
9.18 General radio specifications	232
10. Human body communications PHY specification	235
10.1 General	235
10.2 HBC packet structure	235
10.3 HBC transmitter	236
10.4 PLCP Preamble	237
10.5 Start frame delimiter and rate indicator	239
10.6 PHY Header	242
10.7 PSDU	244
10.8 Transmitter specifications	247
10.9 Receiver specifications	248
10.10 General requirements	248
10.11 PHY layer timing	249
Annex A (informative) Bibliography	250
Annex B (informative) Coexistence applicability guide	251
Annex C (informative) Ultra wideband	252
Annex D (informative) Features of human body communication	256

IEEE Standard for Local and metropolitan area networks— Part 15.6: Wireless Body Area Networks

IMPORTANT NOTICE: IEEE Standards documents are not intended to ensure safety, 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 is a standard for short-range, wireless communication in the vicinity of, or inside, a human body (but not limited to humans). It uses existing industrial scientific medical (ISM) bands as well as frequency bands approved by national medical and/or regulatory authorities. Support for quality of service (QoS), extremely low power, and data rates up to 10 Mbps is required while simultaneously complying with strict non-interference guidelines where needed. This standard considers effects on portable antennas due to the presence of a person (varying with male, female, skinny, heavy, etc.), radiation pattern shaping to minimize specific absorption rate (SAR) into the body, and changes in characteristics as a result of the user motions.

1.2 Purpose

The purpose is to provide an international standard for a short-range (i.e., about human body range), low power, and highly reliable wireless communication for use in close proximity to, or inside, a human body. Data rates, typically up to 10Mbps, can be offered to satisfy an evolutionary set of entertainment and healthcare services. Current *personal area networks* (PANs) do not meet the medical (proximity to human tissue) and relevant communication regulations for some application environments. They also do not support the combination of reliability, QoS, low power, data rate, and noninterference required to broadly address the breadth of body area network (BAN) applications.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used, so each referenced document is cited in 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.

ETSI EN 301 839-1, V1.3.1 (2009), Electromagnetic Compatibility and Radio Spectrum Matters (ERM); Short Range Devices (SRD); Ultra Low Power Active Medical Implants (ULP-AMI) and Peripherals (ULP-AMI-P) operating in the frequency range 402 MHz to 405 MHz; Part 1: Technical characteristics and test methods.¹

FIPS Pub 186-3 (2009), Digital Signature Standard (DSS).²

FIPS Pub 197 (2001), Advanced Encryption Standard (AES).

IEEE Std 802[®]-2001, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.^{3, 4}

IEEE Std 1363[™]-2000, IEEE Standard Specifications for Public-Key Cryptography.

ISO/IEC 10646:2003, Information Technology—Universal Multiple-Octet Coded Character Set (UCS). Amendment 1, November 2005; Amendment 2, July 2006; Amendment 3, February 2008.⁵

NIST Special Publication 800-38B (2005), Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication.⁶

NIST Special Publication 800-38C (2004), Recommendation for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality.

3. Definitions, acronyms, and abbreviations

For the purposes of this document, the following terms and definitions apply. The *IEEE Standards Dictionary: Glossary of Terms & Definitions* [B1] should be consulted for terms not defined in this clause.^{7, 8}

3.1 Definitions

abbreviated address: A one-octet identifier (ID) selected as an address for a node, hub, or body area network (BAN) in frame exchanges.

active state: An internal power management state that is ready for frame reception and transmission.

¹ ETSI publications are available from the European Telecommunications Standards Institute (<http://www.etsi.org>).

² FIPS publications are available from the National Technical Information Service (<http://www.ntis.gov/>).

³ IEEE publications are available from The Institute of Electrical and Electronics Engineers (<http://standards.ieee.org>).

⁴ The IEEE standards or products referred to in Clause 2 are trademarks owned by The Institute of Electrical and Electronics Engineers, Incorporated.

⁵ ISO/IEC publications are available from the ISO Central Secretariat (<http://www.iso.org/>). ISO publications are also available in the United States from the American National Standards Institute (<http://www.ansi.org/>).

⁶ NIST publications are available from the National Institute of Standards and Technology (<http://www.nist.gov/>).

⁷ The *IEEE Standards Dictionary: Glossary of Terms & Definitions* is available at <http://shop.ieee.org/>.

⁸ The numbers in brackets correspond to those of the bibliography in Annex A.

active superframe: A superframe in which frame transmission typically occurs within the body area network (BAN) of the hub announcing such a superframe.

allocation: One or more time intervals that a node or a hub obtains using an access method for initiating one or more frame transactions. An allocation comprises one or more allocation intervals. Reference to allocation of a node means that the node is the sender or recipient in the allocation.

allocation interval: A continuous time interval in an allocation, comprising one or more consecutive allocation slots. Reference to allocation interval of a node means that the node is the sender or recipient in the allocation interval.

allocation slot: A time unit used to designate the lengths of medium access related time intervals, such as beacon period (superframe) and allocation interval.

beacon: A frame transmitted by a hub to facilitate network management, such as the coordination of medium access and power management of the nodes in the body area network (BAN) of the hub, and to facilitate clock synchronization therein.

beacon period: A repetitive time interval to which medium access is referenced and in which a beacon is transmitted when appropriate, comprising the same number of time units (allocation slots) of equal duration.

bilink: A communications link for transfer of management and data traffic from a hub to a node or/and vice versa.

bilink allocation: An allocation with allocation interval(s) in which a hub or a node initiates one or more frame transactions to transmit management and data traffic to a node or a hub, respectively, and the recipient returns acknowledgment if required, with the provision that the node initiates frame transaction(s) only after receiving a poll from the hub.

connected node: A node that has a connection with a hub.

connection: A relationship between a node and a hub in a body area network (BAN), substantiated by a connected node identification assigned to the node by the hub and by a wakeup arrangement between them, and optionally by one or more scheduled allocations or unscheduled bilink allocations between them.

contended allocation: A non-reoccurring time interval, within a random access phase (RAP) or a contention access phase (CAP), that a node obtains using random access for initiating a frame transaction. A contended allocation is an uplink allocation, suitable for servicing “unpredictable” uplink traffic (for example, due to data rate variations and/or channel impairments).

contention access: An access method, based on carrier sense multiple access with collision avoidance (CSMA/CA) or slotted Aloha access but not both, whereby a node obtains a time interval in a contention access phase (CAP) for initiating one or more frame transactions. As an access method, contention access is synonymous with random access.

contention access phase (CAP): A time span set aside by a hub and announced via a preceding non-beacon frame for contention access to the medium by the nodes in the body area network (BAN) of the hub.

downlink: A communications link for transfer of management and data traffic from a hub to a node.

downlink allocation: An allocation with allocation interval(s) in which a hub initiates one or more frame transactions to transmit management and data traffic to a node and the node returns acknowledgment if required.

entity authentication: Corroboration of the identity of the node or the hub in a security association procedure.

exclusive access phase (EAP): A time span set aside by a hub in a beacon period (superframe) for transfer of the traffic of the highest user priority (UP) (for emergency or medical implant event report).

frame: An uninterrupted sequence of octets delivered by the medium access control (MAC) sublayer to the physical (PHY) layer, or vice versa, within a node or a hub.

hub: An entity that possesses a node's functionality and coordinates the medium access and power management of the nodes in its body area network (BAN).

hub identifier (HID): An abbreviated address of a hub.

improvised access: An access method, based on impromptu polling or posting by a hub, whereby a hub grants to a node or itself a polled or posted allocation, typically outside scheduled downlink and uplink allocations, for initiating one or more frame transactions by the node or hub, respectively.

inactive state: An internal power management state that is not ready for frame reception and transmission.

inactive superframe: A superframe in which no frame transmission occurs within the body area network (BAN) of the hub announcing such a superframe.

information element (IE): An optional part, with variable but self-identifiable length, of some frames.

managed access phase (MAP): A time span set aside by a hub for improvised access, scheduled access, and unscheduled access to the medium by the hub and the nodes in the body area network (BAN) of the hub.

master key (MK): A secret bit string activated or established between a node and a hub in a security association procedure and used to create a pairwise temporal key (PTK) shared between them.

message authentication: Corroboration of the origin of a message in a message transfer.

multi-periodic (m-periodic) allocation: A scheduled allocation or an unscheduled bilink allocation that has allocation intervals reoccurring in every m th beacon period (superframe) with m being an integer larger than one. An m -periodic scheduled allocation is an uplink allocation, a downlink allocation, or a bilink allocation, suitable for servicing low duty cycle periodic or quasi-periodic traffic on a committed schedule. An m -periodic unscheduled allocation is a bilink allocation, suitable for servicing low duty cycle periodic or quasi-periodic traffic on a best-effort basis.

node: An entity that contains a medium access control (MAC) sublayer, a physical (PHY) layer, and that optionally provides security services.

node identifier (NID): An abbreviated address of a single node or of a logical group of nodes.

nonce: A number that is unique per instantiation of a cryptography protocol as part of a measure to thwart cryptanalytic and other cryptographic attacks.

non-secure frame: A term that is interchangeable with unsecured frame.

one-periodic (1-periodic) allocation: A scheduled allocation that has allocation intervals reoccurring in every beacon period (superframe), or an unscheduled bilink allocation that has allocation intervals reoccurring in every beacon period (superframe) or in round-robin together with the allocation intervals of other 1-periodic unscheduled bilink allocations. A 1-periodic scheduled allocation is an uplink allocation, a

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

downlink allocation, or a bilink allocation, suitable for servicing high duty cycle periodic or quasi-periodic traffic on a committed schedule. A 1-periodic unscheduled bilink allocation is a bilink allocation, suitable for servicing high duty cycle periodic or quasi-periodic traffic on a best-effort basis.

pairwise temporal key (PTK): A secret bit string shared between a node and a hub and used to secure frames transferred between them.

pairwise temporal key (PTK) creation: A procedure used to create a PTK between a node and a hub based on a master key (MK) shared between them, and to confirm possession of a shared MK between the node and the hub.

poll: A control type frame or its variant sent by a hub to grant an immediate polled allocation to the addressed node or to inform the node of a future poll or post.

polled allocation: A non-reoccurring time interval that a hub grants to a node using polling access for initiating one or more frame transactions by the node. A polled allocation is an uplink allocation interval, suitable for servicing “ordinary,” “unexpected,” or “extra” uplink traffic (for example, due to data rate variations and/or channel impairments). A polled allocation is also called a *polled allocation interval*.

polling access: An access method, based on impromptu or scheduled polling by a hub, whereby a hub grants to a node a polled allocation for initiating one or more frame transactions by the node.

post: A management or data type frame sent by a hub to a node within its body area network (BAN). A post starts a posted allocation.

posted allocation: A non-reoccurring time interval that a hub grants to itself using posting access for initiating a frame transaction. A posted allocation is a downlink allocation interval, suitable for servicing “unexpected” or “extra” downlink traffic (for example, due to network management needs, data rate variations, and/or channel impairments).

posting access: An access method, based on impromptu or scheduled posting by a hub, whereby a hub grants to itself a posted allocation, typically outside scheduled uplink allocations, for initiating one or more frame transactions by the hub.

random access: An access method, based on carrier sense multiple access with collision avoidance (CSMA/CA) or slotted Aloha access but not both, whereby a node obtains a time interval in a random access phase (RAP) for initiating one or more frame transactions.

random access phase (RAP): A time span set aside by a hub and announced via a beacon frame for random access to the medium by the nodes in the body area network (BAN) of the hub.

relayed node: A node that communicates with a hub through another node.

relaying node: A node through which another node communicates with a hub.

scheduled access: An access method, based on advance reservation and committed scheduling, whereby a node and a hub obtain scheduled reoccurring time intervals for initiating frame transactions.

scheduled allocation: One or more scheduled reoccurring time intervals that a node and a hub obtains using scheduled access for initiating frame transactions. A scheduled allocation is an uplink allocation, a downlink allocation, or a bilink allocation, suitable for servicing high or low duty cycle periodic or quasi-periodic traffic on a committed schedule.

scheduled-polling access: A combination of scheduled access and polling access, whereby a node and a hub obtain scheduled reoccurring time intervals, wherein the hub grants to the node and/or itself non-reoccurring time intervals for initiating frame transactions on uplink and/or downlink.

secure frame: A term that is interchangeable with secured frame.

secured communication: Exchange of secured frames.

secured frame: A frame that is secured with authenticity, integrity, confidentiality if required, and replay protection.

security association: A procedure used to identify a node and a hub to each other and establish a new master key (MK) shared between them or activate an existing MK pre-shared between them.

star network: A logical network partition comprising a hub and zero or more nodes whose medium access and power management are coordinated by the hub.

superframe: A term that is interchangeable with beacon period used especially when no beacons are transmitted.

type-I polled allocation: A polled allocation the length of which is specified in terms of the duration of time granted for transmission.

type-I polling access: Polling access that provides type-I polled allocations.

type-II polled allocation: A polled allocation the length of which is specified in terms of the number of frames granted for transmission.

type-II polling access: Polling access that provides type-II polled allocations.

unconnected node: A node that does not have a connection with a hub in a body area network (BAN).

unscheduled access: A combination of best-effort scheduled access and polling access, whereby a node and a hub obtain unscheduled reoccurring time intervals, wherein the hub grants to the node and/or itself non-reoccurring time intervals for initiating frame transactions in an uplink and/or downlink.

unscheduled bilink allocation: One or more unscheduled reoccurring time intervals that a node and a hub obtains using unscheduled access for initiating frame transactions. An unscheduled bilink allocation is a bilink allocation, suitable for servicing high or low duty cycle periodic or quasi-periodic traffic in an uplink and/or downlink on a best-effort basis.

unsecured communication: Exchange of unsecured frames.

unsecured frame: A frame that is not secured with authenticity, integrity, confidentiality, or replay protection.

uplink allocation: An allocation with allocation interval(s) in which a node initiates one or more frame transactions to transmit management and data traffic to a hub and the hub returns acknowledgment if required.

uplink: A communications link for transfer of management and data traffic from a node to a hub.

3.2 Special terms

association: A term synonymous to security association in absence of the qualifier “security.”

appropriate: Subject to the rules specified in Clause 5, Clause 6, and Clause 7.

frame transaction: All management or data type frames of the same frame subtype, and an acknowledgment frame if required, that are separated in time by an appropriate interframe space and transmitted between a node and a hub. Examples of a frame transaction: (a) a management type frame and an immediate acknowledgment frame; (b) all data type frames of the same frame subtype after a block acknowledgment frame, and the block acknowledgment frame that follows; (c) a management or data type frame, which is followed by another management or data type frame of a different frame subtype; (d) a management or data type frame, which is separated by more than an appropriate interframe space by another management or data type frame of the same frame subtype.

implant: An entity that is placed inside a human body for medical purposes.

low power low duty cycle (LP/LDC): Power, duty cycle, and transmission (transmissions/per hour) limits defined by standards and regulations for implants transmitting within the band 403.5 MHz to 403.8 MHz without coordination with a hub.

medical implant event: A term referenced in regulatory documents governing the use of the Medical Device Radiocommunication Service (MedRadio), which includes the Medical Implant Communications Service (MICS) from 402 MHz to 405 MHz.

3.3 Acronyms and abbreviations

AES	Advanced Encryption Standard
ARQ	automatic repeat request
AWGN	additive white Gaussian noise
B-Ack	block acknowledgment
BAN	body area network
BCH code	Bose, Ray-Chaudhuri, Hocquenghem code
BPSK	binary phase shift keying
CAP	contention access phase
CBC	cipher block chaining
CCA	clear channel assessment
CCM	counter mode for message encryption and cipher block chaining (CBC) mode for message authentication
CMAC	(block) cipher-based message authentication code algorithm
CP	contention probability
CP-BFSK	continuous-phase binary frequency shift keying
CRC	cyclic redundancy check

CSMA/CA	carrier sense multiple access with collision avoidance
CW	contention window (for carrier sense multiple access with collision avoidance)
D8PSK	differential 8-phase-shift keying
DBPSK	differential (or differentially encoded) binary phase-shift keying
DQPSK	differential (or differentially encoded) quadrature phase-shift keying
DRF	data rate field
EAP	exclusive access phase
ED	energy detection
EFC	electric field communication
EIRP	equivalent isotropically radiated power
EUI	extended unique identifier
EVM	error vector magnitude
FCS	frame check sequence
FEC	forward error correction code
FM-UWB	frequency modulation ultra-wideband
FSDT	frequency selective digital transmission
FS-Spreader	frequency selective spreader
G-Ack	group acknowledgment
GF	Galois field
GFSK	Gaussian frequency-shift keying
GPPM	group pulse position modulation
GT	guard time
GTK	group temporal key
HARQ	hybrid automatic repeat request
HBC	human body communications
HCS	header check sequence
HID	hub identifier
HME	hub management entity
I-Ack	immediate acknowledgment
IE	information element
IR-UWB	impulse radio ultra-wideband
ISM	industrial scientific medical
ISO	International Organization for Standardization

IEEE Std 802.15.6-2012
 IEEE Standard for Local and metropolitan area networks—
 Part 15.6: Wireless Body Area Networks

KCK	key confirmation key
KMAC	key message authentication code
L-Ack	late acknowledgment (acknowledgment later)
LFSR	linear feedback shift register
LP/LDC	low power low duty cycle
LSB	least significant bit
MAC	media access control
MAP	managed access phase
MCU	micro controller unit
MIC	message integrity code
MICS	medical implant communications service
MIFS	minimum interframe space
MK	master key
MPDU	medium access control protocol data unit
MSB	most significant bit
MSDU	media access control service data unit
MUX	multiplexer
N-Ack	no acknowledgment
NID	node identifier
NME	node management entity
OSI	open systems interconnection
PER	packet error rate
PHR	physical layer header
PHY	physical or physical layer
PLCP	physical layer convergence protocol
PN	pseudo-random noise
PPDU	physical layer protocol data unit
PPRF	peak pulse repetition frequency
PRF	pulse-repetition-frequency
PSD	power spectral density
PSDU	physical layer service data unit
PSK	phase shift keying
PTK	pairwise temporal key

QPSK	quadrature phase shift keying
RAP	random access phase
RI	rate indicator
RX	receive or reception
S2P	serial-to-parallel
SAP	service access point
SAR	specific absorption rate
SF	spreading factor
SFD	start-of-frame delimiter
SHR	synchronization header
SIFS	short interframe spacing
SRRC	square-root raised cosine
TK	temporal key
TX	transmit or transmission
UP	user priority
UWB	ultra wideband

4. General framework elements

4.1 General

This clause provides the basic framework of nodes and hubs. The framework serves as a prerequisite to supporting the functions of nodes and hubs and their interactions specified later in detail. It covers the following aspects: the network topology used for medium access, the reference model used for functional partitioning, the time base used for access scheduling, the state diagrams used for frame exchange, and the security paradigm used for message protection.

4.2 Network topology

All nodes and hubs are to be organized into logical sets, referred to as *body area networks* (BANs) in this specification, and coordinated by their respective hubs for medium access and power management as illustrated in Figure 1. There is to be one and only one hub in a BAN, whereas the number of nodes in a BAN is to range from zero to mMaxBANSize. In a one-hop star BAN, frame exchanges are to occur directly between nodes and the hub of the BAN. In a two-hop extended star BAN, the hub and a node are to exchange frames optionally via a relay-capable node.

Access coordination at the MAC sublayer between BANs is not specified in this standard. Optional mechanisms for coexistence and interference mitigating between adjacent or overlapping BANs are provided (in 6.13). Nodes referenced in this standard are in the context of a given BAN, unless noted otherwise.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 1—Network topology

4.3 Reference model

All nodes and hubs are internally partitioned into a physical (PHY) layer and a medium access control (MAC) sublayer, in accordance with the IEEE 802[®] reference model, as shown in Figure 2. Direct communications between a node and a hub are to transpire at the PHY layer and MAC sublayer as specified in this standard; the PHY layer and MAC sublayer of a node or a hub are to use only one operating channel at any given time. Message security services are to occur at the MAC sublayer, and security key generations are to take place inside and/or outside the MAC sublayer.




Figure 2—Reference model

Within a node or a hub, the MAC provides its service to the MAC client (higher layer) through the MAC service access point (SAP) located immediately above the MAC sublayer, while the PHY provides its service to the MAC through the PHY SAP located between them. On transmission, the MAC client passes MAC service data units (MSDUs) to the MAC sublayer via the MAC SAP, and the MAC sublayer passes MAC frames (also known as MAC protocol data units or MPDUs) to the PHY layer via the PHY SAP. On reception, the PHY layer passes MAC frames to the MAC sublayer via the PHY SAP, and the MAC sublayer passes MSDUs to the MAC client via the MAC SAP. Both MAC SAP and PHY SAP are not exposed and their specifications are beyond the scope of this standard.

There may be a logical node management entity (NME) or hub management entity (HME) that exchanges network management information with the PHY and MAC as well as with other layers. The HME is a superset of the NME in terms of the management functionality they each support. However, the presence of the NME or HME and the partitioning between the NME or HME and the MAC or the PHY is not mandated, nor is the behavior of the NME or HME specified, in this standard.

4.4 Time base

All nodes and hubs are to establish a time reference base, as shown in Figure 3, if their medium access is to be scheduled in time, where the time axis is divided into beacon periods (superframes) of equal length and each beacon period (superframe) is composed of allocation slots of equal length and numbered from 0, 1, ..., s , where $s \leq 255$. An allocation interval may be referenced in terms of the numbered allocation slot comprising it, and a point of time may be referenced in terms of the numbered allocation slots preceding or following it as appropriate.




Figure 3—Time reference base

If time reference is needed for access scheduling in its BAN, the hub is required to choose the boundaries of beacon periods (superframes) and hence of the allocation slots therein. In beacon mode operation for which beacons are transmitted, the hub needs to communicate such boundaries by transmitting beacons at the start or other specified locations of beacon periods (superframes), and optionally timed frames (T-Poll frames) containing their transmit time relative to the start time of current beacon period (superframe). In non-beacon mode operation for which beacons are not transmitted but time reference is needed, the hub is required to communicate such boundaries by transmitting timed frames (T-Poll frames) also containing their transmit time relative to the start time of current superframe.

A node requiring a time reference in the BAN needs to derive and recalibrate the boundaries of beacon periods (superframes) and allocation slots from reception of beacons or/and timed frames (T-Poll frames).

A frame transmission may span more than one allocation slot, starting or ending not necessarily on an allocation slot boundary.

4.5 MAC and security state diagrams

All nodes and hubs are to go through certain stages, i.e., states, at the MAC level before they exchange user (MAC client) data, as shown in Figure 4, where frames permitted or required to exchange between a node and a hub at each state are also indicated. State classification and transition are defined with respect to a pair of a node and a hub, but is often referenced in the name of the node only. Broadcast or multicast frames, such as beacons, are frame exchanges between a hub and a group of nodes but not between a hub and a node; their transmissions are not dictated by the pairwise state diagrams specified in this subclause.

4.5.1 Secured communication

A node and a hub are to follow the MAC state diagram Figure 4(a) for secured communication if either of them requires secured frame exchanges with the other.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 4—MAC and security state diagrams

4.5.1.1 Orphan state

At this state, the node does not have any relationship with the hub for secured communication. It is the state that the node initially enters in relation with the hub. The node and the hub are not allowed to transmit any frames to each other—other than Security Association and control unsecured frames. The node is allowed to exchange Security Association frames with the hub to establish a security association, i.e., to activate a pre-shared master key (MK) or generate a new MK, and to authenticate with each other if so required, thereby transitioning to next state, *Associated* state. However, if the node and the hub fail to activate or establish a shared MK, they are not allowed to advance to *Associated* state.

4.5.1.2 Associated state

At this state, the node is associated, i.e., holds a shared MK, with the hub for their pairwise temporal key (PTK) creation. The node and the hub are not allowed to transmit any frames to each other—other than Security Disassociation and PTK unsecured frames, as well as control unsecured frames. The node is allowed to exchange PTK frames with the hub to establish a secured relationship, i.e., to confirm possession of a shared MK and to create a PTK, thereby transitioning to next state, *Secured* state. However, if the node and the hub fail to create a PTK, they are not allowed to advance to *Secured* state. If their MK is missing or invalid during their PTK creation, they are required to move back to *Orphan* state. To repeal the security association and hence the current MK, either the node or the hub is allowed to send a Security Disassociation frame to the other, thus moving back to *Orphan* state.

4.5.1.3 Secured state

At this state, the node is secured, i.e., holds a PTK, with the hub for message security, i.e., for secured frame exchanges. The node and the hub are not allowed to transmit any frames to each other—other than Security Disassociation, Connection Request, and Connection Assignment secured frames, as well as control secured or unsecured frames depending on whether authentication of control type frames was selected during their association. The node is allowed to exchange Connection Request and Connection Assignment frames with the hub to establish a connection, thereby transitioning to the next and final state, *Connected* state. However, if the node and the hub fail to establish a connection, they are not allowed to advance to *Connected* state. If their PTK is missing or invalid or the node's *Connected_NID* is lost, they are required to move back to *Associated* state. To repeal the security association and hence the current MK, either the node or the hub is allowed to send a Security Disassociation frame to the other, thus moving back to *Orphan* state.

4.5.1.4 Connected state

At this state, the node is connected, i.e., holds an assigned *Connected_NID*, a wakeup arrangement, and optionally one or more scheduled and unscheduled allocations with the hub for abbreviated node addressing, desired wakeup, and optionally scheduled and unscheduled access. The node and the hub are allowed to transmit any secured frames to each other—other than Security Association secured frames, but are not allowed to transmit any unsecured frames to each other—other than control unsecured frames if authentication of control type frames was not selected during their association. If their MK is missing or invalid during a PTK recreation, the node and the hub are required to move back to *Orphan* state. If their PTK is missing or invalid, the node and the hub are required to move back to *Associated* state. To repeal the security association and hence the current MK, either the node or the hub is allowed to send a Security Disassociation frame to the other, thus moving back to *Orphan* state. To repeal the connection and hence the node's *Connected_NID*, wakeup arrangement, and scheduled and unscheduled allocations if any, either the node or the hub is allowed to send a Disconnection frame to the other, thereby moving back to *Associated* state.

4.5.2 Unsecured communication

A node and a hub are required to follow the MAC state diagram Figure 4(b) for unsecured communication if neither of them requires secured frame exchanges with the other.

4.5.2.1 Orphan state

At this state, the node does not have any relationship with the hub for unsecured communication. It is the state that the node initially enters in relation with the hub. The node and the hub are not allowed to transmit any frames to each other—other than Connection Request, Connection Assignment, and control unsecured frames. The node is allowed to exchange Connection Request and Connection Assignment frames with the hub to establish a connection, thereby transitioning to next and final state, *Connected* state. However, if the node and the hub fail to establish a connection, they are not allowed to advance to *Connected* state.

4.5.2.2 Connected state

At this state, the node is connected, i.e., holds an assigned *Connected_NID*, a wakeup arrangement, and optionally one or more scheduled and unscheduled allocations with the hub for abbreviated node addressing, desired wakeup, and optionally scheduled and unscheduled access. The node and the hub are allowed to transmit any unsecured frames to each other—other than Security Association, Security Disassociation, and PTK frames, but are not allowed to transmit secured frames to each other. To repeal the

connection and hence the node's *Connected_NID*, wakeup arrangement, and scheduled and unscheduled allocations if any, either the node or the hub is allowed to send a Disconnection frame to the other, thereby moving back to *Orphan* state.

To change to secured communications between them, the node and the hub are required to disconnect from each other, thus moving back to *Orphan* state and then following the state diagram Figure 4(a) for secured communication.

4.6 Security paradigm

All nodes and hubs are to choose three security levels in this standard, as follows:

- Level 0—unsecured communication. At this level, messages are transmitted in unsecured frames, which provide no measures for message authenticity and integrity validation, confidentiality and privacy protection, and replay defense.
- Level 1—authentication but not encryption. At this level, messages are transmitted in secured authenticated but not encrypted frames, which provide measures for message authenticity and integrity validation and replay defense but not confidentiality and privacy protection.
- Level 2—authentication and encryption. At this level, messages are transmitted in secured authenticated and encrypted frames, which provide measures for message authenticity and integrity validation, confidentiality and privacy protection, and replay defense.

During association, a node and a hub need to jointly select a security level suitable for their subsequent secured frame exchanges, and whether to require authentication of control type frames, based on their respective security requirements and certain information specific to each other.

For unicast secured communication, the node and the hub need to activate a pre-shared MK or establish a new MK via an unauthenticated or authenticated association, and create a PTK via a PTK creation procedure. For multicast secured communication, the hub needs to distribute a GTK to the corresponding multicast group in a unicast manner.

The node and the hub are to follow the security structures shown in Figure 5 to perform security key generations and provide message security services. A “session” indicated in this figure refers to a time span in which a temporal key (TK) remains valid. The length of the session is determined by the security policy governing data transfers between the two communicating parties. It is further limited by the technical restrictions on the reuse of the same TK over successive messages (i.e., MAC frames in this specification).




Figure 5—Security hierarchy

5. MAC frame formats

5.1 Conventions

A MAC frame is an ordered sequence of fields delivered to or from the physical layer service access point (PHY SAP). Each figure in Clause 5 depicts the fields contained in a MAC frame, or a part thereof, from left to right in the transmit order, with fields that are optional or selectively absent drawn in dashes where possible. The transmit order starts from top to bottom if the fields are depicted in multiple rows, symbolically linked by dashes extended to the right of one row and to the left of next row, respectively. Also indicated is the number of octets contained in each field along with the corresponding octet transmit order, on top of the field, as illustrated in Figure 6.




Figure 6—Example of fields aligned with octet boundaries

Unless otherwise noted, an atomic field, i.e., a field that is not in turn comprised of other fields, denotes a numerical value encoded in unsigned binary. If such a field contains F octets ($F > 1$), octet 0 is the octet containing the least significant bits (LSBs) of that field and is the first octet transmitted of the field, whereas octet $F-1$ is the octet containing the most significant bits (MSBs) of that field and is the last octet transmitted of the field. The octet order is indicated as “L-R,” i.e., from left to right, above a multi-octet non-atomic field.

In a field encoding an EUI-48, it is set to a string of six octets per Clause 9 of IEEE Std 802-2001.⁹ Octets 0–2 of the field are set sequentially to the first three octets of the string that specify an organizationally unique identifier (OUI), with bit 0 of octet 0 being the individual/group (I/G) address bit. Octets 3–5 are set to the last three octets selected as an extension identifier (ID) by the organization identified by the OUI.

In a figure that depicts certain fields not aligned with octet boundaries, the number of bits and the corresponding bit order of encoding are shown instead for each field in the figure, as illustrated in Figure 7. Bits are ordered continually across the fields that are not aligned with octet boundaries, from left to right, starting from bit 0, i.e., the LSB of the bits comprising these fields. Bit numbering restarts from zero in fields located on octet boundaries.




Figure 7—Example of fields not aligned with octet boundaries

Each field is defined, or set, based on the perspective of the node or the hub that is sending the frame containing that field, referred to as the *sender*. It is parsed based on its definition by the hub or the node intended to receive the frame containing it, referred to as the *recipient* or the *intended recipient*.

⁹ Information on references can be found in Clause 2.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Reserved fields are set to zero on transmission and ignored on reception. If some values in a field are reserved, the field is not set to any of those reserved values on transmission. Unless otherwise noted, fields that are set to reserved values or defined based on other fields that are set to reserved values are ignored on reception.

MAC constants are referenced through parameters as listed in 6.15. These parameters are denoted with a preceding “m” if they are PHY independent or with a preceding “p” if they are PHY dependent.

5.2 General format

A MAC frame consists of a fixed-length MAC header, a variable-length MAC frame body, and a fixed-length Frame Check Sequence (FCS) field as shown in Figure 8. The MAC frame body has an octet length L_{FB} such that $0 \leq L_{FB} \leq pMaxFrameBodyLength$, and is present only if it has a nonzero length.




Figure 8—MAC frame format

5.2.1 MAC Header

The MAC Header is formatted as shown in Figure 9.




Figure 9—MAC Header format

5.2.1.1 Frame Control

The Frame Control is formatted as shown in Figure 10.




Figure 10—Frame Control format

5.2.1.1.1 Protocol Version

The Protocol Version field is set to zero for this revision of the standard. All other values are reserved. This field is invariant in size and place across all revisions of this standard.

5.2.1.1.2 Acknowledgment (Ack) Policy

The Ack Policy field is set according to Table 1 to indicate the acknowledgement requirement of the current frame.

Table 1—Acknowledgement Policy field encoding

Field value b2 b1	Acknowledgment requirement
00	No acknowledgment (N-Ack) or group acknowledgment (G-Ack)
01	Immediate acknowledgment (I-Ack)
11	Block acknowledgment later (L-Ack)
10	Block acknowledgment (B-Ack)

The group acknowledgment (G-Ack) value is applicable to frames sent to a hub and of frame type set to data and frame subtype set to mG-AckDataSubtype.

5.2.1.1.3 Security Level

The Security Level field is set according to Table 2 such that it indicates the security level of the current frame.

Table 2—Security Level field encoding

Field value b4 b3	Security level of current frame
00	Level 0—frame not secured
01	Level 1—frame authenticated but not encrypted
10	Level 2—frame authenticated and encrypted
11	Reserved

5.2.1.1.4 Temporal key (TK) Index

The TK Index field is set as follows to indicate the PTK or GTK being used to secure the current frame:

- In frames secured by a PTK, it is set to the value of the PTK Index field in the PTK frames that were exchanged in creating the PTK.
- In frames secured by a GTK, it is set to the value of the GTK Index field in the GTK frame that was exchanged in distributing the GTK.

The TK Index field is reserved in unsecured frames.

5.2.1.1.5 BAN Security/Relay

The BAN Security/Relay field is set as follows:

- a) In beacon, Poll, and T-Poll frames sent by a hub, it is used as a BAN Security field, which is set to one if this hub accepts only secured communication with it as described in 4.5.1, or is set to zero if this hub accepts either secured or unsecured communication with it as described in 4.5.1 or 4.5.2.
- b) In frames sent to or from a relaying node in a two-hop extended star network communication, it is used as a Relay field, which is set to one.
- c) In all other frames, it is reserved.

5.2.1.1.6 Ack Timing/EAP Indicator/First Frame On Time

The Ack Timing/EAP Indicator/First Frame On Time field is set as follows:

- a) In beacon frames, it is used as an EAP Indicator field, which is set to one if exclusive access phase 1 (EAP 1) in the current (with beacon shifting not enabled) or next (with beacon shifting enabled) beacon period has a nonzero length, or is set to zero otherwise.
- b) In non-beacon management or data type frames or Poll frames sent by a hub to a node, it is used as a First Frame On Time field, which is set to one if this is the first frame sent by the hub to the node at the start of an allocation interval of a scheduled allocation assigned to the node, or is set to zero otherwise.
- c) In non-beacon management or data type frames with the Ack Policy field of the MAC header set to I-Ack or B-Ack sent by a node to a hub, it is used as an Ack Timing field, which is set to one if the following acknowledgment (I-Ack, B-Ack, I-Ack+Poll, or B-Ack+Poll) frame is to include a timestamp in its frame payload, or is set to zero otherwise. The timestamp encodes the start time of the acknowledgment frame transmission based on the hub's clock.
- d) In I-Ack, B-Ack, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as an Ack Timing field, which is set to one if the frame includes a timestamp in the frame payload, or is set to zero otherwise.
- e) In all other frames, it is reserved.

5.2.1.1.7 Frame Subtype

The Frame Subtype field is set to indicate the subtype of the current frame of a given type according to Table 3. The name of the frame subtype of a frame is used as the name of the frame throughout Clause 2 through Clause 7.

5.2.1.1.8 Frame Type

The Frame Type field is set to indicate the type of the current frame according to Table 3.

Table 3—Frame Type and Frame Subtype field encoding

Frame Type value b5 b4	Frame Type name	Frame Subtype value b3 b2 b1 b0	Frame Subtype name
00	Management	0000	Beacon
00	Management	0001	Reserved
00	Management	0010	Security Association
00	Management	0011	Security Disassociation
00	Management	0100	PTK
00	Management	0101	GTK
00	Management	0110–0111	Reserved
00	Management	1000	Connection Request
00	Management	1001	Connection Assignment
00	Management	1010	Disconnection
00	Management	1011–1110	Reserved
00	Management	1111	Command
01	Control	0000	I-Ack
01	Control	0001	B-Ack
01	Control	0010–0011	Reserved
01	Control	0100	I-Ack+Poll
01	Control	0101	B-Ack+Poll
01	Control	0110	Poll
01	Control	0111	T-Poll
01	Control	1000–1101	Reserved
01	Control	1110	Wakeup
01	Control	1111	B2
10	Data	0000	User Priority 0 or Allocation Mapped Data Subtype
10	Data	0001	User Priority 1 or Allocation Mapped Data Subtype
10	Data	0010	User Priority 2 or Allocation Mapped Data Subtype
10	Data	0011	User Priority 3 or Allocation Mapped Data Subtype
10	Data	0100	User Priority 4 or Allocation Mapped Data Subtype
10	Data	0101	User Priority 5 or Allocation Mapped Data Subtype
10	Data	0110	User Priority 6 or Allocation Mapped Data Subtype
10	Data	0111	Emergency
10	Data	1000–1111	Allocation Mapped Data Subtype
11	Reserved	0000–1111	Reserved

5.2.1.1.9 More Data

The More Data field is set as follows:

- a) In management or data type frames sent by a node to a hub,
 - 1) it is set to zero if the node has no management or data type frame pending for transmission to the hub, except for a possible retransmission of the current frame, or
 - 2) it is set to one if the node has at least a management or data type frame pending for transmission or retransmission to the hub.
- b) In I-Ack and B-Ack frames sent by a node to a hub,
 - 1) it is set to zero if the node has no management or data type frame pending for transmission or retransmission to the hub, or
 - 2) it is set to one if the node has at least one management or data type frames pending for transmission or retransmission to the hub.
- c) In non-beacon management or data type frames sent by a hub to a node,
 - 1) it is set to zero if the hub has no poll or post pending for transmission to the node, except for a possible retransmission of the current frame, or
 - 2) it is set to one if the hub has at least a poll or post pending for transmission to the node.
- d) In I-Ack and B-Ack frames sent by a hub to a node,
 - 1) it is set to zero if the hub is not to send a poll or post to the node immediately after the current allocation interval is reclaimed or ended, or
 - 2) it is set to one if the hub is to send a poll or post to the node immediately after the current allocation interval is reclaimed or ended.
- e) In Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node,
 - 1) if Access Mode = 0 indicating that the hub is operating in beacon or non-beacon mode with superframes,
 - i) it is set to zero if via the current frame the hub grants a type-I (immediate) polled allocation that starts pSIFS thereafter and ends at a time as encoded in the Poll-Post Window and Next fields therein, or
 - ii) it is set to one if via the current frame the hub grants no polled allocation but is to send a (future) poll or post to the node at a time as encoded in the Poll-Post Window and Next fields thereof;
 - 2) if Access Mode = 1 indicating that the hub is operating in non-beacon mode without superframes, and that via the current frame the hub grants a type-II (immediate) polled allocation that starts pSIFS thereafter and allows for the node to send up to a specified number of frames as encoded in the Poll-Post Window field thereof, it is reserved.
- f) In all other frames, it is reserved.

5.2.1.1.10 Last Frame/Access Mode/B2

The Last Frame/Access Mode/B2 field is set as follows:

- a) In beacon frames, it is used as a B2 field, which is set to one if a B2 frame is to be transmitted in the current beacon period, or is set to zero otherwise.

- b) In non-beacon management or data type frames, it is used as a Last Frame field, which is
 - 1) set to zero if the sender is likely to send another frame in the current allocation interval after the current frame, or
 - 2) set to one if the sender is definitely not to send another frame in the current allocation interval after the current frame.
- c) In I-Ack, B-Ack, I-Ack+Poll, B-Ack+Poll, Poll, and T-Poll frames sent by a hub to a node, it is used as an Access Mode field, which is
 - 1) set to zero if the hub that is the sender or recipient of the current frame is operating in beacon or non-beacon mode with superframes, or
 - 2) set to one if the hub is operating in non-beacon mode without superframes.
- d) In all other frames, it is reserved.

5.2.1.1.11 Sequence Number/Poll-Post Window

The Sequence Number/Poll-Post Window field is set as follows:

- a) In beacon frames, it is used as a Sequence Number field, which is
 - 1) incremented by one, modulo 256, from its value applicable to the last beacon period (superframe), active or inactive;
 - 2) also referred to as the sequence number of the current beacon (superframe) or of the current beacon period (superframe).
- b) In non-beacon management type frames, it is used as a Sequence Number field, which is
 - 1) set to zero if the frame is the first non-beacon management type frame sent to a recipient or a group of recipients;
 - 2) incremented by one, modulo 256, from its value in the frame that was of the same frame type and addressed to the same recipient(s); or
 - 3) kept with the same value if the frame is being retransmitted to the same recipient(s).
- c) In Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as a Poll-Post Window field, which is set as follows:
 - 1) If Access Mode = 0 indicating that the hub is operating in beacon or non-beacon mode with superframes, and
 - i) if More Data = 0 indicating that via the current frame the hub grants a type-I (immediate) polled allocation that starts pSIFS thereafter, it is set to E such that the polled allocation ends at the end of the allocation slot that is numbered E and located in a beacon period (superframe) encoded in the Next field thereof; or
 - ii) if More Data = 1 indicating that via the current frame the hub grants no polled allocation, it is set to S such that the hub is to send a (future) poll or post to the node at the start of the allocation slot that is numbered S and located in a beacon period (superframe) encoded in the Next field thereof.
 - 2) If Access Mode = 1 indicating that the hub is operating in non-beacon mode without superframes, and that via the current frame the hub grants a type-II (immediate) polled allocation that starts pSIFS thereafter, it is set to M such that the node is allowed to send up to M frames in the polled allocation.
- d) In Wakeup frames, it is used as a Poll-Post Window field, which is set to $Poll_Post_Window$ such that this hub is to send a (future) poll after $2 \times (1 + Scale) \times (256 \times Next + Poll_Post_Window)$

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

milliseconds of the current frame, where *Next* and *Scale* are the values of the Next and Scale fields, respectively.

- e) In B2 frames, it is used as a Poll-Post Window field, which is set to *C* such that the current frame starts a contention access phase (CAP) that ends at the end of the allocation slot numbered *C* in the current beacon period (superframe).
- f) In data type frames, it is used as a Sequence Number field, which is
 - 1) set to zero if the frame is the first data type frame of this frame subtype sent to a recipient or a group of recipients;
 - 2) incremented by one, modulo 256, from its value in the data type frame that was of the same frame subtype and addressed to the same recipient(s) and that contained the previous MSDU or final fragment thereof;
 - 3) kept with the same value in frames containing fragments of the same MSDU; or
 - 4) kept with the same value if the frame is being retransmitted to the same recipient(s).
- g) In all other frames, it is reserved.

5.2.1.1.12 Fragment Number/Next/Coexistence

The Fragment Number/Next/Coexistence field is set as follows:

- a) In beacon and B2 frames, it is used as a Coexistence field, which is formatted as shown in Figure 11.
 - 1) The Beacon Shifting field is set to one if beacon shifting is currently enabled, or is set to zero otherwise.
 - 2) The Channel Hopping field is set to one if channel hopping is currently enabled, or is set to zero otherwise.
 - 3) The Superframe Interleaving field is set to one if the sender (a hub) supports active superframe interleaving and Command frames, or is set to zero otherwise.




Figure 11—Coexistence format

- b) In non-beacon management type frames, it is used as a Fragment Number field, which is
 - 1) set to zero if the current frame contains a frame payload not fragmented or the first fragment of a fragmented frame payload;
 - 2) incremented by one from its value in the frame containing the previous fragment of the frame payload if the current frame contains a non-first fragment of a fragmented frame payload; or
 - 3) kept with the same value if the frame is being retransmitted to the same recipient(s).
- c) In data type frames, it is used as a Fragment Number field, which is
 - 1) set to zero if the current frame contains no MSDU, an MSDU not fragmented, or the first fragment of a fragmented MSDU;
 - 2) incremented by one from its value in the frame containing the previous fragment of the MSDU if the current frame contains a non-first fragment of a fragmented MSDU; or

- 3) kept with the same value if the frame is being retransmitted to the same recipient(s).
- d) In Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as a Next field, which is set as follows:
 - 1) If Access Mode = 0 indicating that the hub is operating in beacon or non-beacon mode with superframes, and
 - i) if More Data = 0 indicating that via the current frame the hub grants a type-I (immediate) polled allocation that starts pSIFS thereafter and ends at the end of an allocation slot as encoded in the Poll-Post Window field thereof, it is set to N such that the allocation slot is the one located in the current beacon period (superframe) for $N = 0$ or in the next N th beacon period (superframe) not counting the current one for $N > 0$ (a case possible if no beacon is to be sent in the polled allocation); or
 - ii) if More Data = 1 indicating that via the current frame the hub grants no polled allocation but is to send a (future) poll or post to the node at the start of an allocation slot as encoded in the Poll-Post Window field thereof, it is set to F such that the allocation slot is the one located in the current beacon period (superframe) for $F = 0$ or in the next F th beacon period (superframe) not counting the current one for $F > 0$.
 - 2) If Access Mode = 1 indicating that the hub is operating in non-beacon mode without superframes, and that via the current frame the hub grants a type-II (immediate) polled allocation that starts pSIFS thereafter and allows for the node to send up to a specified number of frames as encoded in the Poll-Post Window field thereof, it is reserved.
- e) In Wakeup frames, it is used as a Next field, which is set to $Next$ such that this hub is to send a (future) poll after $2 \times (1 + Scale) \times (256 \times Next + Poll_Post_Window)$ milliseconds of the current frame, where $Poll_Post_Window$ and $Scale$ are the values of the Poll-Post Window and Scale fields, respectively.
- f) In all other frames, it is reserved.

5.2.1.1.13 Non-final Fragment/Cancel/Scale/Inactive

The Non-final Fragment/Cancel/Scale/Inactive field is set as follows:

- a) In beacon and B2 frames, it is used as an Inactive field, which is set to one if one or more inactive superframes are enabled (starting) at the end of the current beacon period (superframe), or is set to zero otherwise.
- b) In non-beacon management type frames, it is used as a Non-final Fragment field, which is
 - 1) set to zero if the current frame contains no frame payload, a frame payload not fragmented, or the final fragment of a fragmented frame payload; or
 - 2) set to one if the frame contains a non-final fragment of a fragmented frame payload.
- c) In data type frames, it is used as a Non-final Fragment field, which is
 - 1) set to zero if the current frame contains no MSDU, an MSDU not fragmented, or the final fragment of a fragmented MSDU; or
 - 2) set to one if the frame contains a non-final fragment of a fragmented MSDU.
- d) In I-Ack, B-Ack, Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames sent by a hub to a node, it is used as a Cancel field, which is
 - 1) set to zero if the current frame does not cancel any future polls/posts previously improvised by this hub; or
 - 2) set to one if the current frame cancels all future polls/posts previously improvised by this hub.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

- e) In Wakeup frames, it is used as a Scale field, which is set to *Scale* such that this hub is to send a (future) poll after $2 \times (1 + \text{Scale}) \times (256 \times \text{Next} + \text{Poll_Post_Window})$ milliseconds of the current frame, where *Poll_Post_Window* and *Next* are the values of the Poll-Post Window and Next fields, respectively.
- f) In all other frames, it is reserved.

5.2.1.2 Recipient ID

The Recipient ID field is set to the abbreviated address (i.e., NID or HID) of the recipient of the current frame.

5.2.1.3 Sender ID

The Sender ID field is set to the abbreviated address (i.e., NID or HID) of the sender of the current frame.

5.2.1.4 BAN ID

The BAN ID field is set to the abbreviated address of the BAN in which the current frame is transferred.

5.2.2 MAC Frame Body

The MAC Frame Body, when it has a nonzero length, is formatted as shown in Figure 12. The length of the MAC Frame Body *L_{FB}* is not to exceed *pMaxFrameBodyLength*.




Figure 12—MAC Frame Body format

The Low-Order Security Sequence Number and Message Integrity Code (MIC) fields are not present in unsecured frames, as indicated by the Security Level field of the MAC header of the current frame.

5.2.2.1 Low-Order Security Sequence Number

The Low-Order Security Sequence Number field is set as follows for message freshness as needed for nonce construction and replay detection:

- a) It is set to zero, if the current frame is secured with a PTK or GTK that has never been used.
- b) It is incremented by one from its value in the last frame secured with the same PTK or GTK, containing a valid MIC value, and transmitted by the same sender, even if the current frame transmission is a retransmission of the last frame or an earlier frame. Incrementing by one from the maximum value of the field wraps around to zero, causing the High-Order Security Sequence Number maintained internally to increment by one.

The value of the Low-Order Security Sequence Number field increments in frames secured with the same PTK or GTK, rather than in frames of the same frame type or frame subtype. It increments even if the current frame transmission is a retransmission of the last frame or an earlier frame. It is set independently between the frames sent from a hub to a node and the frames sent from the node to the hub, although the same PTK applies to secured frames sent in either direction.

5.2.2.2 Frame Payload

The Frame Payload field is set as follows:

- a) In management type frames, prior to encryption (if any) it is set to a sequence of fields to be communicated to the recipient(s), with the fields defined in 5.3.
- b) In control type frames, it is set to a sequence of fields to be communicated to the recipient(s), with the fields defined in 5.4.
- c) In data type frames, prior to encryption (if any) it is set to a sequence of octets passed as an MSDU through the MAC SAP to the MAC, without altering the order of the sequence.
- d) In data type frames with the Relay field in their MAC header set to one, prior to encryption (if any) it is set to the relayed MAC frame that would otherwise be sent directly between a node and a hub without relay.

If a frame payload is fragmented and carried in multiple frames, the Frame Payload field is set to a fragment of the otherwise unfragmented frame payload.

The length of the Frame Payload field, denoted as L_{FP} in Figure 12, needs to be such that the length of the MAC frame body does not exceed $pMaxFrameBodyLength$. If the Frame Payload has a zero length, i.e., if a frame does not have a frame payload, then the frame does not have a MAC frame body if it is not secured, but still has a MAC frame body containing the Low-Order Security Sequence Number and MIC fields if it is secured.

5.2.2.3 Message Integrity Code (MIC)

The MIC field is set to a message authentication code for preserving the authenticity and integrity of the MAC header and the MAC frame body of the current secured frame, as further specified in 7.3.1.5.

5.2.3 Frame Check Sequence (FCS)

The FCS field is formatted as shown in Figure 13, where its transmit order is defined such that a_{15} is the LSB of the field, and a_0 is the MSB. The bits a_{15} , a_{14} , ..., a_0 are the binary coefficients of a cyclic redundancy check (CRC) polynomial of degree 15 denoted, as shown in Equation (1):

$$R(x) = a_{15}x^{15} + a_{14}x^{14} + \dots + a_1x + a_0 \quad (1)$$




Figure 13—FCS format

The CRC polynomial is calculated over a calculation field using the following CRC-16-CCITT standard generator polynomial of degree 16, shown in Equation (2):

$$G(x) = x^{16} + x^{12} + x^5 + 1 \quad (2)$$

The calculation field is the transmitted MAC frame except the FCS field for this specification. It is mapped to a message polynomial $M(x)$ of degree $k-1$, where k is the number of bits in the calculation field. The LSB of the first octet presented to the PHY SAP is the coefficient of the x^{k-1} term, the next LSB is the coefficient of the x^{k-2} term, and finally the MSB of the last octet transmitted is the coefficient of the x^0 term.

The CRC polynomial is the remainder $R(x)$ resulting from the (modulo 2) division of $[x^{16} \times M(x)]$ by $G(x)$, shown in Equation (3):

$$R(x) = x^{16} \times M(x) \bmod G(x) \quad (3)$$

The initial remainder of the division is set to zero, and the final remainder is not inverted as is the case in some other standards.

5.3 Management type frames

A management type frame contains certain mandatory fixed-length fields and some optional variable length components referred to as *information elements* (IEs).

5.3.1 Beacon

A beacon frame contains a Frame Payload that is formatted as shown in Figure 14. It is locally broadcast by a hub in every beacon period (superframe).




Figure 14—Frame Payload format for Beacon frames

5.3.1.1 Sender Address

The Sender Address field is set to the EUI-48 of the hub sending the current beacon.

5.3.1.2 Beacon Period Length

The Beacon Period Length field is set to the length of the current beacon period (superframe), in units of allocation slots. It is set to zero to encode a value of 256 allocation slots.

5.3.1.2.1 Allocation slot numbering

If beacon shifting is not enabled, the allocation slots in a beacon period (superframe) are numbered $0, 1, \dots$, starting from the allocation slot that starts at the beacon transmission time of the beacon period (superframe) consecutively to the allocation slot that ends at the end of the beacon period (superframe), as shown in Figure 15.




Figure 15—Allocation slot ordering and numbering in beacon periods (superframes)

5.3.1.2.2 Allocation slot numbering with beacon shifting

If beacon shifting is enabled, a beacon period (superframe) needs to have $4N$ allocation slots in length, where N is an integer. The allocation slots in a beacon period (superframe) are numbered $0, 1, \dots, 4N-1$, starting from the allocation slot that starts at the beacon transmission time of the beacon period (superframe) to the allocation slot that ends at the end of the beacon period (superframe) and, if the beacon transmission time is not at the start of the beacon period (superframe), wrapping back to the allocation slot that starts at the start of the beacon period (superframe) and finally to the allocation slot that ends at the beacon transmission time, as shown in Figure 16. A beacon period (superframe) has four quarters, which are comprised of allocation slots 0 to $N-1$, N to $2N-1$, $2N$ to $3N-1$, and $3N$ to $4N-1$, respectively.




Figure 16—Allocation slot ordering and numbering in beacon periods (superframes) with exemplary beacon shifting

5.3.1.3 Allocation Slot Length

The Allocation Slot Length field is set to L such that the length of an allocation slot is equal to $p_{\text{AllocationSlotMin}} + L \times p_{\text{AllocationSlotResolution}}$.

5.3.1.4 RAP1 Start

The RAP1 Start field is present only if EAP1 has a nonzero length as indicated by the EAP Indicator field of the MAC header of the current beacon frame. When present, it is set to *S1* such that random access phase 1 (RAP1) starts at the beginning of the allocation slot that is numbered *S1* and located in the current beacon period ends EAP1, and it occurs after the PHY Capability field.

Both EAP1 (if of nonzero length) and RAP1 (if of nonzero length) occur in the current beacon period if beacon shifting is not enabled or in the next beacon period if beacon shifting is enabled. If EAP1 has a nonzero length, it immediately follows the preceding beacon, and is immediately followed by RAP1. If EAP1 has a zero length, RAP1 immediately follows the preceding beacon.

5.3.1.5 RAP1 End

The RAP1 End field is set to *E1* such that RAP1 ends at the end of the allocation slot that is numbered *E1* and located in the current beacon period. The value of this field is not to be smaller than the value of the Earliest RAP1 End field in any Connection Assignment frame sent by the hub transmitting this beacon.

5.3.1.6 RAP2 Start

The RAP2 Start field is set to *S2* such that random access phase 2 (RAP2) starts at the beginning of the allocation slot that is numbered *S2* and located in the current beacon period if either exclusive access phase 2 (EAP2) or RAP2 is of nonzero length, or is set to zero otherwise. If EAP2 is of nonzero length, it ends at the time indicated by this field. The start time of EAP2 is encoded in Connection Assignment frames defined in 5.3.6.15.

Both EAP2 (if of nonzero length) and RAP2 (if of nonzero length) occur in the current beacon period if beacon shifting is not enabled or in the next beacon period if beacon shifting is enabled.

5.3.1.7 RAP2 End

The RAP2 End field is set to *E2* such that RAP2 ends at the end of the allocation slot that is numbered *E2* and located in the current beacon period if RAP2 is of nonzero length, or is set to zero otherwise.

5.3.1.8 MAC Capability

The MAC Capability is as defined in 5.6.1.

5.3.1.9 PHY Capability

The PHY Capability is as defined in 5.6.2.

5.3.1.10 Beacon Shifting Sequence

The Beacon Shifting Sequence is present only if beacon shifting is currently enabled as encoded according to Figure 11. When present, it is formatted as shown in Figure 17 to indicate the beacon transmission time in the current beacon period (superframe).

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 17 —Beacon Shifting Sequence format

5.3.1.10.1 Beacon Shifting Sequence Index

The Beacon Shifting Sequence Index field is set according to Table 4 to the index m of the PN sequence $PN_m(n)$ governing the beacon transmission time pattern.

Table 4—Beacon Shifting Sequence field encoding

Beacon Shifting Sequence Index m in decimal value	Beacon Shifting Sequence as function of Beacon Shifting Sequence Phase $n = 0, 1, 2, \dots, 15$	Beacon Shifting Sequence pattern (“...” denotes pattern repeat)
0	$PN_0(n) = n \bmod 2$	$PN_0(n) = 0, 1, 0, 1, \dots$
1	$PN_1(n) = 2 \times PN_0(n)$	$PN_1(n) = 0, 2, 0, 2, \dots$
2	$PN_2(n) = n \bmod 4$	$PN_2(n) = 0, 1, 2, 3, \dots$
3	$PN_3(n) = [PN_0(n) + PN_2(n)]/2 \bmod 2 + [PN_0(n) + PN_1(n) + PN_2(n)] \bmod 4$	$PN_3(n) = 0, 1, 3, 2, \dots$
4	$PN_4(n) = [PN_0(n) + PN_1(n) + PN_2(n)]/2$	$PN_4(n) = 0, 2, 1, 3, \dots$
5	$PN_5(n) = \{PN_2(n) + [PN_0(n) + PN_2(n)]/2\} \bmod 4$	$PN_5(n) = 0, 2, 3, 1, \dots$
6	$PN_6(n) = PN_1(n) + \{[PN_0(n) + PN_2(n)]/2 \bmod 2\}$	$PN_6(n) = 0, 3, 1, 2, \dots$
7	$PN_7(n) = [PN_1(n) + PN_2(n)] \bmod 4$	$PN_7(n) = 0, 3, 2, 1, \dots$
8–15	Reserved	Reserved

5.3.1.10.2 Beacon Shifting Sequence Phase

The Beacon Shifting Sequence Phase field is incremented by one, modulo 16, from its value applicable to the last beacon period (superframe), active or inactive.

5.3.1.11 Channel Hopping State

The Channel Hopping State field is present only if channel hopping is currently enabled as encoded according to Figure 11. When present, it is set to the current state of a 16-bit maximum-length linear feedback shift register (LFSR) used to generate the channel hopping sequence by the hub sending this beacon, as further specified in 6.13.2.

5.3.1.12 Next Channel Hop

The Next Channel Hop field is present only if channel hopping is currently enabled as encoded according to Figure 11. When present, it is set to the sequence number of the beacon period (superframe) in which the hub sending the current beacon is to hop to another channel according to its channel hopping sequence. Channel hopping is independent of inactive duration as defined in 5.3.1.13.

5.3.1.13 Inactive Duration

The Inactive Duration field is present only if one or more inactive superframes are starting at the end of the current beacon period (superframe) as encoded by the Non-final Fragment/Cancel/Scale/Inactive field according to 5.2.1.1.13. When present, it is set to the number of inactive superframes after each active superframe.

5.3.2 Security Association

A Security Association frame contains a Frame Payload that is formatted as shown in Figure 18. It is exchanged between a node and a hub during the execution of a security association protocol to activate a pre-shared MK or generate a new shared MK.




Figure 18—Frame Payload format for Security Association frames

5.3.2.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame, or is set to zero if such an EUI-48 is yet unknown.

5.3.2.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.2.3 Security Suite Selector

The Security Suite Selector is formatted as shown in Figure 19.




Figure 19—Security Suite Selector format

5.3.2.3.1 Security Association Protocol

The Security Association Protocol field is set according to Table 5 to select a desired security association protocol.

Table 5—Security Association Protocol field encoding

Field value in decimal	Security association protocol
0	Master key pre-shared association
1	Unauthenticated association
2	Public key hidden association
3	Password authenticated association
4	Display authenticated association
5–7	Reserved

5.3.2.3.2 Security Level

The Security Level field is set to the security level required by this sender according to Table 6.

Table 6—Security Level field encoding

Field value in decimal	Security level required
0	Level 0—unsecured communication
1	Level 1—authentication but not encryption
2	Level 2—authentication and encryption
3	Reserved

5.3.2.3.3 Control Frame Authentication

The Control Frame Authentication field is set to one if control type frames sent from or to this sender needs to be authenticated but not encrypted when they are required to have security level 1 or 2. It is set to zero if control type frames sent from or to this sender needs to be neither authenticated nor encrypted even when they are otherwise required to have security level 1 or 2.

5.3.2.3.4 Cipher Function


The Cipher Function field is set according to Table 7 to indicate the underlying cipher function selected by this sender for performing security services as specified in Clause 7.

Table 7—Cipher Function field encoding

Field value in decimal	Cipher function
0	AES-128 forward cipher function
1	Camellia-128 forward cipher function
2–15	Reserved

5.3.2.4 Association Control

The Association Control is formatted as shown in Figure 20.

**Figure 20—Association Control format****5.3.2.4.1 Association Sequence Number**

The Association Sequence Number field is set to the number (i.e., position) of the current Security Association frame in the execution of the selected security association protocol. In particular, it is set to one in the first Security Association frame of the protocol, two in the second Security Association frame, etc. The first Security Association frame is the Security Association frame transmitted or retransmitted by the node initializing the security association, the second Security Association frame is the next Security Association frame transmitted or retransmitted by the responding hub, and the like.

5.3.2.4.2 Association Status

The Association Status field in the second Security Association frame of a security association protocol (procedure) is set according to Table 8 to indicate the status of the current security association procedure. It is reserved in other Security Association frames.

Table 8—Association Status field encoding

Field value in decimal	Association status
0	Joining the security association procedure
1	Aborting the security association procedure with a different security suite selector
2	Aborting the security association procedure due to lack of needed security credential (no MK pre-shared with the initiator of the security association procedure for MK pre-shared association; no public key of the initiator of the security association procedure for public key hidden association; no password shared with the initiator of the security association procedure for password authenticated association)
3	Aborting the security association procedure due to temporary lack of resources
4	Aborting the security association procedure due to security policy restrictions as imposed by the administrator/owner of this hub on the communications in its BAN
5–15	Reserved

5.3.2.5 Security Association Data

The Security Association Data field is specific to the selected security association protocol.

For MK pre-shared association, the Security Association Data field is not present.

For unauthenticated association, public key hidden association, password authenticated association, and display authenticated association, the Security Association Data field is formatted as shown in Figure 21.




Figure 21—Security Association Data format for security association protocols 1 to 4

5.3.2.5.1 Sender Nonce

The Sender Nonce field is set to a statistically unique number per sender and per security association procedure, except otherwise indicated, as follows:

- a) For unauthenticated association, public key hidden association, and password authenticated association
 - 1) In the first Security Association frame of the security association procedure, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders.
 - 2) In the second Security Association frame of the procedure, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders if the sender of the frame is to join the security association procedure, or is set to zero otherwise.
 - 3) In the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
 - 4) In the fourth Security Association frame, the field is set to its value contained in the first Security Association frame of the procedure.
- b) For display authenticated association
 - 1) In the first Security Association frame of the security association procedure, the field is set to a witness (message authentication code) of its value contained in the fourth Security Association frame of the procedure.
 - 2) In the second Security Association frame of the procedure, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders if the sender of the frame is to join the security association procedure, or is set to zero otherwise.
 - 3) In the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
 - 4) In the fourth Security Association frame, the field is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders.

5.3.2.5.2 Sender PK_x

The Sender PK_x field is set to the X-coordinate of the sender's elliptic curve public key, except otherwise indicated, as follows:

- a) For unauthenticated association and display authenticated association,
 - 1) in the first Security Association frame of the current security association procedure, the field is set to the X-coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$;
 - 2) in the second Security Association frame of the procedure, the field is set to the X-coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

- 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure;
 - 4) in the fourth Security Association frame, the field is set to its value contained in the first Security Association frame of the procedure.
- b) For public key hidden association,
- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the second Security Association frame, the field is set to the X -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
- c) For password authenticated association,
- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to the X -coordinate of the sender's password-scrambled elliptic curve public key $PK' = (PK'_x, PK'_y)$;
 - 2) in the second Security Association frame, the field is set to the X -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.

5.3.2.5.3 Sender PK_y

The Sender PK_y field is set to the Y -coordinate of the sender's elliptic curve public key, except otherwise indicated, as follows:

- a) For unauthenticated association and display authenticated association,
 - 1) in the first Security Association frame of the current security association procedure, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$;
 - 2) in the second Security Association frame of the procedure, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure;
 - 4) in the fourth Security Association frame, the field is set to its value contained in the first Security Association frame of the procedure.
- b) For public key hidden association,
 - 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the second Security Association frame, the field is set to the Y -coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
 - 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.
- c) For password authenticated association,

- 1) in the first and fourth Security Association frames of the current security association procedure, the field is set to the Y-coordinate of the sender's password-scrambled elliptic curve public key $PK' = (PK'_x, PK'_y)$;
- 2) in the second Security Association frame, the field is set to the Y-coordinate of the sender's elliptic curve public key $PK = (PK_x, PK_y)$ if the sender of the frame is to join the security association procedure, or is set to zero otherwise;
- 3) in the third Security Association frame, the field is set to its value contained in the second Security Association frame of the procedure.

5.3.2.5.4 MK_KMAC

The MK_KMAC field is set to a key message authentication code (KMAC) for certain fields of the frame payloads of the Security Association frames of the current security association procedure, except otherwise indicated, as follows:

- a) For unauthenticated association and display authenticated association,
 - 1) in the first and second Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the third Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the first and second Security Association frames of the procedure;
 - 3) in the fourth Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the second and fourth Security Association frames.
- b) For public key hidden association,
 - 1) in the first and second Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the third Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the first and second Security Association frames of the procedure;
 - 3) in the fourth Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the second and fourth Security Association frames.
- c) For password authenticated association,
 - 1) in the first and Security Association frames of the current security association procedure, the field is set to zero;
 - 2) in the third Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the first and second Security Association frames of the procedure;
 - 3) in the fourth Security Association frame, the field is set to a KMAC for certain fields of the frame payloads of the second and fourth Security Association frames.

5.3.3 Security Disassociation

A Security Disassociation frame contains a Frame Payload that is formatted as shown in Figure 22. It is transmitted by either an associated node or a hub to repeal an existing security association, i.e., the shared MK.




Figure 22—Frame Payload format for Security Disassociation frames

5.3.3.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.3.3.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.3.3 DA_KMAC

The DA_KMAC field is set to a KMAC for the fields of the frame payload of this Security Disassociation frame.

5.3.4 Pairwise Temporal Key (PTK)

A PTK frame contains a Frame Payload that is formatted as shown in Figure 23. It is exchanged between a node and the hub with which the node is associated to create a PTK based on a shared MK.




Figure 23—Frame Payload format for PTK frames

5.3.4.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.3.4.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.4.3 Sender Nonce

The Sender Nonce field is set to a statistically unique number per sender and per PTK creation procedure, as follows:

- In the first PTK frame of the current PTK creation procedure, it is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders.
- In the second PTK frame of the procedure, it is set afresh to an integer randomly drawn from $\{1, 2, \dots, 2^{128}-1\}$ and independent of the nonces of other senders if the sender of the frame is to join the PTK creation procedure, or is set to zero otherwise.
- In the third PTK frame of the procedure, it is set to its value contained in the first PTK frame.

5.3.4.4 PTK Control

The PTK Control is formatted as shown in Figure 24.




Figure 24—PTK Control format

5.3.4.4.1 Message Number

The Message Number field is set to the number (i.e., position) of the current PTK frame in the current PTK creation procedure. In particular, it is set to one in the first PTK frame of the procedure, two in the second PTK frame, and three in the third. The first PTK frame is the PTK frame transmitted or retransmitted by the node or hub initializing the procedure, the second PTK frame is the PTK frame transmitted or retransmitted by the responding hub or node, and the third PTK frame is the PTK frame transmitted or retransmitted by the initiating node or hub again. The other values of the field are reserved.

5.3.4.4.2 PTK Index

The PTK Index field in the first PTK frame transmitted or retransmitted by the node or hub initiating the PTK creation procedure is set as follows to identify the PTK being created, as follows:

- If no PTK was previously created with the responding node, it is set to zero.
- Otherwise, it is set to the modulo-2 sum of one and its value used in successfully creating the last PTK between the sender and the recipient.

5.3.4.4.3 PTK Creation Status

The PTK Creation Status field in the second PTK frame of a security association protocol is set according to Table 9 to indicate the status of the current PTK creation procedure. It is reserved in other PTK frames.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 9—PTK Creation Status field encoding

Field value in decimal	PTK creation status
0	Joining the PTK creation procedure
1	Reserved
2	Aborting the PTK creation procedure due to lack of shared master key
3	Aborting the PTK creation procedure due to temporary lack of resources
4	Aborting the PTK creation procedure due to security policy restrictions as imposed by the administrator/owner of this hub on the communications in its BAN
5–15	Reserved

5.3.4.5 PTK_KMAC


The PTK_KMAC field is set to a KMAC for certain fields of the frame payloads of the PTK frames of the current PTK creation procedure, except otherwise indicated:

- In the first PTK frame of the current PTK creation procedure, it is set to zero.
- In the second PTK frame, it is set to a KMAC for certain fields of the frame payloads of the first and second PTK frames of the procedure if the sender of this frame is to join the PTK creation procedure, or it is set to zero otherwise.
- In the third PTK frame, it is set to a KMAC for certain fields of the frame payloads of the second and third PTK frames.

The PTK Index field takes on a value of either zero or one.

5.3.5 Group Temporal Key (GTK)

A GTK frame contains a Frame Payload that is formatted as shown in Figure 25. It is transmitted by a hub to distribute a GTK to a secured node for securing multicast traffic.

**Figure 25—Frame Payload format for GTK frames****5.3.5.1 Recipient Address**

The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.3.5.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.5.3 GTK NID

The GTK NID field is set according to Table 17 to the broadcast or multicast NID that is to appear in the Recipient ID field of the MAC header of the frames secured by the GTK distributed in the current frame.

5.3.5.4 GTK Control

The GTK Control is formatted as shown in Figure 26.




Figure 26—GTK Control format

The GTK Index field is set as follows to identify the GTK being distributed:

- If no GTK was previously distributed by this hub for the GTK NID indicated in the current frame, it is set to zero.
- Otherwise, it is set to the modulo-2 sum of one and its value used in successfully distributing the last GTK by this hub for the indicated GTK NID.

The GTK Index field takes on a value of either zero or one.

5.3.5.5 GTK SSN

The GTK SSN field is set to the security sequence number, comprising a low-order security sequence number as its LSBs and a high-order security sequence number as its MSBs, of the last frame secured with the GTK distributed in the current frame and addressed to the GTK NID indicated in the current frame.

5.3.5.6 GTK

The GTK field is set to the bit string representing the GTK being distributed in the current frame.

5.3.6 Connection Request

A Connection Request frame contains a Frame Payload that is formatted as shown in Figure 27. It is transmitted by a node to request creation or modification of a connection with a hub.




Figure 27 —Frame Payload format for Connection Request frames

5.3.6.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame, or is set to zero if such an EUI-48 is yet unknown.

5.3.6.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.6.3 MAC Capability

The MAC Capability is as defined in 5.6.1.

5.3.6.4 PHY Capability

The PHY Capability is as defined in 5.6.2.

5.3.6.5 Connection Change Indicator

The Connection Change Indicator is formatted as shown in Figure 28. It indicates certain fields that follow in the current frame have been newly provided or changed in value since their last exchange between the sender and the recipient.




Figure 28—Connection Change Indicator format

5.3.6.5.1 Ack Data Rates Change

The Ack Data Rates Change field is set to one if the value of the Requested Ack Data Rates or Assigned Ack Data Rates field has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.2 Wakeup Phase Change

The Wakeup Phase Change field is set to one if the value of the Requested Wakeup Phase or Assigned Wakeup Phase field has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.3 Wakeup Period Change

The Wakeup Period Change field is set to one if the value of the Requested Wakeup Period or Assigned Wakeup Period field has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.4 Uplink Request/Assignment IE Change

The Uplink Request/Assignment IE Change field is set to one if the value of the Uplink Request IE or Uplink Assignment IE has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.5 Downlink Request/Assignment IE Change

The Downlink Request/Assignment IE Change field is set to one if the value of the Downlink Request IE or Downlink Assignment IE has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.6 Bilink Request/Assignment IE Change

The Bilink Request/Assignment IE Change field is set to one if the value of the Bilink Request IE or Bilink Assignment IE has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.7 Unscheduled Bilink Request/Assignment IE Change

The Unscheduled Bilink Request/Assignment IE Change field is set to one if the value of the Type-I Unscheduled Bilink Request IE, Type-II Unscheduled Bilink Request IE, Type-I Unscheduled Bilink

Assignment IE, or Type-II Unscheduled Bilink Assignment IE has been newly provided or changed, or is set to zero otherwise.

5.3.6.5.8 Channel Order IE Change

The Channel Order IE Change field in Connection Request frames is reserved.

The Channel Order IE Change field in Connection Assignment frames is set to one if the value of the Nibble Encoded Channel Order IE or Channel Hopping and Ordering IE has been newly provided or changed, or is set to zero otherwise.

5.3.6.6 Requested Ack Data Rates

The Requested Ack Data Rates is formatted as shown in Figure 29. It defines the data rates requested for use to send I-Ack and B-Ack data frames between the sender and recipient of the current frame while they are exchanging data type frames.




Figure 29—Requested Ack Data Rates and Assigned Ack Data Rates format

5.3.6.6.1 Node Ack Data Rate Control

The Node Ack Data Rate Control field is set to one if the sender or recipient of this frame (a node) is to send its I-Ack and B-Ack frames at the same data rate as used to send the last frame it received, or is set to zero if the node is to send its I-Ack and B-Ack frames at a data rate indicated in the following Node Ack Data Rate field.

5.3.6.6.2 Node Ack Data Rate

The Node Ack Data Rate field is set to R such that the sender or recipient of this frame (a node) is to send its I-Ack and B-Ack frames at the information data rate as encoded by $R = R2R1R0$ of the Data Rate field defined in the corresponding physical layer (PHY) clause, if the preceding Node Ack Data Rate Control is set to zero, or is reserved otherwise. Here, bit $R0$ denotes the LSB of R , and bit $R2$ denotes the MSB.

5.3.6.6.3 Hub Ack Data Rate Control

The Hub Ack Data Rate Control field is set to one if the sender or recipient of this frame (a hub) is to send its I-Ack and B-Ack frames at the same data rate as used to send the last frame it received, or is set to zero if the hub is to send its I-Ack and B-Ack frames at a data rate indicated in the following Hub Ack Data Rate field.

5.3.6.6.4 Hub Ack Data Rate

The Hub Ack Data Rate field is set to R such that the sender or recipient of this frame (a hub) is to send its I-Ack and B-Ack frames at the information data rate as encoded by $R = R2R1R0$ of the Data Rate field defined in the corresponding PHY clause, if the preceding Hub Ack Data Rate Control is set to zero, or is reserved otherwise. Here, bit $R0$ denotes the LSB of R , and bit $R2$ denotes the MSB.

5.3.6.6.7 Requested Wakeup Phase

The Requested Wakeup Phase field is set to the sequence number of the next beacon period (superframe) in which the sender (a node) plans to wake up for frame reception and transmission, with the sequence number of a beacon period (superframe) treated as incremented by one modulo 2^{16} , instead of modulo 2^8 , from that of the previous beacon period (superframe). The value of this field is calculated as $S+D$ modulo 2^{16} , where S is the one-octet sequence number of the current beacon period (superframe) and D is such that the node is to wake up D beacon periods (superframes) later after receiving a Connection Assignment frame. With a length of two octets, this field allows a node's next wakeup to be scheduled up to 2^{16} beacon periods (superframes) away from the current one.

The Requested Wakeup Phase field is reserved in non-beacon mode without superframes.

5.3.6.6.8 Requested Wakeup Period

The Requested Wakeup Period field is set to the length, in units of beacon periods (superframes), between the start of successive wakeup beacon periods (superframes) in which the sender (a node) plans to wake up for reception and transmission, starting from the one indicated in the preceding Requested Wakeup Phase field. It is set to zero to encode a value of 2^{16} beacon periods (superframes). With a length of two octets, this field allows a node's wakeup period to be up to 2^{16} beacon periods (superframes).

The value of this field determines whether the IEs in this frame denote 1-periodic or m-periodic allocations, as follows:

- If Requested Wakeup Period = 1, these IEs denote 1-periodic allocations.
- If Requested Wakeup Period $\neq$ 1, these IEs denote m-periodic allocations.

The Requested Wakeup Period field is reserved in non-beacon mode without superframes.

5.3.6.6.9 Max Sync Interval/Clock PPM

The Max Sync Interval/Clock PPM is present only if this node is requiring centralized guard time provisioning as indicated in the MAC Capability field of the current frame. When present, it is formatted as shown in Figure 30.




Figure 30—Max Sync Interval/Clock PPM format

5.3.6.9.1 Node Max Sync Interval

The Node Max Sync Interval field is set to the length of this node's maximum synchronization interval, in units of the Requested Wakeup Period field value in the current frame, over which this node is to synchronize with its hub at least once. It is set to zero to encode a value of 8 such units.

5.3.6.9.2 Node Clock PPM

The Node Clock PPM field is set to the PPM of this node's MAC clock encoded according to Table 10.

Table 10—Node Clock PPM field encoding

Field value in decimal	Clock accuracy (ppm)
0	20
1	40
2	50
3	100
4	200
5	300
6	400
7	500

5.3.6.10 Uplink Request IE

The Uplink Request IE is as defined in 5.7.2.

5.3.6.11 Downlink Request IE

The Downlink Request IE is as defined in 5.7.3.

5.3.6.12 Bilink Request IE

The Bilink Request IE is as defined in 5.7.4.

5.3.6.13 Unscheduled Bilink Request IE

The Unscheduled Bilink Request IE, when present, is either Type-I Unscheduled Bilink Request IE as defined in 5.7.5 or Type-II Unscheduled Bilink Request IE as defined in 5.7.6.

5.3.6.14 Former Hub Address IE

The Former Hub Address IE is as defined in 5.7.14.

5.3.6.15 Application Specific IE

The Application Specific IE is as defined in 5.7.15.

5.3.7 Connection Assignment

A Connection Assignment frame contains a Frame Payload that is formatted as shown in Figure 31. It is transmitted by a hub to respond to a connection request or to initiate or change a connection assignment.




Figure 31—Frame Payload format for Connection Assignment frames

5.3.7.1 Recipient Address


The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.3.7.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.7.3 Mode/Status

The Mode/Status field is formatted as shown in Figure 32.

**Figure 32—Mode/Status format****5.3.7.3.1 Access Indicator**

The Access Indicator field is set to the access mode of this hub encoded according to Table 11.

Table 11—Access Indicator field encoding

Field value in decimal	Status
0	Beacon mode with superframes, and CSMA/CA for random access
1	Beacon mode with superframes, and slotted aloha for random access
2	Non-beacon mode with superframes, and no random access
3	Non-beacon mode without superframes, and CSMA/CA for random access

5.3.7.3.2 Connection Status

The Connection Status field is set to the status of the connection assignment encoded according to Table 12.

5.3.7.4 Current Superframe Number

The Current Superframe Number field is set to the sequence number of the current superframe (beacon period) as follows:

- a) It is set to the sequence number of the beacon that is or was transmitted in the current beacon period.
- b) It is set to the sequence number of the beacon that would otherwise have been or be transmitted in the current superframe if no beacon was or is to be transmitted in this superframe.

The Current Superframe Number field is reserved in non-beacon mode without superframes.

5.3.7.5 Earliest RAP1 End

The Earliest RAP1 End field is set to $E > 0$ such that random access phase 1 (RAP1) is guaranteed not to end before the start of the allocation slot numbered E in any beacon period (superframe) if RAP1 is of nonzero minimum length, or is set to 0 otherwise.

Table 12—Connection Status field encoding

Field value in decimal	Status
0	Connection request accepted
1	Connection request rejected—due to access policy restrictions as imposed by the administrator/owner of this hub on the communications in its BAN
2	Connection request rejected—invalid or unsupported frame format
3	Connection request rejected—no unsecured communication with this hub
4	Connection request rejected—no more channel bandwidth for a new connection
5	Connection request rejected—no more Connected_NID for a new connection
6	Connection request rejected—no more internal resources for a new connection
7	Connection request rejected—node's maximum synchronization interval too long to support
8	Connection request rejected—node's clock ppm too large to support
9	Connection request rejected—beacon shifting enabled but not supported by requestor
10	Connection request rejected—channel hopping enabled but not supported by requestor
11–15	Reserved
16	Connection assignment modified
17–31	Reserved

5.3.7.6 EAP2 Start

The EAP2 Start field is set to the number of the allocation slot whose start time starts EAP2 if EAP2 is of nonzero length, or is set to zero otherwise. If EAP 2 is of nonzero length, it ends at the start time of random access phase 2 (RAP2) defined in the beacon.

5.3.7.7 Minimum CAP Length

The Minimum CAP Length field is set to the least length guaranteed for CAP, in units of allocation slots.

5.3.7.8 MAC Capability

The MAC Capability is as defined in 5.6.1.

5.3.7.9 PHY Capability

The PHY Capability is as defined in 5.6.2.

5.3.7.10 Connection Change Indicator

The Connection Change Indicator is as defined in 5.3.6.5.

5.3.7.11 Assigned Ack Data Rates

The Assigned Ack Data Rates is as formatted in Figure 29 and as encoded in 5.3.6.6. It defines the data rates assigned for use to send I-Ack and B-Ack data frames between the sender and recipient of the current

frame while they are exchanging data type frames. The field in this Connection Assignment frame supersedes the Requested Ack Data Rates field in the Connection Request frame previously exchanged between the sender and the recipient.

5.3.7.12 Assigned Wakeup Phase

The Assigned Wakeup Phase field is set to the sequence number of the next beacon period (superframe) in which the recipient (a node) needs to wake up for frame reception and transmission. It is as encoded in 5.3.6.7. The field in this Connection Assignment frame supersedes the Requested Wakeup Phase field in the Connection Request frame previously exchanged between the sender and the recipient.

The Assigned Wakeup Phase field is reserved in non-beacon mode without superframes.

5.3.7.13 Assigned Wakeup Period

The Assigned Wakeup Period field is set to the length, in units of beacon periods (superframes), between the start of successive wakeup beacon periods (superframes) in which the recipient (a node) needs to wake up for reception and transmission, starting from the one indicated in the preceding Assigned Wakeup Phase field. It is set to zero to encode a value of 2^{16} beacon periods (superframes). It is as encoded in 5.3.6.8. The field in this Connection Assignment frame supersedes the Requested Wakeup Phase field in the Connection Request frame previously exchanged between the sender and the recipient.

The value of this field determines whether the IEs in this frame denote 1-periodic or m-periodic allocations as follows:

- a) If Assigned Wakeup Period = 1, these IEs denote 1-periodic allocations.
- b) If Assigned Wakeup Period $\neq$ 1, these IEs denote m-periodic allocations.

The Assigned Wakeup Period field is reserved in non-beacon mode without superframes.

5.3.7.14 Superframe Parameters IE

The Superframe Parameters IE is as defined in 5.7.1.

5.3.7.15 Uplink Assignment IE

The Uplink Assignment IE is as defined in 5.7.7.

5.3.7.16 Downlink Assignment IE

The Downlink Assignment IE is as defined in 5.7.8.

5.3.7.17 Bilink Assignment IE

The Bilink Assignment IE is as defined in 5.7.9.

5.3.7.18 Unscheduled Bilink Assignment IE

The Unscheduled Bilink Assignment IE is either Type-I Unscheduled Bilink Assignment IE as defined in 5.7.10 or Type-II Unscheduled Bilink Assignment IE as defined in 5.7.11.

5.3.7.19 Channel Order IE

The Channel Order IE is either a Nibble Encoded Channel Order IE or a Channel Hopping and Ordering IE as defined in 5.7.12 and 5.7.13, respectively.

5.3.7.20 Application Specific IE

The Application Specific IE is as defined in 5.7.15.

5.3.8 Disconnection

A Disconnection frame contains a Frame Payload that is formatted as shown in Figure 33. It is transmitted by a hub to repeal the connection with a node or by a node to repeal the connection with a hub.




Figure 33—Frame Payload format for Disconnection frames

5.3.8.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.3.8.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.8.3 New Hub Address

The New Hub Address field in Disconnection frames sent by a node is set as follows:

- a) If the node is not newly connected with another hub, it is null, i.e., not present.
- b) Otherwise, it is set to the EUI-48 of the hub with which the node is newly connected.

The *i*th New hub Address field in Disconnection frames sent by a hub is set as follows:

- c) If the hub does not have a suggested *i*th preferred new hub for the addressed node, it is null, i.e., not present.
- d) Otherwise, it is set to the EUI-48 of the suggested *i*th preferred new hub for the addressed node.

5.3.9 Command

A Command frame contains a Frame Payload that is formatted as shown in Figure 34. It is optionally transmitted by a hub or node.




Figure 34—Frame Payload format for Command frames

5.3.9.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.3.9.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.3.9.3 Command ID

The Command ID field is set according to Table 13 such that it identifies the specific command of the current frame.

5.3.9.4 Command Data

The Command Data field is specific to the command conveyed in the current frame.

Table 13—Command ID Field encoding

Field value decimal	Command name
0	Command—Active Superframe Interleaving Request
1	Command—Active Superframe Interleaving Response
2–255	Reserved

5.3.9.5 Command—Active Superframe Interleaving Request

The Command Data field is formatted as shown in Figure 35 for a Command—Active Superframe Interleaving Request frame, which is optionally transmitted by a hub to another hub to request for channel sharing through active superframe interleaving.




Figure 35—Command Data format for Command—Active Superframe Interleaving Request frames

5.3.9.5.1 HID

The HID field is set to the abbreviated address of the sender (a hub) of the current frame.

5.3.9.5.2 BAN ID

The BAN ID field is set to the abbreviated address of the BAN of this hub.

5.3.9.5.3 BAN Priority

The BAN Priority field is set according to Table 14 to indicate the priority of the services provided to the BAN of the sender of the current frame. The higher the value of this field, the higher the priority of the BAN services.

Table 14—BAN Priority field encoding

Field value in decimal	BAN services
0	Non-medical services
1	Mixed medical and non-medical services
2	General health services
3	Highest priority medical services

5.3.9.5.4 Requested Beacon Period Length

The Requested Beacon Period Length field is set to the length of the beacon period (superframe), in units of allocation slots, as requested by the sender of the current frame. It is set to zero to encode a value of 256 allocation slots.

5.3.9.5.5 Requested Allocation Slot Length

The Requested Allocation Slot Length field is set to L such that the length of an allocation slot, as requested by the sender of the current frame, is equal to $p\text{AllocationSlotMin} + L \times p\text{AllocationSlotResolution}$.

5.3.9.5.6 Requested Active Superframe Offset

The Requested Active Superframe Offset field is set to the length, in units of requested beacon periods (superframes) defined in the current frame, as requested by the sender of the current frame, between the end of an active superframe of the recipient of the current frame and the start of the next active superframe of the sender of the current frame.

5.3.9.5.7 Requested Inactive Duration

The Requested Inactive Duration field is set to the number of inactive superframes of the sender of the current frame after each active superframe of the sender, as requested by the sender.

5.3.9.6 Command—Active Superframe Interleaving Response

The Command Data field is formatted as shown in Figure 36 for a Command—Active Superframe Interleaving Response frame, which is optionally transmitted by a hub to another hub in response to a request for channel sharing through active superframe interleaving.

5.3.9.6.1 HID

The HID field is set to the abbreviated address of the sender (a hub) of the current frame.




**Figure 36—Command Data format for
Command—Active Superframe Interleaving Response frames**

5.3.9.6.2 BAN ID

The BAN ID field is set to the abbreviated address of the BAN of this hub.

5.3.9.6.3 BAN Priority

The BAN Priority field is as defined in 5.3.9.5.3.

5.3.9.6.4 Request Status

The Request Status field is set to one if the request for active superframe interleaving is accepted, or is set to zero otherwise.

5.3.9.6.5 Offered Beacon Period Length

The Offered Beacon Period Length field is set to the length of the beacon period (superframe), in units of allocation slots, as offered by the sender of the current frame. It is set to zero to encode a value of 256 allocation slots.

5.3.9.6.6 Offered Allocation Slot Length

The Offered Allocation Slot Length field is set to L such that the length of an allocation slot, as offered by the sender of the current frame, is equal to $p\text{AllocationSlotMin} + L \times p\text{AllocationSlotResolution}$.

5.3.9.6.7 Offered Active Superframe Offset

The Offered Active Superframe Offset field is set to the length, in units of offered beacon periods (superframes) defined in the current frame, as offered by the sender of the current frame, between the end of an active superframe of the sender of the current frame and the start of the next active superframe of the recipient of the current frame.

5.3.9.6.8 Offered Inactive Duration

The Offered Inactive Duration field is set to the number of inactive superframes of the recipient of the current frame after each active superframe of the recipient, as offered by the sender of the current frame.

5.3.9.6.9 Current Allocation Slot Number

The Current Allocation Slot Number field is set to S such that the sender of the current frame starts sending this frame during the offered allocation slot numbered S .

5.3.9.6.10 Current Allocation Slot Offset

The Current Allocation Slot Offset field is set to F in units of $\lceil \text{Offered Allocation Slot Length in microseconds}/65536 \rceil$ microseconds such that the sender of the current frame starts sending this frame at F after the start of the offered allocation slot indicated in the preceding field. Here, the function $\lceil x \rceil$ is defined to be the least integer not smaller than x .

5.4 Control type frames

A control type frame contains no frame payload or a frame payload of a fixed or variable length.

5.4.1 Immediate Acknowledgement (I-Ack)

An I-Ack frame transmitted by a node to a hub contains no Frame Payload. An I-Ack frame transmitted by a hub to a node selectively contains a Frame Payload that is formatted as shown in Figure 37. An I-Ack frame is transmitted by a node or a hub to acknowledge receipt of the preceding frame, while optionally providing a timestamp by the hub in terms of a Current Allocation Slot Number and a Current Allocation Slot Offset in the Frame Payload for the node's clock synchronization.




Figure 37—Frame Payload format for I-Ack frames

The two fields shown in Figure 37 are either both present or both absent.

5.4.1.1 Current Allocation Slot Number

The Current Allocation Slot Number field is present only if the Ack Timing field of the MAC header of the current frame is set to one. When present, it is as defined in 5.4.6.1.

5.4.1.2 Current Allocation Slot Offset

The Current Allocation Slot Offset field is present only if the Ack Timing field of the MAC header of the current frame is set to one. When present, it is as defined in 5.4.6.2.

5.4.2 Block Acknowledgement (B-Ack)

A B-Ack frame selectively contains a Frame Payload that is formatted as shown in Figure 38. A timestamp comprising a Current Allocation Slot Number and a Current Allocation Slot Offset field is selectively present only in the Frame Payload of B-Ack frames transmitted by a hub to a node. A B-Ack frame is transmitted by a node or a hub to acknowledge the reception status of certain preceding data type frames each containing a whole MSDU, while optionally providing a timestamp by the hub in terms of a Current Allocation Slot Number and a Current Allocation Slot Offset in the Frame Payload for the node's clock synchronization.




Figure 38—Frame Payload format for B-Ack frames

The first two fields shown in Figure 38 are either both present or both absent. The final three fields shown in Figure 38 are either all present or all absent.

5.4.2.1 Current Allocation Slot Number

The Current Allocation Slot Number field is present only if the Ack Timing field of the MAC header of the current frame is set to one. When present, it is as defined in 5.4.6.1.

5.4.2.2 Current Allocation Slot Offset

The Current Allocation Slot Offset field is present only if the Ack Timing field of the MAC header of the current frame is set to one. When present, it is as defined in 5.4.6.2.

5.4.2.3 Oldest Frame Expected

The Oldest Frame Expected field is not present if only one new frame is expected and allowed. When present, it is set as follows:

- If one or more frames that are of the same frame subtype as, and older than, the frame preceding this B-Ack frame are still expected to be, but not yet, received, it is set to the sequence number of the oldest frame.
- Otherwise, it is set to the next expected sequence number, i.e., one plus SN modulo 256, where SN is the sequence number of the frame preceding this B-Ack frame.

5.4.2.4 Next Block Size

The Next Block Size field is not present if only one new frame is expected and allowed. When present, it is set to the maximum number of data type frames permitted in the next block transmission from the acknowledged node or hub to the acknowledging hub or node, where the next block transmission is a transmission of data type frames whose reception status will be provided in the next B-Ack frame and whose frame subtype is the same as that of the data type frame preceding this B-Ack frame.

5.4.2.5 Frame Status Bitmap

The Frame Status Bitmap field is not present if only one new frame is expected and allowed. When present, it is set as follows to indicate the reception status of up to mBackLimit frames that are of the same frame subtype as the frame preceding this B-Ack frame, and bounded in sequence between but not including the frame indicated in the Oldest Frame Expected field and the frame preceding this B-Ack frame:

- a) The LSB, i.e., bit 0, of the field denotes the oldest of these frames, i.e., the frame that is immediately subsequent in sequence to the frame indicated in the Oldest Frame Expected field.
- b) Each successive bit, up to and including bit $F-1$, denotes a successive frame, i.e., a frame with a successive sequence number, of these frames, where F is the number of these frames.
- c) A bit is set to one if it denotes a corresponding frame and the corresponding frame is received, or is set to zero otherwise.
- d) The field is set to zero if there are no such frames, i.e., if no previous frames are still expected.

5.4.3 Immediate Acknowledgement + Poll (I-Ack+Poll)

An I-Ack+Poll frame selectively contains a Frame Payload that is formatted as defined in 5.4.1. It is transmitted by a hub to acknowledge receipt of the preceding frame and to send a poll to the addressed node, while optionally providing a timestamp in terms of a Current Allocation Slot Number and a Current Allocation Slot Offset in the Frame Payload for the node's clock synchronization. An I-Ack+Poll frame is equivalent in function to an I-Ack frame followed by a Poll or T-Poll frame.

5.4.4 Block Acknowledgement + Poll (B-Ack+Poll)

A B-Ack+Poll frame selectively contains a Frame Payload that is formatted as defined in 5.4.2. It is transmitted by a hub to acknowledge the reception status of certain preceding data type frames and to send a poll to the addressed node, while optionally providing a timestamp in terms of a Current Allocation Slot Number and a Current Allocation Slot Offset in the Frame Payload for the node's clock synchronization. A B-Ack+Poll frame is equivalent in function to a B-Ack frame followed by a Poll or T-Poll frame.

5.4.5 Poll

A Poll frame contains no Frame Payload. It is transmitted by a hub to grant to the addressed node an immediate polled allocation that starts pSIFS after the end of the frame or to inform the node of a future poll or post.

5.4.6 Timed-Poll (T-Poll)

A T-Poll frame contains a Frame Payload that is formatted as shown in Figure 39. Except as stated otherwise, it is transmitted by a hub to grant to the addressed node(s) an immediate polled allocation that

starts pSIFS after the end of the frame or to inform the node of a future poll or post, while providing a timestamp by the hub in terms of a Current Allocation Slot Number, a Current Allocation Slot Offset, and selectively a Current Allocation Slot Length or a Current Superframe Number in the Frame Payload for the node's clock synchronization. A T-Poll frame is equivalent in function to a Poll frame expanded by a frame payload containing a transmit timestamp for superframe and allocation slot boundary synchronization and optionally a relay link quality for relay selection.




Figure 39—Frame Payload format for T-Poll frames

5.4.6.1 Current Allocation Slot Number

The Current Allocation Slot Number field is set to S such that the hub starts sending this frame during the allocation slot numbered S .

5.4.6.2 Current Allocation Slot Offset

The Current Allocation Slot Offset field is set to F in units of $\lceil \text{Allocation Slot Length in microseconds} / 65536 \rceil$ microseconds such that the hub starts sending this frame at F after the start of the allocation slot indicated in the preceding field. Here, the function $\lceil x \rceil$ is defined to be the least integer not smaller than x .

5.4.6.3 Current Allocation Slot Length

The Current Allocation Slot Length field is present only in T-Poll frames with the Recipient ID field of the MAC header set to the Unconnected_Broadcast_NID or Broadcast_NID per Table 17. When present, it is as defined in 5.3.1.3.

5.4.6.4 Current Superframe Number

The Current Superframe Number field is present only in T-Poll frames with the Recipient ID field of the MAC header set to the Broadcast_NID per Table 17. When present, it is as defined in 5.3.7.4.

5.4.6.5 Relay Link Quality

The Relay Link Quality field is optionally present in T-Poll frames with the Recipient ID field of the MAC header set to the Broadcast_NID per Table 17, which are transmitted by a node willing to support relay. When present, the Relay Link Quality field is set as follows:

- A bit of the field is set to one if in the active superframe (beacon period) designated by the bit, at least one frame was transmitted or received by the sender of the current frame, and more than half of the frames transmitted between this sender and its hub were received, or is set to zero otherwise. An expected beacon that was not received by this sender was considered to have been

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

transmitted. A frame sent by this sender was considered to have been received by the hub if its expected acknowledgment was received by this sender.

- b) The LSB, b₀, of the field designates the last active superframe (beacon period), and each successively more significant bit designates a successively earlier active superframe (beacon period).

5.4.7 Wakeup

A Wakeup frame contains a Frame Payload that is formatted as shown in Figure 40. It is optionally transmitted by a hub to wake up a node operating in the medical implant communications service (MICS) band.




Figure 40—Frame Payload format for Wakeup frames

5.4.7.1 Recipient Address

The Recipient Address field is set to the EUI-48 of the recipient of the current frame.

5.4.7.2 Sender Address

The Sending Address field is set to the EUI-48 of the sender of the current frame.

5.4.8 B2 Frame

A B2 frame contains a Frame Payload that is formatted as shown in Figure 41. It is optionally broadcast by a hub to announce B2-aided time-sharing information and/or provide group acknowledgment.




Figure 41—Frame Payload format for B2 frames

The B2-aided time-sharing information consists of the Beacon Period Length, Allocation Slot Length, Current Allocation Slot Number, Current Allocation Slot Offset, and CAP End fields, and is present only if the Coexistence field of the MAC header of the current B2 frame is set to one.

5.4.8.1 Beacon Period Length

The Beacon Period Length field, when present, is as defined in 5.3.1.2.

5.4.8.2 Allocation Slot Length

The Allocation Slot Length field, when present, is as defined in 5.3.1.3.

5.4.8.3 Current Allocation Slot Number

The Current Allocation Slot Number field, when present, is as defined in 5.4.6.1.

5.4.8.4 Current Allocation Slot Offset

The Current Allocation Slot Offset field, when present, is as defined in 5.4.6.2.

5.4.8.5 CAP End

The CAP End field, when present, is set to E such that the CAP ends at the end of the allocation slot that is numbered E and located in the current beacon period. The CAP starts at the end of this B2 frame. A B2 frame that does not contain the CAP End field defines a CAP that starts at the end of the B2 frame and ends at the end of the current beacon period.

5.4.8.6 Inactive Duration

The Inactive Duration field is as defined in 5.3.1.13.

5.4.8.7 NID

One or more NID fields are optionally present. When present, each NID field is set to the NID of a node from which this hub received a frame requiring group acknowledgment since the last transmitted B2 frame.

5.5 Data type frames

A data type frame contains a full, a fragmented, or no MSDU.

A data type frame of Emergency subtype, i.e., an Emergency frame, is transmitted to indicate an emergency or medical implant event report.

A User Priority UP frame ($UP = 0, 1, \dots$, or 6) is transmitted to indicate that the frame payload contained in the frame is identified by the frame subtype of the frame and has a user priority UP, if no IE of the Connection Request or Connection Assignment frame previously transmitted or received by the sender of the current frame contained an Allocation ID comprising the frame subtype value and the user priority UP.

An Allocation Mapped Data Subtype frame is transmitted to indicate that the frame payload contained in the frame is identified by the frame subtype of the frame and has a user priority UP, if an information element (IE) of the Connection Request or Connection Assignment frame previously transmitted or received by the sender of the current frame contained an Allocation ID comprising the frame subtype value and the user priority UP.

5.6 MAC/PHY Capability fields

A MAC Capability field and a PHY Capability field are included in beacons and some other management type frames. A hub and a node typically have different functional requirements for the capability of supporting a given function as listed in these fields, especially in the MAC Capability field. Indication of being capable of supporting a function implies being willing to enable and perform the function as well.

5.6.1 MAC Capability

The MAC Capability is formatted as shown in Figure 42.




Figure 42—MAC Capability format

5.6.1.1 CSMA/CA

The CSMA/CA field is set to one if the sender supports contention allocations obtained by using CSMA/CA in exclusive access phase 1 (EAP1), random access phase 1 (RAP1), exclusive access phase 2 (EAP2), random access phase 2 (RAP2), and contention access phase (CAP), or is set to zero otherwise.

5.6.1.2 Slotted Aloha Access

The Slotted Aloha Access field is set to one if the sender supports contention allocations obtained by using slotted Aloha access in exclusive access phase 1 (EAP1), random access phase 1 (RAP1), exclusive access phase 2 (EAP2), random access phase 2 (RAP2), and contention access phase (CAP), or is set to zero otherwise.

5.6.1.3 Type-I Polling Access

The Type-I Polling Access field is set to one if the sender supports type-I polled allocations, or is set to zero otherwise.

5.6.1.4 Type-II Polling Access

The Type-II Polling Access field is set to one if the sender supports type-II polled allocations, or is set to zero otherwise.

5.6.1.5 Scheduled Access

The Scheduled Access field is set to one if the sender supports scheduled allocations, or is set to zero otherwise. The sender supports scheduled bilink allocations if and only if it supports both scheduled allocations and type-I polled allocations.

5.6.1.6 Unscheduled Access

The Unscheduled Access field is

- a) set to one in beacon or non-beacon mode with superframes, if the sender supports unscheduled bilink allocations and type-I polled allocations and will be always in active state (abbreviated as always active) ready to receive and transmit frames during time intervals wherein polls and posts are allowed to be sent;
- b) set to one in non-beacon mode without superframes, if the sender supports unscheduled bilink allocations and type-II polled allocations; or
- c) is set to zero otherwise.

5.6.1.7 Fragmentation/Reassembly

The Fragmentation/Reassembly field is set to one if the sender supports fragmentation and reassembly, or is set to zero otherwise.

5.6.1.8 Command Frames

The Command Frames field is set to one if the sender supports the processing and functionality of Command frames, or is set to zero otherwise.

5.6.1.9 Node Always Active/Hub Clock PPM

The Node Always Active/Hub Clock PPM field is set as follows:

- a) In frames sent by a node, it is used as a Node Always Active field, which is set to one if the node will be always in active state (abbreviated as always active) ready to receive and transmit frames during time intervals wherein polls and posts are allowed to be sent, or is set to zero if the node will not be always in active state.
- b) In frames sent by a hub, it is used as a Hub Clock PPM field, which is set to one if the hub has a clock with a minimum accuracy of $\text{ppm} = \text{mHubClockPPMLimit}/2$, or is set to zero if the hub has a clock with a minimum accuracy of $\text{ppm} = \text{mHubClockPPMLimit}$.

5.6.1.10 Guard Time Provisioning

The Guard Time Provisioning field is set as follows:

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

- a) In frames sent by a node, it is set to one if the node supports and requires centralized guard time provisioning, or is set to zero if the node supports and requires distributed guard time provisioning.
- b) In frames sent by a hub, it is reserved.

5.6.1.11 L-Ack/B-Ack

The L-Ack/B-Ack field is set to one if the sender supports both L-Ack and B-Ack acknowledgment, or is set to zero otherwise.

5.6.1.12 G-Ack

The G-Ack field is set to one if the sender supports group acknowledgment, or is set to zero otherwise.

5.6.1.13 Relaying Node

The Relaying Node field is set to one if the sender is a node that supports the functionality required of a relaying node in a two-hop extended star BAN, or is set to zero if the sender is a node that does not support such a functionality. It is reserved if the sender is a hub.

5.6.1.14 Relayed Hub/Node

The Relayed Hub/Node field is set to one if the sender supports the functionality required of a relayed hub or node in a two-hop extended star BAN, or is set to zero otherwise.

5.6.1.15 Beacon Shifting

The Beacon Shifting field is set to one if the sender supports beacon shifting, or is set to zero otherwise.

5.6.1.16 Channel Hopping

The Channel Hopping field is set to one if the sender supports channel hopping, or is set to zero otherwise.

5.6.1.17 Data Subtypes

The Data Subtypes field is set to the maximum number of data subtypes supported by the sender for data type frames received from the recipient of the current frame. It is set to zero to encode a value of 16.

5.6.2 PHY Capability

The PHY Capability is formatted as shown in Figure 43.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 43—PHY Capability format

The Data Rate i field is set to one if the sender supports the information data rate for both transmission and reception in the operating frequency band, as encoded by $R2R1R0 = i$ of the Data Rate field defined in the corresponding PHY clause, or is set to zero otherwise. Table 15 is compiled based on the data rate specifications provided in the PHY clauses (where rsvd means *reserved* or *undefined*).

Table 15—Data Rate field representation

PHY	Frequency band (MHz), center frequency (MHz), or modulation	Data rate 0 (kb/s)	Data rate 1 (kb/s)	Data rate 2 (kb/s)	Data rate 3 (kb/s)	Data rate 4 (kb/s)	Data rate 5 (kb/s)	Data rate 6 (kb/s)	Data rate 7 (kb/s)
Narrow band (NB)	402 to 405	75.9	151.8	303.6	455.4	Rsvd	Rsvd	Rsvd	Rsvd
	420 to 450	75.9	151.8	187.5	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd
	863 to 870	101.2	202.4	404.8	607.1	Rsvd	Rsvd	Rsvd	Rsvd
	902 to 928	101.2	202.4	404.8	607.1	Rsvd	Rsvd	Rsvd	Rsvd
	950 to 958	101.2	202.4	404.8	607.1	Rsvd	Rsvd	Rsvd	Rsvd
	2360 to 2400	121.4	242.9	485.7	971.4	Rsvd	Rsvd	Rsvd	Rsvd
	2400 to 2483.5	121.4	242.9	485.7	971.4	Rsvd	Rsvd	Rsvd	Rsvd
Ultra wideband (UWB)	Non-coherent	394.8	789.7	1579	3159	6318	12 636	Rsvd	Rsvd
	Differentially coherent	487	975	1950	3900	7800	15 600	557	1114
	FM	202.5	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd	Rsvd
Human body communications (HBC)	21	164	328	656	1312.5	Rsvd	Rsvd	Rsvd	Rsvd

In calculating the transmission time of a MAC frame, care needs to be taken to not directly use a data rate given in this table, if block coding is applied to the MAC frame at the PHY before transmission. In particular, such a data rate does not account for the following factors: The information bits resulting from the MAC frame likely is not an integral number of the information bits contained in a codeword for the block coding, yet each codeword adds the same number of parity bits. The information or coded bits are not necessarily an integral number of the bits of each PHY symbol, and pad bits are added for symbol boundary alignment. Calculation of the frame transmission time is given in the corresponding PHY clause.

5.7 Information elements

An information element (IE) is formatted as shown in Figure 44. It is optionally contained in certain management type frames.




Figure 44—IE format—general

The Element ID field is set to the value that identifies the information element according to Table 16.

The Length field is set to the length, in octets, of the IE-specific Information field that follows.

The Information field is set based on the Element ID as defined in the remainder of this subclause.

Table 16—Information elements


Element ID in decimal value	IE name	Description
0	Superframe Parameters IE	Specifies superframe (beacon period) operation parameters
1	Uplink Request IE	Specifies allocation slot-based requirements by a node for scheduled uplink allocation(s) in beacon or non-beacon mode with superframes
2	Downlink Request IE	Specifies allocation slot-based requirements by a node for scheduled downlink allocation(s) in beacon or non-beacon mode with superframes
3	Bilink Request IE	Specifies allocation slot-based requirements by a node for scheduled bilink allocation(s) in beacon or non-beacon mode with superframes
4	Type-I Unscheduled Bilink Request IE	Specifies allocation slot-based requirements by a node for unscheduled bilink allocation(s) in beacon or non-beacon mode with superframes
5	Type-II Unscheduled Bilink Request IE	Specifies frame count-based requirements by a node for unscheduled bilink allocation(s) in non-beacon mode without superframes
6	Reserved	Reserved
7	Uplink Assignment IE	Specifies allocation slot-based scheduled uplink allocation(s) assigned to a node in beacon or non-beacon mode with superframes
8	Downlink Assignment IE	Specifies allocation slot-based scheduled downlink allocation(s) assigned to a node in beacon or non-beacon mode with superframes

Table 16—Information elements (continued)

Element ID in decimal value	IE name	Description
9	Bilink Assignment IE	Specifies allocation slot-based scheduled bilink allocation(s) assigned to a node in beacon or non-beacon mode with superframes
10	Type-I Unscheduled Bilink Assignment IE	Specifies allocation slot-based unscheduled bilink allocation(s) assigned to a node in beacon or non-beacon mode with superframes
11	Type-II Unscheduled Bilink Assignment IE	Specifies frame count-based unscheduled bilink allocation(s) assigned to a node in non-beacon mode without superframes
12	Reserved	Reserved
13	Nibble Encoded Channel Order IE	Specifies a list of 4-bit encoded channels in an operating band containing no more than 15 channels in the order of their selection by a hub as the operating channel
14	Channel Hopping and Ordering IE	Specifies a subset of channels included in channel hopping in the operating frequency band and/or a list of 8-bit encoded channels in the operating band in the order of their selection by a hub as the operating channel
15	Former Hub Address IE	Specifies the EUI-48 of the last hub with which the node was connected
16–244	Reserved	Reserved
255	Application Specific IE	Provides user-defined application-specific information

5.7.1 Superframe Parameters IE

The Superframe Parameters IE is formatted as shown in Figure 45. It is optionally contained in Connection Assignment frames to convey the values of chosen superframe (beacon period) operation parameters.

**Figure 45—Superframe Parameters IE format**

5.7.1.1 Superframe Parameters Bitmap

The Superframe Parameters Bitmap is formatted as shown in Figure 46. It indicates what superframe (beacon period) operation parameters are present in this IE in the order shown in Figure 45.




Figure 46—Superframe Parameters Bitmap format

5.7.1.1.1 Period/Slot Lengths Indicator

The Period/Slot Lengths Indicator field is set to one if both Beacon Period Length and Allocation Slot Length fields are present in this IE, or is set to zero if both fields are absent.

5.7.1.1.2 Beacon Shifting Sequence Indicator

The Beacon Shifting Sequence Indicator field is set to one if beacon shifting is currently enabled and the Beacon Shifting Sequence field is present in this IE, or is set to zero if the Beacon Shifting Sequence field is absent.

5.7.1.1.3 Channel Dwell Duration Indicator

The Channel Dwell Duration Indicator field is set to one if channel hopping is currently enabled and the Channel Dwell Duration field is present in this IE, or is set to zero if the Channel Dwell Duration field is absent.

5.7.1.1.4 Channel Hopping Indicator

The Channel Hopping Indicator field is set to one if channel hopping is currently enabled and both Channel Hopping State and Next Channel Hop fields are present in this IE, or is set to zero if both fields are absent.

5.7.1.1.5 Active/Inactive Superframe Indicator

The Active/Inactive Superframe Indicator field is set to one if inactive superframes are periodically provided and both Inactive Duration and Next Active Superframe fields are present in this IE, or is set to zero if both fields are absent.

5.7.1.2 Beacon Period Length

The Beacon Period Length field, when present, is as defined in 5.3.1.2.

5.7.1.3 Allocation Slot Length

The Allocation Slot Length field, when present, is as defined in 5.3.1.3.

5.7.1.4 Current Superframe Number

The Current Superframe Number field, when present, is as defined in 5.4.6.4.

5.7.1.5 Beacon Shifting Sequence

The Beacon Shifting Sequence field, when present, is as defined in 5.3.1.10.

5.7.1.6 Channel Dwell Duration

The Channel Dwell Duration field, when present, is set to the number of beacon periods (superframes) over which the hub sending the current frame is to dwell in any chosen operating channel before hopping to another one. When present, it is set to zero to encode a value of 256 beacon periods (superframes).

5.7.1.7 Channel Hopping State

The Channel Hopping State field, when present, is as defined in 5.3.1.11.

5.7.1.8 Next Channel Hop

The Next Channel Hop field, when present, is as defined in 5.3.1.12.

5.7.1.9 Inactive Duration

The Inactive Duration field, when present, is as defined in 5.3.1.13.

5.7.1.10 Next Active Superframe

The Next Active Superframe field, when present, is set to the sequence number of the next active beacon period (superframe).

5.7.2 Uplink Request IE

The Uplink Request IE is formatted as shown in Figure 47 (where N is the number of Allocation Request fields contained in the IE). It is optionally contained in Connection Request frames to request, using allocation slot-based requirements, for creation or modification of one or more scheduled uplink allocations in beacon or non-beacon mode with superframes.




Figure 47 —Uplink Request, Downlink Request, Bilink Request, or Type-I Unscheduled Bilink Request IE format

One or more Allocation Request fields are present. Each Allocation Request is formatted as shown in Figure 48 to describe the allocation slot-based requirements of an allocation for servicing the data belonging to a given user priority (UP).




Figure 48—Allocation Request format

5.7.2.1 Allocation ID

The Allocation ID identifies an allocation requested by the node. It is comprised of the Frame Subtype and User Priority fields as defined in 5.7.2.1.1 and 5.7.2.1.2.

5.7.2.1.1 Frame Subtype

The Frame Subtype field is set to the frame subtype of the data type frames to be transferred in this requested allocation.

5.7.2.1.2 User Priority

The User Priority field is set to the UP of the frame payloads to be transferred in this requested allocation according to 6.2.3.

5.7.2.2 Maximum Gap

The Maximum Gap field is set to the largest length, in units of allocation slots, between the end of an allocation interval and the start of the next allocation interval of this requested allocation in the same beacon period (superframe) or across beacon periods (superframes), if the Requested Wakeup Period field in the current frame has a value of one. It is reserved otherwise.

5.7.2.3 Minimum Gap

The Minimum Gap field is set to the smallest length, in units of allocation slots, between the end of an allocation interval and the start of the next allocation interval of this requested allocation in the same beacon period (superframe) or across beacon periods (superframes).

5.7.2.4 Minimum Length

The Minimum Length field is set to the smallest length, in units of allocation slots, of any of the allocation intervals of this requested allocation.

5.7.2.5 Allocation Length

The Allocation Length field is set to the overall length, in units of allocation slots, of the allocation intervals of this requested allocation in each wakeup beacon period (superframe) of this node.

5.7.2.6 Scaling Down Factor

The Scaling Down Factor field is set to D such that $\lceil (1 - D/16) \times \text{Allocation Length} \rceil$ is the smallest overall length, in units of allocation slots, of the allocation intervals this node is willing to accept for this requested allocation in each of its wakeup beacon periods (superframes). Here, $\lceil x \rceil$ is the least integer that is not smaller than x .

5.7.2.7 Scaling Up Factor

The Scaling Up Factor field is set to U such that $\lceil (1 + U/8) \times \text{Allocation Length} \rceil$ is the largest overall length, in units of allocation slots, of the allocation intervals this node is willing to accept for this requested allocation in each of its wakeup beacon periods (superframes). Here, $\lceil x \rceil$ is the least integer that is not smaller than x .

5.7.3 Downlink Request IE

The Downlink Request IE is as formatted in Figure 47 in conjunction with Figure 48 and as encoded in 5.7.2. It is optionally contained in Connection Request frames to request, using allocation slot-based requirements, for creation or modification of one or more scheduled downlink allocations in beacon or non-beacon mode with superframes.

5.7.4 Bilink Request IE

The Bilink Request IE is as formatted in Figure 47 in conjunction with Figure 48 and as encoded in 5.7.2. It is optionally contained in Connection Request frames to request, using allocation slot-based requirements, for creation or modification of one or more scheduled bilink allocations in beacon or non-beacon mode with superframes.

5.7.5 Type-I Unscheduled Bilink Request IE

The Type-I Unscheduled Bilink Request IE is formatted as shown in Figure 47 in conjunction with Figure 48. It is optionally contained in Connection Request frames to request, using allocation slot-based requirements, for creation or modification of one or more unscheduled bilink allocations in beacon or non-beacon mode with superframes.

5.7.6 Type-II Unscheduled Bilink Request IE

The Type-II Unscheduled Bilink Request IE is formatted as shown in Figure 49 (where M is the number of Type-II Unscheduled Allocation Request fields contained in the IE). It is optionally contained in Connection Request frames to request, using frame count-based requirements, for creation or modification of one or more unscheduled bilink allocations in non-beacon mode without superframes.




Figure 49—Type-II Unscheduled Bilink Request IE format

One or more Type-II Unscheduled Allocation Request fields are present. Each Type-II Unscheduled Allocation Request is formatted as shown in Figure 50 to describe the frame count-based requirements of an allocation for servicing the data belonging to a given UP.




Figure 50—Type-II Unscheduled Allocation Request format

5.7.6.1 Allocation ID

The Allocation ID is as defined in Figure 48.

5.7.6.1.1 Frame Subtype

The Frame Subtype field is as defined in 5.7.2.1.1.

5.7.6.1.2 User Priority

The User Priority field is as defined in 5.7.2.1.2.

5.7.6.2 Minimum Allocation Length Requested

The Minimum Allocation Length Requested field is set to the minimum number of non-control type frames to be transferred between the node and the hub in each allocation interval of this requested allocation subject to round-robin scheduling policy, whereby each allocation has one allocation interval among the allocation intervals of other allocations per round-robin cycle.

5.7.6.3 Nominal Allocation Length Requested

The Nominal Allocation Length Requested field is set to the expected number of non-control type frames to be transferred between the node and the hub in each allocation interval of this requested allocation subject to round-robin scheduling policy.

5.7.6.4 Maximum Allocation Length Requested

The Maximum Allocation Length Requested field is set to the maximum number of non-control type frames to be transferred between the node and the hub in each allocation interval of this requested allocation subject to round-robin scheduling policy.

5.7.7 Uplink Assignment IE

The Uplink Assignment IE is formatted as shown in Figure 51 (where J is the number of Allocation Assignment fields contained in the IE). It is optionally contained in Connection Assignment frames to assign or reassign one or more allocation slot-based scheduled uplink allocations to the addressed node in beacon or non-beacon mode with superframes.




Figure 51 —Uplink Assignment, Downlink Assignment, Bilink Assignment, or Type-I Unscheduled Bilink Assignment IE format

One or more Allocation Assignment fields are present. Each Allocation Assignment is formatted as shown in Figure 52 to specify an allocation interval of an assigned allocation for the data belonging to a given UP. One or more Allocation Assignment fields are included in this IE to specify an allocation. Still more Allocation Assignment fields are included in the IE to specify additional allocations.




Figure 52 —Allocation Assignment format

5.7.7.1 Allocation ID

The Allocation ID identifies the allocation being assigned or reassigned to the node. It is set to the Allocation ID used to identify the allocation requested earlier by the node.

5.7.7.1.1 Frame Subtype

The Frame Subtype field is set to the frame subtype of the data type frames to be transferred in this assigned allocation.

5.7.7.1.2 User Priority

The User Priority field is set to the UP of the frame payloads to be transferred in this assigned allocation.

5.7.7.2 Interval Start

The Interval Start field is set to S such that an allocation interval of this assigned allocation starts at the beginning of the allocation slot that is numbered S .

5.7.7.3 Interval End

The Interval End field is set to E such that the allocation interval whose start time is specified in the preceding field ends at the end of the allocation slot that is numbered E .

The Interval Start field is set to 255 and the Interval End field is set to zero, if the allocation is being ended.

5.7.8 Downlink Assignment IE

The Downlink Assignment IE is as formatted in Figure 51 in conjunction with Figure 52 and as encoded in 5.7.7. It is optionally contained in Connection Assignment frames to assign one or more allocation slot-based scheduled downlink allocations to the addressed node in beacon or non-beacon mode with superframes.

5.7.9 Bilink Assignment IE

The Bilink Assignment IE is as formatted in Figure 51 in conjunction with Figure 52 and as encoded in 5.7.7. It is optionally contained in Connection Assignment frames to assign one or more allocation slot-based scheduled Bilink allocations to the addressed node in beacon or non-beacon mode with superframes.

5.7.10 Type-I Unscheduled Bilink Assignment IE

The Type-I Unscheduled Bilink Assignment IE is as formatted in Figure 51 in conjunction with Figure 52 and as encoded in 5.7.7. It is optionally contained in Connection Assignment frames to assign or reassign one or more allocation slot-based unscheduled bilink allocations to the addressed node in beacon or non-beacon mode with superframes.

The values of the Interval Start and Interval End fields contained in this IE, i.e., the locations and lengths of the allocation intervals, are assigned on a tentative basis and subject to change within the wakeup beacon period (superframe) of the node. If the Assigned Wakeup Period field contained in the current frame is set to one, the availability of allocation intervals is further subject to round-robin scheduling policy, i.e., each allocation has one allocation interval among the allocation intervals of other allocations per round-robin cycle.

5.7.11 Type-II Unscheduled Bilink Assignment IE

The Type-II Unscheduled Bilink Assignment IE is formatted as shown in Figure 53 (where L in the Length field is the number of Type-II Unscheduled Allocation Assignment fields contained in the IE). It is optionally contained in Connection Assignment frames to assign or reassign one or more frame count-based unscheduled bilink allocations to the addressed node in non-beacon mode without superframes.




Figure 53—Type-II Unscheduled Bilink Assignment IE format

One or more Type-II Unscheduled Allocation Assignment fields are present. Each Type-II Unscheduled Allocation Assignment is formatted as shown in Figure 54 to specify an assigned allocation for the data belonging to a given UP.




Figure 54—Type-II Unscheduled Allocation Assignment format

5.7.11.1 Allocation ID

The Allocation ID identifies the allocation being assigned or reassigned to the node. It is set to the Allocation ID used to identify the allocation requested earlier by the node.

5.7.11.1.1 Frame Subtype

The Frame Subtype field is set to the frame subtype of the data type frames to be transferred in this assigned allocation.

5.7.11.1.2 User Priority

The User Priority field is set to the UP of the frame payloads to be transferred in this assigned allocation.

5.7.11.2 Minimum Allocation Length Assigned

The Minimum Allocation Length Assigned field is set to the minimum number of data type frames to be transferred between the node and the hub in each allocation interval of this assigned allocation subject to round-robin scheduling policy, whereby each allocation has one allocation interval per round-robin cycle.

5.7.11.3 Nominal Allocation Length Assigned

The Nominal Allocation Length Assigned field is set to the expected number of data type frames to be transferred between the node and the hub in each allocation interval of this assigned allocation subject to round-robin scheduling policy, whereby each allocation has one allocation interval per round-robin cycle.

5.7.11.4 Maximum Allocation Length Assigned

The Maximum Allocation Length Assigned field is set to the maximum number of data type frames to be transferred between the node and the hub in each allocation interval of this assigned allocation subject to round-robin scheduling policy, whereby each allocation has one allocation interval per round-robin cycle.

5.7.12 Nibble Encoded Channel Order IE

The Nibble Encoded Channel Order IE is formatted as shown in Figure 55. It is optionally contained in Connection Assignment frames to indicate some or all channels included in the operating channel selection in the operating frequency band that has no more than 15 operating channels and the order in which the operating channel is selected.




Figure 55—Nibble Encoded Channel Order IE format

The Nibble Encoded Channel List field is set as follows:

- The value of the four LSBs is set to the channel number of the channel that will be the first candidate to each instance of the operating channel selection.
- The value of each successive four bits is set to the channel number of the channel that will be the next candidate to the instance of the operating channel selection.
- If the list conveys an odd number of channels, four bits with a binary value of 1111 are padded as the MSBs to the field.

5.7.13 Channel Hopping and Ordering IE

The Channel Hopping and Ordering IE is formatted as shown in Figure 56. It is optionally contained in Connection Assignment frames to indicate the channels included in channel hopping in the operating frequency band or the order in which the operating channel is selected.




Figure 56—Channel Hopping and Ordering IE format

5.7.13.1 Channel Control

The Channel Control is formatted as shown in Figure 57.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 57 —Channel Control format

5.7.13.1.1 Channel Bitmap Length

The Channel Bitmap Length field is set to the length, a nonzero value, of the Channel Bitmap field in octets if channel hopping is currently enabled, or to be enabled later, in selected channels of the operating frequency band, or is set to zero otherwise. Whether channel hopping is currently enabled is indicated by the presence of the Channel Hopping State and Next Channel Hop fields in the last transmitted beacon or/and Connection Assignment frames.

5.7.13.1.2 Channel Separation

The Channel Separation field is set to the minimum number of channels separated between two consecutive hops when channel hopping is enabled. It is reserved if channel hopping is neither currently enabled nor to be enabled later. This field supersedes the default value of pChannelSeparation of Table 25.

5.7.13.2 Channel Bitmap

The Channel Bitmap field is present only if the Channel Bitmap Length field has a nonzero value. When present, it is set as follows:

- The LSB, bit 0, of the field denotes the lowest numbered channel of the operating frequency band.
- Each successive bit, up to and including bit $N-1$, denotes the next higher numbered channel of the operating frequency band, where N is the number of the channels in the operating frequency band.
- A bit is set to one if it denotes a corresponding channel and the corresponding channel is included in channel hopping, or is set to zero otherwise.


5.7.13.3 Octet Encoded Channel List

The Octet Encoded Channel List field is present only if an order in which the operating channel is selected is to be followed when channel hopping is not enabled. When present, it is set as follows with the channel number specified in the corresponding PHY clause:

- The value of the eight LSBs is set to the channel number of the channel that will be the first candidate to each instance of the operating channel selection.
- The value of each successive eight bits is set to the channel number of the channel that will be the next candidate to the instance of the operating channel selection.

5.7.14 Former Hub Address IE


The Former Hub Address IE is formatted as shown in Figure 58. It is optionally contained in Connection Request frames to convey the EUI-48 of the last hub with which this node was connected.

**Figure 58—Former Hub Address IE format**

The Former Hub Address field is set to the EUI-48 of the last hub with which the node was connected.

5.7.15 Application Specific IE

The Application Specific IE is formatted as shown in Figure 59. It is optionally contained in some management type frames to convey application-specific information.

**Figure 59—Application Specific IE format**

5.7.15.1 OUI-24

The OUI-24 field is set to the 24-bit Organizationally Unique Identifier (OUI) assigned by the IEEE Registration Authority to the vendor or manufacturer that defines this IE.

5.7.15.2 Application Specific Data

The Application Specific Data field is set by the owner (“assignee” in IEEE terms) of the OUI.

6. MAC functions

6.1 General

This clause specifies MAC sublayer functionality. It starts with the ground rules for preparing frame transmission and performing frame reception in 6.2. Subclause 6.3 then presents an access umbrella encompassing the access modes and structures provided in this standard for medium access. Before treating individual access methods, 6.4 describes the creation of a BAN by a hub, the connection of a node with a hub, and the disconnection of a node from a hub. Subclauses 6.5, 6.6, and 6.7 specify the operation of a variety of selectable access schemes for obtaining and using allocation intervals of preferred attributes. Subclause 6.8 provides access continuation, termination, and timeout rules in allocation intervals. Additional access steps are given in 6.9 for communications in the tightly regulated MICS band. Two-hop access extension of star BAN topology is optionally made available in 6.10.

Supplemental MAC functions follow in the remaining subclauses. Clock synchronization and guard time provisioning are addressed in 6.11. Power management is expounded in 6.12. Coexistence support and

interference mitigation are optionally offered in 6.13. MAC/PHY optional capability handling and Application Specific IE usage are elucidated in 6.14. MAC sublayer parameters are listed in 6.15.

6.2 Frame processing

This subclause provides fundamental rules on preparing MAC frames for transmission and processing them on reception.

6.2.1 Abbreviated addressing

A hub shall select a one-octet body area network identifier (BAN ID) from an integer between 0x00 and 0xFF, inclusive, as its BAN's abbreviated address contained in the MAC header of all frames sent from or to the hub. The hub should select a BAN ID that is currently not being used by neighbor BANs.

A hub shall select a one-octet hub identifier (HID) from an integer in the Connected_NID subset as specified in Table 17 as its abbreviated address contained in the MAC header of all frames sent from or to the hub. The hub shall not reselect an HID equal to a Connected_NID that is currently being used for a node connected with it. The hub should select an HID that is currently not being used by neighbor hubs.

A one-octet node identifier (NID) selected in accordance with Table 17 shall be used as a node's abbreviated address contained in the MAC header of all frames sent (unicast) from or to the node, or shall be used as a group of nodes' common abbreviated address contained in the MAC header of all frames sent (multicast or broadcast) to the nodes by a hub. The Broadcast_NID shall be the value of the Recipient ID field of the MAC header of beacon frames.

Table 17—NID selection

NID value in hex	NID subtotal	NID notation	NID usage
0x00	1	Unconnected_Broadcast_NID	For broadcast to unconnected nodes
0x01	1	Unconnected_NID	For unicast from/to unconnected nodes in a BAN
0x02–0xF5	244	Connected_NID	For unicast from/to connected nodes in a BAN
0xF6	1	Reserved	Reserved
0xF7–0xFD	7	Multicast_NID	For multicast to connected nodes in a BAN
0xFE	1	Local_Broadcast_NID	For broadcast to all nodes in a BAN
0xFF	1	Broadcast_NID	For broadcast to all nodes and hubs

An unconnected node without a Connected_NID shall choose the Unconnected_NID as its NID in sending to a hub a non-command management type frame and retries thereof. The node shall treat the Unconnected_NID or any Connected_NID as its NID in receiving the I-Ack frame expected to follow, as shown in Figure 60.

Upon receiving a non-command management type frame with the Sender ID field of the MAC header set to the Unconnected_NID, a hub shall keep the Unconnected_NID as the node's NID or shall assign as the node's NID a Connected_NID that is not being used for itself or another node in its BAN.

IEEE Std 802.15.6-2012
 IEEE Standard for Local and metropolitan area networks—
 Part 15.6: Wireless Body Area Networks

- If the hub does not assign a Connected_NID to the node, it shall use the Unconnected_NID as the Recipient ID of the MAC header in its I-Ack frame sent to the node, and shall not send any management or data type frames to the node.
- If the hub assigns a Connected_NID to the node, it shall use that Connected_NID as the Recipient ID of the MAC header in its I-Ack and all subsequent frames sent to the node. The hub may later decide to send no more management type frames to the node before sending a Connection Assignment frame to the node, as appropriate, such as due to a failed security association, a failed PTK creation, or an EUI-based access control.

The unconnected node shall treat the Unconnected_NID or any Connected_NID as its NID in receiving the I-Ack frame expected to follow, as shown in Figure 60.

- If the node receives the I-Ack frame with the Recipient ID field of the MAC header set to the Unconnected_NID, indicating that the hub is not at a position to assign a Connected_NID to the node, the node shall keep the Unconnected_NID as its NID expecting no management or data type frames from the hub.
- If the node receives the I-Ack frame with the Recipient ID field of the MAC header set to a Connected_NID, the node shall treat the Connected_NID as its NID in receiving the next management type frame from the hub, with the Connected_NID yet to be confirmed through the matching of the node's EUI-48 with the Recipient Address field of the frame payload. If the frame is received and contains the node's EUI-48 in the Recipient Address field of the frame payload, the node shall treat the Connected_NID as its NID, now confirmed, in subsequent frame exchanges with the hub. However, the node shall not send an I-Ack frame after the received frame, if it would not be able to send the I-Ack frame in time after checking the Recipient Address field of the received management type frame.
- If the node does not receive the I-Ack frame but instead receives a management type frame from the hub that contains its EUI-48 in the Recipient Address field of the frame payload, it shall treat the Connected_NID in the Recipient ID field of the MAC header as its NID, which is confirmed, in subsequent frame exchanges with the hub. As noted before, the node shall not send an I-Ack frame after the received frame, if it would not be able to send the I-Ack frame in time after checking the Recipient Address field of the received management type frame.

A node with an Unconnected_NID as its NID shall not send any I-Ack frames. A node with an unconfirmed Connected_NID as its NID shall not send any frames and shall be treated as unconnected with the hub providing the Connected_NID. A node with a confirmed Connected_NID as its NID shall be treated as unconnected with the hub until it has received a Connection Assignment frame from the hub and sent an I-Ack frame.

A node with an unconfirmed Connected_NID as its NID shall consider itself to have lost the Connected_NID after one of the following events occur:

- The node receives a management type frame with the Recipient Address of the frame payload not set to its EUI-48.
- The node has not received a management type frame with the Recipient Address of the frame payload set to its EUI-48 within an implementation-dependent expected timeout.

A node with a confirmed Connected_NID as its NID shall consider itself to have lost the Connected_NID and not connected with the hub after one of the following events occurs:

- The node receives a management type frame with the Recipient Address of the frame payload not set to its EUI-48.

- The node is not at *Connected* state per Figure 4 and has not received a Connection Assignment frame with the Recipient Address of the frame payload set to its EUI-48 within an expected time.
- The node receives a Disconnection frame with the Recipient Address of the frame payload set to its EUI-48.

In sending an I-Ack or B-Ack frame, a node shall set the Sender ID field of the MAC header to the Recipient ID of the MAC header of the frame that immediately preceded the I-Ack or B-Ack frame.




Figure 60—Node NID transition

6.2.2 Full addressing

A separate EUI-48 shall be used to uniquely identify a sender or a recipient when desired. In particular, the EUI-48 of a hub sending a beacon is included in the frame payload of the beacon. The EUI-48 values of both the sender and the recipient of other management type frames are included in the frame payload of those frames.

Exceptionally, an unconnected node may set to zero the Recipient Address field of the first management type frame it sends to a hub, if the node has not yet known the EUI-48 of the hub.

A recipient shall check the Recipient Address and Sender Address fields of the frame payload of a received management type frame to determine if the frame was indeed addressed to it from an expected sender, taking the aforementioned exception into account.

6.2.3 Priority mapping

UP values, when referenced in prioritizing medium access of data and management type frames, shall be determined based on the designation of frame payloads (traffic) contained in the frames according to Table 18. The traffic designation for background (BK), best effort (BE), excellent effort (EE), video (VI), voice (VO), and network control is based on some traffic types defined in Annex G.1 of IEEE Std 802.1D™-2004 [B2].

Table 18—User priority mapping

Priority	User priority	Traffic designation	Frame type
Lowest	0	Background (BK)	Data
	1	Best effort (BE)	Data
	2	Excellent effort (EE)	Data
	3	Video (VI)	Data
	4	Voice (VO)	Data
	5	Medical data or network control	Data or management
Highest	6	High-priority medical data or network control	Data or management
	7	Emergency or medical implant event report	Data

6.2.4 Frame reception

A node shall receive a frame if the following conditions are met unless specified otherwise:

- The Recipient ID field of the MAC header of the frame is set to its own NID, any applicable Unconnected_Broadcast_NID, multicast_NID, Local_Broadcast_NID, or Broadcast_NID.
- The Sender ID field of the MAC header of the frame is set to the HID of the desired hub with which to exchange frames.
- The BAN ID field of the MAC header of the frame is set to an expected value.
- The Protocol Version of the MAC header of the frame is set to a value it supports.
- The FCS of the frame is valid, i.e., equal to the FCS value it calculates over the applicable fields received.

A hub shall receive a frame if the following conditions are met unless specified otherwise:

- The Recipient ID field of the MAC header of the frame is set to its own HID.
- The Sender ID field of the MAC header of the frame is set to the NID of an expected sender or the Unconnected_NID.
- The BAN ID field of the MAC header of the frame is set to an expected value.
- The Protocol Version of the MAC header of the frame is set to a value it supports.
- The FCS of the frame is valid.

The node or the hub shall ignore a received frame, aside from performing applicable acknowledgment, whose frame payload has a Sender Address field that is not set to the EUI-48 of the expected sender or whose frame payload has a Recipient Address field that is not set to its own EUI-48.

The node or the hub shall ignore a received frame, aside from performing applicable acknowledgment that is detected to be a duplicate, as described in 6.2.10.

6.2.5 Frame sequencing

6.2.5.1 Management type frames

A sender shall exchange non-beacon management type frames with a recipient in the sequence as specified in Figure 4, with frame exchanges for security association, security disassociation, PTK creation, GTK distribution, connection, and disconnection further specified in the respective subclauses.

After sending a management type frame (denoted frame n) and expecting an I-Ack frame and then a management type frame (denoted frame $n+1$) from the recipient, if a node or a hub does not receive the expected I-Ack frame but instead receives frame $n+1$, the node or the hub shall consider frame n has been received and shall process frame $n+1$.

When fragmenting a frame payload that could otherwise be contained in a management type frame without length limitation, the sender shall extract the first fragment, the second fragment, and so on, in sequential octet order. The sender shall transmit the first fragment, then the second fragment, and so on, accordingly.

6.2.5.2 Data type frames

A sender may send a data type frame without a frame payload in a frame transaction, setting the Sequence Number field of the MAC header as if the frame contained a new MSDU.

A sender shall transmit MSDUs contained in data type frames of the same frame subtype and addressed to the same recipient(s) in the octet order in which they arrived at the local MAC service access point (SAP).

When fragmenting an MSDU, the sender shall extract the first fragment, the second fragment, and so on, in sequential octet order. The sender shall transmit the first fragment, then the second fragment, and so on, accordingly.

A sender may transmit an MSDU earlier than another MSDU, even if the former arrived at the local MAC SAP later than the latter, so long as the two MSDUs are contained in data type frames not of the same frame subtype or not addressed to the same recipient(s).

A recipient shall release to the MAC client MSDUs that were transmitted by the same sender and contained in data type frames of the same frame subtype in the octet order in which they were received.

6.2.6 Frame retry

A node or a hub may retry a frame, i.e., may retransmit a frame that was previously transmitted but not necessarily received, to the same recipient(s), as appropriate, taking into consideration such factors as delay requirements, fairness policies, channel conditions, and medium availability.

6.2.7 Frame timeout

A node or a hub shall treat an expected frame, such as an I-Ack or B-Ack frame, as not arriving, if mTimeOut after the end of the PHY preamble of the expected frame, it has not received a PHY preamble of a frame. It shall estimate the end of an expected I-Ack, B-Ack, I-Ack+Poll, or B-Ack+Poll frame determined not arriving by assuming that such a frame was of the length expected for the frame and was transmitted at the data rate currently applicable to that frame. It should estimate the end of any other expected frame determined not arriving by assuming that such an expected frame was of the length

expected for the frame and was transmitted at the highest mandatory data rate of the operating frequency band as specified in the corresponding PHY clause.

6.2.8 Frame separation

If a sender—a node or a hub—is to send a frame pSIFS or pMIFS after (the end of) the previous frame, the frame shall occur between pSIFS and pSIFS+pExtraIFS or between pMIFS and pMIFS+pExtraIFS, respectively, after the end of the previous frame, where the start of a frame and the end of a frame are defined in the specification of the underlying PHY. If the previous frame is an expected frame but determined not arriving per 6.2.7, its end time shall be estimated according to 6.2.7.

If a recipient—a node or a hub—is to receive a frame pSIFS or pMIFS after (the end of) the previous frame, it shall be ready to receive a frame no later than pSIFS or pMIFS, respectively, after the end of the previous frame, and shall not exit receive state earlier than mTimeout after the end of the PHY preamble of the expected frame.

In determining if a new frame transaction will fit into an allocation interval, a sender shall treat the value of the pSIFS or pMIFS involved in the frame transaction, if any, as no less than $pSIFS + 0.5 \times pExtraIFS$ or $pMIFS + 0.5 \times ExtraIFS$, respectively.

6.2.9 Frame acknowledgement

A node or a hub shall set the Ack Policy field of the MAC header of a frame to be transmitted according to Table 19. A recipient shall acknowledge a received frame, by sending an immediate acknowledgment (I-Ack) or block acknowledgment (B-Ack) frame, if the criteria in 6.2.4 for qualifying a frame as received are met and if required by the acknowledgment policy set in the frame as further described in the remainder of this subclause. In sending an I-Ack or B-Ack frame, the recipient may instead send an I-Ack+Poll or B-Ack+Poll frame as appropriate, not only providing frame acknowledgment but also granting an immediate polled allocation or announcing a future poll or post as specified in 6.6.1.

To send an I-Ack or B-Ack frame, a node or a hub

- shall use the lowest mandatory data rate of the operating frequency band as specified in the corresponding PHY clause, if the sender and the recipient are not at *Connected* state per Figure 4, or
- shall use the data rate indicated in the Assigned Ack Data Rates field of the last Connection Assignment frame exchanged with the recipient, if the sender and the recipient are at *Connected* state per Figure 4.

To send an I-Ack+Poll or B-Ack+Poll frame conveying no immediate polled allocation, a hub shall use the same data rate as currently applicable for an I-Ack or B-Ack frame.

Table 19—Acknowledgement (ACK) Policy field setting

Frame type name	Frame subtype name	Ack Policy field
Management	Beacon	N-Ack
Management	Security Association	I-Ack
Management	Security Disassociation	I-Ack
Management	PTK	I-Ack
Management	GTK	I-Ack
Management	Connection Request	I-Ack
Management	Connection Assignment	I-Ack
Management	Disconnection	I-Ack
Management	Command	I-Ack
Control	I-Ack	N-Ack
Control	B-Ack	N-Ack
Control	I-Ack+Poll	N-Ack
Control	B-Ack+Poll	N-Ack
Control	Poll	N-Ack
Control	T-Poll	N-Ack
Control	Wakeup	N-Ack
Control	B2	N-Ack
Data	Data subtype set to mG-AckDataSubtype	G-Ack
Data	User-defined data subtype other than mG-AckDataSubtype	N-Ack, I-Ack, L-Ack, or B-Ack

6.2.9.1 No acknowledgment (N-Ack)

A node or a hub may send a frame with the Ack Policy field set to N-Ack to require no acknowledgment at all, as permitted by Table 19. The sender may retry a data type frame with the Ack Policy field set to N-Ack, as appropriate.

A recipient shall not acknowledge a received frame containing an Ack Policy field set to N-Ack, either immediately or later.

6.2.9.2 Group acknowledgment (G-Ack)

A node, but not a hub, may send to a hub a frame with the Ack Policy field set to G-Ack if the following two conditions are satisfied:

- The frame is a data type frame with the frame subtype set to mG-AckDataSubtype.
- The hub supports G-Ack as indicated in its last transmitted MAC Capability field.

A hub shall acknowledge data type frames with the Ack Policy field set to G-Ack and the Frame Subtype field set to mG-AckDataSubtype that were received since its last B2 transmission by including the NIDs of the nodes from which those frames were received in the frame payload of the next B2 frame. The hub should send such a B2 frame as soon as permitted.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

The node may retry the frame requiring G-Ack upon failing to receive the expected B2 frame or the expected acknowledgment in the next received B2 frame.

An example for group acknowledgment (G-Ack) is shown in Figure 61.




Figure 61—Group acknowledgment (G-ACK)

6.2.9.3 Immediate acknowledgment (I-Ack)

A node or a hub may send a frame with the Ack Policy field set to I-Ack to require an immediate acknowledgment.

A recipient shall acknowledge a received frame with the Ack Policy field set to I-Ack by unconditionally sending back an I-Ack frame pSIFS after the end of the received frame.

Examples for immediate acknowledgment are shown in Figure 62.




Figure 62—Immediate acknowledgment (I-Ack) illustration

6.2.9.4 Block acknowledgment later (L-Ack) and block acknowledgment (B-Ack)

A source, either a node or a hub, may send a data type frame with the Ack Policy field set to B-Ack if the following two conditions are satisfied:

— The frame contains a whole MSDU.

— The recipient supports L-Ack/B-Ack as indicated in the latter's MAC Capability field.

The recipient shall acknowledge a received frame with the Ack Policy field set to B-Ack by unconditionally sending back a B-Ack frame pSIFS after the end of the received frame. The B-Ack frame shall contain a frame payload as defined in 5.4.2 unless the following two conditions are both true:

- No older frames of the same frame subtype as the last received frame are still expected to be received.
- Only one frame in the next block transmission is allowed.

A block transmission starts from the first frame sent after the last B-Ack frame received and ends with the next earliest frame with the Ack Policy field set to B-Ack.

The source may send a frame with the Ack Policy field set to L-Ack if the following conditions are all satisfied:

- The frame contains a whole MSDU.
- It has sent a frame of the same frame subtype with the Ack Policy field set to B-Ack and received a B-Ack frame acknowledging that frame and containing a frame payload.
- That B-Ack frame was the last B-Ack frame received from the recipient.

The source shall not transmit more frames in a block transmission than allowed as specified in the last B-Ack frame received. The source shall end a block transmission with a frame with the Ack Policy field set to B-Ack.

The source shall separate the frames in a block transmission within an allocation interval by pMIFS or pSIFS, depending on whether it is setting the block transmission as a burst mode transmission defined in the specification of the underlying PHY.

The source shall send frames in a block transmission in the order of increasing sequence number values, which are not necessarily consecutive if the block transmission contains retransmitted frames, considering that sequence number wraparound is also increasing the sequence number value. The source shall not retransmit frames that are older than the frame indicated in the Oldest Frame Expected field of the last B-Ack frame received. It should retransmit frames that were not received, but shall not retransmit frames that were received, as indicated in the Frame Status Bitmap field of that B-Ack frame, starting with the oldest frame expected or the next oldest frame still buffered. The source may discard frames if permitted by the application generating those frames, for example, due to buffer constraints or aging considerations.

The source, once starting a block transmission, shall not transmit frames of another frame type or subtype other than control type frames until it has sent a frame with the Ack Policy field set to B-Ack. Subject to this restriction, a block transmission may span more than one allocation interval or beacon period (superframe).

The recipient shall not acknowledge immediately a received frame with the Ack Policy field set to L-Ack. Rather, it shall indicate the reception status of the frames newer than the oldest frame still expected through the B-Ack frame it returns at the end of the current block transmission.

The source may retransmit in an appropriate time the last frame that had the Ack Policy field set to B-Ack or send a later frame of the same frame subtype with the Ack Policy field set to B-Ack, but shall not send any earlier frame of the same frame subtype, after failing to receive an expected B-Ack frame.

The recipient may implement a timeout that enables it to stop waiting for missing old frames if appropriate, hence allowing new MSDUs to be released to the MAC client and some B-Ack buffer resources to be freed without receiving those missing frames.

Examples for late acknowledgment (L-Ack) and block acknowledgment (B-Ack) are shown in Figure 63.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 63—Late acknowledgement (L-ACK) and block acknowledgement (B-Ack) illustration

6.2.10 Duplicate detection

A recipient shall record the sequence number and fragment number of the MAC header of the last non-beacon management type frame received from each sender. It shall treat the next management type frame received from the same sender with the same sequence number and fragment number of the MAC header as a duplicate frame and discard it upon acknowledging it if required by the Ack Policy field.

A recipient shall record the sequence number and fragment number of the MAC header of the last data type frame of each frame subtype received from each sender, for up to the maximum number of data subtypes it supports as indicated in its last transmitted MAC Capability field. It shall treat the next data type frame of the same frame subtype received from the same sender with the same sequence number and fragment number of the MAC header as a duplicate frame and discard it upon acknowledging it if required by the Ack Policy field.

6.2.11 Fragmentation and reassembly

A sender may fragment a frame payload that could otherwise be contained in a management type frame without length limitation, if the recipient supports fragmentation/reassembly as indicated in the latter's last transmitted MAC Capability.

A sender may fragment only MSDUs to be transferred in data type frames with the Ack Policy field set to N-Ack, G-Ack, or I-Ack. The sender shall not fragment MSDUs into data type frames if the recipient does not support fragmentation/reassembly as indicated in the latter's last transmitted MAC Capability field.

The sender shall not divide a frame payload of a management type frame or any MSDU into more than `mMaxFragmentCount` fragments. All fragments of the same frame payload of a management type frame or of the same MSDU shall have the same length, except the last one, which may be shorter but shall not be longer. The sender shall not alter the length of the fragments of the frame payload of a management type frame or of an MSDU, by refragmentation or recombination, in retransmitting them.

The sender shall complete the transmission of a fragmented frame payload of a non-beacon management type frame before transmitting another non-beacon management type frame or a fragment thereof. The sender may discard fragments of a frame payload of a non-beacon management type frame before they are transmitted or received, but shall not subsequently transmit any remaining fragments of the frame payload.

The sender shall complete the transmission of a fragmented MSDU before transmitting another MSDU or a fragment thereof, regardless of whether these MSDUs are contained in data type frames with the same or different frame subtypes. The sender may discard fragments of an MSDU before they are transmitted or received, but shall not subsequently transmit any remaining fragments of the MSDU.

A recipient shall completely reassemble an MSDU in the correct order before delivery to the MAC client. The recipient shall discard any MSDU with missing fragments.

6.3 Access classification and division

To provide or support time referenced allocations in its BAN, a hub shall establish a time base as specified in 4.4, which divides the time axis into beacon periods (superframes) regardless of whether it is to transmit beacons. In such cases, the hub shall transmit a beacon in each beacon period (superframe), except in inactive superframes, or shall not transmit a beacon in any superframe (beacon period). The hub may shift (rotate) its beacon transmission time, as specified in 6.13.1, from one offset from the start of current beacon period (superframe) to another offset from the start of next beacon period (superframe), thereby shifting the time reference for all scheduled allocations, to prevent large-scale repeated transmission collisions between its BAN and neighbor BANs.

In cases where a hub is not to provide or support time referenced allocations in its BAN, it may operate without a time base or superframes and hence without transmitting beacons at all.

Equivalently, a hub shall operate in beacon mode transmitting a beacon in every beacon period other than in inactive superframes to enable time referenced allocations; or shall operate in non-beacon mode transmitting no beacons, with superframes and allocation slots established if access to the medium in its BAN involves time referencing, or without superframes or allocation slots if access to the medium in its BAN involves no time referencing.

In summary, a hub shall operate in one of the following three access modes:

- Beacon mode with beacon periods (superframes);
- Non-beacon mode with superframes;
- Non-beacon mode without superframes.

6.3.1 Beacon mode with beacon periods (superframes)

In this mode, a hub shall organize applicable access phases in each active beacon period (superframe) as illustrated in Figure 64, where B stands for beacon (B). The hub may maintain I inactive superframes (beacon periods) after each active superframe (beacon period), if there are no allocation intervals scheduled in the inactive superframes, where I is a positive integer chosen by the hub. In an active superframe (beacon period), a hub shall transmit a beacon and may provide access phases. In an inactive superframe (beacon period), a hub shall not transmit any beacon and shall not provide any access phases.




Figure 64—Layout of access phases in a beacon period (superframe) for beacon mode

The hub shall place the access phases—exclusive access phase 1 (EAP1), random access phase 1 (RAP1), managed access phase (MAP), exclusive access phase 2 (EAP2), random access phase 2 (RAP2), another managed access phase (MAP), and contention access phase (CAP)—in the order stated and shown above. The hub may set to zero the length of any of these access phases, but shall not have RAP1 end before the guaranteed earliest time as communicated in Connection Assignment frames sent to nodes that are still connected with it. To provide a non-zero length CAP, the hub shall transmit a preceding B2 frame. The hub shall not transmit a B2 frame if the CAP that follows has a zero length, unless it needs to announce B2-aided time-sharing information and/or provide group acknowledgment.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

A node may obtain, and initiate frame transactions, in contended allocations in EAP1, RAP1, EAP2, RAP2, and CAP in any active superframe using CSMA/CA or slotted Aloha based random access as specified in 6.5.

Only in a MAP, as shown in Figure 65, may the hub

- arrange scheduled uplink allocation intervals, scheduled downlink allocation intervals, and scheduled bilink allocation intervals;
- provide unscheduled bilink allocation intervals; and
- improvise type-I, but not type-II, immediate polled allocation intervals and posted allocation intervals starting in this MAP.

In an EAP, RAP, or CAP, or MAP, as shown in Figure 64, the hub may also improvise future polls or posts starting and ending in a MAP as shown in Figure 65 (through Poll, T-Poll, I-Ack+Poll, and B-Ack+Poll frames as specified in Table 21).

These allocation intervals along with the corresponding access methods whereby they are obtained are illustrated in Figure 65.




Figure 65—Allocation intervals and access methods permitted in a managed access phase

6.3.2 Non-beacon mode with superframes

In this mode, a hub may have only a managed access phase (MAP) as described in 6.3.1 and depicted in Figure 65, in any superframe (beacon period) as illustrated in Figure 66.




Figure 66—Layout of access phases in a superframe (beacon period) for non-beacon mode

6.3.3 Non-beacon mode without superframes

In this mode, a hub may provide unscheduled bilink allocation intervals comprising type-II polled allocations and/or posted allocations, as illustrated in Figure 67. After determining that the hub for the next frame exchange is operating in non-beacon mode without superframe boundaries, a node may treat any time interval as a portion of EAP1 or RAP1 and employ CSMA/CA based random access to obtain a contended allocation as specified in 6.5.1.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 67—Allocation intervals and access methods permitted for non-beacon mode without superframes

6.4 BAN creation/operation and node connection/disconnection

6.4.1 BAN creation/operation

A hub shall choose an operating channel to start a BAN, based on policy regulations, channel conditions, application requirements, coexistence considerations, etc. The hub shall choose a new channel when required by regulations as applicable in the MICS band. The hub may hop to new channels periodically, as specified in 6.13.2, to effect frequency diversity and interference mitigation.

A node shall find the operating channel of the hub it needs to communicate with before sending a frame to the hub, unless otherwise stated (such as for an emergency event report in the MICS band specified in 6.9.3).

The hub shall choose and enable an applicable access mode as described in 6.3 to support desired access methods.

If the hub selected non-beacon mode without superframes, it shall transmit Poll frames each addressed to Unconnected_Broadcast_NID and providing a type-II polled allocation to enable unconnected nodes' connection or reconnection with it.

If the hub selected non-beacon mode with superframes, it shall transmit T-Poll frames each addressed to Unconnected_Broadcast_NID and providing a type-I polled allocation to enable unconnected nodes' connection or reconnection with it.

If the hub selected beacon mode with superframes, it may transmit T-Poll frames each addressed to Unconnected_Broadcast_NID and providing a type-I polled allocation to facilitate unconnected nodes' connection or reconnection with it.

An unconnected polled allocation is a type-I or type-II polled allocation granted by a hub via a Poll or T-Poll frame addressed to Unconnected_Broadcast_NID.

6.4.2 Node connection

An unconnected node shall send a Connection Request frame to a hub, and the hub shall send a Connection Assignment frame to the node, for the node to be connected with the hub per Figure 4 and as further shown in Figure 68.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

In view of the node NID transition procedure specified in 6.2.1, the hub may need to retry its Connection Assignment frame before receiving an expected I-Ack frame, if the node needs to take time to confirm its newly assigned Connected_NID through a Connection Assignment frame, such as in the case where the node has lost its connection (and hence its Connected_NID) or is at Orphan state attempting to start unsecured communication with the hub.




Figure 68—Connection procedure

The unconnected node may send at most a management type frame to the hub in an unconnected polled allocation with contention probability (CP) $P = 1/2$. If the node does not receive an expected acknowledgment, it may retry it in another unconnected polled allocation with CP $P = \max(1/8, (1/2) / \lceil (R+1)/2 \rceil)$, where R counts the retries of the frame, i.e., R equals 1 for the first retry of the frame, 2 for the second retry, and so on. The function $\lceil x \rceil$ is defined to be the least integer not smaller than x . With CP P , the node shall transmit if $z \leq P$ or shall not otherwise, where z is a value the node has newly drawn at random from the interval $[0, 1]$.

In beacon or non-beacon mode with superframes, the hub shall not provide unconnected polled allocations with a time duration shorter than `pUnconnectedPolledAllocationMin`, which is the longest time required for an unconnected node to send a management type frame containing no information elements (IEs) at the highest mandatory data rate of the operating frequency band as specified in the corresponding PHY clause and for the hub to send an I-Ack frame at the lowest mandatory data rate of the operating frequency band as also specified in the corresponding PHY clause.

In beacon or non-beacon mode with superframes, the unconnected node, if allowed to send, may send a management type frame with or without IEs at any mandatory data rate of the operating frequency band as specified in the corresponding PHY clause, so long as the combined transmission time of this frame and the expected I-Ack frame, plus pSIFS, plus an appropriate guard time, fits into the provided polled allocation. The unconnected node may later send another management type frame of the same frame subtype, such as Connection Request, with updated fields and/or additional IEs (e.g., for Requested Wakeup Phase and Requested Wakeup Period fields, and Uplink Request, Downlink Request, and Bilink Request IEs), if it needed more time to calculate the values of these fields and/or IEs after processing the timing information contained in the last received T-Poll frame addressed to the `Unconnected_NID`.

In non-beacon mode without superframes, the unconnected node, if allowed to send, may send a management type frame with or without IEs at any mandatory data rate of the operating frequency band as specified in the corresponding PHY clause.

After the hub assigns a `Connected_NID` to an unconnected node through a management type frame sent to the node, it should provide polled allocations specifically addressed (unicast) to this node to facilitate the transmission by the unconnected node. The unconnected node should stay in active state until it does not need polled or posted allocations that are not announced in advance.

Figure 69 illustrates how to instantiate or obtain unconnected and unicast polled allocations for or by an unconnected node.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 69—Unconnected and unicast polled allocations for unconnected nodes

A node may also use CSMA/CA or slotted Aloha access as described in 6.5, if capable, to transmit its frames before and after it is connected with a hub.

6.4.3 Node disconnection

A node or a hub may send a Disconnection frame as shown in Figure 70 to end their connection, i.e., to void the node's Connected_NID, wakeup arrangement, and any scheduled and unscheduled allocations with the hub.




Figure 70—Disconnection procedure

The hub should send a Disconnection frame to the node after receiving an I-Ack frame from the node following a Connection Assignment frame sent to the node with a connection status indicating connection request rejected for some reason per Table 12, so as to reclaim the Unconnected_NID from the node, which would not be able to communicate with the hub anyway.

6.5 Random access

In exclusive access phase 1 (EAP1), random access phase 1 (RAP1), exclusive access phase 2 (EAP2), random access phase 2 (RAP2), and contention access phase (CAP), as depicted in 6.3, allocations may

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

only be contended allocations, which are non-reoccurring time intervals valid per instance of access. The access method for obtaining the contended allocations shall be

- CSMA/CA as specified in 6.5.1 if pRandomAccess is set to CSMA/CA, or
- slotted Aloha access as specified in 6.5.2 if pRandomAccess is set to Slotted Aloha.

A hub or a node may obtain contended allocations in EAP1 and EAP2, only if it needs to send data type frames of the highest UP (i.e., containing emergency information) as defined in Table 18. The hub may obtain such a contended allocation pSIFS after the start of EAP1 or EAP2 without actually performing the CSMA/CA or slotted Aloha access procedure. Only nodes may obtain contended allocations in RAP1, RAP2, and CAP, to send management or data type frames.

To obtain contended allocations in EAP1, RAP1, EAP2, or RAP2 of a beacon period in beacon mode with superframes, a node shall first receive the beacon that specifies the start and end times of these access phases.

To send data type frames of the highest UP based on CSMA/CA, a hub or a node may treat the combined EAP1 and RAP1 as a single EAP1, and the combined EAP2 and RAP2 as a single EAP2, so as to allow continual invocation of CSMA/CA and improve channel utilization. To send data type frames of the highest UP based on slotted Aloha access, a hub or node may treat RAP1 as another EAP1 but not a continuation of EAP1, and RAP2 as another EAP2 but not a continuation of EAP2, due to the time slotted attribute of slotted Aloha access.

Prioritized access for traffic of differing user priorities (UPs) shall be attained through the predefined relationships in Table 18 and Table 20 between contention window (CW) bounds CWmax and CWmin and UP for CSMA/CA and between contention probability (CP) thresholds CPmax and CPmin and UP for slotted Aloha access.

**Table 20—Contention window bounds for CSMA/CA
and contention probability thresholds for slotted Aloha access**

User Priority	CSMA/CA		Slotted Aloha access	
	CWmin	CWmax	CPmax	CPmin
0	16	64	1/8	1/16
1	16	32	1/8	3/32
2	8	32	1/4	3/32
3	8	16	1/4	1/8
4	4	16	3/8	1/8
5	4	8	3/8	3/16
6	2	8	1/2	3/16
7	1	4	1	1/4

6.5.1 CSMA/CA

To employ CSMA/CA, a node shall maintain a back-off counter and contention window to determine when it obtains a new contended allocation as described in 6.5.1.1, and shall initialize the back-off counter to zero.

6.5.1.1 Starting a contended allocation

To obtain a new contended allocation, a node shall set its back-off counter to a sample of an integer random variable uniformly distributed over the interval $[1, CW]$, when its back-off counter has a value of zero and the node has at least one frame of user priority UP or higher to transmit or retransmit, where CW is a contention window chosen as follows.

- If the node did not obtain any contended allocation previously, it shall set the CW to $CW_{min}[UP]$.
- If the node succeeded, i.e., if the node received an expected acknowledgment, I-Ack or B-Ack frame, to its last frame transmission, in the last contended allocation it had obtained, it shall set the CW to $CW_{min}[UP]$.
- If the node transmitted a frame requiring no acknowledgment, late acknowledgment, or group acknowledgment at the end of its last contended allocation, it shall keep the CW unchanged.
- If the node failed, i.e., if the node did not receive an expected acknowledgment to its last frame transmission, in the last contended allocation it had obtained,
 - 1) it shall keep the CW unchanged if this was the m th time the node had failed consecutively, where m is an odd number;
 - 2) it shall double the CW if this was the n th time the node had failed consecutively, where n is an even number.
- If doubling the CW would have the new CW exceed $CW_{max}[UP]$, the node shall set the CW to $CW_{max}[UP]$.

The node shall lock the back-off counter when any of the following events occurs:

- The back-off counter is reset upon decrementing to zero.
- The channel is busy. If the channel is busy because the node detected a frame transmission, the channel remains busy until at least the end of the frame transmission without the node having to re-sense the channel.
- The current time is outside any RAP or CAP and $UP < 7$ (i.e., not for an emergency according to Table 18), or is outside any EAP, RAP, or CAP and $UP = 7$ (i.e., for an emergency).
- The current time is at the start of a CSMA slot within an EAP, RAP, or CAP, but the time between the end of the slot and the end of the EAP, RAP, or CAP is not long enough for completing a frame transaction.

The node shall unlock the back-off counter when both of the following conditions are met:

- The channel has been idle for pSIFS within a RAP or CAP and $UP < 7$ (i.e., not for an emergency according to Table 18), or within an EAP, RAP, or CAP and $UP = 7$ (i.e., for an emergency).
- The time duration between the current time plus a CSMA slot and the end of the EAP, RAP, or CAP is long enough for completing a frame transaction.

Upon unlocking its back-off counter, the node shall decrement its back-off counter by one for each idle CSMA slot that follows, as shown in Figure 71. In particular, the node shall treat a CSMA slot to be idle if it determines that the channel has been idle between the start of the CSMA slot and pCCATime later, decrementing the back-off counter effectively pCCATime after the start of the CSMA slot, so that the node will transmit a frame to the transport medium (i.e., air) at the end of the CSMA slot in case its back-off counter reaches zero, as further described in the remainder of this subclause. Each CSMA slot shall have a fixed duration of pCSMASlotLength.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 71 —CSMA slot structure

If the back-off counter reaches zero in the current CSMA slot, the node shall have obtained a contended allocation that starts at the end of the current CSMA slot and ends at or by the end of the current RAP or CAP. Figure 72 illustrates how to start and use contended allocations based on CSMA/CA.




Figure 72 —CSMA/CA Illustration

6.5.1.2 Using a contended allocation

A node may transmit a frame for which it obtained the contended allocation at the start of the allocation. A hub shall be ready to receive frames in EAP1, RAP1, EAP2, RAP2, and CAP it provides, accounting for an appropriate guard time.

After receiving an expected acknowledgment frame, in the current contended allocation the node may retransmit an old frame or transmit a new frame, with a UP not lower than the UP used to obtain the current contended allocation, pSIFS after the end of the acknowledgment frame. After sending a frame with the Ack Policy field of the MAC header not set to I-Ack or B-Ack, in the current contended allocation the node may retransmit the previous frame or transmit a new frame with a UP not lower than the UP used to obtain the allocation, pSIFS or pMIFS as appropriate after the end of the previous frame. However, after failing to receive an expected acknowledgment frame, the node shall not retransmit an old frame or transmit a new frame in the current contended allocation.

The node shall end its transmission in the allocation such that the last transmission in the allocation completes by the end of the allocation.

The node shall not send more than mCSMATxLimit management or data type frames in a contended allocation.

6.5.1.3 Modifying a contended allocation

The node shall not modify a contended allocation to a different location without performing another instance of CSMA/CA.

6.5.1.4 Ending a contended allocation

A node may at any time end a contended allocation by not following with another frame transmission in the allocation.

A node shall treat a contended allocation to have been ended after failing to receive an expected acknowledgment frame in the allocation.

A node may start a new contended allocation procedure as specified in 6.5.1 to obtain another contended allocation.

6.5.2 Slotted Aloha access

To employ slotted Aloha access, a node shall maintain a CP to determine if it obtains a new contended allocation in an Aloha slot as described in 6.5.2.1.

6.5.2.1 Starting a contended allocation

To obtain a new contended allocation for the transmission or retransmission of a frame of user priority *UP* or higher, a node shall set its CP as follows.

- If the node did not obtain any contended allocation previously, it shall set the CP to $CP_{max}[UP]$.
- If the node succeeded, i.e., if the node received an expected acknowledgment, I-Ack or B-Ack frame, to its last frame transmission, in the last contended allocation it had obtained, it shall set the CP to $CP_{max}[UP]$.
- If the node transmitted a frame requiring no acknowledgment, late acknowledgment, or group acknowledgment at the end of its last contended allocation, it shall keep the CP unchanged.
- If the node failed, i.e., if the node did not receive an expected acknowledgment to its last frame transmission, in the last contended allocation it had obtained,
 - 1) it shall keep the CP unchanged if this was the m th time the node had failed consecutively, where m is an odd number;
 - 2) it shall halve the CP if this was the n th time the node had failed consecutively, where n is an even number.

If halving the CP would make the new CP smaller than $CP_{min}[UP]$, the node shall set the CP to $CP_{min}[UP]$.

With the CP set to CP as stated above, the node shall have obtained a contended allocation delimited by the current Aloha slot if $z \leq CP$ or shall not have otherwise, where z is a value the node has newly drawn at random from the interval $[0, 1]$. Each Aloha slot shall be of equal length predetermined to be $pAlohaSlotLength$ as shown in Figure 73.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 73—Aloha slot structure

An Aloha slot available to the node for contention shall start at the beginning of a RAP or CAP if $UP < 7$ (i.e., not for an emergency according to Table 18), or after the start of an EAP, RAP, or CAP if $UP = 7$ (i.e., for an emergency). Successive Aloha slots may be available to the node for contention if they are fully located within the EAP, RAP, or CAP. Figure 74 illustrates how to start and use contended allocations based on slotted Aloha access.




Figure 74—Slotted Aloha access illustration

6.5.2.2 Using a contended allocation

A node may transmit a frame for which it obtained the contended allocation in the allocation. A hub shall be ready to receive frames in EAP1, RAP1, EAP2, RAP2, and CAP it provides, accounting for an appropriate guard time

After receiving an or no expected acknowledgment frame, in the current contended allocation the node may retransmit an old frame or transmit a new frame, with a UP not lower than the UP used to obtain the current contended allocation, pSIFS after the end of the acknowledgment frame. After sending a frame with the Ack Policy field of the MAC header not set to I-Ack or B-Ack, in the current contended allocation the node may retransmit the previous frame or transmit a new frame with a UP not lower than the UP used to obtain the allocation, pSIFS or pMIFS as appropriate after the end of the previous frame.

The node shall end its transmission in the allocation such that the last transmission in the allocation completes by the end of the allocation.

6.5.2.3 Modifying a contended allocation

The node shall not modify a contended allocation obtained in an Aloha slot to a different Aloha slot without performing another instance of slotted Aloha access.

6.5.2.4 Ending a contended allocation

A node may at any time end a contended allocation by not following with another frame transmission in the allocation.

A node may start a new contended allocation procedure as specified in 6.5.2.1 to obtain another contended allocation in a different Aloha slot.

6.6 Improvised access and unscheduled access

A hub may employ improvised access

- as an independent access method to send polls or posts on a best-effort basis, without advance reservation and assignment via Connection Request and Connection Assignment frames;
- as a supplemental access method to scheduled access and unscheduled access to send additional polls and posts outside scheduled allocations and unscheduled bilink allocations; and
- as an enabling access method for scheduled-polling access and unscheduled access to send polls or posts inside scheduled bilink allocations and unscheduled bilink allocations.

A hub and a node that support unscheduled access as indicated in their last exchanged MAC Capability field may employ unscheduled access to initiate frame transactions in a downlink and/or uplink on a best-effort basis, with advance reservation and tentative assignment via Connection Request and Connection Assignment frames. To support unscheduled access in beacon or non-beacon mode with superframes, a node shall be always active during time intervals wherein polls and posts are allowed to be sent.

6.6.1 Improvised access

As characterized in 6.3, in beacon or non-beacon mode with superframes, a hub may employ improvised access to send polls and posts to a node without advance notice or at previously announced times according to Table 21 and as illustrated in Figure 75, thus granting polled or posted allocations in any access mode.

A poll is a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame, whereas a post is a management or data type frame. A polled or posted allocation is started by a poll and bounded by an explicit or implicit time interval that does not reoccur subsequently without the hub invoking another instance of improvised access.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 21 —Polls for improvised access initialization

Beacon or non-beacon mode with superframes (Access Mode = 0)				
More Data	Poll-Post Window	Next	Upon sending to a node a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame with the fields of the MAC header listed on the left and set as given below, a hub	Upon receiving from a hub such a frame with the fields of the MAC header listed on the left and set as given below, a node
0	E	N	shall grant to the node a type-I (immediate) polled allocation that starts pSIFS thereafter and ends at the end of the allocation slot that is numbered <i>E</i> and located in the current beacon period (superframe) for <i>N</i> = 0 or in the next <i>N</i> th beacon period (superframe) not counting the current one for <i>N</i> > 0 (a case possible if no beacon is to be sent in the polled allocation).	shall initiate a frame transaction starting pSIFS later, if it supports type-I polling access, if it has management or data type frames to send, and if the frame transaction and an appropriate guard time fit into the allocation.
1	S	F	shall grant to the node no immediate polled allocation, but shall send to the node a (future) poll or post at the start of the allocation slot that is numbered <i>S</i> and located in the current beacon period (superframe) for <i>F</i> = 0 or in the next <i>F</i> th beacon period (superframe) not counting the current one for <i>F</i> > 0 (a case possible if no beacon is to be sent in the polled allocation), unless the future poll or post is cancelled by an I-Ack, B-Ack, Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame with the Cancel field of the MAC header set to one that is subsequently sent to the node before the start of the future poll or post.	shall be ready to receive a poll or post from the hub at the preannounced time, taking into account an appropriate guard time, unless the (future) poll or post is cancelled by an I-Ack, B-Ack, Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame with the Cancel field of the MAC header set to one that is subsequently received from the hub before the start of the future poll or post..
Non-beacon mode without superframes (Access Mode = 1)				
Poll-Post Window	Upon sending to a node a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame with the field of the MAC header listed on the left and set as given below, a hub		Upon receiving from a hub a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame with the field of the MAC header listed on the left and set as given below, a node that supports type-II polling access and has data or management type frames to send	
<i>M</i>	shall grant a type-II (immediate) polled allocation that starts pSIFS thereafter and in which the node is allowed to send up to <i>M</i> .		shall send up to <i>M</i> management or data type frames starting pSIFS later, but shall not send any more frames (except to retransmit the last frame) in the current allocation after sending a frame with the Ack Policy field of the MAC header set to I-Ack or B-Ack.	


IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

M/D = management or data type frame


Ack = I-Ack or B-Ack frame Ack+Poll = I-Ack+Poll or B-Ack+Poll frame Poll = Poll or T-Poll frame

M = More Data L = Last Frame A = Access Mode P = Poll-Post Window N = Next

(N1 = N2 = N3 = 0 for A = 1)



(a) For immediate polled allocations



(b) For future polls and posts (and immediate polled allocations)

Figure 75—Example frame transactions for improvised access

6.6.1.1 Polling access

A hub may send polls and grant type-I or type-II polled allocations to a node only if both of them support polling access of the corresponding type as indicated in their last exchanged MAC Capability field (i.e., with Type-I or Type-II Polling Access field set to one). A hub may send to a node a poll conveying no polled allocations but a future poll or post, even if the node does not support polling access, treating the future poll or post as being sent in a posted allocation. A hub may send polls addressed to Unconnected_Broadcast_NID, granting type-I or type-II polled allocations as described in 6.4 for management type frame transmissions by unconnected nodes without a Connected_NID, without regard to the polling access settings in the MAC Capability indicated by the nodes in its BAN.

6.6.1.1.1 Starting a polled allocation

To obtain a polled allocation for initiating a frame transaction with a hub, a node shall set to one the More Data field of the MAC header of the frame it is transmitting to the hub.

The hub shall send to the node a poll conveying a polled allocation as soon as feasible and according to 6.8. The hub may send an I-Ack+Poll or B-Ack+Poll frame through a required acknowledgment in an uplink allocation interval of the node, if the newly granted polled allocation interval extends the existing uplink allocation interval, as illustrated in Figure 75(a).

To grant a polled allocation to a node that is not always active, a hub shall send a poll to the node at a preannounced time as specified in Table 21 or through access continuation or termination as specified in Table 22 and Table 23.

A hub may send Poll or T-Poll frames to a node that is always active as indicated in the node's last transmitted MAC Capability field at times not indicated earlier to the node in any time intervals wherein such frames may be sent as specified in 6.3.

6.6.1.1.2 Using a polled allocation

Following the last frame transaction in a polled allocation in beacon or non-beacon mode with superframe boundaries, the node may initiate another frame transaction pSIFS later, regardless of whether it received the acknowledgment frame if immediate or block acknowledgment was expected, if the current frame transaction and an appropriate guard time fit into the allocation, and if appropriate as specified in 6.8.1.

The management or data type frame(s) in any frame transaction may be new or old (retried), as appropriate.

6.6.1.1.3 Modifying a polled allocation

A hub may modify a polled allocation of a node by sending to the node within the polled allocation an I-Ack+Poll or B-Ack+Poll frame that extends the allocation interval, effectively granting a new polled allocation interval in place of the remaining allocation interval.

A hub shall not cancel a polled allocation granted via a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame by sending another frame, except for extending the existing polled allocation interval.

6.6.1.2 Posting access

A hub may send posts to a node while granting posted allocations, even if the node does not support polling access, so long as it has informed the node of the transmit times of the posts through previously transmitted frames or if the node has indicated to be always active in its last transmitted MAC Capability field (i.e., with Always Active field set to one).

6.6.1.2.1 Starting a posted allocation

To obtain a posted allocation for receiving a post from a hub while not in a scheduled downlink or bilink allocation interval, a node that is not always active may send to the hub in an uplink allocation interval a management or data type frame with the Ack Policy field of the MAC header set to I-Ack or B-Ack, thus enabling the hub to improvise a future post to the node, through a corresponding I-Ack+Poll or B-Ack+Poll frame as described in Table 21 and illustrated in Figure 75(b).

Alternatively, the node may send to the hub in an uplink allocation interval a non-Emergency data type frame without a frame payload and with the Ack Policy, More Data, and Last Frame fields of the MAC header set to N-Ack, zero, and one, respectively, thus relinquishing the current allocation interval for the hub to reclaim the interval as described in Table 22 and send a post to the node pSIFS later if appropriate.

After receiving such a data type frame from the node, the hub should send to the node a post, if any, pSIFS later. The node should be ready to receive an expected post from the hub at this time.

To send to a node that is not always active a post in a posted allocation, a hub shall send a poll to the node at a preannounced time as specified in Table 21 or through access continuation or termination as specified in Table 22 and Table 23

A hub may send posts to a node that is always active or supports unscheduled access at times not indicated earlier to the node in any time intervals wherein posted allocations may be granted as specified in 6.3.

6.6.1.2.2 Using a posted allocation

Following the last frame transaction in a posted allocation, a hub may initiate another frame transaction pSIFS later, regardless of whether it received the acknowledgment frame if immediate or block acknowledgment was expected, if the current frame transaction and an appropriate guard time fit into the allocation, and if appropriate as specified in 6.8.2.

The management or data type frame(s) in any frame transaction may be new or old (retried), as appropriate.

A node shall transmit an I-Ack or B-Ack frame, when required, pSIFS after the end of the previous frame it received in a posted allocation.

If a node does not receive any required I-Ack, B-Ack, I-Ack+Poll, or B-Ack+Poll frame in a scheduled uplink allocation interval of its own, it should be ready to receive a post from the hub at the end of the allocation interval accounting for an appropriate guard time, if it has no other frame transmission or reception pending. It may be in inactive state from mTimeout after the end of the PHY preamble of the expected post if at this time it has not received any portion (such as a PHY preamble) of a frame, unless it has other frame transmission or reception pending.

6.6.1.2.3 Modifying a posted allocation


A hub may modify a posted allocation of a node by initializing another frame transaction with the node as described in 6.6.1.2.1, effectively extending the remaining posted allocation interval.

A hub may cancel a future poll or post announced earlier for a node by subsequently sending to the node a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame or an I-Ack or B-Ack frame (further specified in 6.8.1), with the Cancel field of the MAC header set to one, before the start of the future poll or post, as illustrated in Figure 76. After announcing a future poll or post for a node, a hub may cancel or retain that future poll or post while announcing another or no future poll or post for the node.


A node shall treat a future poll or post announced earlier for it as cancelled upon receiving a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame or an I-Ack or B-Ack frame (further specified in 6.8.1), with the Cancel field of the MAC header set to one, as illustrated in Figure 76.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

M/D = management or data type frame data = data type frame
Ack = I-Ack or B-Ack frame Ack+Poll = I-Ack+Poll or B-Ack+Poll frame Poll = Poll or T-Poll frame
M = More Data L = Last Frame A = Access Mode P = Poll-Post Window N = Next C = Cancel



(a) Cancellation with replacement and cancellation without replacement



(b) No cancellation with addition and no cancellation without addition

Figure 76—Example cancellations or no cancellations of future polls and posts

6.6.2 Unscheduled access

As indicated in 6.3, in any access mode, a node that supports unscheduled access as indicated in its last transmitted MAC Capability field and a hub may employ unscheduled access to obtain unscheduled bilink allocations and polled and posted allocations therein.

In beacon or non-beacon mode with superframes, unscheduled bilink allocations may be 1-periodic/round-robin or m-periodic allocations, except that a node shall not have both 1-periodic and m-periodic allocations in the same BAN. In non-beacon mode without superframes, unscheduled bilink allocations may only be round-robin.

To have a 1-periodic/round-robin unscheduled bilink allocation, which has one or more allocation intervals occurring in every beacon period (superframe) on a best-effort basis or occurring in round robin along with the allocation intervals of other 1-periodic unscheduled bilink allocations, a node shall treat all beacon periods (superframes) as its wakeup beacon periods (superframes), as illustrated in Figure 77(a1) through Figure 77(a5).

To have an m-periodic allocation, which has one or more allocation intervals occurring in every $m > 1$ beacon periods (superframes) on a best-effort basis, the node shall treat the beacon periods (superframes) containing its tentatively assigned allocation intervals as its wakeup beacon periods (superframes), as illustrated in Figure 77(b1) through Figure 77(b3).

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 77—Example unscheduled bilink allocations

6.6.2.1 Starting unscheduled bilink allocations

To obtain one or more new unscheduled bilink allocations, a node shall send a Connection Request frame to the hub when permitted to do so, setting the Requested Wakeup Period field in the frame to one for 1-periodic or round-robin allocations and to $m > 1$ for m -periodic allocations in beacon or non-beacon mode with superframes, or keeping the Requested Wakeup Phase and Requested Wakeup Period fields reserved in non-beacon mode without superframes. The node shall include the following in the frame:

- Type-I Unscheduled Bilink Request IE if its hub is operating on beacon or non-beacon mode with superframes; or
- Type-II Unscheduled Bilink Request IE if its hub is operating on non-beacon mode without superframes.

In either IE, the Minimum Length and Allocation Length fields, or the Nominal Allocation Length Requested and Maximum Allocation Length Requested fields, when present, shall have nonzero values.

To grant unscheduled bilink allocations, i.e., best-effort scheduled bilink allocations, requested by the node or initiated by itself, a hub shall send a Connection Assignment frame to the node, including the following:

- Type-I Unscheduled Bilink Assignment IE if the hub is operating on beacon or non-beacon mode with superframes; or
- Type-II Unscheduled Bilink Assignment IE if the hub is operating on non-beacon mode without superframes.

In either IE, the Interval Start and Interval End fields, or the Minimum Allocation Length Assigned, Nominal Allocation Length Assigned, and Maximum Allocation Length Assigned fields, may be all set to zero to convey no unscheduled bilink assignment, or some of them may be set to nonzero values to convey a tentative bilink assignment.

6.6.2.2 Using unscheduled bilink allocations

Upon successfully sending the Connection Assignment frame granting the node unscheduled bilink allocation intervals, the hub should provide the node with the following:

- Unscheduled bilink allocation intervals in each of the node's wakeup beacon periods (superframes), as defined by the values of the Interval Start and Interval End fields, if the hub is operating in beacon or non-beacon mode with superframes, as illustrated in Figure 77(a1) and Figure 77(b1); if this is not possible, then
- Unscheduled bilink allocation intervals in each of the node's wakeup beacon periods (superframes), as defined by the values of the Interval Start and Interval End fields but shifted in the same beacon period (superframe), if the hub is operating in beacon or non-beacon mode with superframes, as illustrated in Figure 77(a2) and Figure 77(b2); if this is not possible, then
- Unscheduled bilink allocation intervals in each of the node's wakeup beacon periods (superframes), as defined by the values of the Interval Start and Interval End fields but shifted in the same beacon period (superframe) and reduced or expanded in length, if the hub is operating in beacon or non-beacon mode with superframes, as illustrated in Figure 77(a3) and Figure 77(b3); if this is not possible for 1-periodic allocations, then
- Unscheduled bilink allocation intervals, as defined by the values of the Interval Start and Interval End fields, not necessarily occurring in every beacon period (superframe) but reoccurring across beacon periods (superframes) in round robin, i.e., sequentially along with the unscheduled bilink

allocation intervals of the other assigned 1-periodic unscheduled bilink allocations of this node and other connected nodes, if the hub is operating in beacon or non-beacon mode with superframes, as illustrated in Figure 77(a4); or

- Unscheduled bilink allocation intervals, as defined by the values of the Minimum Allocation Length Assigned, Nominal Allocation Length Assigned, and Maximum Allocation Length Assigned fields, and reoccurring over time in round robin, i.e., sequentially along with the unscheduled bilink allocation intervals of the other assigned unscheduled bilink allocations of this node and other connected nodes, if the hub is operating in non-beacon mode without superframes, as illustrated in Figure 77(a5).

To provide the node with an unscheduled bilink allocation interval, at the start of the allocation interval, the hub may initiate a frame transaction with the node, or the hub may send to the node a poll granting a polled allocation for the node to initiate one or more frame transactions, as illustrated in Figure 78. The node shall not initiate a frame transaction, until it receives a Poll or T-Poll frame, in an unscheduled bilink allocation interval. The recipient, the node or the hub, shall be ready to receive the frames transmitted by the sender and return appropriate frames during the provided unscheduled bilink allocation intervals.

In the remaining provided unscheduled bilink allocation interval, pSIFS after the end of the preceding frame transaction initiated by the hub or after the end of the preceding polled allocation interval, the hub may initiate one or more frame transactions with the node, or the hub may send to the node another poll granting a polled allocation for the node to initiate one or more frame transactions therein. However, the hub shall not continue its transmission to the node in a provided unscheduled bilink allocation interval after sending to the node `mUnscheduledNoResponseLimit` consecutive frames each requiring a response but receiving no response from the node, thus effectively relinquishing and reclaiming the allocation interval.

Frame transactions with the hub, including acknowledgment frames if required, shall fit in the provided unscheduled bilink allocation intervals, accounting for an appropriate guard time.




Figure 78—Example frame transactions in unscheduled bilink allocation intervals

6.6.2.3 Modifying unscheduled bilink allocations

A node may modify existing unscheduled bilink allocations by sending another Connection Request frame specifying the new requirements for the allocations of the same Allocation ID values. The hub shall treat this request as a new one, except that it shall set the Connection Change Indicator field in its responding Connection Assignment frame with reference to the last Connection Assignment frame it sent to the node. In particular, if the hub rejects the modifications but maintains the existing allocations, it shall respond with a Connection Assignment frame with the Connection Change Indicator field set to zero and the other fields kept to the respective values contained in the last Connection Assignment frame sent to the node.

A hub may, but should not where possible, modify unscheduled bilink allocations of a node on its own by sending the node an unsolicited Connection Assignment frame specifying the new tentative assignments of those allocations, setting the Connection Change Indicator field in the frame with reference to the last Connection Assignment frame it sent to the same node.

6.6.2.4 Aborting unscheduled bilink allocations

A node or a hub shall treat an existing unscheduled bilink allocation to have been aborted after failing to receive any frame in the last `mUnscheduledAllocationAborted` assigned allocation intervals of the allocation if the hub is operating in beacon or non-beacon mode with superframes, or in the last `mUnscheduledAllocationAborted` seconds if the hub is operating in non-beacon mode without superframes. Subsequently, the hub may reclaim the unscheduled bilink allocation by not providing the node with the corresponding unscheduled bilink allocation intervals.

A node or a hub should transmit at least one frame requiring an immediate return of a frame, such as an I-Ack or B-Ack frame, or a poll if applicable, in every allocation interval of a provided unscheduled bilink allocation so as to reduce the chance of experiencing an abortion of the unscheduled bilink allocation.

A node and a hub may start a new unscheduled bilink allocation procedure as specified in 6.6.2.1 to reinstate their lost unscheduled bilink allocations or obtain their replacements.

6.6.2.5 Ending unscheduled bilink allocations

A node may at any time end unscheduled bilink allocations by sending a modified Connection Request frame that contains Allocation Request fields with the Allocation ID fields identifying those allocations, and with the corresponding Minimum Length and Allocation Length fields set to zero, or with the Minimum Allocation Length Requested, Nominal Allocation Length Requested, and Maximum Allocation Length Requested fields set to zero.

A hub may, but should not where possible, at any time end any unscheduled bilink allocations of a node by sending the node a modified Connection Assignment frame that contains Allocation Assignment fields with the Allocation ID fields identifying those allocations, and with the Interval Start and Interval End fields set to zero, or with the Minimum Allocation Length Assigned, Nominal Allocation Length Assigned, and Maximum Allocation Length Assigned fields set to zero.

6.7 Scheduled access and scheduled-polling access

As outlined in 6.3, in beacon or non-beacon mode with superframes but not in non-beacon mode without superframes, a node and a hub may employ scheduled access to obtain scheduled uplink allocations and scheduled downlink allocations, and may employ scheduled-polling access to obtain scheduled bilink allocations and polled and posted allocations therein. Scheduled uplink allocations, scheduled downlink allocations, and scheduled bilink allocations, which are collectively referred to as *scheduled allocations*, may be 1-periodic or m-periodic allocations, except that a node shall not have both 1-periodic and m-periodic allocations in the same BAN.

To have a 1-periodic allocation, which has one or more allocation intervals spanning the same allocation slots in every beacon period (superframe), a node shall treat all beacon periods (superframes) as its wakeup beacon periods (superframes), as illustrated in Figure 79(a).

To have an m-periodic allocation, which has one or more allocation intervals spanning the same allocation slots in every $m > 1$ beacon periods (superframes), a node shall treat the beacon periods (superframes) containing its allocation intervals as its wakeup beacon periods (superframes), as illustrated in Figure 79(b).

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 79—Example scheduled uplink, downlink, and bilink allocations

6.7.1 Starting scheduled allocations

To obtain one or more new scheduled allocations, a node shall send a Connection Request frame to the hub when permitted to do so, setting the Requested Wakeup Period field in the frame to one for 1-periodic allocations and to $m > 1$ for m-periodic allocations. The node shall include in the frame the following:

- Uplink Request IE if scheduled uplink allocations are needed;
- Downlink Request IE if scheduled downlink allocations are needed; and
- Bilink Request IE if scheduled bilink allocations are needed.

In these IEs, the Minimum Length and Allocation Length fields, when present, shall have nonzero values.

To grant scheduled allocations, requested by the node or initiated by itself, a hub shall send a Connection Assignment frame to the node, including the following:

- Uplink Assignment IE if scheduled uplink allocations are granted;
- Downlink Assignment IE if scheduled downlink allocations are granted; and
- Bilink Assignment IE if scheduled bilink allocations are granted.

6.7.2 Using scheduled allocations

Upon receiving the Connection Assignment frame granting it scheduled uplink allocation intervals, the node may initiate a frame transaction with the hub at the start of each of the allocation intervals, as illustrated in Figure 80, if the frame transaction and an appropriate guard time fit into the current allocation interval. The hub shall be ready to receive the frames transmitted by the node, taking into account an appropriate guard time, during these allocation intervals.

Upon successfully sending the Connection Assignment frame granting the node scheduled downlink allocation intervals, the hub may initiate a frame transaction with the node at the start of each of the allocation intervals, as also illustrated in Figure 80, if the frame transaction and an appropriate guard time fit into the current allocation interval. The node shall be ready to receive the frames transmitted by the hub, taking into account an appropriate guard time, during these allocation intervals.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Upon successfully sending the Connection Assignment frame granting the node scheduled bilink allocation intervals, at the start of each of the allocation intervals, the hub may initiate a frame transaction with the node, if the frame transaction and an appropriate guard time fit into the current allocation interval. Alternatively, the hub may send to the node a poll granting a polled allocation for the node to initiate one or more frame transactions, as illustrated in Figure 81, if the poll and the polled allocation fit into the current bilink allocation interval. The hub shall send a Poll or T-poll frame granting an immediate polled allocation to the node within each scheduled bilink allocation interval, unless the node indicates it has no more data to send. The poll should be a Poll frame if it is sent at the start of a scheduled bilink allocation interval. The node shall not initiate a frame transaction, until it receives a Poll or T-Poll frame, in a bilink allocation interval. The recipient, the node or the hub, shall be ready to receive the frames transmitted by the sender, taking into account an appropriate guard time, during these allocation intervals.

Following a frame transaction in a scheduled uplink or downlink allocation interval, the node or the hub, respectively, may initiate another frame transaction pSIFS later as also illustrated in Figure 80, regardless of whether it received an acknowledgment frame if immediate or block acknowledgment is expected, if the current frame transaction and an appropriate guard time fit into the allocation interval, and if appropriate as specified in 6.8.1 and 6.8.2, respectively.

Following a frame transaction in a posted allocation, or the final frame transaction in a polled allocation, of a scheduled bilink allocation interval, the hub may initiate another frame transaction or send a poll providing an or no immediate polled allocation pSIFS later as also illustrated in Figure 81, if the frame transaction and an appropriate guard time, or if the poll and the polled allocation (if any), fit into the bilink allocation interval.

The management or data type frame(s) in any frame transaction may be new or old (retried), as appropriate.




Figure 80 —Example frame transactions in scheduled uplink/downlink allocation intervals




Figure 81 —Example frame transactions in scheduled bilink allocation intervals

6.7.3 Modifying scheduled allocations

A node may modify existing scheduled allocations by sending another Connection Request frame specifying the new requirements for the allocations of the same Allocation ID values. The hub shall treat this request as a new one, except that it shall set the Connection Change Indicator field in its responding Connection Assignment frame with reference to the last Connection Assignment frame it sent to the node. In particular, if the hub rejects the modifications but maintains the existing allocations, it shall respond with a Connection Assignment frame with the Connection Change Indicator field set to zero and the other fields kept to the respective values contained in the last Connection Assignment frame sent to the node.

A hub may, but should not where possible, modify scheduled allocations of a node on its own by sending the node an unsolicited Connection Assignment frame specifying the new assignments of those allocations, setting the Connection Change Indicator field in the frame with reference to the last Connection Assignment frame it sent to the same node.

6.7.4 Aborting scheduled allocations

A node or a hub shall treat an existing scheduled allocation to have been aborted after failing to receive any frame in the last `mScheduledAllocationAborted` allocation intervals of the allocation. Subsequently, the hub may reclaim the aborted scheduled allocation.

A node or a hub shall not exchange frames with each other in their aborted scheduled allocation.

A node or a hub should transmit at least one frame requiring an immediate return of a frame, such as an I-Ack or B-Ack frame, or a poll if applicable, in every allocation interval of a scheduled allocation so as to reduce the chance of experiencing an abortion of the scheduled allocation.

A node and a hub may start a new scheduled allocation procedure as specified in 6.7.1 to reinstate their lost scheduled allocations or obtain their replacements.

6.7.5 Ending scheduled allocations

A node may at any time end scheduled allocations by sending a modified Connection Request frame that contains Allocation Request fields with the Allocation ID fields identifying those allocations, and with the corresponding Minimum Length and Allocation Length fields set to zero.

A hub may, but should not where possible, at any time end any scheduled allocations of a node by sending the node a modified Connection Assignment frame that contains Allocation Assignment fields with the Allocation ID fields identifying those allocations, and with the Interval Start field set to 255 and the Interval End field set to zero.

A node or a hub shall not exchange frames with each other in their ended scheduled allocation.

6.8 Access continuation, termination, and timeout

A node or a hub shall treat the access rules specified in Table 22 and Table 23 of this subclause for the continuation, termination, and timeout of frame transmissions and receptions within an obtained uplink or downlink allocation interval as supplementary to, but not a substitute of, those provided in 6.5, 6.6, and 6.7 for obtaining and using an allocation interval.

In particular, in setting the More Data and Last Frame fields of the MAC header of a management or data type frame to be sent, a node or a hub shall determine first if it may initiate another frame transaction in the current allocation interval according to 6.5, 6.6, and 6.7.

An uplink allocation interval is a contended, polled, or scheduled uplink allocation interval. A downlink allocation interval is a posted or scheduled downlink allocation interval. A bilink allocation interval comprises polls, polled allocation intervals (started by polls), and/or posted allocation intervals (started by posts).

6.8.1 In an uplink allocation interval

Table 22 —Access continuation, termination, and timeout in uplink allocation interval

More Data	Last Frame	After sending to a hub a management or data type frame with the Ack Policy field of the MAC header not set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, a node,
0	0	having indicated no more data or management type frames waiting for transmission, other than a potential current frame retransmission, in the current allocation interval, shall not transmit another frame, but may retransmit the current frame, to the hub pSIFS or pMIFS later as appropriate, thus relinquishing the current allocation interval if not retransmitting.
1	0	having indicated data or management type frame(s) waiting for transmission or retransmission in the current allocation interval, should transmit or retransmit a frame to the hub pSIFS or pMIFS later as appropriate.
0	1	having indicated no more data or management type frames waiting for transmission in the current allocation interval, other than a potential frame retransmission in a next allocation interval, shall not transmit or retransmit a frame to the hub pSIFS or pMIFS later, thus relinquishing the current allocation interval.
1	1	having indicated data or management type frame(s) waiting for transmission or retransmission not in the current but a next allocation interval, shall not transmit or retransmit a frame to the hub pSIFS or pMIFS later, most likely due to not enough time remaining in the current allocation interval for completing another frame transaction plus an appropriate guard time, thus relinquishing the current allocation interval.
More Data	Last Frame	After receiving from a node a management or data type frame with the Ack Policy field of the MAC header not set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, a hub,
0	0	expecting no data or management type frame(s) waiting for transmission, other than a potential current frame retransmission, from the node in the current allocation interval, shall be ready to receive the expected frame pSIFS or pMIFS later as appropriate, but may reclaim the current allocation interval mTimeOut after the end of the PHY preamble of the expected frame if at this time it has not received a PHY preamble of a frame.
1	0	expecting data or management type frame(s) waiting for transmission or retransmission from the node in the current allocation interval, shall be ready to receive the expected frame(s) pSIFS or pMIFS later as appropriate and in the remaining allocation interval.
0	1	expecting no more data or management type frames waiting for transmission from the node in the current allocation interval, other than a potential current frame retransmission from the node in a next allocation interval, may reclaim the current allocation interval.
1	1	expecting data or management type frame(s) waiting for transmission or retransmission from the node not in the current but a next allocation interval, may reclaim the current allocation interval.

**Table 22—Access continuation, termination, and timeout in uplink allocation interval
 (continued)**

More Data	Last Frame	After sending to a hub a management or data type frame with the Ack Policy field of the MAC header set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, and receiving from the hub an or no expected I-Ack or B-Ack frame (or I-Ack+Poll or B-Ack+Poll frame with the More Data field of the MAC header set to one in beacon or non-beacon mode with superframes), a node,
0	0	having indicated no more data or management type frames waiting for transmission, other than a potential current frame retransmission, in the current allocation interval, shall not transmit another frame, but may retransmit the current frame, to the hub pSIFS later, thus relinquishing the current allocation interval if not retransmitting.
1	0	in beacon or non-beacon mode with superframe boundaries, having indicated data or management type frame(s) waiting for transmission or retransmission in the current allocation interval, should transmit or retransmit a frame to the hub pSIFS later; in non-beacon mode without superframes, having indicated data or management type frames waiting but not allowed for transmission, and a potential current frame retransmission, in the current allocation interval, shall not transmit another frame, but may retransmit the current frame, to the hub pSIFS later, thus relinquishing the current allocation interval if not retransmitting.
0	1	having indicated no more data or management type frames waiting for transmission in the current allocation interval, other than a potential current frame retransmission, in a next allocation interval, shall not transmit or retransmit a frame to the hub pSIFS later, thus relinquishing the current allocation interval.
1	1	having indicated data or management type frame(s) waiting for transmission or retransmission not in the current but a next allocation interval, shall not transmit or retransmit a frame to the hub pSIFS later, most likely due to not enough time remaining in the current allocation interval for completing another frame transaction plus an appropriate guard time, thus relinquishing the current allocation interval.
More Data	Last Frame	After receiving from a node a management or data type frame with the Ack Policy field of the MAC header set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, and sending to the node a required I-Ack or B-Ack frame (or I-Ack+Poll or B-Ack+Poll frame with the More Data field of the MAC header set to one in beacon or non-beacon mode with superframes), a hub,
0	0	expecting no data or management type frame(s) waiting for transmission, other than a potential current frame retransmission, from the node in the current allocation interval, shall be ready to receive the expected frame pSIFS later, but may reclaim the current allocation interval mTimeOut after the end of PHY preamble of the expected frame if at this time it has not received a PHY preamble of a frame.
1	0	in beacon or non-beacon mode with superframe boundaries, expecting data or management type frame(s) waiting for transmission or retransmission from the node in the current allocation interval, shall be ready to receive the expected frame(s) pSIFS later and in the remaining allocation interval; in non-beacon mode without superframes, expecting data or management type frame(s) waiting but not allowed for transmission, and a potential current frame retransmission, from the node in the current allocation interval, pSIFS later the hub shall be ready to receive the expected frame, but may reclaim the current allocation interval mTimeOut after the end of the PHY preamble of the expected frame if at this time it has not received a PHY preamble of a frame.
0	1	expecting no more data or management type frames waiting for transmission from the node in the current allocation interval, other than a potential current frame retransmission from the node in a next allocation interval, may reclaim the current allocation interval.
1	1	expecting data or management type frame(s) waiting for transmission or retransmission from the node not in the current but a next allocation interval, may reclaim the current allocation interval.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

**Table 22—Access continuation, termination, and timeout in uplink allocation interval
(continued)**

More Data	After sending to a node an expected I-Ack or B-Ack frame with the More Data field of the MAC header set as given below, a hub,
0	having indicated no frames waiting for transmission to the node, should not transmit to the node, but may transmit to another node as appropriate, a poll or post, (a) pSIFS after the current allocation interval is reclaimed if appropriate, or else (b) after the current allocation interval is ended, unless indicated otherwise by another frame received from or sent to the node before the current allocation interval is reclaimed or ended; nor should transmit to the node a future poll or post announced earlier (if any).
1	having indicated frame(s) waiting for transmission or retransmission to the node, shall send to the node a poll or post (a) pSIFS after the current allocation interval is reclaimed if appropriate or else (b) after the current allocation interval is ended, unless indicated otherwise by another I-Ack or B-Ack frame sent to the node before the current allocation interval is reclaimed or ended.
More Data	After receiving from a hub a required I-Ack or B-Ack frame with the More Data field of the MAC header set as given below, a node,
0	expecting no more frames waiting for transmission from the hub, may be in inactive state in the remaining allocation interval, if it is not to send another frame to the hub pSIFS later (as indicated by the preceding frame with the More Data field set to zero or the Last Frame field set to one in any access mode, or with the More Data field set to one and the Last Frame field set to zero in non-beacon mode without superframes), or unless it is to send another frame to the hub pSIFS later (as indicated by the preceding frame with the More Data field set to one and the Last Frame field set to zero) in beacon or non-beacon mode with superframes; and may be in inactive state at the time announced earlier from the hub for sending a future poll or post to the node (if any), given that the future poll or post is cancelled.
1	expecting frame(s) waiting for transmission or retransmission from the hub, shall be ready to receive a poll or post from the hub (a) pSIFS after the current allocation interval is reclaimed if appropriate or else (b) when the current allocation interval is ended taking into account an appropriate guard time, but may be in inactive state in the remaining allocation interval from mTimeOut after the end of the PHY preamble of the expected poll or post if at this time it has not received any portion of a frame. However, the node shall continue to be ready to receive frames in a scheduled bilink allocation interval unless the node and the hub have both indicated no more frames to transmit.
More Data	After sending to a node a Poll, T-Poll, I-Ack+Poll, or B-Ack+Poll frame with the More Data field of the MAC header set as given below, a hub,
0	having granted an immediate polled allocation to the node, expecting one or more frames sent from the node, shall be ready to receive the expected frame(s) pSIFS later, but may reclaim the current allocation interval mTimeOut after the end of the PHY preamble of the expected frame if at this time it has not received a PHY preamble of a frame. However, if the polled allocation interval extends from within an existing uplink allocation interval of the node, the hub may reclaim only the extended portion of the polled allocation interval.

6.8.2 In a downlink allocation interval

Table 23—Access continuation, termination, and timeout in downlink allocation interval

More Data	Last Frame	After sending to a node a management or data type frame with the Ack Policy field of the MAC header not set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, a hub,
0	0	having indicated no more frames waiting for transmission to the node, other than a potential current frame retransmission, in the current allocation interval, shall not transmit another frame, but may retransmit the current frame, to the node pSIFS or pMIFS later as appropriate, thus relinquishing and reclaiming the current allocation interval if not retransmitting.
1	0	having indicated frame(s) waiting for transmission or retransmission to the node in the current allocation interval, shall transmit or retransmit a frame to the node pSIFS or pMIFS later as appropriate.
0	1	having indicated no more frames waiting for transmission to the node in the current allocation interval, other than a potential frame retransmission in a next allocation interval, shall not transmit or retransmit a frame to the node pSIFS or pMIFS later, thus relinquishing and reclaiming the current allocation interval.
1	1	having indicated frame(s) waiting for transmission or retransmission to the node not in the current but a next allocation interval, shall not transmit or retransmit a frame to the node pSIFS or pMIFS later, most likely due to not enough time remaining in the current allocation interval for completing another frame transaction plus an appropriate guard time, thus relinquishing the current allocation interval.
More Data	Last Frame	After receiving from a hub a management or data type frame with the Ack Policy field of the MAC header not set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, a node,
0	0	expecting no frame(s) waiting for transmission, other than a potential current frame retransmission, from the hub in the current allocation interval, shall be ready to receive the expected frame pSIFS or pMIFS later as appropriate, but may be in an inactive state in the remaining allocation interval from mTimeOut after the end of the PHY preamble of the expected frame if at this time it has not received any portion of a frame, thus relinquishing the current allocation interval.
1	0	expecting frame(s) waiting for transmission or retransmission from the hub in the current allocation interval, shall be ready to receive the expected frame(s) pSIFS or pMIFS later as appropriate and in the remaining allocation interval.
0	1	expecting no more frames waiting for transmission from the hub in the current allocation interval, other than a potential current frame retransmission from the hub in a next allocation interval, may be in inactive state in the remaining allocation interval, thus relinquishing the current allocation interval.
1	1	expecting frame(s) waiting for transmission or retransmission from the hub not in the current but a next allocation interval, may be in inactive state in the remaining allocation interval, thus relinquishing the current allocation interval.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

**Table 23—Access continuation, termination, and timeout in downlink allocation interval
(continued)**

More Data	Last Frame	After sending to a node a management or data type frame with the Ack Policy field of the MAC header set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, and receiving from the node an or no expected I-Ack or B-Ack frame, a hub,
0	0	having indicated no more frames waiting for transmission, other than a potential current management or data type frame retransmission, to the node in the current allocation interval, shall not transmit another frame, but may retransmit the current management or data type frame, to the node pSIFS later, thus relinquishing and reclaiming the current allocation interval if not retransmitting. However, if the hub received from the node an I-Ack or B-Ack with the More Data field of the MAC header set to one (and if both the hub and the node support polling access), pSIFS later the hub should send to the node a Poll or T-Poll frame conveying an immediate polled allocation or a future poll or post for the node, if there is a corresponding time interval available.
1	0	having indicated frame(s) waiting for transmission or retransmission to the node in the current allocation interval, shall transmit or retransmit a frame to the node pSIFS later.
0	1	having indicated no more frames waiting for transmission to the node in the current allocation interval, other than a potential current management or data type frame retransmission to the node in a next allocation interval, shall not transmit or retransmit a frame to the node pSIFS later, thus relinquishing and reclaiming the current allocation interval.
1	1	having indicated frame(s) waiting for transmission or retransmission to the node not in the current but a next allocation interval, shall not transmit or retransmit a frame to the node pSIFS later, most likely due to not enough time remaining in the current allocation interval for completing another frame transaction plus an appropriate guard time, thus relinquishing and reclaiming the current allocation interval.
More Data	Last Frame	After receiving from a hub a management or data type frame with the Ack Policy field of the MAC header set to I-Ack or B-Ack, and with the More Data and Last Frame fields of the MAC header set as given below, and sending to the hub a required I-Ack or B-Ack frame, a node,
0	0	expecting no frame(s) waiting for transmission, other than a potential current frame retransmission, from the hub in the current allocation interval, should be ready to receive the expected frame pSIFS later, but may be in inactive state in the remaining allocation interval from mTimeOut after the end of the PHY preamble of the expected frame if at this time it has not received any portion of a frame, thus relinquishing the current allocation interval. However, if the node sent to the hub an I-Ack or B-Ack with the More Data field of the MAC header set to one (and if both the hub and the node support polling access), pSIFS later the node should be ready to receive a Poll or T-Poll frame from the hub potentially providing an (uplink) polled allocation to the node, but may be in inactive state in the remaining allocation interval from mTimeOut after the end of the PHY preamble of the expected frame if at this time it has not received any portion of a frame, thus relinquishing the current allocation interval.
1	0	expecting frame(s) waiting for transmission or retransmission from the hub in the current allocation interval, shall be ready to receive the expected frame(s) pSIFS later and in the remaining allocation interval.
0	1	expecting no more frames waiting for transmission from the hub in the current allocation interval, other than a potential current management or data type frame retransmission from the hub in a next allocation interval, may be in inactive state in the remaining allocation interval, thus relinquishing the current allocation interval.
1	1	while knowing frame(s) waiting for transmission or retransmission from the hub not in the current but a next allocation interval, may be in inactive state in the remaining allocation interval, thus relinquishing the current allocation interval.

6.9 MICS band communication

A hub or a node may support no mechanisms described solely in this subclause if it does not communicate in the medical implant communications service (MICS) band.

In the MICS band, a hub shall operate with or without superframes as specified in 6.3. The hub may choose a new channel only when required by, and in compliance with, applicable considerations, regulations, and standards including subclause 8.6 or Clause 10 of ETSI EN 301 839-1. An implant shall communicate as a node with a hub, taking into account applicable considerations, regulations, and standards including subclause 8.6 or Clause 10 of ETSI EN 301 839-1. The hub and the node may perform a mutual discovery procedure described in 6.9.1 and 6.9.2 before their exchange of data or management type frames.

6.9.1 Unconnected mutual discovery

A hub may facilitate mutual discovery and connection with unconnected nodes as follows and in compliance with applicable regulations and standards including Clause 10 of ETSI EN 301 839-1.

If the hub is to wake up an unconnected node with a known EUI-48, it should send one or more Wakeup frames that have the Ack Policy field of the MAC header set to N-Ack and contain the EUI-48 in the frame payload, prior to transmitting one or more unconnected T-Poll frames separated by pMICSPollSpace, as illustrated in Figure 82(a). The hub shall listen for a frame arrival after each transmission and before the next transmission of such an unconnected T-Poll frame. It may repeat this sequence of Wakeup and unconnected T-Poll frame transmissions until it receives a response from the node.

Every pMICSUnconnectedPollPeriod or shorter, the hub shall transmit a group of up to pMICSUnconnectedPolls unconnected T-Poll frames separated by pMICSPollSpace, as illustrated in Figure 82(b). The hub shall listen for a frame arrival after each transmission and before the next transmission of such an unconnected T-Poll frame. An unconnected T-Poll frame is a T-Poll frame addressed to Unconnected_Broadcast_NID as defined in Table 17 and providing an unconnected polled allocation as specified in 6.4.

The hub shall transmit unconnected T-Poll frames at the highest mandatory data rate of the MICS band as specified in the corresponding PHY subclause, 8.1.1.

An unconnected node in need of a connection may discover a hub as follows:

The unconnected node should cyclically tune to each MICS band channel for pMICSUnconnectedPollRxTime, unless or until it receives a Wakeup or unconnected T-Poll frame and hence discovers a hub. Once it starts receiving a frame, it should stay on the receive mode until it receives the whole frame. The unconnected node should not further change the channel unless recommended otherwise.

A hub may facilitate mutual discovery with connected nodes immediately prior to frame exchanges with them, based on any or a combination of the following discovery procedures.

As illustrated in Figure 83(a), prior to more frame exchanges with a connected node, the hub may send one or more Wakeup frames that have the Ack Policy field in the MAC header set to N-Ack and contain the node's EUI-48 in the frame payload, prior to transmitting Poll frames to the node. The hub shall transmit a group of up to pMICSPolls Poll frames separated by pMICSPollSpace, each addressed to the node and providing an immediate polled allocation, as illustrated in Figure 83(b). The hub shall transmit these Poll frames at the highest mandatory data rate of the MICS band as specified in the corresponding PHY subclause, 8.1.1. The hub shall listen for a frame arrival after each transmission and before the next transmission of such a Poll frame.

The connected node shall be in active mode in anticipation of pending frame exchanges with the hub by scheduling or other means, accounting for an appropriate guard time. It should tune to the channel in which it last received a frame with a valid FCS from the hub for a time equal to $(pMICSPollTxTime + pMICSPollSpace + pMICSPollRxTime)$ to receive a Poll frame. If it does not receive a Poll frame, it should change to another channel. After dwelling on the current channel for $pMICSPollRxTime$, it should switch to yet another channel in accordance with the channel order last provided by the hub, unless or until it receives a Poll frame addressed to it. The node should not further change the channel unless recommended otherwise.



IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Now that the hub and the unconnected node are in the same channel, they are in a position to exchange more frames with each other using appropriate access methods as specified in 6.6, 6.7, and 6.8 in view of 6.3.

6.9.2.2 Multicast poll aided discovery

As illustrated in Figure 84, prior to more frame exchanges with a group of connected nodes, the hub shall transmit a group of up to pMICSMcastPolls Poll frames forming a lockup phase and separated by pMIFS, each addressed to the Multicast_NID of the group and providing no immediate polled allocation but announcing a future poll starting at the intended beginning of the first individual communication phase with a node of the group. In particular, the time of the future poll is encoded according to Table 21. The hub shall transmit these Poll frames at the highest mandatory data rate of the MICS band as specified in the corresponding PHY subclause, 8.1.1.

The hub shall transmit a poll addressed to a node of the group at the indicated future poll time and providing an immediate polled allocation, thus starting the individual phase for the node as also illustrated in Figure 84.

The connected nodes of the group should be in active mode in anticipation of pending frame exchanges with the hub by scheduling or other means, accounting for an appropriate guard time. Each of them should tune to the channel in which it last received a frame with a valid FCS from the hub. After dwelling in the current channel for pMICSMcastPollRxTime, each should switch to another channel in accordance with the channel order last provided by the hub, unless or until it receives a Poll frame addressed to the Multicast_NID of its group. It should not further change the channel unless recommended otherwise. It may enter sleep state until the start of the first individual phase as announced in the received Poll frame, when it should be in active state to receive a unicast Poll frame.




Figure 84—Multicast poll aided connected discovery and frame exchange in MICS band

Now that the hub and the connected nodes of the group are in the same channel, they are in a position to exchange more frames with each other using appropriate access methods as specified in 6.6, 6.7, and 6.8 in view of 6.3.

6.9.3 Medical implant event report

When not transmitting, a hub should stay in receive mode in the channel selected according to its channel order list communicated to the nodes connected with it.

A node connected with a hub may transmit frames reporting a medical implant event in its next scheduled bilink allocation interval, if available, following a Poll or T-Poll frame conveying a polled allocation to it by the hub, using the connected mutual discovery procedure as specified in 6.9.2.

The node may also transmit such frames anytime as illustrated in Figure 85. Before discovering the hub's operating channel, the node should send Emergency frames without a frame payload. In particular, the node should transmit an Emergency frame with no frame payload and with the Ack Policy field set to I-Ack, in the first channel of the channel order last communicated to it by the hub. It should retry the frame for up to $pMICSHubEmergencyRetries$ times in this channel upon failing to receive an expected acknowledgment. If it still receives no acknowledgment, it should similarly send and retry the frame in the next MICS channel in the channel order list, and again in another channel, until it receives an expected I-Ack frame or pauses its transmission. If the node did not receive an expected I-Ack frame after sending an Emergency frame in each of the channels in the channel order list, it may pause its transmission or restart with the list for the Emergency frame transmission. The node shall space an Emergency frame and the next as if an I-Ack frame were received in between.

After receiving an I-Ack frame, it should proceed to transmit the Emergency frames (with incremental sequence numbers) containing frame payloads generated from the medical implant event. The node shall set to one the More Data field in the MAC header of these Emergency frames except the last one, and shall set to zero the More Data field in the MAC header of that last frame to indicate the end of the medical implant event report transfer.

On receiving an Emergency frame with More Data field set to one from the node, the hub should not initiate its own frame transactions with this node or another one until it has received all Emergency frames as indicated by the More Data field value.




Figure 85—Medical implant event report outside scheduled allocations in MICS band

To prevent a prolonged collision between overlapped transmissions by the hub and its nodes, after retrying a frame for up to $pMICSHubMaxRetries$ without receiving an expected response, the hub should enter the receive mode to allow for transmission and reception of possible Emergency frames.

6.9.4 Low-power low duty cycle (LP/LDC) operation

A hub shall stay in receive mode while in a channel selected for LP/LDC operation as specified in applicable considerations, regulations, and standards including subclause 8.6 of ETSI EN 301 839-1.

A node may transmit a data type frame to a hub at any time in a channel, with a duty cycle, and subject to a limit on the number of transmissions within an hour, as specified for LP/LDC operation in applicable regulations and standards including subclause 8.6 of ETSI EN 301 839-1. The node shall set the Ack Policy field of the MAC header of the data type frame to N-Ack.

6.10 Two-hop star topology extension

Except in the MICS band, a node and a hub may use a two-hop extension to exchange frames through another node that is connected and capable of direct communication with both of them, as illustrated in Figure 86, turning the terminal and intermediate nodes into the relayed and relaying nodes, respectively, and the hub into the target hub of the relayed node.

Either the relayed node or the target hub may initiate a two-hop extension at times determined fit by the initiator, regardless of whether they have been never, or are no longer, in direct communication with each other.

The relaying node may also exchange its own frames with the hub directly just as in a one-hop star network.

H = hub N = node
 H – N1c – N1d = two-hop extension between H and N1d
 H – N1f – N1e = two-hop extension between H and N1e
 N1d = relayed node N1e = relayed node
 N1c = relaying node N1f = relaying node




Figure 86—Two-hop extended star network topology

6.10.1 Exchanging frames for a two-hop extension

The relayed node and the target hub may exchange unicast management or data, but not control, type frames through the relaying node by frame encapsulation as illustrated in Figure 87 and Figure 88 and described in the remainder of this subclause. The relayed node, the relaying node, and the target hub shall not apply frame encapsulation to control type frames, such as I-Ack frames.




Figure 87—General frame encapsulation format

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 88—Management or data type frame exchanges on a two-hop extension

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

- The relayed node and the target hub shall follow the state diagram specified in 4.5 for a one-hop star network in exchanging encapsulated frames through the relaying node, as if they were in direct communication, applying an appropriate security level for both transmission and reception of the encapsulated frames.
- The relaying node and the target hub shall follow the state diagram specified in 4.5 for a one-hop star network in exchanging frames, encapsulated or not, as if the relaying node were a non-relaying node, applying an appropriate security level for both transmission and reception of the frames.
- The relayed node and the relaying node shall follow the state diagram specified in 4.5 for a one-hop star network in exchanging frames, encapsulated or not, as if the relaying node were a hub, applying an appropriate security level for both transmission and reception of the frames, with the followings exception: To exchange frames for a two-hop extension, they shall not be connected with each other in the way a node and a hub would be in a one-hop star network, i.e., they shall not exchange with each other Connection Request and Connection Assignment frames not encapsulating another frame.

6.10.1.1 Frame transmission from relayed node through relaying node to target hub

To send a management or data type frame, designated as an encapsulated “X” frame, through the relaying node to the target hub as shown in Figure 88(a), the relayed node shall send to the relaying node an encapsulating “X” frame, wherein

- the Recipient ID is set to the NID of the relaying node;
- the Relay field of the MAC header is set to one;
- the other fields of the MAC header are set as if the relaying node were a hub and the relayed node were sending the encapsulated “X” frame to that hub without frame encapsulation in a one-hop star network, e.g., the Frame Type and Frame Subtype fields are set to the corresponding values of their counterparts in the encapsulated “X” frame;
- the frame payload is set to the encapsulated “X” frame, which is formatted as if the relayed node were sending the encapsulated “X” frame directly to the target hub for the first time in a one-hop star network, with the Recipient ID field of the MAC header set to zero if the relayed node does not yet know the HID of the target hub.

If the relayed node does not have a Connected_NID yet, it shall treat the Unconnected_NID as its NID in receiving an expected I-Ack frame from the relaying node.

Upon receiving an encapsulating “X” frame, i.e., a management or data type frame with the Relay field of the MAC header set to one, the relaying node shall process the frame according to 6.2 specified for a one-hop star network with the following additional considerations:

- The Sender ID of the MAC header is set to the NID of the relayed node.
- The frame payload, after appropriate security is applied, is an encapsulated “X” frame.
- The Relay field of the MAC header of the required I-Ack frame is set to one if relay is feasible, or is set to zero otherwise.

If relay is feasible, i.e., if the relaying node is capable of providing relay between the relayed node and the target hub and if the target hub is capable of being a relayed hub as indicated in its MAC Capability field, the relaying node shall send to the target hub an encapsulating “X” frame, wherein

- the Relay field of the MAC header is set to one;

- the other fields of the MAC header are set as if the relaying node were sending the encapsulated “X” frame to the hub without frame encapsulation in a one-hop star network;
- the frame payload is set to the encapsulated “X” frame to be next relayed to the target hub.

Upon receiving an encapsulating “X” frame, i.e., a management or data type frame with the Relay field of the MAC header set to one, the target hub shall process the frame according to 6.2 specified for a one-hop star network with the following additional considerations:

- The frame is not a duplicate of another frame with the Relay field of the MAC header set to zero even if it would be otherwise treated as a duplicate according to 6.2.10 specified for a one-hop star network.
- The frame payload, after appropriate security is applied, is an encapsulated “X” frame.
- The Relay field of the MAC header of the required I-Ack frame is set to one if it is capable of being a relayed hub as announced in its MAC Capability field, or is set to zero otherwise.

If the relayed node does not have a Connected_NID yet, i.e., if the Sender ID field of the MAC header of the encapsulated “X” frame is set to the Unconnected_NID, the target hub shall keep the Unconnected_NID as the relayed node’s NID or shall assign as the relayed node’s NID a Connected_NID according to the NID selection rules of 6.2.1.

6.10.1.2 Frame transmission from target hub through relaying node to relayed node

To send a management type frame, other than Connection Assignment, or a data type frame, designated as an encapsulated “Y” frame, through the relaying node to the relayed node as shown in Figure 88(b), the target hub shall send to the relaying node an encapsulating “Y” frame, wherein

- the Relay field of the MAC header is set to one;
- the other fields of the MAC header are set as if the target hub were sending the encapsulated “Y” frame to the relaying node without frame encapsulation in a one-hop star network;
- the frame payload is set to the encapsulated “Y” frame, which is formatted as if the target hub were sending the encapsulated “Y” frame directly to the relayed node for the first time in a one-hop star network;
- the Requested Ack Data Rates field of the frame payload of the encapsulated “Y” frame is set as if the node and the hub referenced in the definition of the field in 5.3.6.6 were the relayed node and the relaying node, respectively.

Upon receiving an encapsulating “Y” frame, i.e., a management or data type frame with the Relay field of the MAC header set to one, the relaying node shall process the frame according to 6.2 specified for a one-hop star network with the following additional considerations:

- The frame is not a duplicate of another frame with the Relay field of the MAC header set to zero even if it would be otherwise treated as a duplicate according to 6.2.10 specified for a one-hop star network.
- The frame payload, after appropriate security is applied, is an encapsulated “Y” frame.
- The Relay field of the MAC header of the required I-Ack frame is set to one if relay is feasible or is set to zero otherwise.

If relay is feasible, i.e., if the relaying node is capable of providing relay between the target hub and the relayed node, the relaying node shall send to the relayed node an encapsulating “Y” frame, wherein

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

- the Recipient ID field of the MAC header is set to the NID of the relayed node, i.e., the Recipient ID of the MAC header of the encapsulated “Y” frame;
- the Sender ID field of the MAC header is set to the NID of the relaying node;
- the Relay field of the MAC header is set to one;
- the other fields of the MAC header are set as if the relaying node were a hub and sending the encapsulated “Y” frame to the relayed node without frame encapsulation in a one-hop star network;
- the frame payload is set to the encapsulated “Y” frame to be next relayed to the relayed node.

If the relayed node does not have a Connected_NID yet, it shall treat the Unconnected_NID or any Connected_NID as its NID in receiving an expected encapsulating “Y” frame.

Upon receiving an encapsulating “Y” frame, i.e., a management or data type frame with the Relay field of the MAC header set to one, the relayed node shall process the frame according to 6.2 specified for a one-hop star network with the following additional considerations:

- The Sender ID field of the MAC header is set to the NID of the relaying node.
- The frame payload, after appropriate security is applied, is an encapsulated “Y” frame.
- The Relay field of the MAC header of the required I-Ack frame is set to one if it is capable of being a relayed node as announced in its MAC Capability field, or is set to zero otherwise.

6.10.1.3 Connection assignment from target hub to relaying node and relayed node

To specify connection assignment for a two-hop extension involving a relaying node and a relayed node as shown in Figure 88(c), the target hub shall send two encapsulated Connection Assignment frames 1 and 2 in two encapsulating Connection Assignment frames 1 and 2 to the relayed node and the relaying node, respectively, as further described in the remainder of this subclause.

6.10.1.3.1 Connection assignment from target hub through relaying node to relayed node

To send encapsulated Connection Assignment frame 1 through the relaying (capable) node to the relayed node as shown in the upper diagram of Figure 88(c), the target hub and the relaying node shall follow the frame transmission procedure for “Y” frame = Connection Assignment frame 1 as described in Figure 88(b) and Figure 88(c), with the following modifications made to the frame payload of encapsulated Connection Assignment frame 1:

- the MAC Capability and PHY Capability fields are set to those for the relaying node;
- the Assigned Ack Data Rates field is set as if the node and the hub referenced in the definition of the field in 5.3.7.11 were the relayed node and the relaying node, respectively;
- the Connection Change Indicator field is set such that it accounts for all the included IEs defining both hops;
- after each Uplink Assignment, Downlink Assignment, or Bilink Assignment IE defining the assignment of scheduled allocations applicable between the target hub and the relaying node for the two-hop extension, another Uplink Assignment, Downlink Assignment, or Bilink Assignment IE is inserted defining the assignment of scheduled allocations applicable between the relaying node and the relayed node for the corresponding two-hop extension, with the relaying node considered as a hub for link direction referencing purposes.

Upon processing the assignment of scheduled allocations applicable between the target hub and the relaying node, the relayed node may listen in those scheduled allocations to determine if it can be in a one-hop communication with the target hub.

6.10.1.3.2 Connection assignment from target hub to relaying node

To send encapsulated Connection Assignment frame 2 to the relaying (capable) node as shown in the lower diagram of Figure 88(c), the target hub shall send to the relaying node encapsulating Connection Assignment frame 2, wherein

- the Relay field of the MAC header is set to one;
- the other fields of the MAC header are set as if the target hub were sending encapsulated Connection Assignment frame 2 to the relaying node without frame encapsulation in a one-hop star network;
- the frame payload is set to encapsulated Connection Assignment frame 2, which is formatted as if the target hub were sending the frame directly to the relaying node for the first time in a one-hop star network, with some modifications to support the two-hop extension.

In particular, encapsulated Connection Assignment frame 2 is formatted with the following modifications:

- In the MAC header, the Sender ID field is set to the NID of the relayed node to indicate that the connection assignment is for a two-hop extension to the relayed node.
- In the frame payload,
 - the Assigned Wakeup Phase and Assigned Wakeup Period fields are set to those for the relayed node;
 - the MAC Capability and PHY Capability fields are set to those for the relayed node if they are known to the target hub, or are set to zero otherwise;
 - the Assigned Ack Data Rates field is set as if the node and the hub referenced in the definition of the field in 5.3.7.11 were the relayed node and the relaying node, respectively;
 - the Connection Change Indicator field is set such that it accounts for all the included IEs defining both hops;
 - after each Uplink Assignment, Downlink Assignment, or Bilink Assignment IE defining the assignment of scheduled allocations applicable between the target hub and the relaying node for the two-hop extension, another Uplink Assignment, Downlink Assignment, or Bilink Assignment IE is inserted defining the assignment of scheduled allocations applicable between the relaying node and the relayed node for the corresponding two-hop extension, with the relaying node considered as a hub for link direction referencing purposes.

Upon receiving encapsulating Connection Assignment frame 2 as defined above, the relaying node shall process the frame according to 6.2 specified for a one-hop star network, taking into account the above modifications.

The relaying node shall not send to the relayed node another encapsulating Connection Assignment frame with the frame payload set to encapsulated Connection Assignment frame 2, after it has already sent to the relayed node an encapsulating Connection Assignment frame with the frame payload set to encapsulated Connection Assignment frame 1.

6.10.1.4 Control type frame transmission without frame encapsulation

The target hub may send a control type frame to the relaying node so long as a hub may send the frame to a node in a one-hop star network, with the modification that in the MAC header the Relay field is set to one or zero depending on whether relay is involved and feasible.

The relaying node may send a control type frame to the target hub so long as a node may send the frame to a hub in a one-hop star network, with the modification that in the MAC header the Relay field is set to one or zero depending on whether relay is involved and feasible.

The relaying node may send a control type frame to the relayed node so long as a hub may send the frame to a node in a one-hop star network, with the modification that in the MAC header

- the Relay field is set to one or zero depending on whether relay is involved and feasible;
- the Recipient ID field is set to the NID of the relayed node;
- the Sender ID field is set to the NID of the relaying node.

Upon receiving a control type frame, the recipient shall process the frame according to 6.2 specified for a one-hop star network with the additional considerations as noted above in defining the frame.

6.10.2 Selecting a relaying node for a two-hop extension

Either the relayed node or the target hub may select their relaying node through prearrangement.

The relayed node may also select node B as its relaying node if it recently received acknowledgment frames sent from node B to the target hub. The relayed node may receive such acknowledgment frames based on the frame reception rules specified in 6.2.4 with appropriate exceptions given to the values of the Recipient ID and Sender ID fields of the MAC header of the frames.

The relayed node may alternatively select node C as its relaying node if it recently received T-Poll frames broadcast by node C. The relayed node shall not select more than one node as its relaying node at any given time.

In selecting the relaying node, the relayed node should take into account the quality of the links between itself and the relaying node and between the relaying node and the target hub.

A relaying node in emergency, i.e., with its own pending Emergency frames to send to the target hub, may recognize whether the relayed nodes to which it is providing relay are in emergency through receipt of Emergency frames from them. Such a relaying node may continue or discontinue its relay service to relayed nodes that are not in emergency and to which it is providing relay, by taking into account such factors as quality of service (QoS), level of traffic load, and availability of power. The relayed nodes relaying through a relaying node in emergency may attempt to relay through an alternative relaying capable node not in emergency to reduce the load on the relaying node. A relaying node in emergency may accept or decline to provide relay to nodes newly requesting for relay, depending on whether the requesting nodes are in emergency, the number of relayed nodes that are in emergency and to which it is providing relay, etc.

6.10.3 Using broadcast T-Polls for a two-hop extension

To facilitate a two-hop extension for relayed nodes, a relaying node may obtain a scheduled uplink allocation in accordance with 6.7 to be used as described in the remainder of this subclause and illustrated in Figure 89, setting the UP of the allocation request to that for network control as defined in Table 18.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 89—Scheduled uplink allocation for relaying node to send broadcast T-Polls and to relay frames

In the scheduled uplink allocation, the relaying node shall broadcast T-Poll frames formatted as if it were sending them as a hub, with the modifications that in the MAC header

- the Recipient ID field is set to the Broadcast_NID;
- the Sender ID field is set to the NID of the relaying node.

Via such a T-Poll frame, the relaying node may provide either no or an immediate (shared) polled allocation within the scheduled uplink allocation, facilitating the selection of a relaying node and the synchronization with the target hub by potential relayed nodes, as well as offering the latter a frame transmission opportunity for frame relay to the target hub.

In the scheduled uplink allocation, the relaying node may also send to the hub or relayed nodes frames it has received from relayed nodes or the hub for further transmission to the hub or relayed nodes, respectively.

A relayed node that does not directly receive beacons from the target hub should indirectly synchronize with the hub through reception of T-Poll frames sent by a relaying node in the same BAN.

A relayed node may send at most a frame to the relaying node in a shared polled allocation, initially with contention probability (CP) $P = 1/2$. If the relayed node does not receive an expected acknowledgment from the relaying node, it may retry it in the shared polled allocation conveyed by another such T-Poll frame broadcast by the relaying node, with CP $P = \max(1/8, (1/2) / \lceil (R+1)/2 \rceil)$, where R counts the retries of the frame, i.e., R equals 1 for the first retry of the frame, 2 for the second retry, and so on. The function $\lceil x \rceil$ is defined to be the least integer not smaller than x . With CP P , the relayed node shall transmit if $z \leq P$ or shall not otherwise, where z is a value the relayed node has newly drawn at random from the interval $[0, 1]$.

A relayed node shall not send its frames to the relaying node in contention access phase (RAP) or contention access phase (CAP) provided by the target hub.

A relayed node shall not send a frame in a polled allocation if the frame transmission and the expected acknowledgment plus pSIFS would not be located within the allocation. A relaying node should not provide an immediate shared polled allocation that is not adequately long.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

A relayed node should send only management type or control type frames in shared polled allocations. To transmit data type frames, it should obtain scheduled allocations for its two-hop extension.

6.10.4 Using improvised access for a two-hop extension

The target hub and the relaying node may obtain improvised polled and posted allocations according to 6.6.1, as if they were a hub and a node, respectively, in a one-hop star network, to exchange data or management type frames originated from or destined to the relayed node.

The relaying node and the relayed node may obtain improvised polled and posted allocations in the scheduled uplink allocations applicable between the target hub and the relaying node according to 6.6.1, as if they were a hub and a node, respectively, in a one-hop star network, to exchange data or management type frames originated from or destined to the target hub.

6.10.5 Starting scheduled allocations for a two-hop extension

Either the relayed node or the target hub may initiate scheduled allocations for a two-hop extension as illustrated in Figure 90, which also shows equivalent one-hop scheduled allocations, regardless of whether the latter have been obtained.




Figure 90—Equivalent one-hop and two-hop scheduled allocations

6.10.5.1 Allocation request for a two-hop extension

To obtain scheduled allocations for a two-hop extension, the relayed node shall send an encapsulated Connection Request frame through the relaying node to the target hub, as illustrated in Figure 88(a) and specified in 6.10.1.1.

In this frame, the relayed node shall include an Uplink Request IE, a Downlink Request IE, or/and a Bilink Request IE specifying equivalent scheduled allocation(s) applicable between the relayed node and the target hub in a one-hop star network, setting the Allocation ID field of each Allocation Request field of these IEs to the Allocation ID value that identifies the corresponding one-hop scheduled allocation if it currently holds the latter.

6.10.5.2 Allocation assignment for a two-hop extension

To grant scheduled allocations for a two-hop extension, requested by the relayed node or initiated by itself, the target hub shall send an encapsulated Connection Assignment frame to the relaying node, which shall subsequently send the frame upon some modifications to the relayed node, as illustrated in Figure 88(c) and specified in 6.10.1.3.

In this frame, the target hub shall include appropriate Uplink Assignment IEs, Downlink Assignment IEs, or/and Bilink Assignment IEs specifying the two-hop scheduled allocations with the following matches to the corresponding one-hop scheduled allocations:

- Each allocation applicable between the target hub and the relaying node for the two-hop extension has the same allocation direction (i.e., uplink for uplink, downlink for downlink, and bilink for bilink), and approximately the same total allocation length, as the equivalent allocation applicable between the target hub and the relayed node in a one-hop star network.
- Each allocation applicable between the relaying node and the relayed node for the two-hop extension has a bilink allocation direction and more total allocation length than the equivalent allocation applicable between the target hub and the relayed node in a one-hop star network, with the additional allocation length provided for T-Poll frame transmissions by the relaying hub.

The target hub should take into account the overall latency that results from a two-hop extension in specifying the two-hop scheduled allocation intervals.

6.10.6 Using scheduled allocations for a two-hop extension

Upon successful transmission of the encapsulating Connection Assignment frame from the target hub to the relaying node, the target hub and the relaying node shall use their scheduled allocations conveyed in the encapsulated Connection Assignment frame according to 6.7.2, as if they were a hub and a node, respectively, in a one-hop star network, to exchange data, and occasionally management, type frames originated from or destined to the relayed node.

Upon successful transmission of the encapsulating Connection Assignment frame from the relaying node to the relayed node, the relaying node and relayed node shall use their scheduled bilink allocations conveyed in the encapsulated Connection Assignment frame according to 6.7.2, as if they were a hub and a node, respectively, in a one-hop star network, to exchange data type frames originated from or destined to the target hub. The relaying node should send one or more T-Poll frames to the relayed node in each allocation interval of the bilink allocations, and the relayed node should perform its clock synchronization based on received T-Poll frames.

6.10.7 Modifying scheduled allocations for a two-hop extension

The relayed node and the target hub may modify their two-hop scheduled allocations via another exchange of Connection Request or/and Connection Assignment frames according to 6.7.3, as if they were a node and a hub in a one-hop star network, with additional considerations based on the modified definition in 6.10.1.3 of the Connection Assignment frames for a two-hop extension.

6.10.8 Aborting scheduled allocations for a two-hop extension

The target hub and the relaying node shall abort their two-hop scheduled allocations according to 6.7.4, as if they were a hub and a node in a one-hop star network. Subsequently, the hub may reclaim the aborted scheduled allocations.

Once the relaying node aborts its two-hop scheduled allocations applicable with the target hub, it shall abort its two-hop scheduled allocations applicable with the relayed node as well.

The relaying node and the relayed node shall abort their two-hop scheduled allocations according to 6.7.4, as if they were a hub and a node, respectively, in a one-hop star network.

The target hub, the relaying node, and the relayed node shall respectively transmit at least one frame requiring an immediate return of a frame in every allocation interval of their two-hop scheduled allocations allowing for such a transmission, so as to reduce the chance of experiencing an abortion of their two-hop scheduled allocations, as also specified in 6.7.4 for one-hop star network.

6.10.9 Ending scheduled allocations for a two-hop extension

The relayed node or the target hub may initiate to end their two-hop scheduled allocations at any time the initiator determines as appropriate, in exchange for or without regaining equivalent one-hop scheduled allocations.

The relaying node may end the two-hop scheduled allocations applicable between a relayed node and the target hub by setting to zero the Relay field of the MAC header of a required I-Ack frame in response to a frame received from the relayed node or the target hub, when it determines that its relay between them is no longer feasible. The relaying may, and should, keep the two-hop scheduled allocations, if any, applicable between another relayed node and the target hub, so long as its relay between them is feasible and desirable.

A relayed node, a relaying node, or a target hub shall not send a frame in an already ended scheduled allocation.

6.10.9.1 In exchange for equivalent one-hop scheduled allocations

To request for replacing the two-hop scheduled allocations with equivalent one-hop scheduled allocations, the relayed node shall send a Connection Request frame directly to the target hub as in a one-hop star network, where the frame specifies the one-hop scheduled allocations using the Allocation IDs for the two-hop scheduled allocations. The relayed node should send the frame in a scheduled allocation applicable between itself and the relaying node.

To replace the two-hop scheduled allocations with equivalent one-hop scheduled allocations, in response to a request from the relayed node as described in the above or to its own decision, the target hub shall send to

- the relaying node an encapsulated Connection Assignment frame formatted to end the two-hop scheduled allocations as described in 6.7.5, and
- the relayed node directly a Connection Assignment frame formatted to specify the granted one-hop scheduled allocations in a one-hop star network.

The target hub should send the encapsulated Connection Assignment frame to the relaying node in a scheduled allocation applicable between itself and the relaying node. The target hub should send the Connection Assignment frame directly to the relayed node in a scheduled allocation currently or previously applicable between the relaying node and the relayed node.

6.10.9.2 Without regaining equivalent one-hop scheduled allocations

To request for ending the two-hop scheduled allocations without regaining equivalent one-hop scheduled allocations, the relayed node shall send a Connection Request frame directly to the target hub as in a

one-hop star network, where the frame is formatted with the Allocation IDs for the two-hop scheduled allocations to end the equivalent one-hop scheduled allocations as described in 6.7.5. The relayed node should send the frame in a scheduled allocation applicable between itself and the relaying node.

To end the two-hop scheduled allocations without regaining equivalent one-hop scheduled allocations, in response to a request from the relayed node as described in the above or to its own decision, the target hub shall send two encapsulated Connection Assignment frames 1 and 2 in two encapsulating Connection Assignment frames 1 and 2 to the relayed node and the relaying node, respectively, as described in 6.10.1.3, with the encapsulated Connection Assignment frames formatted to end the two-hop scheduled allocations.

6.11 Clock synchronization and guard time provisioning

A node or a hub shall maintain a MAC clock with a minimum resolution of $mClockResolution$ and with a minimum accuracy of $mHubClockPPMLimit$ to time its frame transmission and reception, except that a node may use a MAC clock with a PPM higher than $mHubClockPPMLimit$ subject to certain restrictions as stated later in this subclause. The node or the hub shall time its transmission and reception in any of their allocation intervals according to its local clock.

The node may request the hub to include a timestamp in an acknowledgment (I-Ack, B-Ack, I-Ack+Poll, or B-Ack+Poll) frame by setting to one the Ack Timing field of a management or data type frame being sent with the Ack Policy field of the MAC header set to I-Ack or B-Ack. The timestamp encodes the start time of the acknowledgment frame transmission based on the hub's clock. The hub shall include such a timestamp in the acknowledgment frame if and only if requested by the node.

The node shall synchronize to the hub through the beacons, T-Poll frames, acknowledgment frames containing a timestamp, or the first frames (on-time frames) in scheduled allocation intervals received from the hub. In particular, the node shall advance or delay its clock by a total amount as shown in Equation (4) and Equation (5), respectively:

$$D = T_S - T_L, \text{ if } T_S > T_L \quad (4)$$

$$D = T_L - T_S, \text{ if } T_S < T_L \quad (5)$$

where T_S is the time when such a frame started to be transmitted on the transport medium (i.e., air), and T_L is the time when the frame started to be received according to the local clock.

A node may rely on itself or a hub to track and set aside appropriate guard times in its allocation intervals. A hub shall be ready to accommodate either choice, referred to as *distributed* or *centralized* guard time provisioning, respectively, as indicated in the node's last transmitted MAC Capability field.

6.11.1 Distributed guard time provisioning

For distributed guard time provisioning, the node and the hub shall include appropriate guard times in the scheduled allocation intervals they requested or assigned, respectively. The hub shall also include appropriate guard times in the polled allocation intervals granted to the node.

6.11.1.1 Distributed guard time computation

If the node and the hub have the same clock accuracy designated as $HubClockPPM$ in terms of PPM, as shown in Figure 91, the node and the hub shall compute a nominal guard time GT_n to compensate for their

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

clock drifts over an interval not longer than a nominal synchronization interval SI_n , as shown in Equation (6), Equation (7), and Equation (8):

$$GT_n = GT_0 + 2 \times D_n \quad (6)$$

$$GT_0 = pSIFS + pExtraIFS + mClockResolution \quad (7)$$

$$D_n = SI_n \times HubClockPPM, SI_n = mNominalSynchInterval \quad (8)$$

The parameter GT_0 comprises the receive-to-transmit or transmit-to-receive turnaround time $pSIFS$, the synchronization error tolerance $pExtraIFS$, and the timing uncertainty $mClockResolution$, which are all of fixed values that are independent of clock drifts. The parameter D_n represents the maximum clock drift of the node or the hub relative to an ideal (nominal) clock over SI_n . The parameter SI_n delimits a nominal synchronization interval over which the clock drifts of the node and the hub are accounted for in the nominal guard time GT_n .

The node shall further compute an additional guard time GT_a to compensate for additional clock drifts of itself and the hub over an interval SI_a beyond SI_n , as shown in Equation (9):

$$GT_a = 2 \times D_a, D_a = SI_a \times HubClockPPM \quad (9)$$

The parameter SI_a denotes the length of the time interval that has accrued in addition to SI_n since the node's last synchronization with the hub. The corresponding additional clock drift D_a is a function of SI_a and accounts for the required additional guard time GT_a . The values of D_a and SI_a are specific to the node and time of concern.

A node may time its frame transmission and reception with a clock accuracy $NodeClockPPM$ larger than $HubClockPPM$, provided it reduces its nominal synchronization interval to SI_n such that, as shown in Equation (10):

$$SI_n \times NodeClockPPM = mNominalSynchInterval \times HubClockPPM \quad (10)$$

If the time interval length SI since its last synchronization with the hub exceeds the reduced SI_n by SI_a , i.e., if $SI = SI_n + SI_a$, the node shall calculate the required additional guard time GT_a as shown in Equation (11):

$$GT_a = SI_a \times NodeClockPPM + \min[0, (SI - mNominalSynchInterval) \times HubClockPPM] \quad (11)$$

An illustration of clock drifts and guard times for the case of a hub and nodes operating with the same clock accuracy is given in Figure 91, with the following legend:

N_f = fast node N_s = slow node H = slow hub in (a) and fast hub in (b)

tm_H = position of ideal (nominal) clock when N_H 's local clock is at tm , $m = 1, \dots$, or 4

tm_f = position of ideal (nominal) clock when N_f 's local clock is at tm , $m = 1, \dots$, or 4

tm_s = position of ideal (nominal) clock when N_s 's local clock is at tm , $m = 1, \dots$, or 4

SI_n = nominal synchronization interval GT_n = nominal guard dtime

D_n = maximum clock drift over SI_n relative to ideal clock

SI_a = additional synchronization interval GT_a = additional guard time

D_a = maximum clock drift over SI_a relative to ideal clock

allocation interval of H = allocation interval in which H controls the timing for frame transactions

allocation interval of N = allocation interval in which N controls the timing for frame transactions

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 91—Analysis of clock drifts and guard times for distributed provisioning

6.11.1.2 Distributed guard time compensation

With reference to Figure 91 and Figure 92, and with GT_n given in Equation (6), GT_0 in Equation (7), SI_n in Equation (8) or Equation (10) as appropriate, and GT_a in Equation (11), the node and the hub shall account for clock drifts and guard times in their frame transmission and reception as follows:

- The hub shall commence its beacon transmission at the nominal start of the beacon.
- The hub shall commence its transmission in the node's next scheduled downlink or bilink allocation interval at the nominal start of the interval, and shall end its transmission in the interval

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

early enough such that the last transmission in the interval completes at least GT_n prior to the nominal end of the interval.

- The hub shall commence its transmission of the node's next future poll or post at the nominal start of the poll or post.
- The hub shall commence its reception in the node's next scheduled uplink allocation interval up to $GT_n - GT_0$ earlier than the nominal start of the interval to account for pertinent clock drifts.
- If the node's last synchronization to the hub was less than SI_n ago at the nominal end of its next scheduled uplink or polled allocation interval, the node shall commence its transmission in the interval at the nominal start of the interval, and the node shall end its transmission in the interval early enough such that the last transmission in the interval completes at least GT_n prior to the nominal end of the interval.
- If the node's last synchronization to the hub was less than SI_n ago at the nominal start of the next beacon transmission, the node shall commence its reception of the beacon up to $GT_n - GT_0$ earlier than the nominal start of the beacon to account for pertinent clock drifts.
- If the node's last synchronization to the hub was less than SI_n ago at the nominal start of its next future poll or post, the node shall commence its reception of the poll or post up to $GT_n - GT_0$ earlier than the nominal start of the poll or post to account for pertinent clock drifts.
- If the node's last synchronization to the hub was less than SI_n ago at the nominal start of its next scheduled downlink or bilink allocation interval, the node shall commence its reception in the interval up to $GT_n - GT_0$ earlier than the nominal start of the interval to account for pertinent clock drifts. The node may commence its reception up to $GT_n - GT_0$ later than the start of the interval based on its estimate of the relative clock drift with respect to the hub since its last synchronization with the hub.
- If the node's last synchronization to the hub was $SI_n + SI_a$ ago at the nominal end of its next scheduled uplink allocation interval, the node shall commence its transmission in the interval GT_a later than that nominal start time, and shall end its transmission in the interval early enough such that the last transmission in the interval completes at least $GT_n + GT_a$ prior to the nominal end of the interval.
- If the node's last synchronization to the hub was $SI_n + SI_a$ ago at the nominal end of its next polled allocation interval, the node shall commence its transmission in the interval at the nominal start of the interval, and shall end its transmission in the interval early enough such that the last transmission in the interval completes at least $GT_n + GT_a$ prior to the nominal end of the interval.
- If the node's last synchronization to the hub was less than $SI_n + SI_a$ ago at the nominal start of the next beacon transmission, the node shall commence its reception of the beacon up to $GT_n + GT_a - GT_0$ earlier than the nominal start of the beacon to account for pertinent clock drifts.
- If the node's last synchronization to the hub was less than $SI_n + SI_a$ ago at the nominal start of its next future poll or post, the node shall commence its reception of the poll or post up to $GT_n + GT_a - GT_0$ earlier than the nominal start of the poll or post to account for pertinent clock drifts.
- If the node's last synchronization to the hub was $SI_n + SI_a$ ago at the nominal start of its next scheduled downlink or bilink allocation interval, the node shall commence its reception in the interval up to $GT_n + GT_a - GT_0$ earlier than the nominal start of the interval to account for pertinent clock drifts. The node may commence its reception up to $GT_n + GT_a - GT_0$ later than the start of the interval based on its estimate of the relative clock drift with respect to the hub since its last synchronization with the hub.




Figure 92—Distributed provisioning of guard times for frame transmissions

6.11.1.3 Distributed guard time allocation

The node and the hub shall include a nominal guard time GT_n as given in Equation (6) and, if applicable, twice an additional guard time GT_a as given in Equation (11) in each of the scheduled allocation intervals they request or assign. The hub shall also include the nominal guard time GT_n in each of the polled allocation intervals granted to the node.

6.11.1.4 Clock synchronization for distributed guard time provisioning

The node shall synchronize with the hub at least once within the nominal synchronization interval SI_n given in Equation (8) or Equation (10) as appropriate, if only the nominal guard time GT_n as given in Equation (6) is accounted for per 6.11.1.3. The node shall synchronize with the hub at least once within the nominal synchronization interval SI_n given in Equation (8) or Equation (10) as appropriate, plus the additional synchronization interval SI_a given in Equation (10), if both the nominal guard time GT_n as given in Equation (6) and the additional guard time GT_a as given in Equation (11) are accounted for per 6.11.1.3.

6.11.2 Centralized guard time provisioning

For centralized guard time provisioning, the node shall not include clock drifts or guard times in the scheduled allocation intervals it requests, but the hub shall include appropriate clock drifts in the downlink or bilink scheduled allocation intervals it assigns to the node. The hub shall also provision an appropriate guard time between two neighboring allocation intervals one or both of which are assigned to the node requiring centralized guard time provisioning.

6.11.2.1 Centralized guard time computation

As shown in Figure 93, the hub shall compute a centralized guard time GT_c between two neighboring allocation intervals (with beacon treated as an allocation interval), both of which do not include a guard time, to compensate for pertinent clock drifts, as follows:

For case (a) where each of the two allocation intervals is a beacon or an allocation interval in which the hub controls the timing for frame transactions, as shown in Equation (12):

$$GT_c = GT_0 \quad (12)$$

For case (b) where one of the two allocation intervals is a beacon or an allocation interval in which the hub controls the timing for frame transactions, and the other is an allocation interval in which the node controls the timing for frame transactions, given the node's maximum synchronization interval SI_N and its clock accuracy P_N in terms of PPM, and the hub's clock accuracy P_H in terms of PPM, as shown in Equation (13):

$$GT_c = GT_0 + SI_N \times (P_H + P_N) \quad (13)$$

For case (c) where one of the two allocation intervals is an allocation interval in which the node controls the timing for frame transactions, and the other is an allocation interval in which another node controls the timing for frame transactions, given the node's maximum synchronization interval SI_{N1} and its clock accuracy P_{N1} in terms of PPM, the other node's maximum synchronization interval SI_{N2} and its clock accuracy P_{N2} in terms of PPM, and the hub's clock accuracy P_H in terms of PPM, with the other node also requiring centralized guard time provisioning, as shown in Equation (14):

$$GT_c = GT_0 + P_{N1} \times SI_{N1} + P_{N2} \times SI_{N2} + P_H \times |SI_{N1} - SI_{N2}| \quad (14)$$

The parameter GT_0 is a fixed value independent of clock drifts as given in Equation (7).

In Figure 93(a), there are no relative clock drifts since it is the same hub that controls the timing for frame transactions in both allocation intervals. In Figure 93(b), since the node last synchronized to the hub SI_N ago, the hub's clock has drifted by D_H toward the other allocation interval, and the node's clock has drifted by D_N toward the other direction, both relative to an ideal clock. In Figure 93(c), since the two nodes last synchronized to the hub SI_{N1} and SI_{N2} ago, their clocks have drifted by D_{N1} and D_{N2} in opposite directions, respectively; between the times of the nodes' last synchronization, the hub's clock has also drifted by D_H in the same direction as the clock of the node that synchronized with the hub later, all relative to an ideal clock.

Of the two neighboring allocation intervals, in case the earlier one is provided for distributed guard time provisioning and thus includes a nominal guard time GT_n as given in Equation (6) at the end, the hub may deduct GT_n from GT_c in inserting a centralized guard time between the two intervals. Further, if the earlier one is a scheduled uplink or polled allocation interval provided to a node for distributed guard time provisioning, the hub shall set SI_N or SI_{N1} to SI_n as given in Equation (8) in computing GT_c according to Equation (13) or Equation (14).

On the other hand, in case the later one is a scheduled downlink, bilink, or uplink allocation interval assigned to a node requiring distributed guard time provisioning, the hub shall treat such an interval as one assigned for centralized guard time provisioning in inserting a centralized guard time between the two intervals. Further, if such an interval is a scheduled uplink allocation interval, the hub shall set SI_N or SI_{N2} to SI_n as given in Equation (8) in computing GT_c according to Equation (13) or Equation (14), respectively.

An illustration of clock drifts and guard times for the case of both neighboring allocation intervals (with beacon treated as an allocation interval) not including guard times is given in Figure 91, with the following legend:

H = hub N = node N1 = node 1 N2 = node 2

P_H = PPM of H's clock P_N = PPM of N's clock

P_{N1} = PPM of N1's clock P_{N2} = PPM of N2's clock

SI_N = maximum synchronization interval of N

SI_{N1} = maximum synchronization interval of N1

SI_{N2} = maximum synchronization interval of N2

D_H = clock drift of H relative to ideal clock D_N = clock drift of N, ideal clock

D_{N1} = clock drift of N1 relative to ideal clock D_{N2} = clock drift of N2 relative to ideal clock

GT_c = centralized guard time

allocation interval of H = allocation interval in which H controls the timing for frame transactions

allocation interval of N = allocation interval in which N controls the timing for frame transactions




Figure 93—Analysis of clock drifts and guard times for centralized provisioning

6.11.2.2 Centralized guard time compensation

With reference to Figure 93 and Figure 94, and with GT_0 given in Equation (7), and GT_c in Equation (12), Equation (13), or Equation (14) as appropriate, the node and the hub shall account for clock drifts in their frame transmission and reception as follows, where the node applies Equation (13) to calculate GT_c for its reception time:

- The hub shall commence its beacon transmission at the nominal start of the beacon.
- The hub shall commence its transmission in the node's next scheduled downlink or bilink allocation interval at the nominal start of the interval, and shall end its transmission in the interval early enough such that the last transmission in the interval completes by the nominal end of the interval.
- The hub shall commence its transmission of the node's next future poll or post at the nominal start of the poll or post.
- The hub shall commence its reception in the node's next scheduled uplink allocation interval up to $GT_c - GT_0$ earlier than the nominal start of the interval to account for pertinent clock drifts since the node last synchronized with it.
- The node shall commence its transmission in a scheduled uplink allocation interval at the nominal start of the interval, and shall end its transmission in the interval early enough such that the last transmission in the interval completes by the nominal end of the interval.
- The node shall commence its reception of the beacon up to $GT_c - GT_0$ earlier than the nominal start of the beacon to account for pertinent clock drifts since it last synchronized with the hub.
- The node shall commence its reception in its next scheduled downlink or bilink allocation interval up to $GT_c - GT_0$ earlier or later than the nominal start of the interval to account for pertinent clock drifts since it last synchronized with the hub.
- The node shall commence its reception of its next poll or post up to $GT_c - GT_0$ earlier than the nominal start of the poll or post to account for pertinent clock drifts, where the node's last synchronization interval is measured up to the nominal start of the poll or post.




Figure 94—Centralized provisioning of guard times for frame transmissions

6.11.2.3 Centralized guard time allocation

The node shall not include clock drifts or guard times in the scheduled allocation intervals it requests. The hub shall include $2 \times (GT_c - GT_0)$ with GT_c given in Equation (13) in each of the scheduled downlink or bilink allocation intervals it assigns to the node. The hub shall also provision at least a centralized guard time GT_c given in Equation (12), Equation (13), or Equation (14) as appropriate, between two neighboring allocation intervals, minus a nominal guard time GT_n given in Equation (6) if the earlier one of the allocation intervals is provided to a node requiring distributed guard time provisioning and hence includes GT_n in the end, treating a beacon as an allocation interval that does not include GT_n .

6.11.2.4 Clock synchronization for centralized guard time provisioning

The node shall synchronize with the hub at least once within its maximum synchronization interval SI_N as indicated in its last transmitted Connection Request frame.

6.12 Power management

A node not indicating to be always active in its last transmitted MAC Capability field, referred to as a *node in short* in this subclause, may hibernate, i.e., be in inactive state across its non-wakeup beacon periods (superframes). It may sleep, i.e., be in inactive state over some time intervals in its wakeup beacon periods (superframes).

6.12.1 Hibernation—macroscopic power management

To be in inactive state across more than one beacon period (superframe), a node shall set the Requested Wakeup Period field in its last Connection Request frame sent to a hub to an integer larger than one, while setting the Requested Wakeup Phase field in the frame to a value specifying its intended next wakeup beacon period (superframe).

To wake up, i.e., be ready for receiving or transmitting frames, in every beacon period (superframe), a node shall set the Requested Wakeup Period field in its last Connection Request frame sent to a hub to one, while setting the Requested Wakeup Phase field in the frame to a value identifying the next beacon period (superframe).

The hub should honor the values of the Requested Wakeup Period and, to a less extent, Requested Wakeup Phase fields whenever possible, but may set its Assigned Wakeup Period and Assigned Wakeup Phase fields of its responding Connection Assignment frame to different values if need be. The hub may later modify these values by sending to the node another Connection Assignment frame if warranted by new operating conditions.

If the hub sets the Assigned Wakeup Period field in its responding frame to an integer not equal to one, it may grant only m-periodic allocations to the node, with the allocation intervals being in the node's wakeup beacon periods (superframes), in accordance with the node's last connection request whenever possible, but shall not grant to the node any 1-periodic allocations.


If the hub sets the Assigned Wakeup Period field in its responding frame to one, it may grant only 1-periodic allocations to the node, with the allocation intervals being in every beacon period (superframe), in accordance with the node's last connection request whenever possible, but shall not grant to the node any m-periodic allocations.

If the Assigned Wakeup Period value in the Connection Assignment frame last received from the hub is not equal to one, the node shall wake up in each of its wakeup beacon periods (superframes) based on the latest Assigned Wakeup Period and Assigned Phase Wakeup values provided in that frame by the hub, to transmit or receive frames in the granted m-periodic allocation intervals, and to receive the beacon if needed.


If the Assigned Wakeup Period value in the Connection Assignment frame last received from the hub is one, the node shall wake up in every beacon period (superframe), to transmit or receive frames in the granted 1-periodic allocation intervals, and to receive the beacon as appropriate.

Figure 95 illustrates macroscopic power management across beacon periods (superframes).

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks



(a) Assigned Wakeup Period $\neq 1$



(b) Assigned Wakeup Period = 1

Figure 95—Macroscopic power management

6.12.2 Sleep—microscopic power management


In a wakeup beacon period (superframe), in beacon mode with superframes, a node may be in inactive state during the beacon transmission time, if it does not need to receive a beacon during this time. The node may be in inactive state in exclusive access phase 1 (EAP1), random access phase 1 (RAP1), exclusive access phase 2 (EAP2), or random access phase 2 (RAP2), if it does not need to transmit a management or data type frame in the corresponding access phase.

In a MAP, a node may be in inactive state outside its scheduled allocation intervals and polled or posted allocation intervals (a future poll/post is sent in a posted allocation). It may also be in inactive state during its scheduled or polled allocation intervals on determining that the remaining interval has been relinquished by itself or reclaimed by its hub due to no more pending transmissions as specified in 6.8.


In beacon mode with superframes, to receive a B2 frame or to transmit in a CAP, a node shall be active in the time interval wherein a B2 frame may be sent as specified in 6.3.

Figure 96 illustrates microscopic power management in a wakeup beacon period (superframe).

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks



(a) Beacon mode with superframes



(b) Non-beacon mode without superframes

Figure 96—Microscopic power management

6.13 Coexistence and interference mitigation

A hub may employ one or more of the optional mechanisms described in this subclause for coexistence and/or interference mitigation between its BAN and neighbor BANs.

6.13.1 Beacon shifting

A hub may transmit its beacons at different time offsets relative to the start of the beacon periods by including a Beacon Shifting Sequence field in its beacons as defined in 5.3.1.10. A hub should choose a beacon shifting sequence that is not being used by its neighbor hubs to mitigate potential repeated beacon collisions and scheduled allocation conflicts between overlapping or adjacent BANs operating in the same channel.

As shown in Figure 97, the hub shall transmit a beacon out of its PHY at a time $t = PN_m(n) \times BP/4$ relative to the start of beacon period n . Here, m is the beacon shifting sequence index that the hub has chosen for its BAN, BP is the length of its beacon period, and n is the phase of the chosen sequence ($n = 0, 1, \dots$) for this beacon period.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 97—Beacon shifting illustration with $PN_4(n)$

As defined in 5.3.1.2, the allocation slots in a beacon period shift around with the beacon transmit time. The access phases—exclusive access phase 1 (EAP1), random access phase 1 (RAP1), exclusive access phase 2 (EAP2), random access phase 2 (RAP2), and contention access phase (CAP)—shown in Figure 64 are referenced to numbered allocation slots and shift around with the beacon in the beacon period accordingly. The RAP1 and RAP2 related fields contained in the beacon of the current beacon period now refer to the EAP1, RAP1, EAP2, and RAP2 in the next beacon period. A node does not know nor use these access phases in the beacon period in which it received its first beacon indicating beacon shifting is enabled, but it may use a MAP through polls or posts.

In choosing access phases and the beacon shifting sequence, the beacon shift shall not result in a split of any of the aforementioned access phases into two parts.

Scheduled allocation intervals are also referenced to numbered allocation slots and shift around with the beacon transmit time accordingly in the beacon period as shown in Figure 98. A scheduled allocation interval in a beacon period may be split into two portions as a result of shifting around the beacon period, but the aggregate length remains the same.




Figure 98—Beacon and scheduled allocation interval shifting

6.13.2 Channel hopping

A hub may enable channel hopping only if its PHY is a narrow band (NB) PHY not operating in the MICS band or a frequency modulation ultra-wideband (FM-UWB) PHY. In such cases, the hub may change its operating channel in the operating frequency band periodically by including the Channel Hopping State and Next Channel Hop fields in its beacons as defined in 5.3.1.11 and 5.3.1.12, respectively, or/and in the

Superframe Parameters IE of its Connection Assignment frames as defined in 5.7.1. A hub should choose a channel hopping sequence that is not being used by its neighbor hubs.

The hub shall hop to another channel after dwelling in the current channel for a fixed number of beacon periods (superframes) as communicated to the nodes connected with the hub through Connection Assignment frames. To hop to a new channel, the hub shall start switching to the new channel at pChannelSwitchTime prior to the start of the beacon period (superframe) that begins with the new channel, neither sending nor receiving frames during the channel switch. A node should not send a frame during this transition.

A hub shall generate a channel hopping sequence based on the maximum-length Galois linear feedback shift register (LFSR) defined in Figure 99 and by the generator polynomial shown in Equation (15):

$$g(x) = x^{16} + x^{14} + x^{13} + x^{11} + 1 \quad (15)$$

The state of the LFSR at stage k is given by Equation (16):

$$Y_k = 2^0 \times r_{k,0} + 2^1 \times r_{k,1} + \dots + 2^{15} \times r_{k,15} \quad (16)$$

Y_k represents the binary value read from the bits $r_{k,0}, r_{k,1}, \dots, r_{k,15}$ of the individual registers at stage k , with $r_{k,0}$ being the LSB and $r_{k,15}$ being the MSB.




Figure 99—16-bit Galois LFSR for channel hopping sequence generation

Given the current state Y_k of the LFSR, the hub shall generate the next state Y_{k+1} of the LFSR, i.e., the state of the LFSR at the next stage, stage $k+1$. Accordingly, the hub shall generate the channel index C_{k+1} identifying the next channel from the channel index C_k identifying the current channel and the next state Y_{k+1} of the LFSR as shown in Equation (17), Equation (18), and Equation (19):

$$N_{reduced} = N_{ch} - 2N_{sep} + 1 \quad (17)$$

$$Z_{k+1} = Y_{k+1} \bmod N_{reduced} \quad (18)$$

$$C_{k+1} = (C_k + Z_{k+1} + N_{sep}) \bmod N_{ch} \quad (19)$$

In the equations, the notation “mod” denotes modular operation. N_{ch} = pChannelsTotal is the number of total channels in the operating frequency band as listed in Table 25. N_{sep} = pChannelSeparation is the minimum number of channels separated between two consecutive hops as illustrated in Figure 100. $N_{reduced}$ is the number of channels available for each hop on account of the channel separation constraint N_{sep} . The channel index is the channel number as specified in the corresponding PHY clause.

If channel hopping is currently enabled, or to be enabled later, in selected channels of the operating frequency band, the hub shall include in its Connection Assignment frames a Channel Hopping and Ordering IE indicating those channels. In this case, N_{ch} is the number of channels included in channel hopping in the operating frequency band as indicated in the IE, and N_{sep} is the value of the Channel Separation field as also indicated in the IE. Moreover, the channel index = C identifies a channel designated

by the Cth of all the bits set to one in the Channel Bitmap of the IE, with such bits numbered at zero from the LSB that is set to one.

Given the current channel number N_k , the next channel number N_{k+1} is such that $|N_{k+1} - N_k| \geq N_{sep}$. All channels are selected with equal probability, as also shown in Figure 100.




Figure 100 —Channel separation in consecutive hops

The hub shall set the initial state Y_0 of its LFSR to the 16 LSBs of its EUI-48, with $r_{0,0}$ corresponding to the LSB. The hub shall select the channel number C_0 of its initial channel as shown in Equation (20):


$$C_0 = Y_0 \bmod N_{ch} \quad (20)$$

To communicate with a hub, a node shall hop to the same channel as the hub. If the hub does not enable channel hopping, as indicated by exclusion of the Channel Hopping State and Next Channel Hop fields from its last transmitted beacon or/and Connection Assignment frames, the node should find the hub's operating channel according to the channel order list provided by the hub, if any.


6.13.3 Active superframe interleaving

A BAN may share the same operating channel with one or more other BANs with or without interleaving their active superframes.

A BAN, denoted BAN 1, may at any time share the same operating channel with another BAN, denoted BAN 2, through interleaving their active superframes as illustrated in Figure 101. A hub that supports active superframe interleaving and operates in non-beacon mode with superframes shall send a B2 frame in every active superframe.



(a) Without active superframe adjustment



(b) With active superframe adjustment

Figure 101 —Active superframe interleaving illustration

Regardless of whether BAN 1 is operating with a beacon period (superframe) length and inactive duration that are suitable for interleaving BAN 1 and BAN 2's active superframes, BAN 2's hub, denoted hub 2, may send to BAN 1's hub, denoted hub 1, a Command–Active Superframe Interleaving Request frame to request for active superframe interleaving between the two BANs, after receiving a beacon or B2 frame of hub 1 with the Superframe Interleaving field of the MAC header set to one indicating support for active superframe interleaving.

Upon acknowledging receipt of the frame, hub 1 should send to hub 2 a Command–Active Superframe Interleaving Response frame to indicate whether it accepts or rejects the request, if it supports active superframe interleaving as indicated in the MAC header of its last transmitted beacon or B2 frame. If hub 1 accepts the request,

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

- in some cases as illustrated in Figure 101(a), it may continue with its current beacon period (superframe) length and inactive duration to enable the offered active superframe interleaving;
- in other cases as illustrated in Figure 101(b), it shall adjust its beacon period (superframe) length and inactive duration to enable the offered active superframe interleaving before sending its response.

Hub 1 should accept the request if it may continue with its current beacon period (superframe) length and inactive duration to enable the requested superframe interleaving, as in the cases illustrated in Figure 101(a), unless its inactive duration has been mostly taken by other hubs also for active superframe interleaving. Hub 1 should also accept the request from hub 2, if hub 1 has a lower BAN priority than hub 2.

If hub 1 rejects the request, it may continue with its current beacon period (superframe) length and inactive duration, even if it offers alternative beacon period (superframe) and inactive duration values in its response for active superframe interleaving between the two BANs.

If hub 1 accepts hub 2's request, hub 2 should set up or/and adjust its beacon period (superframe) boundary and inactive duration to attain active superframe interleaving as it has requested, as illustrated in Figure 101, once hub 1 makes its own adjustment if required as in the cases illustrated in Figure 101(b).

If hub 1 rejects hub 2's request, hub 2 may send another request for active superframe interleaving based on the alternative offer in hub 1's response, or may start or continue BAN 2 operation in the same channel without regard to active superframe interleaving.

Hub 2 may send to hub 1 another Command–Active Superframe Interleaving Request frame even if it has previously sent such a frame containing the same or different requested field values. If the new request is accepted, it shall supersede the previous request. If the new request is rejected, the last accepted request, if any, shall remain valid.

If hub 2 previously sent to hub 1 a request for active superframe interleaving and the request was accepted by hub 1, hub 2 should send to hub 1 another request when hub 2 needs fewer or no active superframes.

Hub 1 may send to hub 2 a Command–Active Superframe Interleaving Request frame for active superframe interleaving any time as well, following the procedure specified in the above with hub 1 and hub 2 swapping their roles.

To send a Command–Active Superframe Interleaving Request or a Command–Active Superframe Interleaving Response frame, the sender shall send the frame as if it were an unconnected node of the recipient's BAN, for both medium access and MAC header setting. The transmission and setting of I-Ack frames for acknowledging receipt of the above two frames shall be the same as for acknowledging receipt of any other frame.

6.14 MAC/PHY capability handling/interaction and Application Specific IE usage

6.14.1 MAC/PHY optional capability support

A node or a hub shall be ready to carry out a function that it has indicated to be capable of supporting in its last transmitted MAC Capability or PHY Capability field. A node or a hub shall not initiate or execute a function that its communication partner has indicated to be incapable of supporting through the last MAC Capability or PHY Capability field sent by the partner.

To transmit beacons, Poll and T-Poll frames addressed to Unconnected_Broadcast_NID, Local_Broadcast_NID, or Broadcast_NID as appropriate, a hub shall use a mandatory data rate in the operating frequency band for the underlying PHY.

6.14.2 MAC/PHY interaction for hybrid ARQ

A node and a hub may employ type-II hybrid automatic repeat request (HARQ) in transmitting and receiving a frame as specified in Clause 9 for the UWB PHY if pHybridARQ is set to TRUE.

The sender shall send the frame with the Ack Policy field of the MAC header set to I-Ack. Unless the sender receives an expected I-Ack frame or aborts the current frame retransmission at the PHY, it shall transmit a frame containing alternately either the parity bits of the MAC frame constructed at the PHY per the HARQ scheme or the MAC frame itself, pSIFS after the estimated end of the expected I-Ack frame.

The recipient shall receive and acknowledge the frame according to 6.2.4 and 6.2.9.

When a sender invokes a type-II HARQ operation for a frame transmission, the transmissions of the frame containing a MAC frame or the parity bits thereof, the expected I-Ack frames, and an appropriate guard time shall all fit within an allocation interval it has obtained.

6.14.3 Application Specific IE usage

A hub may include one or more Application Specific IEs at the end of its beacon.

A node may include one or more Application Specific IEs at the end of its Connection Request frame. A hub may also include one or more Application Specific IEs at the end of its Connection Assignment frame.

A recipient shall ignore unrecognized Application Specific IEs.

6.15 MAC sublayer parameters

Table 24 provides the values for the MAC sublayer parameters.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 24—MAC sublayer parameters

Parameter	Value
mBackLimit	8
mCSMATxLimit	2 for $UP \leq 5$ or 4 for $UP \geq 6$
mHubClockPPMLimit	40 ppm
mClockResolution	4 μ s
mG-AckDataSubtype	1111 (binary)
mMaxFragmentCount	8
mMaxBANSize	64
mNominalSynchInterval	$8 \times$ Beacon Period (Superframe) Length
mScheduledAllocationAborted	32
mTimeOut	30 μ s
mUnscheduledAllocationAborted	32
mUnscheduledNoResponseLimit	3

Table 25, Table 26, and Table 27 provide the values of the PHY dependent parameters used by the MAC sublayer.

Table 25—PHY-dependent MAC sublayer parameters pertaining to narrowband PHY

Parameter	Value
pAllocationSlotMin	500 μ s
pAllocationSlotResolution	500 μ s
pCCATime	63 / Symbol Rate (see Table 29 to Table 35 for Symbol Rate)
pChannelSeparation	2
pChannelsTotal	See Table 45
pChannelSwitchTime	100 μ s
pCSMAMACPHYTime	40 μ s
pCSMASlotLength	$pCCATime + pCSMAMACPHYTime$
pExtraIFS	10 μ s
pHybridARQ	FALSE
pMaxFrameBodyLength	255 octets
pMICSCchannelsTotal	10
pMICSCchannelSwitchTime	100 μ s
pMICSHubMaxRetries	10
pMICSMcastPollRxTime	$pMICSPollTxTime + pMIFS + pMICSPreambleTxTime = 1567 \mu$ s
pMICSMcastPolls	$\lceil pMICSCchannelsTotal \times (pMICSMcastPollRxTime + pMICSCchannelSwitchTime) / (pMICSPollTxTime + pMIFS) \rceil = 16$
pMICSCnodeEmergencyRetries	2
pMICSPollRxTime	$pMICSPollTxTime + pMICSPollSpace + pMICSPreambleTxTime = 2157 \mu$ s
pMICSPollSpace	$2 \times pSIFS + pMICSPreambleTxTime + mTimeOut = 610 \mu$ s

**Table 25—PHY-dependent MAC sublayer parameters pertaining to narrowband PHY
 (continued)**

Parameter	Value
pMICSPollTxTime	$pMICSPreambleTxTime + pMICSPLCHeaderTxTime + \{(7+2) \times 8 + 12 \times \lceil (7+2) \times 8/51 \rceil\} / 187.5 \text{ ms} = 1323 \text{ } \mu\text{s}$
pMICSPreambleTxTime	$90 / 187.5 \text{ ms} = 480 \text{ } \mu\text{s}$
pMICSPLCHeaderTxTime	$2 \times 31 / 187.5 \text{ ms} = 331 \text{ } \mu\text{s}$
pMICSUnconnectedPollPeriod	$> (pMICSUnconnectedPollTxTime + pMICSPollSpace) \times pMICSUnconnectedPolls = 25 \text{ } 130 \text{ } \mu\text{s}$
pMICSUnconnectedPollRxTime	$pMICSUnconnectedPollTxTime + pMICSPollSpace + pMICSPreambleTxTime = 2275 \text{ } \mu\text{s}$
pMICSUnconnectedPolls	$\lceil pMICSChannelsTotal \times (pMICSUnconnectedPollRxTime + pMICSCChannelSwitchTime) / (pMICSUnconnectedPollTxTime + pMICSPollSpace) \rceil = 14$
pMICSUnconnectedPollTxTime	$pMICSPreambleTxTime + pMICSPLCHeaderTxTime + \{(7+4+2) \times 8 + 12 \times \lceil (7+4+2) \times 8/51 \rceil\} / 187.5 \text{ ms} = 1558 \text{ } \mu\text{s}$
pMIFS	$20 \text{ } \mu\text{s}$
pRandomAccess	CSMA/CA
pSIFS	$75 \text{ } \mu\text{s}$
pUnconnectedPolledAllocationMin	$\geq$ transmission time of two PHY packets containing a MAC frame of 7+104+2 and 7+2 octets, respectively, both transmitted at the highest mandatory data rate of the operating frequency band specified in Clause 8.

Table 26—PHY-dependent MAC sublayer parameters pertaining to UWB PHY

Parameter	Value
pAllocationSlotMin	$16 \text{ } \mu\text{s}$
pAllocationSlotResolution	$16 \text{ } \mu\text{s}$
pAlohaSlotLength	pUnconnectedPolledAllocationMin
pCCATime	$252 \text{ } \mu\text{s}$
pCSMAMACPHYTime	$40 \text{ } \mu\text{s}$
pCSMASlotLength	$pCCATime + pCSMAMACPHYTime$
pExtraIFS	$10 \text{ } \mu\text{s}$
pHybridARQ	TRUE
pMaxFrameBodyLength	255 octets
pMIFS	$20 \text{ } \mu\text{s}$
pRandomAccess	CSMA/CA or Slotted Aloha
pSIFS	$75 \text{ } \mu\text{s}$
pUnconnectedPolledAllocationMin	$\geq$ transmission time of two PHY packets containing a MAC frame of 7+104+2 and 7+2 octets, respectively, both transmitted at the mandatory data rate of the operating frequency band specified in Clause 9.

Table 27 —PHY-dependent MAC sublayer parameters pertaining to HBC PHY

Parameter	Value
pAllocationSlotMin	500 μ s
pAllocationSlotResolution	500 μ s
pAlohaSlotLength	pUnconnectedPolledAllocationMin
pExtraIFS	10 μ s
pHybridARQ	FALSE
pMaxFrameBodyLength	255 octets
pMIFS	20 μ s
pRandomAccess	Slotted Aloha
pSIFS	75 μ s
pUnconnectedPolledAllocationMin	$\geq$ transmission time of two PHY packets containing a MAC frame of 7+104+2 and 7+2 octets, respectively, both transmitted at the highest mandatory data rate of the operating frequency band specified in Clause 10.

7. Security services

This clause expounds on the elements of the security hierarchy introduced in Figure 5. Security in this standard starts with a negotiation of the desired security suite between the two communicating parties, a node and a hub. The security selection in turn sets off a security association between the two parties for activating a pre-shared or generating a new shared master key (MK). Several security association protocols suitable for a variety of use cases are provided in 7.1, which finishes off with security disassociation for legitimately repealing a shared MK between the two parties. Pairwise temporal key (PTK) creation and group temporal key (GTK) distribution are then described in 7.2.

Treated in 7.3 is message security at the MAC level, i.e., message authentication and encryption, based on the Advanced Encryption Standard (AES) forward cipher function for 128-bit keys operating on counter mode and cipher block chaining (CBC) mode, respectively. As part of message security, replay protection is also provided in this subclause.

Support for mandatory and optional cipher functions is clarified in 7.4.

7.1 Security association and disassociation

The security association protocols specified in 7.1.2, 7.1.3, 7.1.4, and 7.1.5 shall be based on the Diffie-Hellman key exchange employing the elliptic curve public key cryptography. The elliptic curve, characterized as shown in Equation (21):

$$y^2 \equiv x^3 + ax + b \pmod{p}, \text{ with } a, b \in GF(p), 4a^3 + 27b^2 \neq 0 \quad (21)$$

where $GF(p)$ is a prime finite field, shall have the following values for its coefficients and domain parameters, as specified for Curve P-256 in FIPS Pub 186-3, with p (an odd prime), r (order of base point G), and a (a coefficient) given in decimal form, and coefficient b and base point $G = (G_x, G_y)$ given in hex:

$$p = 2^{256} - 2^{224} + 2^{192} + 2^{96} - 1$$

$$= 115792089210356248762697446949407573530086143415290314195533631308867097853951$$

$$r = 115792089210356248762697446949407573529996955224135760342422259061068512044369$$

$$a = p - 3$$

$$b = 5ac635d8 aa3a93e7 b3ebbd55 769886bc 651d06b0 cc53b0f63bce3c3e 27d2604b$$

$$G_x = 6b17d1f2 e12c4247 f8bce6e5 63a440f2 77037d81 2deb33a0f4a13945 d898c296$$

$$G_y = 4fe342e2 fe1a7f9b 8ee7eb4a 7c0f9e16 2bce3357 6b315ececbb64068 37bf51f5$$

The private keys (also called *secret keys*) SK_A and SK_B of the elliptic curve public key cryptography for the two communicating parties, a node (party A) and a hub (party B), respectively, shall be each statically unique 256-bit integers chosen independently and at random from the set of integers $\{1, 2, r-1\}$. The corresponding 256-bit public keys PK_A and PK_B shall be computed as shown in Equation (22):

$$PK_A = SK_A \times G, PK_B = SK_B \times G \quad (22)$$

where $\times$ denotes scalar multiplication of the base point $G = (G_x, G_y)$ by an integer as described in A.9.2 of IEEE Std 1363-2000. A received public key, denoted by a pair of X -coordinate and Y -coordinate values, shall be treated valid only if it is a non-infinity point on the elliptic curve defined in the above, i.e., that its X and Y coordinates satisfy the elliptic curve equation given above.

In the security association and disassociation procedures, except that for MK pre-shared association, that are specified in this subclause, the cipher-based message authentication code (CMAC) algorithm as specified in the NIST Special Publication 800-38B, with the AES forward cipher function under a 128-bit key as specified in FIPS Pub 197, is used to compute key message authentication codes (KMAC) and the desired shared MK. Specifically, the functional notation $CMAC(K, M, L)$ represents the L -bit output of the CMAC applied under key K to message M based on the AES forward cipher function.

Moreover, the bit string truncation functions $LMB_n(S)$ and $RMB_n(S)$ designate the n leftmost and the n rightmost bits of the bit string S , respectively. The sign $\parallel$ denotes concatenation of bit strings that are converted according to IEEE Std 1363-2000 from certain fields of the frames of concern.

7.1.1 Master key pre-shared association

A node and a hub shall each have a secret pre-shared MK prior to running the MK pre-shared association protocol to activate their pre-shared MK as their shared MK for their PTK creation, with the benefit of keeping third parties not possessing the secret MK from launching impersonation attacks via the PTK creation procedure.

The node, but not the hub, may initiate a security association procedure to run the MK pre-shared association protocol, by sending to the hub the first Security Association frame of the procedure.

Upon receiving the first Security Association frame, the hub shall send to the node the second Security Association frame, joining or aborting the security association procedure.

If the node receives the second Security Association frame indicating the hub is aborting the security association procedure, it shall abort the current security association procedure. It may initiate a new security association procedure if the hub aborted the procedure with a different security suite selector. It may later initiate a new security association procedure if the hub aborted the procedure due to temporary lack of resources.

Upon successfully sending the second Security Association frame indicating it is joining the security association procedure, the hub shall activate the pre-shared MK as its shared MK with the node, treating the node's true identity unauthenticated but the security association procedure completed. Upon receiving the second Security Association frame indicating the hub is joining the security association procedure, the node

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

shall also activate the pre-shared MK as its shared MK with the hub, treating the hub's true identity unauthenticated but the security association procedure completed as well. The node shall proceed to the PTK creation procedure to create a PTK with the hub, meanwhile performing mutual authentication of each other based on the claimed pre-shared MK.

The MK pre-shared association procedure is illustrated in Figure 102, where

Address_A is the Sender Address field of the frame payload of the first Security Association frame.

Address_B is the Recipient Address field of the frame payload of the first Security Association frame.

Security_Suite_Selector is the Security Suite Selector field of the frame payload of the first Security Association frame.

Association_Control is the Association Control field of the frame payload of the Security Association frame containing the field.




Figure 102 —MK pre-shared association procedure

7.1.2 Unauthenticated association

A node and a hub shall each require no authentication credentials such as a shared secret or human intervention prior to running the unauthenticated association protocol to generate their shared MK for their PTK creation, without the benefit of keeping third parties from launching impersonation attacks.

The node, but not the hub, may initiate a security association procedure to run the unauthenticated association protocol, by sending to the hub the first Security Association frame of the procedure.

Upon receiving the first Security Association frame, the hub shall send to the node the second Security Association frame, joining or aborting the security association procedure.

If the node receives the second Security Association frame indicating the hub is aborting the security association procedure, it shall abort the current security association procedure. It may initiate a new security association procedure if the hub aborted the procedure with a different security suite selector. It may later initiate a new security association procedure if the hub aborted the procedure due to temporary lack of resources. However, the node may resume the security association procedure if it subsequently receives the third Security Association frame with the MK_KMAC field set to MK_KMAC_3 as calculated via Equation (25), treating the earlier received second Security Association frame to have been sent by an impersonator of the hub.

The hub shall abort the security association procedure if the node's public key (PK_{AX} , PK_{AY}) contained in the first Security Association frame is not a valid public key. Otherwise, upon successfully sending the second Security Association frame indicating it is joining the security association procedure, the hub shall send to the node the third Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_3 as calculated via Equation (25).

The node shall abort the security association procedure if the hub's public key (PK_{BX} , PK_{BY}) contained in the second Security Association frame is not a valid public key. Otherwise, upon receiving the third

Security Association frame with the MK_KMAC equal to MK_KMAC_3 , the node shall send to the hub the fourth Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_4 as calculated via Equation (26).

Upon successfully sending the fourth Security Association frame, the node shall compute the shared MK as given in Equation (27), treating the hub's true identity unauthenticated but the association procedure completed. Upon receiving the fourth Security Association frame with the MK_KMAC equal to MK_KMAC_4 , the hub shall also compute the shared MK, treating the node's true identity unauthenticated but the association procedure completed as well.

The node and the hub shall each compute $DHKey$ and extract $Temp_1$ as shown in Equation (23) and Equation (24):

$$DHKey = X(SK_A \times PK_B) = X(SK_B \times PK_A) = X(SK_A \times SK_B \times G) \quad (23)$$

$$Temp_1 = RMB_128(DHKey) \quad (24)$$

The node and the hub shall each derive MK_KMAC_3 and MK_KMAC_4 as shown in Equation (25) and Equation (26):

$$MK_KMAC_3 = CMAC(Temp_1, Address_A || Address_B || Nonce_A || Nonce_B || Security_Suite_Selector || Association_Control, 64) \quad (25)$$

$$MK_KMAC_4 = CMAC(Temp_1, Address_B || Address_A || Nonce_B || Nonce_A || Security_Suite_Selector || Association_Control, 64) \quad (26)$$

After the aforementioned verifications have passed, the node and the hub shall each derive their shared MK as shown in Equation (27):

$$MK = CMAC(Temp_2, Nonce_A || Nonce_B, 128), \text{ where } Temp_2 = LMB_128(DHKey) \quad (27)$$

In the above, $X(P) = X(P_X, P_Y) = P_X = X$ -coordinate of P , which is computed from $SK_A \times PK_B$ at the node and from $SK_B \times PK_A$ at the hub, respectively.

SK_A is the node's 256-bit private key (an integer) kept secret by the node.

SK_B is the hub's 256-bit private key (an integer) kept secret by the hub.

PK_A is the node's 256-bit public key (a pair of X and Y coordinates) transmitted by the node.

PK_B is the hub's 256-bit public key (a pair of X and Y coordinates) transmitted by the hub.

$Address_A$ is the Sender Address field of the frame payload of the first or third Security Association frame.

$Address_B$ is the Recipient Address field of the frame payload of the first or third Security Association frame.

$Nonce_A$ is the Sender Nonce field of the frame payload of the first Security Association frame.

$Nonce_B$ is the Sender Nonce field of the frame payload of the second Security Association frame.

$Security_Suite_Selector$ is the Security Suite Selector field of the frame payload of the first Security Association frame.

$Association_Control$ is the Association Control field of the frame payload of the Security Association frame containing the field.

The unauthenticated association procedure is illustrated in Figure 103.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 103 —Unauthenticated association or public key hidden association procedure

7.1.3 Public key hidden association

A node and a hub shall have a secured, secret transfer of the node's public key to the hub, typically through an out-of-band channel, prior to running the public key hidden association protocol to generate their shared MK for their PTK creation, with the benefit of assisting in keeping third parties from launching impersonation attacks.

The node, but not the hub, may initiate a security association procedure to run the public key hidden association protocol, by sending to the hub the first Security Association frame of the procedure.

Upon receiving the first Security Association frame, the hub shall send to the node the second Security Association frame, joining or aborting the security association procedure.

If the node receives the second Security Association frame indicating the hub is aborting the security association procedure, it shall abort the current security association procedure. It may initiate a new security association procedure if the hub aborted the procedure with a different security suite selector. It may later initiate a new security association procedure if the hub aborted the procedure due to temporary lack of resources. However, the node may resume the security association procedure if it subsequently receives the third Security Association frame with the MK_KMAC field set to MK_KMAC_3 as calculated via Equation (30), treating the earlier received second Security Association frame to have been sent by an impersonator of the hub.

The hub shall abort the security association procedure if the node's public key (PK_{AX} , PK_{AY}) transferred through an out-of-band channel is not a valid public key. Otherwise, upon successfully sending the second Security Association frame indicating it is joining the security association procedure, the hub shall send to the node the third Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_3 as calculated via Equation (30).

The node shall abort the security association procedure if the hub's public key (PK_{BX} , PK_{BY}) contained in the second Security Association frame is not a valid public key. Otherwise, upon receiving the third Security Association frame with the MK_KMAC equal to MK_KMAC_3 , the node shall send to the hub the fourth Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_4 as calculated via Equation (31).

Upon successfully sending the fourth Security Association frame, the node shall compute the shared MK as given in Equation (32), treating the hub's true identity authenticated and the association procedure completed. Upon receiving the fourth Security Association frame with the MK_KMAC equal to MK_KMAC_4 , the hub shall also compute the shared MK, treating the node's true identity authenticated and the association procedure completed as well.

The node and the hub shall each compute $DHKey$ and extract $Temp_1$ as shown in Equation (28) and Equation (29):

$$DHKey = X(SK_A \times PK_B) = X(SK_B \times PK_A) = X(SK_A \times SK_B \times G) \quad (28)$$

$$Temp_1 = RMB_128(DHKey) \quad (29)$$

The node and the hub shall each derive MK_KMAC_3 and MK_KMAC_4 as shown in Equation (30) and Equation (31):

$$MK_KMAC_3 = CMAC(Temp_1, Address_A || Address_B || Nonce_A || Nonce_B || Security_Suite_Selector || Association_Control, 64) \quad (30)$$

$$MK_KMAC_4 = CMAC(Temp_1, Address_B || Address_A || Nonce_B || Nonce_A || Security_Suite_Selector || Association_Control, 64) \quad (31)$$

After the aforementioned verifications have passed, the node and the hub shall each derive their shared MK as shown in Equation (32):

$$MK = CMAC(Temp_2, Nonce_A || Nonce_B, 128), \text{ where } Temp_2 = LMB_128(DHKey) \quad (32)$$

In the above, $X(P) = X(P_X, P_Y) = P_X = X$ -coordinate of P , which is computed from $SK_A \times PK_B$ at the node and from $SK_B \times PK_A$ at the hub, respectively.

SK_A is the node's 256-bit private key (an integer) kept secret by the node.

SK_B is the hub's 256-bit private key (an integer) kept secret by the hub.

PK_A is the node's 256-bit public key (a pair of X and Y coordinates) transferred only to the hub by a secure out-of-band channel.

PK_B is the hub's 256-bit public key (a pair of X and Y coordinates) transmitted by the hub.

$Address_A$ is the Sender Address field of the frame payload of the first or third Security Association frame.

$Address_B$ is the Recipient Address field of the frame payload of the first or third Security Association frame.

$Nonce_A$ is the Sender Nonce field of the frame payload of the first Security Association frame.

$Nonce_B$ is the Sender Nonce field of the frame payload of the second Security Association frame.

$Security_Suite_Selector$ is the Security Suite Selector field of the frame payload of the first Security Association frame.

Association Control is the Association Control field of the frame payload of the Security Association frame containing the field.

The public key hidden association procedure is also illustrated in Figure 103 as well.

7.1.4 Password authenticated association

A node and a hub shall each have a secret shared password prior to running the password authenticated association protocol to generate their shared MK for their PTK creation, with the benefit of assisting in keeping third parties not possessing the secret password from launching impersonation attacks.

The node, but not the hub, may initiate a security association procedure to run the password authenticated association protocol, by sending to the hub the first Security Association frame of the procedure.

Upon receiving the first Security Association frame, the hub shall send to the node the second Security Association frame, joining or aborting the security association procedure.

If the node receives the second Security Association frame indicating the hub is aborting the security association procedure, it shall abort the current security association procedure. It may initiate a new security association procedure if the hub aborted the procedure with a different security suite selector. It may later initiate a new security association procedure if the hub aborted the procedure due to temporary lack of resources. However, the node may resume the security association procedure if it subsequently receives the third Security Association frame with the MK_KMAC field set to MK_KMAC_3 as calculated via Equation (38), treating the earlier received second Security Association frame to have been sent by an impersonator of the hub.

The hub shall abort the security association procedure if the node's password-scrambled public key (PK'_{AX} , PK'_{AY}) contained in the first Security Association frame is not a valid public key. Otherwise, upon successfully sending the second Security Association frame indicating it is joining the security association procedure, the hub shall send to the node the third Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_3 as calculated via Equation (38).

The node shall abort the security association procedure if the hub's public key (PK_{BX} , PK_{BY}) contained in the second Security Association frame is not a valid public key. Otherwise, upon receiving the third Security Association frame with the MK_KMAC equal to MK_KMAC_3, the node shall send to the hub the fourth Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_4 as calculated via Equation (39).

Upon successfully sending the fourth Security Association frame, the node shall compute the shared MK as given in Equation (40), treating the hub's true identity authenticated and the association procedure as completed. Upon receiving the fourth Security Association frame with the MK_KMAC equal to MK_KMAC_4, the hub shall also compute the shared MK, treating the node's true identity authenticated and the association procedure completed as well.

The node shall compute its password-scrambled public key $PK'_A = (PK'_{AX}, PK'_{AY})$ from its public or private key and the password shared with the hub as shown in Equation (33) and Equation (34):

$$PK'_A = PK_A - Q(PW) = SK_A \times G - Q(PW) \quad (33)$$

$$Q(PW) = (Q_X = 2^{32} \times PW + M_X, Q_Y = \text{even positive integer}) \quad (34)$$

The hub shall recover the node's public key from the received password-scrambled public key $PK'_A = (PK'_{AX}, PK'_{AY})$ for the subsequent DHKey computation as shown in Equation (35):

$$PK_A = PK'_A + Q(PW), Q(PW) = (Q_X = 2^{32} \times PW + M_X, Q_Y = \text{even positive integer}) \quad (35)$$

The parameters involved in these equations are defined as follows:

PW is a positive integer converted according to IEEE Std 1363-2000 from the UTE-16BE representation specified in ISO/IEC 10646:2003 of the shared password by treating the leftmost octet as the octet containing the MSBs.

M_X is the smallest nonnegative integer such that $Q_X = 2^{32} \times PW + M_X$ is the X -coordinate of a point on the elliptic curve defined earlier.

$Q(PW)$ is the point on the elliptic curve with X -coordinate = Q_X and Y -coordinate = Q_Y of an even positive integer.

The node shall choose a private key SK_A such that the X -coordinate of PK_A is not equal to the X -coordinate of $Q(PW)$.

The node and the hub shall each compute $DHKey$ and extract $Temp_1$ as shown in Equation (36) and Equation (37):

$$DHKey = X(SK_A \times PK_B) = X(SK_B \times PK_A) = X(SK_A \times SK_B \times G) \quad (36)$$

$$Temp_1 = RMB_128(DHKey) \quad (37)$$

The node and the hub shall each derive MK_KMAC_3 and MK_KMAC_4 as shown in Equation (38) and Equation (39):

$$MK_KMAC_3 = CMAC(Temp_1, Address_A \parallel Address_B \parallel Nonce_A \parallel Nonce_B \parallel Security_Suite_Selector \parallel Association_Control, 64) \quad (38)$$

$$MK_KMAC_4 = CMAC(Temp_1, Address_B \parallel Address_A \parallel Nonce_B \parallel Nonce_A \parallel Security_Suite_Selector \parallel Association_Control, 64) \quad (39)$$

After the aforementioned verifications have passed, the node and the hub shall each derive their shared MK as shown in Equation (40):

$$MK = CMAC(Temp_2, Nonce_A \parallel Nonce_B, 128), \text{ where } Temp_2 = LMB_128(DHKey) \quad (40)$$

In the above, $X(P) = X(P_X, P_Y) = P_X = X$ -coordinate of P , which is computed from $SK_B \times PK_A$ at the hub and from $SK_A \times PK_B$ at the node, respectively.

SK_A is the node's 256-bit private key (an integer) kept secret by the node.

SK_B is the hub's 256-bit private key (an integer) kept secret by the hub.

PK_A is the node's 256-bit public key (a pair of X and Y coordinates) kept secret by the node.

PK_B is the hub's 256-bit public key (a pair of X and Y coordinates) transmitted by the hub.

$Address_A$ is the Sender Address field of the frame payload of the first or third Security Association frame.

$Address_B$ is the Recipient Address field of the frame payload of the first or third Security Association frame.

$Nonce_A$ is the Sender Nonce field of the frame payload of the first Security Association frame.

$Nonce_B$ is the Sender Nonce field of the frame payload of the second Security Association frame.

$Security_Suite_Selector$ is the Security Suite Selector field of the frame payload of the first Security Association frame.

$Association_Control$ is the Association Control field of the frame payload of the Security Association frame containing the field.

The password authenticated association procedure is illustrated in Figure 104.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 104 — Password authenticated association procedure

7.1.5 Display authenticated association

A node and a hub shall each have a display of a 5-digit decimal number prior to running the display authenticated association protocol to generate their shared MK for their PTK creation, with the benefit of assisting in keeping third parties from launching man-in-the middle attacks.

The node, but not the hub, may initiate a security association procedure to run the display authenticated association protocol, by sending to the hub the first Security Association frame of the procedure.

Upon receiving the first Security Association frame, the hub shall send to the node the second Security Association frame of the procedure, joining or aborting the security association procedure.

If the node receives the second Security Association frame indicating the hub is aborting the security association procedure, it shall abort the current security association procedure. It may initiate a new security association procedure if the hub aborted the procedure with a different security suite selector. It may later initiate a new security association procedure if the hub aborted the procedure due to temporary lack of resources. However, the node may resume the security association procedure if it subsequently

receives the third Security Association frame with the MK_KMAC field set to MK_KMAC_3 as calculated via Equation (44), treating the earlier received second Security Association frame to have been sent by an impersonator of the hub.

The hub shall abort the security association procedure if the node's public key (PK_{AX}, PK_{AY}) contained in the first Security Association frame is not a valid public key. Otherwise, upon successfully sending the second Security Association frame indicating it is joining the security association procedure, the hub shall send to the node the third Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_3 as calculated via Equation (44).

The node shall abort the security association procedure if the hub's public key (PK_{BX}, PK_{BY}) contained in the second Security Association frame is not a valid public key. Otherwise, upon receiving the third Security Association frame with the MK_KMAC equal to MK_KMAC_3 , the node shall send to the hub the fourth Security Association frame, setting the MK_KMAC field of the Security Association Data thereof to MK_KMAC_4 as also calculated via Equation (45).

Upon successfully sending the fourth Security Association frame, the node shall display a 5-digit decimal number $Display_A$ as derived via Equation (47). Upon receiving the fourth Security Association frame with the MK_KMAC equal to MK_KMAC_4 , the hub shall display a 5-digit decimal number $Display_B$ as also derived via Equation (47). However, the hub shall display a 5-digit decimal number of zero if $Witness_A$ as contained in the received first Security Association frame is not equal to $Witness_A$ as derived via Equation (41) from the received first and fourth Security Association frames of the procedure.

After the node and the hub have been verified to display the same 5-digit number, they shall each be informed through their respective user interfaces that their mutual authentication has succeeded. Otherwise, they shall each be informed that their mutual authentication has failed.

Upon determining that their mutual authentication has succeeded, the node and the hub shall each compute the shared MK as given in Equation (48), treating their association procedure completed.

The node and the hub shall each derive $Witness_A$ before sending the first Security Association frame and after receiving the fourth Security Association frame, respectively, as shown in Equation (41):

$$Witness_A = CMAC(Nonce_A, Address_A || Address_B || PK_{AX} || PK_{AY}, 128) \quad (41)$$

The node and the hub shall each compute $DHKey$ and extract $Temp_1$ as shown in Equation (42) and Equation (43):

$$DHKey = X(SK_A \times PK_B) = X(SK_B \times PK_A) = X(SK_A \times SK_B \times G) \quad (42)$$

$$Temp_1 = RMB_{128}(DHKey) \quad (43)$$

The node and the hub shall each derive MK_KMAC_3 and MK_KMAC_4 as shown in Equation (44) and Equation (45):

$$MK_KMAC_3 = CMAC(Temp_1, Address_A || Address_B || Witness_A || Nonce_B || Security_Suite_Selector || Association_Control, 64) \quad (44)$$

$$MK_KMAC_4 = CMAC(Temp_1, Address_B || Address_A || Nonce_B || Witness_A || Security_Suite_Selector || Association_Control, 64) \quad (45)$$

The node and the hub shall also compute $Display_A$ and $Display_B$, respectively, as shown in Equation (46) and Equation (47):

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

$$D = \text{CMAC}(\text{Nonce_A} \parallel \text{Nonce_B}, \text{Nonce_B} \parallel \text{Nonce_A} \parallel \text{Temp_1}, 16) \quad (46)$$

$$\text{Display_A} = \text{BS2DI}(D), \text{Display_B} = \text{BS2DI}(D) \quad (47)$$

After the aforementioned verifications have passed, the node and the hub shall each derive their shared MK as shown in Equation (48):

$$\text{MK} = \text{CMAC}(\text{Temp_2}, \text{Nonce_A} \parallel \text{Nonce_B}, 128), \text{ where } \text{Temp_2} = \text{LMB_128}(\text{DHKey}) \quad (48)$$

In the above, $X(P) = X(P_X, P_Y) = P_X = X$ -coordinate of P , which is computed from $SK_A \times PK_B$ at the node and from $SK_B \times PK_A$ at the hub, respectively. $\text{BS2DI}(BS)$ converts the bit string BS to a positive decimal integer for display by treating the leftmost bit of the string as the MSB of the equivalent binary integer.

SK_A is the node's 256-bit private key (an integer) kept secret by the node.

SK_B is the hub's 256-bit private key (an integer) kept secret by the hub.

PK_A is the node's 256-bit public key (a pair of X and Y coordinates) transmitted by the node.

PK_B is the hub's 256-bit public key (a pair of X and Y coordinates) transmitted by the hub.

Address_A is the Sender Address field of the frame payload of the first or third Security Association frame.

Address_B is the Recipient Address field of the frame payload of the first or third Security Association frame.

Witness_A is the Sender Nonce field of the frame payload of the first Security Association frame.

Nonce_B is the Sender Nonce field of the frame payload of the second Security Association frame.

Nonce_A is the Sender Nonce field of the frame payload of the fourth Security Association frame.

$\text{Security_Suite_Selector}$ is the Security Suite Selector field of the frame payload of the first Security Association frame.

$\text{Association_Control}$ is the Association Control field of the frame payload of the Security Association frame containing the field.

The display authenticated association procedure is illustrated in Figure 105.

7.1.6 Security disassociation

Either the node or the hub may initiate a security disassociation procedure to nullify an existing security association and hence the shared MK and PTK with a hub or a node, by unilaterally sending a Security Disassociation frame, setting the DA_KMAC field of the frame payload depicted in Figure 22 to DA_KMAC as calculated via Equation (49).

Upon successfully sending the Security Disassociation frame, the sender shall erase the MK and the corresponding PTK materials from its internal storage. Upon receiving a Security Disassociation frame with the DA_KMAC equal to DA_KMAC , the recipient shall also erase the MK and the corresponding PTK materials from its internal storage.

The node and the hub shall compute DA_KMAC as shown in Equation (49):

$$\text{DA_KMAC} = \text{CMAC}(\text{MK}, \text{Address_B} \parallel \text{Address_A}, 64) \quad (49)$$

The input fields to the computation above are defined as follows:

MK is the shared MK to be repealed.

Address_B is the Recipient Address field of the frame payload of the Security Disassociation frame.

Address_A is the Sender Address field of the frame payload of the Security Disassociation frame.

The security disassociation procedure is illustrated in Figure 106.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 105 —Display authenticated association procedure

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 106 —Security disassociation procedure

7.2 PTK creation and GTK distribution

In the PTK creation procedure specified in this subclause, the CMAC algorithm as specified in the NIST Special Publication 800-38B, with the AES forward cipher function under a 128-bit key as specified in FIPS Pub 197, is used to compute KMACs. Specifically, the functional notation $CMAC(K, M, L)$ represents the L -bit output of the CMAC applied under key K to message M based on the AES forward cipher function.

Moreover, the bit string truncation functions $LMB_n(S)$ and $RMB_n(S)$ designate the n leftmost and the n rightmost bits of the bit string S , respectively. The sign $\parallel$ denotes concatenation of bit strings that are converted according to IEEE Std 1363-2000 from certain fields of the frames of concern.

7.2.1 PTK creation

A node and a hub shall have a 128-bit secret shared MK resulting from a successful run of a security association protocol as provided in 7.1 prior to running a PTK creation procedure to generate a PTK for exchanging secured frames with each other.

Either the node or the hub may initiate a PTK creation procedure, by sending to the hub or the node the first PTK frame of the procedure. Whichever sends the first PTK frame is referred to as the initiator, and the other is referred to as the responder.

Upon receiving the first PTK frame, the responder shall send to the initiator the second PTK frame, joining or aborting the PTK creation procedure, and setting the PTK_KMAC field of the frame payload depicted in Figure 23 to PTK_KMAC_2 as calculated via Equation (53) if joining.

If the initiator receives the second PTK frame indicating the responder is aborting the PTK creation procedure, it shall abort the current PTK creation procedure. It may later initiate a new PTK creation procedure if the responder aborted the procedure due to temporary lack of resources. However, the initiator may resume the PTK creation procedure if it subsequently receives another second PTK frame indicating the responder is joining the procedure and containing the PTK_KMAC field set to PTK_KMAC_2 calculated via Equation (53), treating the earlier received second PTK frame to have been sent by an impersonator of the responder.

Upon receiving the second PTK creation frame, the initiator shall send to the responder the third PTK frame, setting the PTK_KMAC field of the frame payload to PTK_KMAC_3 as also calculated via Equation (53). The initiator shall send this PTK frame only after it has verified that the PTK_KMAC contained in the second PTK frame is equal to PTK_KMAC_2 .

Upon successfully sending the third PTK frame, the initiator shall compute a new PTK as given in Equation (50), treating the responder's true identity authenticated and the PTK creation procedure completed. Upon receiving the third PTK frame with the PTK_KMAC equal to PTK_KMAC_3, the responder shall also compute the new PTK, treating the initiator's true identity authenticated and the PTK creation procedure completed as well.

After the respective appropriate verifications have passed, the initiator and the responder shall each derive the PTK, KCK, PTK_KMAC_2, and PTK_KMAC_3 as shown in Equation (50) through Equation (53):

$$PTK = CMAC(MK, Address_I \parallel Address_R \parallel Nonce_I \parallel Nonce_R \parallel PTK_Control, 128) \quad (50)$$

$$KCK = CMAC(MK, Address_R \parallel Address_I \parallel Nonce_R \parallel Nonce_I \parallel PTK_Control, 128) \quad (51)$$

$$P = CMAC(KCK, Address_I \parallel Address_R \parallel Nonce_R \parallel Nonce_I \parallel PTK_Control, 128) \quad (52)$$

$$PTK_KMAC_2 = LMB_64(P), PTK_KMAC_3 = RMB_64(P) \quad (53)$$

The fields that form the message of CMAC correspond to the fields in the PTK frames of the current PTK creation procedure and are converted to bit strings according to IEEE Std 1363-2000:

Address_I is the Sender Address field of the frame payload of the first PTK frame.

Address_R is the Recipient Address field of the frame payload of the first PTK frame.

Nonce_I is the Sender Nonce field of the frame payload of the first PTK frame.

Nonce_R is the Sender Nonce field of the frame payload of the second PTK frame.

PTK_Control is the PTK Control field of the frame payload of the second PTK frame.

The PTK creation procedure is illustrated in Figure 107.




Figure 107 —PTK creation procedure

7.2.2 GTK distribution

A node and a hub shall have a PTK prior to running a GTK distribution procedure to transport a GTK to the node for multicasting secured frames to the node and others.

The hub, but not the node, may send a GTK to the node, by sending to the node a GTK frame containing the GTK and secured with the PTK being used between the hub and the node.

7.3 Message security

A hub shall transmit a beacon, if appropriate, as an unsecured frame or as a secured frame authenticated, but not encrypted, by a GTK distributed to the nodes that are secured with it. Nodes that do not have a secured relationship with the hub may receive and process the beacon without validating the message integrity code (MIC) included therein.

At their *Secured* or *Connected* State, a node and a hub shall exchange and process only secured frames at a security level negotiated during their last association, with the following additional considerations:

- GTK frames shall always be secured, both authenticated and encrypted.
- Poll and T-Poll frames shall never be authenticated or encrypted.
- Control type frames other than Poll and T-Poll frames
 - 1) shall be neither authenticated nor encrypted if the hub and the node have agreed to applying no control type frame authentication;
 - 2) shall be authenticated but not encrypted if the hub and the node have agreed to applying control type frame authentication.

A recipient shall ignore a received frame with an unexpected security level, other than performing acknowledgment if needed. A recipient shall also ignore a received secured frame with an invalid MIC, i.e., the MIC value calculated from the received frame as described in 7.3.1.5 is not the same as the MIC field contained in the received frame, except again for returning an acknowledgment.

7.3.1 Frame authentication, encryption, and decryption

Frames shall be transmitted as secured or unsecured frames according to Figure 4 and Table 28, wherein SL and CFA stand for the Security Level field and the Control Frame Authentication field, respectively, of the Security Suite Selector field contained in the last Association frame exchanged between the sender and the recipient. A node or a hub shall ignore received frames secured or unsecured unexpectedly, except for returning an acknowledgment if required by the acknowledgment policy and for taking an appropriate defensive measure against potential security violations.

Table 28 —Conditions for transmitting unsecured and secured frames

Frame name	Frame (message) security
Beacon	Unsecured or authenticated with a distributed GTK, at sender's discretion
Security Association	Unsecured at Orphan state for secured communication; not sent otherwise
Security Disassociation	Unsecured at Associated state; authenticated but not encrypted at Secured or Connected state if SL = 1; authenticated and encrypted at Secured or Connected state if SL = 2; not sent otherwise
PTK	Unsecured at Associated state; authenticated but not encrypted at Secured or Connected state if SL = 1; authenticated and encrypted at Secured or Connected state if SL = 2; not sent otherwise
GTK	Authenticated but not encrypted at Secured or Connected state if SL = 1; authenticated and encrypted at Secured or Connected state if SL = 2; not sent otherwise
Connection Request, Connection Assignment	Unsecured at Orphan state for unsecured communication; authenticated but not encrypted at Secured or Connected state if SL = 1; authenticated and encrypted at Secured or Connected state if SL = 2; not sent otherwise
Disconnection	Unsecured at Connected state for unsecured communication; authenticated but not encrypted at Connected state if SL = 1; authenticated and encrypted at Connected state if SL = 2; not sent otherwise
Command exchanged within same BAN	Unsecured at Connected state for unsecured communication; authenticated but not encrypted at Connected state if SL = 1; authenticated and encrypted at Connected state if SL = 2; not sent otherwise
Command exchanged between BANs	Unsecured always
I-Ack, B-Ack, I-Ack+Poll, B-Ack+Poll	Unsecured at Orphan or Associated state; unsecured at Secured state if CFA = 0; unsecured at Connected state for unsecured communication or if CFA = 0 for secured communication; authenticated but not encrypted at Secured or Connected state if SL = 1 or 2 and CFA = 1; not sent otherwise
Poll and T-Poll addressed to Connected_NID	
Poll and T-Poll not addressed to Connected_NID	Unsecured always
Wakeup addressed to Unconnected_NID	Unsecured
Wakeup addressed to a Connected_NID	Unsecured at Connected state for unsecured communication; authenticated but not encrypted at Connected state if SL = 1 or 2 and CFA = 1; not sent otherwise
B2	Unsecured or authenticated with a distributed GTK, at sender's discretion
Any data type frame	Unsecured at Connected state for unsecured communication; authenticated but not encrypted at Connected state if SL = 1; authenticated and encrypted at Connected state if SL = 2; not sent otherwise

Secured frames shall be authenticated, and encrypted/decrypted when required, based on AES-128 CCM, i.e., the CCM mode as specified in the NIST Special Publication 800-38C, with the AES forward cipher function for 128-bit keys as specified in FIPS Pub 197 applied as the underlying block cipher algorithm.

Prior to exchanging secured unicast frames, the two communicating parties, a node and a hub, shall have a PTK for use as the AES key applied to these frames, sent from the node to the hub or vice versa. They may have an additional PTK before the current PTK is retired. Once one of them starts using a PTK, both shall no longer use any old PTK. Prior to multicast secured frames to a group, the hub shall have distributed a GTK to the nodes of the group for use as the AES key applied to the multicast frames. Once the hub starts using a GTK for the same group, it shall no longer use any old GTK for that group.

A TK, PTK or GTK, shall be retired no later than when both the Low-Order Security Sequence Number and High-Order Security Sequence Number fields of the last frame secured by the key have reached their respective maximum values supported by the fields. It may be retired earlier as needed.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

The length of what is referred to as the *message authentication code* (MAC) for message (frame) authentication in NIST Special Publication 800-38C but as the *message integrity code* (MIC) in this standard—to be distinguished from another accustomed standing of the term MAC for *medium access control*—shall be four octets. That is, in the NIST Special Publication 800-38C, $t = 4$. Also, $q = 2$ shall be chosen as the octet length of the binary representation of the octet length of the frame payload.

The bit order of each input block to the CCM invocation and AES encryption shall be formatted as illustrated in Figure 108. It is the concatenation of the bits of the ordered octets of the constituent fields of the block, where the octet order of each constituent field is defined in the remainder of this subclause, and the bits of each octet are ordered such that the MSB is the first bit of the octet while the LSB is the last bit of the octet. The first octet or the first bit of a given component is shown on the left, and the last octet or the last bit is shown on the right, in the context of the component. The bit notations $input_0, \dots, input_{127}$ correspond to those used for AES input block formation specified in FIPS Pub 197.




Figure 108 —Bit order for AES input blocks

7.3.1.1 Nonce formation

The Nonce as a required input field to each instance of CCM frame authentication and encryption/decryption is a 13-octet field that is formatted as shown in Figure 109.

Here, the octets of the MAC header and Low-Order Security Sequence Number fields are each ordered from left to right in accordance with their transmit order as defined in 5.1 and 5.2. The octets of the High-Order Security Sequence Number are similarly ordered, i.e., they are ordered with the octet containing the LSBs on the left and the octet containing the most-significant bits on the right.

The High-Order Security Sequence Number field is set as follows:

- It is set to zero if the current frame is secured with a PTK, and if the Low-Order Security Sequence Number field of the frames secured with this PTK, containing a valid MIC value, and transmitted by the sender of the current frame has never wrapped around.
- It is set to H or H + 1, if the current frame is secured with a GTK, and if the Security Sequence Number field of the current frame has a value larger or not larger, respectively, than L, where L and H are the values of the eight LSBs and the 24 most-significant bits, respectively, of the GTK SSN field of the last GTK frame through which this GTK was received.
- It is incremented by one each time the Low-Order Security Sequence Number field of the frames secured with the same PTK or GTK, containing a valid MIC value, and transmitted by the same sender wraps around, i.e., if the Security Sequence Number field of the current frame has a value not larger than the value of the same field of the last frame secured with the same PTK or GTK, containing a valid MIC value, and transmitted by the same sender. The Low-Order Security Sequence Number field of the last frame secured with the same PTK, containing a valid MIC value, and transmitted by the same sender is considered to have a zero value if no such last frame has been transmitted or received from the same sender.

An assumption is made that a recipient is to receive at least one of the last 2^N frames (including retransmitted frames) secured with the same PTK or GTK and transmitted by the same sender, where $N = 8$ is the number of bits of the Low-Order Security Sequence Number field of secured frames.




Figure 109 —Nonce format

7.3.1.2 Initial block B_0 construction

The block B_0 as the first input block to the CBC for frame authentication, i.e., MIC computation, is a 16-octet field that is formatted as shown in Figure 110. Here, Flags = 0x09 if the frame payload is encrypted or 0x49 if the frame payload is not encrypted. The $Q = L_{FP}$ field is set to the length of the frame payload in octets as defined in Figure 12 and its octets are ordered with the octet containing the most-significant bits on the left and the octet containing the LSBs on the right.

Only this block is present if the current frame does not have a frame payload.




Figure 110 —Initial block B_0 format

7.3.1.3 Payload blocks $B_1, \dots, B_m$ construction

The blocks $B_1, \dots, B_m$ as the subsequent input blocks to the CBC frame authentication, and also as the input blocks to the counter mode encryption/decryption, i.e., cipher text computation and plain text recovery, if

the frame payload is to be encrypted/decrypted, are each a 16-octet field that is formatted as shown in Figure 111.

The octets of the frame payload are ordered from left to right in accordance with their transmit order as defined in 5.1 and 5.2.

These blocks are constructed from the unencrypted or decrypted frame payload. The last block contains one or more padded zero octets on the right end if the frame payload is not an integral multiple of 16 octets.

None of these blocks is present if the current frame does not have a frame payload.




Figure 111 —Payload blocks $B_1, \dots, B_m$ format

7.3.1.4 Counter blocks $Ctr_0, \dots, Ctr_m$ formation

The block Ctr_0 as the input block to the counter mode encryption of the CBC output for MIC computation, and each of the blocks $Ctr_1, \dots, Ctr_m$ as the input blocks to the counter mode encryption/decryption if the frame payload is to be encrypted/decrypted, is a 16-octet field that is formatted as shown in Figure 112.

Here, the octets of the Counter i field, with $i = 0, \dots, m$, respectively, are ordered with the octet containing the MSBs on the left and the octet containing the LSBs on the right.




Figure 112 —Counter blocks $Ctr_0, \dots, Ctr_m$ format

7.3.1.5 MIC computation

The MIC field in an authenticated frame is calculated as shown in Figure 113, where, as shown in Equation (54) and Equation (55):

$$MIC = LMB_n(M), M = AES(Ctr_0) \oplus X_m \quad (54)$$

$$X_0 = AES(B_0), X_i = AES(B_i \oplus X_{i-1}), i = 1, \dots, m \quad (55)$$

Here, $LMB_n(M)$ designates the n leftmost bits of the bit string M , the symbol $\oplus$ denotes bitwise exclusive-OR, and $AES(B)$ represents the output of the forward cipher function of the AES block cipher algorithm applied to block B under the AES key PTK or GTK used to secure the frame. The MIC is ordered for transmission from its first octet on the left to its last octet on the right, as also illustrated in Figure 113.

The octet notations $out_0, \dots, out_{15}$ correspond to those used for AES output block formation specified in FIPS Pub 197.

The blocks required for the MIC computation are constructed from the unencrypted version of the frame to be transmitted at the sender side, and from the decrypted version of the received frame at the recipient side if the frame is encrypted.




Figure 113 —MIC calculation and transmit order

7.3.1.6 Frame payload encryption

The encrypted frame payload in an encrypted frame is formatted as shown Figure 114, where, as shown in Equation (56) and Equation (57):

$$B'_i = B_i \oplus \text{AES}(Ctr_i), i = 1, \dots, m-1 \quad (56)$$

$$B'_m = L_n(B_m) \oplus L_n(\text{AES}(Ctr_m)) \quad (57)$$

Here, the symbol $\oplus$ denotes bitwise exclusive-OR, and $L_n(B)$ designates the n leftmost octets of B . Moreover, $\text{AES}(Ctr_i)$ represents the output of the forward cipher function of the AES block cipher algorithm applied to the counter block Ctr_i under the AES key PTK or GTK used to secure the frame. The encrypted frame payload has the same length as the unencrypted frame payload, so that $n \leq 16$ is the number of octets in B_m excluding the zero padding octets if any.

Each encrypted block is ordered for transmission from its first octet on the left to its last octet on the right, as also illustrated in Figure 114. The octet notations $out_0, \dots, out_{15}$ correspond to those used for AES output block formation specified in FIPS Pub 197.




Figure 114 —Encrypted Frame Payload format for encrypted frames

7.3.1.7 Frame payload decryption

The frame payload in an encrypted frame is decrypted as shown in Figure 115, where, as shown in Equation (58) and Equation (59):

$$B_i = B'_i \oplus \text{AES}(\text{Ctr}_i), i = 1, \dots, m - 1 \quad (58)$$

$$B_m = B'_m \oplus L_n(\text{AES}(\text{Ctr}_m)) \quad (59)$$

The decrypted frame payload has the same length as the encrypted frame payload, so that $n \leq 16$ is the number of octets in the last block B'_m of the encrypted frame payload received. The last decrypted block B_m is padded with $16 - n$ zero octets at the right end to form the last block B_m as shown in Figure 110 for MIC calculation over the received frame as described in 7.3.1.5.

Each decrypted block is ordered from its first octet on the left to its last octet on the right, as also illustrated in Figure 115. Again the octet notations $out_0, \dots, out_{15}$ correspond to those used for AES output block formation specified in FIPS Pub 197.




Figure 115 —Decrypted Frame Payload format for encrypted frames

7.3.1.8 MIC validation

Each block as decrypted in 7.3.1.7 is used for MIC calculation as specified in 7.3.1.5. If the calculated MIC value is equal to the received MIC value, the received frame contains a valid MIC value and frame authentication passes. Otherwise, the received frame contains an invalid MIC frame and frame authentication fails.

The recipient shall discard all received frames containing an invalid MIC value. Appropriate steps should be taken with respect to receipt of such frames.

7.3.2 Replay protection

A recipient shall discard any received frames that could or would otherwise result in its High-Order Security Sequence Number wrapping around to zero. In particular, if a recipient has the High-Order Security Sequence Number equal to $2^{24} - 1$ for the last frame secured with a PTK or GTK, containing a valid MIC value, and transmitted by a sender, it shall discard any received frame secured with the same PTK or GTK, containing a valid MIC value, and transmitted by the same sender, if the Low-Order Security Sequence Number field of the received frame has a value not larger than its value found in that last frame. The recipient shall not apply the discarded frame to update either the Low-Order Security Sequence Number or the High-Order Security Sequence Number pertaining to the last frame it received from the same sender and secured with the same PTK or GTK.

7.4 Optional cipher functions

For secured communications, all nodes and hubs shall support the forward cipher function of AES. For secured broadcast and multicast transmissions, the sender shall use the forward cipher function of AES.

If a node and a hub both support a non-AES cipher function as listed in Table 7, in their secured unicast communications with each other, they may use it in place of the AES forward cipher function as specified in 7.1, 7.2, and 7.3.

8. Narrowband PHY specification

This clause specifies an optional narrowband (NB) physical layer (PHY). The PHY is responsible for the following tasks:

- Activation and deactivation of the radio transceiver
- Clear channel assessment (CCA) within the current channel
- Data transmission and reception

This clause also provides a method for transforming a physical-layer service data unit (PSDU) into a physical-layer protocol data unit (PPDU). During the transmission, the PSDU shall be pre-appended with a physical-layer preamble and a physical-layer header in order to create the PPDU. At the receiver, the physical-layer preamble and physical-layer header serve as aids in the demodulation, decoding and delivery of the PSDU.

Figure 116 shows the format for the PPDU, which is composed of three main components: the physical-layer convergence protocol (PLCP) preamble, the PLCP header, and the PSDU. The components are listed in the order of transmission. The PLCP preamble is the first component of the PPDU (see 8.2). The purpose of the preamble is to aid the receiver during timing synchronization and carrier-offset recovery.

The PLCP header is the second main component of the PPDU (see 8.3). The purpose of this component is to convey the necessary information about the PHY parameters to aid in the decoding of the PSDU at the receiver. The PLCP header can be further decomposed into a RATE field, a LENGTH field, a BURST MODE field, a SCRAMBLER SEED field, reserved bits, a header check sequence (HCS), and BCH parity bits. The BCH parity bits are added in order to improve the robustness of the PLCP header. The PLCP header shall be transmitted using the given header data rate in the operating frequency band.

The PSDU is the last component of the PPDU (see 8.4). This component is formed by concatenating the MAC header with the MAC frame body and frame check sequence (FCS). The PSDU may then be encoded and spread/interleaved before being scrambled. The PSDU shall be transmitted using one of the data rates available in the operating frequency band.

When transmitting the packet, the PLCP preamble is sent first, followed by the PLCP header and finally the PSDU. All multiple octet fields shall be transmitted with least significant octet first and each octet shall be transmitted with LSB first.

A compliant device shall be able to support transmission and reception in at least one of the following frequency bands: 402 MHz to 405 MHz, 420 MHz to 450 MHz, 863 MHz to 870 MHz, 902 MHz to 928 MHz, 950 MHz to 958 MHz, 2360 MHz to 2400 MHz, and 2400 MHz to 2483.5 MHz.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks




Figure 116—Standard PPDU structure

8.1 Data-rate-dependent parameters

The data-rate-dependent parameters for each of the possible frequency bands of operation are provided in subsequent subclauses.

8.1.1 402 MHz to 405 MHz

The modulation parameters for this band are defined in Table 29.

Table 29—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	Pulse shape	Information data rate (kbps)	Support
PLCP header	$\pi/2$ -DBPSK ($M = 2$)	187.5	19/31 ^a	2	SRRC	57.5	Mandatory
PSDU	$\pi/2$ -DBPSK ($M = 2$)	187.5	51/63	2	SRRC	75.9	Mandatory
PSDU	$\pi/2$ -DBPSK ($M = 2$)	187.5	51/63	1	SRRC	151.8	Mandatory
PSDU	$\pi/4$ -DQPSK ($M = 4$)	187.5	51/63	1	SRRC	303.6	Mandatory
PSDU	$\pi/8$ -D8PSK ($M = 8$)	187.5	51/63	1	SRRC	455.4	Optional

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.1.2 420 MHz to 450 MHz

The modulation parameters for this band are defined in Table 30.

Table 30—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	BT	Information data rate (kbps)	Support
PLCP header	GMSK (M = 2)	187.5	19/31 ^a	2	0.5	57.5	Mandatory
PSDU	GMSK (M = 2)	187.5	51/63	2	0.5	75.9	Mandatory
PSDU	GMSK (M = 2)	187.5	51/63	1	0.5	151.8	Mandatory
PSDU	GMSK (M = 2)	187.5	1/1	1	0.5	187.5	Optional

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.1.3 863 MHz to 870 MHz

The modulation parameters for this band are defined in Table 31.

Table 31—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	Pulse shape	Information data rate (kbps)	Support
PLCP header	$\pi/2$ -DBPSK (M = 2)	250	19/31 ^a	2	SRRC	76.6	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	250	51/63	2	SRRC	101.2	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	250	51/63	1	SRRC	202.4	Mandatory
PSDU	$\pi/4$ -DQPSK (M = 4)	250	51/63	1	SRRC	404.8	Mandatory
PSDU	$\pi/8$ -D8PSK (M = 8)	250	51/63	1	SRRC	607.1	Optional

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.1.4 902 MHz to 928 MHz

The modulation parameters for this band are defined in Table 32.

Table 32—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	Pulse shape	Information data rate (kbps)	Support
PLCP header	$\pi/2$ -DBPSK (M = 2)	250	19/31 ^a	2	SRRC	76.6	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	250	51/63	2	SRRC	101.2	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	250	51/63	1	SRRC	202.4	Mandatory
PSDU	$\pi/4$ -DQPSK (M = 4)	250	51/63	1	SRRC	404.8	Mandatory
PSDU	$\pi/8$ -D8PSK (M = 8)	250	51/63	1	SRRC	607.1	Optional

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.1.5 950 MHz to 958 MHz

The modulation parameters for this band are defined in Table 33.

Table 33—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	Pulse shape	Information data rate (kbps)	Support
PLCP header	$\pi/2$ -DBPSK (M = 2)	250	19/31 ^a	2	SRRC	76.6	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	250	51/63	2	SRRC	101.2	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	250	51/63	1	SRRC	202.4	Mandatory
PSDU	$\pi/4$ -DQPSK (M = 4)	250	51/63	1	SRRC	404.8	Mandatory
PSDU	$\pi/8$ -D8PSK (M = 8)	250	51/63	1	SRRC	607.1	Optional

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.1.6 2360 MHz to 2400 MHz

The modulation parameters for this band are defined in Table 34.

Table 34—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	Pulse shape	Information data rate (kbps)	Support
PLCP header	$\pi/2$ -DBPSK (M = 2)	600	19/31 ^a	4	SRRC	91.9	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	600	51/63	4	SRRC	121.4	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	600	51/63	2	SRRC	242.9	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	600	51/63	1	SRRC	485.7	Mandatory
PSDU	$\pi/4$ -DQPSK (M = 4)	600	51/63	1	SRRC	971.4	Mandatory

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.1.7 2400 MHz to 2483.5 MHz

The modulation parameters for this band are defined in Table 35.

Table 35—Modulation parameters for PLCP header and PSDU

Packet component	Modulation (M)	Symbol rate = $1/T_s$ (ksps)	Code rate (k/n)	Spreading factor (S)	Pulse shape	Information data rate (kbps)	Support
PLCP header	$\pi/2$ -DBPSK (M = 2)	600	19/31 ^a	4	SRRC	91.9	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	600	51/63	4	SRRC	121.4	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	600	51/63	2	SRRC	242.9	Mandatory
PSDU	$\pi/2$ -DBPSK (M = 2)	600	51/63	1	SRRC	485.7	Mandatory
PSDU	$\pi/4$ -DQPSK (M = 4)	600	51/63	1	SRRC	971.4	Mandatory

^a BCH (31, 19) code is a shortened code derived from a BCH (63, 51) code.

8.2 PLCP preamble

A preamble shall be added prior to the PLCP header in order to aid the receiver in packet detection, timing synchronization and carrier-offset recovery. Two unique preambles are defined in order to mitigate false alarms due to networks operating on adjacent channels. The mapping between channel number and preamble is defined in 8.6.3. Each preamble is constructed by concatenating a length-63 m-sequence with a 01010101010101010101010101010101 extension sequence. The length of the preamble, $N_{preamble}$, is therefore 90 bits. The former sequence can be used to implement packet detection, coarse-timing synchronization, and carrier-offset recovery, while the latter sequence can be used to implement fine-timing synchronization.

The two preamble sequences are defined in Table 36 and Table 37. The preambles shall be transmitted at the symbol rate for the desired band of operation and shall be encoded using the same modulation parameters as defined for the PLCP header in the preceding tables.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 36—Preamble sequence #1

Bit	Bit value	Bit	Bit value	Bit	Bit value	Bit	Bit value
b ₀	0	b ₂₃	0	b ₄₆	1	b ₆₉	0
b ₁	1	b ₂₄	0	b ₄₇	0	b ₇₀	1
b ₂	0	b ₂₅	1	b ₄₈	0	b ₇₁	0
b ₃	1	b ₂₆	1	b ₄₉	0	b ₇₂	1
b ₄	0	b ₂₇	1	b ₅₀	0	b ₇₃	0
b ₅	1	b ₂₈	0	b ₅₁	1	b ₇₄	1
b ₆	1	b ₂₉	0	b ₅₂	0	b ₇₅	1
b ₇	0	b ₃₀	0	b ₅₃	0	b ₇₆	0
b ₈	0	b ₃₁	1	b ₅₄	0	b ₇₇	1
b ₉	1	b ₃₂	0	b ₅₅	0	b ₇₈	1
b ₁₀	1	b ₃₃	1	b ₅₆	0	b ₇₉	0
b ₁₁	0	b ₃₄	1	b ₅₇	1	b ₈₀	1
b ₁₂	1	b ₃₅	1	b ₅₈	1	b ₈₁	1
b ₁₃	1	b ₃₆	1	b ₅₉	1	b ₈₂	0
b ₁₄	1	b ₃₇	0	b ₆₀	1	b ₈₃	1
b ₁₅	0	b ₃₈	0	b ₆₁	1	b ₈₄	1
b ₁₆	1	b ₃₉	1	b ₆₂	1	b ₈₅	0
b ₁₇	1	b ₄₀	0	b ₆₃	0	b ₈₆	1
b ₁₈	0	b ₄₁	1	b ₆₄	1	b ₈₇	1
b ₁₉	1	b ₄₂	0	b ₆₅	0	b ₈₈	0
b ₂₀	0	b ₄₃	0	b ₆₆	1	b ₈₉	1
b ₂₁	0	b ₄₄	0	b ₆₇	0	—	—
b ₂₂	1	b ₄₅	1	b ₆₈	1	—	—

Table 37—Preamble sequence #2

Bit	Bit value	Bit	Bit value	Bit	Bit value	Bit	Bit value
b ₀	0	b ₂₃	1	b ₄₆	1	b ₆₉	0
b ₁	1	b ₂₄	0	b ₄₇	1	b ₇₀	1
b ₂	1	b ₂₅	0	b ₄₈	0	b ₇₁	0
b ₃	0	b ₂₆	1	b ₄₉	0	b ₇₂	1
b ₄	1	b ₂₇	0	b ₅₀	0	b ₇₃	0
b ₅	0	b ₂₈	0	b ₅₁	1	b ₇₄	1
b ₆	0	b ₂₉	1	b ₅₂	1	b ₇₅	1
b ₇	0	b ₃₀	1	b ₅₃	1	b ₇₆	0
b ₈	1	b ₃₁	1	b ₅₄	0	b ₇₇	1
b ₉	0	b ₃₂	1	b ₅₅	1	b ₇₈	1
b ₁₀	0	b ₃₃	0	b ₅₆	0	b ₇₉	0
b ₁₁	0	b ₃₄	0	b ₅₇	1	b ₈₀	1
b ₁₂	0	b ₃₅	0	b ₅₈	1	b ₈₁	1
b ₁₃	1	b ₃₆	0	b ₅₉	1	b ₈₂	0
b ₁₄	0	b ₃₇	0	b ₆₀	1	b ₈₃	1
b ₁₅	1	b ₃₈	1	b ₆₁	1	b ₈₄	1
b ₁₆	1	b ₃₉	1	b ₆₂	1	b ₈₅	0
b ₁₇	0	b ₄₀	0	b ₆₃	0	b ₈₆	1
b ₁₈	0	b ₄₁	1	b ₆₄	1	b ₈₇	1
b ₁₉	1	b ₄₂	1	b ₆₅	0	b ₈₈	0
b ₂₀	0	b ₄₃	1	b ₆₆	1	b ₈₉	1
b ₂₁	1	b ₄₄	0	b ₆₇	0	—	—
b ₂₂	0	b ₄₅	0	b ₆₈	1	—	—

8.3 PLCP header

A PLCP header shall be added after the PLCP preamble to convey information about the PHY parameters that is needed at the receiver in order to decode the PSDU. The length of the PLCP header, N_{header} , is 31 bits, and it shall be constructed for transmission as shown in Figure 117:

- Based on the information provided by the MAC, form the PHY header according to 8.3.1.
- Calculate the 4-bit HCS value over the PHY header using the CRC-4 ITU polynomial: $1 + x + x^4$, according to 8.3.2.
- As shown in Figure 118 and according to 8.3.3, a BCH (31, 19) code, which is a shortened code derived from a BCH (63, 51) code, is applied to the concatenation of the PHY header (15 bits) and HCS (4 bits).
- The encoded bits are spread using a repetition code according to 8.4.3, where $N_{total} = 31$, and then interleaved using a bit interleaver as defined in 8.4.4. The spreading factor is determined by the frequency band of operation (see 8.1).
- The resulting bit stream is then scrambled according to 8.4.5, where the seed of the scrambler is determined by the channel number, i.e., even channels are mapped to scrambler seed zero and odd channels are mapped to scrambler seed one.
- Finally, the resulting scrambled bit stream is then mapped onto the appropriate constellation (see 8.5), which is determined by the frequency band of operation (see 8.1).




Figure 117 —Block diagram of PLCP header construction for transmission




Figure 118 —BCH encoding scheme for PLCP header construction

8.3.1 PHY header

The PHY header contains information about the data rate of the MAC frame body, the length of the MAC frame body (which does not include the MAC header or the FCS) and information about the next packet—whether it is being sent in a burst mode.

The PHY header field shall be composed of 15 bits, numbered from 0 to 14 as illustrated in Figure 119. Bits 0–2 shall encode the RATE field, which conveys the information about the type of modulation, the information data rate, the pulse shaping, the coding rate, and the spreading factor used to transmit the PSDU. Bits 4–11 shall encode the LENGTH field, with the LSB being transmitted first. Bit 13 shall encode whether or not the packet is being transmitted in the burst (streaming) mode. Bit 14 shall encode the scrambler seed. All other bits that are not defined in this clause shall be understood to be reserved for future use and shall be set to zero.




Figure 119 —PHY header bit assignment

8.3.1.1 Data Rate field (RATE)

Depending on the data rate (RATE), bits R0–R2 shall be set according to the values in Table 38.

Table 38 —Rate dependent parameters

	402 to 405 MHz	420 to 450 MHz	863 to 870 MHz, 902 to 928 MHz, 950 to 958 MHz	2360 to 2400 MHz, 2400 to 2483.5 MHz
R0–R2	Data rate (kbps)	Data rate (kbps)	Data rate (kbps)	Data rate (kbps)
000	75.9	75.9	101.2	121.4
100	151.8	151.8	202.4	242.9
010	303.6	187.5	404.8	485.7
110	455.4	Reserved	607.1	971.4
001	Reserved	Reserved	Reserved	Reserved
101	Reserved	Reserved	Reserved	Reserved
011	Reserved	Reserved	Reserved	Reserved
111	Reserved	Reserved	Reserved	Reserved

8.3.1.2 PLCP Length field (LENGTH)

The PLCP Length field shall be an unsigned 8-bit integer that indicates the number of uncoded information octets in the MAC frame body (which does not include the MAC header or the FCS).

8.3.1.3 Burst Mode (BM) field

The MAC shall set the BM bit, as defined in Table 39, to indicate whether the next packet is part of a packet “burst,” i.e., burst mode transmission. In burst mode, the interframe spacing is defined in 8.7.5.

Table 39—Burst Mode field


Burst Mode bit	Next packet status
0	Next packet is <i>not</i> part of burst
1	Next packet is part of burst

8.3.1.4 Scrambler Seed (SS) field

The MAC shall set the SS bit according to the SS identifier value defined in Table 40. This bit value corresponds to the seed value chosen for the data scrambler.

8.3.2 Header Check Sequence

The PHY header shall be protected with a 4-bit (CRC-4 ITU) header check sequence (HCS). The HCS shall be the ones complement of the remainder generated by the modulo-2 division of the PHY header by the polynomial: $1 + x + x^4$. The HCS bits shall be processed in the transmit order. An example schematic of the processing order is shown in Figure 120. The registers shall be initialized to all ones.

**Figure 120—Block diagram of a CRC-4 implementation**

8.3.3 BCH encoder for PLCP header

The PLCP header shall use a systematic BCH (31, 19, $t = 2$) code, which is a shortened code derived from a BCH (63, 51, $t = 2$) code by appending 32 zero (or shortened) bits to the 19 information bits, to improve the robustness of the PLCP header. A description of the BCH (63, 51, $t = 2$) code can be found in 8.4.1.2. The shortened bits are removed prior to transmission.

8.4 PSDU

The PSDU is the last major component of the PPDU and shall be constructed as shown in Figure 121.

- a) Form the PSDU by pre-pending the 7-octet MAC header to the MAC frame body and appending a 2-octet FCS to the result.
- b) If the code rate (k/n) < 1, the PSDU is:
 - 1) Divided into blocks of messages starting with the LSB of the least significant octet of the PSDU and continuing to the MSB of the most significant octet of the PSDU;
 - 2) Shortening bits may then be appended to the messages, which are then encoded into codewords using a BCH (63, 51) encoder to achieve the desired code rate, according to 8.4.1.1;
 - 3) Finally, the shortened bits are removed from each of the codewords.
- c) Pad bits are then added in order to align on a symbol boundary according to 8.4.2.
- d) If the spreading factor is 2 or 4, the resulting uncoded or coded bits are spread using a repetition code, according to 8.4.3, and then interleaved using a bit interleaver defined in 8.4.4.
- e) The resulting bit stream is then scrambled according to 8.4.5.
- f) Finally, the resulting scrambled bit stream is then mapped onto the appropriate constellation (see 8.5), which is determined by the data rate and frequency band of operation (see 8.1).




Figure 121 —Block diagram of PSDU construction for transmission

8.4.1 BCH encoder for PSDU

A code rate of 51/63 shall be supported by a systematic BCH encoder. The information bits shall be encoded using the BCH encoding process defined in 8.4.1.1. The definition for the systematic BCH encoder is given in 8.4.1.2.

8.4.1.1 BCH encoding process

The PSDU shall be encoded using the following procedure:

- a) Compute the number of bits in the PSDU N_{PSDU} as shown in Equation (60):

$$N_{PSDU} = (N_{MACHeader} + N_{MACFrameBody} + N_{FCS}) \times 8 \quad (60)$$

where $N_{MACHeader}$ is the number of octets in the MAC header, $N_{MACFrameBody}$ is the number of octets in the MAC frame body, and N_{FCS} is the number of octets in the FCS.

- b) Calculate the number of BCH codewords N_{CW} as shown in Equation (61):

$$N_{CW} = \left\lceil \frac{N_{PSDU}}{k} \right\rceil \quad (61)$$

where k is the number of message bits for the selected BCH code.

- c) Compute the number of shortening bits, $N_{shorten}$, to be padded to the N_{PSDU} data bits before encoding as shown in Equation (62):

$$N_{shorten} = N_{CW} \times k - N_{PSDU} \quad (62)$$

- d) The shortening bits shall be equally distributed over all N_{CW} codewords with the first $\text{rem}(N_{shorten}, N_{CW})$ codewords being shortened one bit more than the remaining codewords. Let, as shown in Equation (63):

$$N_{spcw} = \left\lfloor \frac{N_{shorten}}{N_{CW}} \right\rfloor \quad (63)$$

where $\lfloor f \rfloor$ is the largest integer not greater than f .

Thus, the first $\text{rem}(N_{shorten}, N_{CW})$ codewords will have $N_{spcw} + 1$ shortened bits (message bits that are set to zero), while the remaining codewords will have N_{spcw} shortened bits. After encoding, the shortened bits shall be discarded prior to transmission, i.e., the shortened bits are never transmitted on-air.

The BCH encoding process is illustrated in Figure 122 for a single codeword.




Figure 122 —BCH encoding process for a single codeword

8.4.1.2 BCH (63, 51) encoder

The generator polynomial for a systematic BCH (63, 51, $t = 2$) code, where t is the number of bit errors that can be corrected, is given by Equation (64):

$$g(x) = 1 + x^3 + x^4 + x^5 + x^8 + x^{10} + x^{12} \quad (64)$$

The parity bits are determined by computing the remainder polynomial $r(x)$ as shown in Equation (65):

$$r(x) = \sum_{i=0}^{11} r_i x^i = x^{12} m(x) \bmod g(x) \quad (65)$$

where $m(x)$ is the message polynomial shown in Equation (66):

$$m(x) = \sum_{i=0}^{50} m_i x^i \quad (66)$$

and $r_i, i = 0, \dots, 11$ and $m_i, i = 0, \dots, 50$ are elements of $GF(2)$.

The message polynomial $m(x)$ is created as follows: m_{50} is the first bit of the message to be transmitted and m_0 is the last bit of the message, which may be a shortened bit. The order of the parity bits is as follows: r_{11} is the first parity bit transmitted, r_{10} is the second parity bit transmitted, and r_0 is the last parity bit transmitted.

8.4.2 Pad bits

Pad bits shall be appended after the BCH encoder to align the bit stream on a symbol boundary. The number of pad bits, N_{pad} , that shall be inserted is a function of the number of PSDU bits N_{PSDU} , the number of codewords N_{CW} , the number of parity bits $(n-k)$, and the modulation constellation size M , see Equation (67):

$$N_{pad} = \log_2(M) \times \left\lceil \frac{N_{PSDU} + N_{CW} \times (n-k)}{\log_2(M)} \right\rceil - [N_{PSDU} + N_{CW} \times (n-k)]. \quad (67)$$

The pad bits shall be appended to the uncoded or coded PSDU and all of the appended pad bits shall be set to zero. In the case of uncoded transmission, N_{CW} is set to zero.

8.4.3 Spreading

For a spreading factor of 2, each input bit is repeated two times [see illustration in Figure 123(a)]. For a spreading factor of 4, each input bit is repeated four times [see illustration in Figure 123(b)].




Figure 123 —Spreading scheme

8.4.4 Bit interleaver

In the case that the spreading factor is equal to 2 or 4, the output of the spreader shall be interleaved prior to modulation to provide robustness against error propagation. The exact structure of the bit interleaver depends on the number of uncoded or coded bits that will be transmitted on-air, which is given by Equation (68):

$$N_{total} = N_{PSDU} + N_{CW} \times (n - k) + N_{pad} \quad (68)$$

where N_{CW} is set to zero in the case of uncoded transmission.

If $\text{rem}(N_{total}, 2) = 0$, the bit interleaving operation is performed by first grouping the spread bits into blocks of $2S$ bits, where S is the spreading factor, and then using a block interleaver of size $S \times 2$ to permute the bits. Let the sequences $a(i)$ and $b(i)$, where $i = 0, 1, \dots, 2S-1$, represent the input and output bits of the $S \times 2$ bit interleaver, respectively. The output of the $S \times 2$ bit interleaver is given by the relationship in Equation (69):

$$b(i) = a \left[S \times \text{rem}(i, 2) + \left\lfloor \frac{i}{2} \right\rfloor \right] \quad i = 0, 1, \dots, 2S-1 \quad (69)$$

If $\text{rem}(N_{total}, 2) = 1$, the bit interleaving operation is performed by grouping the first $3S$ spread bits into a single block and then using a block interleaver of size $S \times 3$ to permute the bits within that single block. Let the sequences $a(i)$ and $b(i)$, where $i = 0, 1, \dots, 3S-1$, represent the input and output bits of the $S \times 3$ bit interleaver, respectively. The output of the $S \times 3$ bit interleaver is given by the relationship in Equation (70):

$$b(i) = a \left[S \times \text{rem}(i, 3) + \left\lfloor \frac{i}{3} \right\rfloor \right] \quad i = 0, 1, \dots, 3S-1 \quad (70)$$

The remaining spread bits are then grouped into blocks of $2S$ bits and interleaved using the block interleaver of size $S \times 2$ shown in Equation (69).

8.4.5 Data scrambler

A side-stream scrambler with polynomial $G(x) = 1 + x^2 + x^{12} + x^{13} + x^{14}$ shall be used to whiten the PSDU. Figure 124 shows a typical implementation of the side-stream scrambler. The output of the scrambler is generated as shown in Equation (71):

$$x[n] = x[n-2] \oplus x[n-12] \oplus x[n-13] \oplus x[n-14] \quad (71)$$

where $\oplus$ denotes modulo-2 addition. For example, when the scrambler seed is set to zero, the first 20 bits out of the scrambler are: 0 0 0 1 1 1 0 0 0 0 1 1 1 1 0 1 1 1 0. Table 40 defines the initialization vector, x_{init} , for the side-stream scrambler as a function of the scrambler seed value.




Figure 124 —Block diagram of a side-stream scrambler

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 40—Scrambler seed selection

Scrambler seed (SS)	Initialization vector
	$x_{init} = x[-1] \ x[-2] \ \dots \ x[-14]$
0	0 0 1 0 1 1 1 1 0 0 1 1 0 1
1	0 0 0 0 0 0 0 1 0 0 1 1 1 1

The MAC shall set the scrambler seed to 0 when the PHY is initialized and the scrambler seed shall be incremented, using a 1-bit rollover counter, for *each* frame sent by the PHY.

At the receiver, the side-stream de-scrambler shall be initialized with the same initialization vector, x_{init} , used by the transmitter. The initialization vector is determined from the SS value in the PHY header of the received frame.

8.5 Constellation mapping

The constellation mapper operates on the binary bit stream $b(n)$, which is the concatenation of the PLCP preamble, the PLCP header, and the PSDU.

8.5.1 Gaussian minimum shift keying (GMSK)

For the GMSK constellation, $b(n)$, $n = 0, 1, \dots, N-1$ shall be mapped onto a corresponding frequency deviation Δf , which shall be the product of the symbol rate and a modulation index of 0.5 divided by 2. The relationship between the bit stream $b(n)$ and the frequency deviation is given in Table 41.

Table 41—GMSK symbol mapping

$b(n)$	Frequency deviation
0	$-\Delta f$
1	$+\Delta f$

8.5.1.1 Gaussian filter pulse shape

The Gaussian pulse shape with bandwidth-time product BT , described in Equation (72), shall be used to filter the symbols and shape the spectrum.

$$h(t) = \frac{\exp\left(\frac{-t^2}{2\delta^2}\right)}{\sqrt{2\pi} \cdot \delta} \quad (72)$$

where

$$\delta = \frac{\sqrt{\ln(2)}}{2\pi BT}$$

The value of BT is defined in Table 30.

8.5.2 Differential phase-shift keying (D-PSK)

For the D-PSK constellations, $b(n), n = 0, 1, \dots, N-1$ shall be mapped onto one of three rotated and differentially encoded constellations: $\pi/2$ -DBPSK, $\pi/4$ -DQPSK, or $\pi/8$ -D8PSK. The encoded information is carried in the phase transitions between symbols. For the PLCP preamble to PLCP header transition, the phase change is relative to the last symbol for the PLCP preamble. For the PLCP header to PSDU transition, the phase change is relative to the last symbol for the PLCP header.

The binary bit stream $b(n), n = 0, 1, \dots, N-1$ shall be mapped onto a corresponding complex-values sequence $S(k), k = 0, 1, \dots, (N/\log_2(\))-1$ as shown in Equation (73):

$$S(k) = S(k-1) \exp(j \varphi_k) \quad k = 0, 1, \dots, (N/\log_2(\))-1 \quad (73)$$

where $S(-1) = \exp(j \pi/2)$ is the reference for the first symbol of the preamble and the relationship between the bit stream $b(n)$ and the phase change φ_k is given in Table 42, Table 43, or Table 44 for $\pi/2$ -DBPSK, $\pi/4$ -DQPSK, or $\pi/8$ -D8PSK, respectively.

Table 42 — $\pi/2$ -DBPSK mapping

$b(n)$	φ_k
0	$\pi/2$
1	$3\pi/2$

Table 43 — $\pi/4$ -DQPSK mapping

$b(2n)$	$b(2n+1)$	φ_k
0	0	$\pi/4$
0	1	$3\pi/4$
1	0	$7\pi/4$
1	1	$5\pi/4$

Table 44 — $\pi/8$ -D8PSK mapping

$b(3n)$	$b(3n+1)$	$b(3n+2)$	φ_k
0	0	0	$\pi/8$
0	0	1	$3\pi/8$
0	1	0	$7\pi/8$
0	1	1	$5\pi/8$
1	0	0	$15\pi/8$
1	0	1	$13\pi/8$
1	1	0	$9\pi/8$
1	1	1	$11\pi/8$

8.5.2.1 SRRC pulse shape

For the D-PSK constellations, the square-root raised cosine (SRRC) pulse shape with roll-off factor β and symbol period T , described in Equation (74), shall be used to filter the symbols and shape the spectrum.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

$$p(t) = \begin{cases} 1 - \beta + 4 \frac{\beta}{\pi} & t = 0 \\ \frac{\beta}{\sqrt{2}} \left[\left(1 + \frac{2}{\pi} \right) \sin \left(\frac{\pi}{4\beta} \right) + \left(1 - \frac{2}{\pi} \cos \left(\frac{\pi}{4\beta} \right) \right) \right] & t = \pm \frac{T_s}{4\beta} \\ \sin \left[\pi \frac{t}{T_s} (1 - \beta) \right] + 4\beta \frac{t}{T_s} \cos \left[\pi \frac{t}{T_s} (1 + \beta) \right] & \text{otherwise} \\ \pi \frac{t}{T_s} \left[1 - \left(4\beta \frac{t}{T_s} \right)^2 \right] & \text{otherwise} \end{cases} \quad (74)$$

The exact value for the roll-off factor β and the duration of the SRRC pulse shape is implementation dependent.

8.6 General requirements

8.6.1 Operating frequency bands

A compliant device shall be able to support transmissions and reception in one or more of the following frequency bands:

- e) 402 MHz to 405 MHz
- f) 420 MHz to 450 MHz
- g) 863 MHz to 870 MHz
- h) 902 MHz to 928 MHz
- i) 950 MHz to 958 MHz
- j) 2360 MHz to 2400 MHz
- k) 2400 MHz to 2483.5 MHz

8.6.2 Channel numbering

The relationship between center frequency, f_c , and channel number, n_c , is shown in Table 45.

Table 45—Relationship between center frequency and channel number

Frequency band (MHz)	Relationship between f_c and n_c	Number of channels (N_{ch})
402 to 405	$f_c = 402.15 + 0.30 \times n_c$ (MHz), $n_c = 0, \dots, 9$	10
420 to 450	$f_c = 420.30 + 0.50 \times g_1(n_c)$ (MHz), $n_c = 0, \dots, 11$	12
863 to 870	$f_c = 863.20 + 0.40 \times g_2(n_c)$ (MHz), $n_c = 0, \dots, 13$	14
902 to 928	$f_c = 903.20 + 0.40 \times n_c$ (MHz), $n_c = 0, \dots, 59$	60
950 to 958	$f_c = 951.10 + 0.40 \times n_c$ (MHz), $n_c = 0, \dots, 15$	16
2360 to 2400	$f_c = 2361.00 + 1.00 \times n_c$ (MHz), $n_c = 0, \dots, 38$	39
2400 to 2483.5	$f_c = 2402.00 + 1.00 \times n_c$ (MHz), $n_c = 0, \dots, 78$	79

The mapping functions $g_1(n_c)$ and $g_2(n_c)$ used in the 420 MHz to 450 MHz and 863 MHz to 870 MHz frequency bands, respectively, are defined as shown in Equation (75) and Equation (76):

$$g_1(n_c) = \begin{cases} n_c & 0 \leq n_c \leq 1 \\ n_c + 6.875 & 2 \leq n_c \leq 4 \\ n_c + 13.4 & n_c = 5 \\ n_c + 35.025 & 6 \leq n_c \leq 7 \\ n_c + 40.925 & 8 \leq n_c \leq 9 \\ n_c + 47.25 & 10 \leq n_c \leq 11 \end{cases} \quad (75)$$

and

$$g_2(n_c) = \begin{cases} n_c & 0 \leq n_c \leq 7 \\ n_c + 0.5 & n_c = 8 \\ n_c + 1 & 9 \leq n_c \leq 12 \\ n_c + 1.5 & n_c = 13 \end{cases} \quad (76)$$

8.6.3 Preamble sequence assignment

The relationship between channel number, n_c , and preamble sequence used in the PLCP preamble is shown in Table 46.

Table 46—Relationship between channel number and preamble

Relationship between preamble sequence and n_c	Preamble sequence
$\text{rem}(n_c, 2) = 0$	1
$\text{rem}(n_c, 2) = 1$	2

8.7 PHY layer timing

The values for the PHY layer timing parameters are defined in Table 47.

Table 47—Physical layer timing parameters

PHY parameter	Value
pSIFS	75 μs
pMIFS	20 μs
pExtraIFS	10 μs
pEDTime	8 preamble symbols
pCCATime	63 preamble symbols
pChannelSwitchTime	100 μs

The values for pEDTime and pCCATime shall be those specified in Table 47 or the values specified by the local regulatory requirements, whichever is lower.

8.7.1 Packet duration

The total duration (in time) of a packet, which comprises the symbols for the PLCP preamble, PLCP header, and PSDU, is given by Equation (77):

$$t_{\text{packet}} = T_s \times \left[N_{\text{preamble}} + N_{\text{header}} \times S_{\text{header}} + \frac{N_{\text{total}}}{\log_2(M)} \times S_{\text{PSDU}} \right] \quad (77)$$

where T_s , S_{header} , S_{PSDU} , and M are defined in Table 29 through Table 35, where N_{preamble} is defined in 8.2, where N_{header} is defined in 8.3, and where N_{total} is defined in 8.4.4. S_{header} refers to the value of S for the PLCP header and S_{PSDU} refers to the value of S for the PSDU data rate.

8.7.2 Start and end of a frame

The start of a frame shall be the time when the first output sample from the transmitter pulse shaping filter that is affected by the first symbol of the PLCP preamble is present on the local air interface. The end of a frame shall be the time when the last output sample from the transmitter pulse shaping filter that is affected by the last symbol of the frame is present on the local air interface.

8.7.3 Receive-to-transmit (RX-to-TX) turnaround time

The RX-to-TX turnaround time for a device shall be between pSIFS and pSIFS + pExtraIFS. The turnaround time is defined as the time elapsed from the end of the received frame at the local air interface to the start of the transmitted frame at the local air interface, where the start and end of the frame are defined in 8.7.2.

8.7.4 Transmit-to-receive turnaround time

The TX-to-RX turnaround time for a device shall not be greater than pSIFS. The turnaround time is defined as the time elapsed from the end of the transmitted frame at the local air interface until the time when the receiver is ready to begin the reception of the start of the next PHY frame, where the start and end of the frame are defined in 8.7.2.

8.7.5 Time between successive transmissions

For burst mode transmissions, the interframe spacing between uninterrupted successive transmissions by a device shall be between pMIFS and pMIFS + pExtraIFS. The interframe spacing is defined as the time elapsed from the end of a frame at the local air interface, to the start of a frame at the local air interface, where the start and end of the frame are defined in 8.7.2.

8.7.6 Center frequency switch time

The center frequency switch time shall not exceed pChannelSwitchTime. The center frequency switch time is defined as the interval from when the PHY transmits or receives the end of a frame on one center frequency until it is ready to transmit or receive the start of a frame on a different center frequency, where the start and end of the frame are defined in 8.7.2.

8.8 Transmitter specifications

8.8.1 Transmit power spectral density (PSD) mask

The transmitted spectral mask shall be less than $-X$ dBr (dB relative to the maximum spectral density of the signal) for $|f - f_c| \geq f_{BW}/2$, where f_c is the channel center frequency and f_{BW} is the channel bandwidth and is a function of the frequency band of operation as defined in Table 48 and illustrated in Figure 125.

Table 48—Channel bandwidth as a function of the frequency band of operation

Frequency (MHz)	$-X$ dBr	f_{BW}
402 to 405	–20	300 kHz
420 to 450	–20	320 kHz
863 to 870	–20	400 kHz
902 to 928	–20	400 kHz
950 to 958	–20	400 kHz
2360 to 2400	–20	1 MHz
2400 to 2483.5	–20	1 MHz




Figure 125—Transmit spectral mask for all frequency bands

The transmitted spectral density also should comply with all regulations defined by local regulatory bodies.

8.8.2 Transmit power

When operating in a low power low duty cycle (LP/LDC) mode, as defined in applicable regulations and standards including subclause 8.3 of ETSI EN 301 839-1, on a center frequency of 403.65 MHz (channel 6), a transmitter shall be capable of transmitting at most -40 dBm *effective isotropic radiated power* (EIRP). When operating in a non-LP/LDC mode in the 402 MHz to 405 MHz frequency band, a transmitter shall be capable of transmitting at most -16 dBm EIRP. When operating in all other frequency bands, a transmitter shall be capable of transmitting at least -10 dBm EIRP.

Devices should transmit lower power when possible in order to reduce interference to other devices and systems and to protect the safety of the human body. The maximum transmit power is limited by local regulatory bodies.

8.8.3 Transmit power-on and power-down ramp

The transmit power-on ramp for 10% to 90% of maximum power shall be no greater than 5 symbols. The transmit power-on ramp is shown in Figure 126. The transmit power-down ramp for 90% to 10% maximum power shall be no greater than 5 symbols. The transmit power-down ramp is shown in Figure 127.




Figure 126 —Transmit power-on ramp




Figure 127 —Transmit power-down ramp

8.8.4 Transmit center frequency tolerance

The transmitted center frequency tolerance shall be ± 20 ppm maximum.

8.8.5 Symbol clock frequency tolerance

The symbol clock frequency tolerance shall be ± 20 ppm maximum.

8.8.6 Clock synchronization

The transmit center frequencies and the symbol clock frequency shall be derived from the same reference oscillator.

8.8.7 Transmitter modulation accuracy**8.8.7.1 D-PSK constellation error**

The modulation accuracy for the D-PSK modulation schemes is determined via an error-vector magnitude (EVM) measurement, which is calculated over N baud-spaced received complex values after differential demodulation $(\hat{I}_k, \hat{Q}_k)$. A decision is then made for each of the received complex values as to the closest ideal position, which is represented by the vector (I_k, Q_k) . The ideal constellation points associated with the various D-PSK modulation schemes after differential demodulation are given in Table 49. The error vector $(\delta I_k, \delta Q_k)$ is defined as the distance from the ideal position to the actual position of the received complex values, i.e., $(\hat{I}_k, \hat{Q}_k) = (I_k, Q_k) + (\delta I_k, \delta Q_k)$.

Table 49—Ideal constellation points for D-PSK after differential demodulation

Constellation	Ideal constellation positions
$\pi/2$ -DBPSK	(0, 1), (0, -1)
$\pi/4$ -DQPSK	(cos($\pi/4$), sin($\pi/4$)), (cos($3\pi/4$), sin($3\pi/4$)), (cos($5\pi/4$), sin($5\pi/4$)), (cos($7\pi/4$), sin($7\pi/4$))
$\pi/8$ -D8PSK	(cos($\pi/8$), sin($\pi/8$)), (cos($3\pi/8$), sin($3\pi/8$)), (cos($5\pi/8$), sin($5\pi/8$)), (cos($7\pi/8$), sin($7\pi/8$)), (cos($9\pi/8$), sin($9\pi/8$)), (cos($11\pi/8$), sin($11\pi/8$)), (cos($13\pi/8$), sin($13\pi/8$)), (cos($15\pi/8$), sin($15\pi/8$))

The EVM is defined in Equation (78):

$$EVM_{dB} = 10 \log_{10} \left(\frac{1}{N} \sum_{k=1}^N (\delta I_k^2 + \delta Q_k^2) \right) \quad (78)$$

A transmitter shall have EVM values less than or equal to those listed in Table 50 when measured for $N = 1000$ symbols. The EVM shall be measured on the baseband I and Q samples at the output of the differential demodulator of a reference receiver. The reference receiver shall perform the following operations: carrier-frequency offset removal, SRRC filtering matched to the transmitter under test, and symbol timing recovery while making the measurements. Due to the noise enhancement of the differential demodulator, the EVM value measured at the demodulator's output will be approximately 3 dB worse when compared to a similar measurement taken at the differential demodulator input.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 50—Permissible EVM numbers as a function of constellation size

Constellation	EVM error (EVM _{dB})
$\pi/2$ -DBPSK	−11 dB
$\pi/4$ -DQPSK	−15 dB
$\pi/8$ -D8PSK	−20 dB

8.8.7.2 GMSK modulation error

The modulation accuracy for the GMSK modulation scheme is determined by measuring the frequency deviation tolerance and the zero crossing error of the eye diagram, as shown in Figure 128. Frequency deviation tolerance is measured as a percentage of the maximum frequency deviation, Δf , which is defined in 8.5.1. The frequency deviation tolerance at $\pm T_s/2$ shall be greater than or equal to 80% and less than or equal to 120%. The zero crossing error is the time difference between the ideal symbol period, T_s , and the measured crossing time. This shall be less than $\pm 1/8$ of T_s . Both measurements shall be performed and satisfied for a sequence of 1000 symbols.




Figure 128—GMSK modulation accuracy

8.8.8 Adjacent channel power ratio

The adjacent channel power ratio (ACPR) is defined as the ratio of the total power in the adjacent channel to the total power in the wanted channel, where the measurement bandwidth in both cases is equal to the channel bandwidth as given in Table 48. The ACPR shall be measured at an output transmit power equal to the maximum possible power output of the device. A compliant device shall have an ACPR that is no greater than values given in Table 51.

Table 51—ACPR as a function of the frequency band of operation

Frequency (MHz)	ACPR node (dB)	ACPR hub (dB)
402 to 405	−26	−32
420 to 450	−26	−26
863 to 870	−26	−26
902 to 928	−26	−26
950 to 958	−26	−26
2360 to 2400	−26	−26
2400 to 2483.5	−26	−26

8.9 Receiver specifications

8.9.1 Receiver sensitivity

For a packet error rate (PER) of less than or equal to 10% with a PSDU of 255 octets in additive white Gaussian noise (AWGN), a compliant device shall achieve receiver sensitivities listed in Table 52, or better. The minimum input levels are measured at the antenna connector, where a noise figure of 13 dB, which may include losses due to external components, and an implementation loss of 6 dB have been assumed.

Table 52—Receiver sensitivity numbers

Frequency band (MHz)	Information data rate (kbps)	Maximum input level at sensitivity (dBm)
402 to 405	75.9	−95
	151.8	−92
	303.6	−89
	455.4	−83
420 to 450	75.9	−90
	151.8	−87
	187.5	−84
	101.2	−94
863 to 870 902 to 928 950 to 958	202.4	−91
	404.8	−87
	607.1	−82
	121.4	−92
2360 to 2400 2400 to 2483.5	242.9	−90
	485.7	−87
	971.4	−83
	121.4	−92

8.9.2 Adjacent channel rejection

The adjacent channel rejection (ACR) is defined as the ratio of the power of the interfering signal in the adjacent channel to the power of the wanted signal, when the desired signal's strength is set to 3 dB above the rate-dependent sensitivity, and the power of the interfering signal has been raised until a 10% PER is reached for a PSDU length of 255 octets. The interfering signal in the adjacent channel shall be a conformant PHY signal at the same information data rate, unsynchronized with the signal in the channel under test, and generated using a suitable test source. A compliant device shall have an ACR that is no less than the values given in Table 53.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 53—ACR as a function of the frequency band of operation

Frequency band (MHz)	Information data rate (kbps)	ACR (dB)
402 to 405	75.9	17
	151.8	14
	303.6	10
	455.4	5
420 to 450	75.9	12
	151.8	9
	187.5	6
863 to 870 902 to 928 950 to 958	101.2	17
	202.4	14
	404.8	10
	607.1	5
2360 to 2400 2400 to 2483.5	121.4	17
	242.9	15
	485.7	13
	971.4	9

8.9.3 Receiver energy detection

The receiver energy detection (ED) measurement is an estimate of the received signal power within the bandwidth of the channel. It is intended for use by upper layers for various tasks, including as part of a channel selection algorithm. No attempt is made to identify or decode signals on the channel.

8.9.3.1 ED threshold

For the frequency band 402 MHz to 405 MHz, the minimum ED value (zero) shall indicate received power less than that which is prescribed by ETSI EN 301 839-1.

For all other frequency bands, the minimum ED value (zero) shall indicate received power less than either

- 10 dB above the receiver sensitivity as defined in Table 52 for the lowest data rate within a given band (see 8.9.1) or
- that which is prescribed by local regulatory requirements, or applicable standards,

whichever is lower.

The range of received power spanned by the ED values shall be at least 40 dB. Within this range, the mapping from the received power in decibels to ED value shall be linear with an accuracy of ± 6 dB.

8.9.3.2 ED measurement time

For the frequency band 402 MHz to 405 MHz, the ED measurement time, to average over, shall be that which is prescribed by ETSI EN 301 839-1.

For all other frequency bands, the ED measurement time, to average over, shall be either

- pEDTime (8 preamble symbol periods) or
- that which is prescribed by local regulatory requirements, or applicable standards,

whichever is longer in duration.

8.9.4 Receiver clear channel assessment

The PHY shall provide the capability to perform CCA according to at least one of the following three methods:

- a) *CCA Mode 1: Energy above threshold.* CCA shall report a busy medium upon detecting any energy above the ED threshold.
- b) *CCA Mode 2: Carrier sense only.* CCA shall report a busy medium only upon the detection of a signal compliant with this standard with the same modulation and characteristics of the PHY that is currently in use by the device. This signal may be above or below the ED threshold. The CCA detection time shall be equal to pCCATime.
- c) *CCA Mode 3: Carrier sense with energy above threshold.* CCA shall report a busy medium using a logical combination of the following:
 - 1) Detection of a signal with the modulation and characteristics of the PHY that is currently in use by the device, and
 - 2) Energy above the ED threshold, where the logical operator may be AND or OR.

The CCA parameters are subject to the following criteria:

- The ED threshold shall correspond to a received signal power as prescribed in 8.9.3.1
- The CCA detection time shall be equal to pCCATime (see Table 47). Any CCA and LBT procedures required by local regulatory requirements shall also be supported.

9. Ultra wideband PHY specification

The ultra wideband (UWB) PHY specification is designed to offer robust performance for BANs and to provide a large scope for implementation opportunities for high performance, robustness, low complexity, and ultra low power operation. Moreover, the interest of UWB lies in the fact that the signal power levels are in the order of those used in the MICS band, therefore providing safe power levels for the human body and low interference to other devices.

The UWB PHY provides a data interface to the MAC layer under the control of the physical layer convergence protocol (PLCP).

The UWB PHY provides three levels of functionality, as follows:

- a) Activation and deactivation of the radio transceivers.
- b) The PLCP constructs the PHY layer protocol data unit (PPDU) by concatenating the synchronization header (SHR), physical layer header (PHR) and physical layer service data unit (PSDU), respectively. Moreover, the PPDU bits are converted into RF signals for transmission in the wireless medium.
- c) The UWB PHY may provide clear channel assessment (CCA) indication to the MAC in order to verify activity in the wireless medium.

9.1 Definition of hubs and devices

There are two different types of UWB technologies included in the specification. Namely, impulse radio UWB (IR-UWB) and wideband frequency modulation (FM-UWB). For coexistence, the interaction of both UWB technologies is defined as follows:

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

In a BAN, the hub shall implement either an IR-UWB transceiver only or shall implement IR-UWB and FM-UWB transceivers in the same hub.

In a BAN, devices shall implement an IR-UWB transceiver or an FM-UWB transceiver or both.

9.2 Modes of operation

There are two modes of operation: default mode and high quality of service (QoS) mode. The default mode shall be used in medical and non-medical applications. The high QoS mode shall be used for high-priority medical applications.

9.2.1 High QoS mode

The high QoS mode shall be defined as user priority 6 in Table 18 of 6.2.3.

9.3 Rules for use of modes and options

For interoperability, a mandatory procedure is required for the default mode and high QoS mode. Therefore, a compliant UWB PHY shall support the following:

9.3.1 Default mode

The default mode shall support IR-UWB as mandatory PHY and FM-UWB as optional PHY according to 9.1.

9.3.1.1 IR-UWB PHY

- One mandatory PPDU (see 9.5).
- One mandatory data rate: 0.4875 Mbps (see 9.11.1).
- One mandatory modulation: on-off signaling (see 9.10.1).
- One mandatory channel in the low band and one mandatory channel in the high band (see 9.12). Implementers shall use at least one mandatory channel.
- One mandatory transmit spectral mask (see 9.13).

9.3.1.2 FM-UWB PHY

- One mandatory PPDU (see 9.5).
- One mandatory data rate: 250 kbps (see 9.16.1).
- Two mandatory modulations: CP-BFSK and wideband FM (see 9.16.3).
- One mandatory channel in the high band (see 9.12).
- One mandatory transmit spectral mask (see 9.13).

9.3.2 High QoS mode

The high QoS mode shall support IR-UWB as mandatory PHY.

- One mandatory PPDU (see 9.5).
- One mandatory data rate: 0.4875 Mbps (see 9.11.1).
- One mandatory modulation: DPSK modulation (see 9.10.2).
- One mandatory channel in the low band and one mandatory channel in the high band (see 9.12). Implementers shall use at least one mandatory channel.
- One mandatory transmit spectral mask (see 9.13).
- One mandatory HARQ (see 9.15).

9.4 Pulse shape option

There is not a mandatory pulse shape for IR-UWB. However, implementers can choose a pulse shape from a pool of pulse shapes (see 9.14). Furthermore, there are two types of pulse waveforms supported and defined as follows:

- 1) Single pulse option shall be defined as a single pulse transmitted per symbol (see 9.9.1).
- 2) Burst pulse option shall be defined as a concatenation of pulses transmitted per symbol (see 9.9.1).

9.5 UWB PHY frame format

The UWB PHY frame format or physical layer protocol data unit (PPDU) is formed by concatenating the synchronization header (SHR), the physical layer header (PHR), and the physical layer service data unit (PSDU), respectively, as illustrated in Figure 129.




Figure 129 —UWB PPDU structure

9.6 PSDU construction

The PSDU contains the MAC protocol data unit (MPDU) plus channel code BCH parity bits in the default mode. In case of high QoS mode operation, the PSDU contains either the MPDU or BCH parity bits.

The MPDU shall be defined as the concatenation of the MAC header, MAC frame body, and FCS as illustrated in Figure 130.




Figure 130 —MPDU structure

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

The bits of the PSDU are formatted for transmission. The PSDU construction process is illustrated in Figure 131.




Figure 131 —PSDU construction for transmission

9.6.1 Scrambler

A scrambler shall be applied in order to eliminate possible long strings of 1s or 0s contained in the MPDU and so eliminating the dependency of the signal's power spectrum upon the actual data.

An additive or synchronous scrambler with generator polynomial $x[n]$ given in Equation (79) shall be employed. Figure 132 shows a typical implementation of the additive or synchronous side-stream scrambler. The output of the scrambler is generated as:

$$x[n] = x[n-2] \oplus x[n-12] \oplus x[n-13] \oplus x[n-14] \quad (79)$$

where $\oplus$ denotes modulo-2 addition.

Table 54 defines the initialization vector for the additive scrambler as a function of the SS value.




Figure 132 —Block diagram of an additive or synchronous scrambler

Table 54 —Scrambler seed selection

Scrambler seed (SS)	Initialization vector $x_{init} = x[-1] \ x[-2] \ \dots \ x[-14]$
0	0 0 1 0 1 1 1 1 0 0 1 1 0 1
1	0 0 0 0 0 0 0 1 0 0 1 1 1 1

The MAC shall set the scrambler seed to SS = 0 in the PHR (see 9.7.1), when the UWB PHY is initialized. The scrambler seed shall be incremented using a 1-bit rollover counter for each frame sent by the UWB PHY.

At the receiver, the additive de-scrambler shall be initialized with the same initialization vector, x_{int} , used by the transmitter. The initialization vector is determined from the SS value in the PHY header (PHR) of the received frame.

9.6.2 BCH encoder

The channel code BCH($n = 63, k = 51$) shall be used for the default mode. In case of high QoS mode operation, the shortened channel code BCH(126,63) shall be used (see 9.6.2.2).

The number of codewords in a frame is given by the following:

$$N_{CW} = \left\lceil \frac{N'_{PSDU}}{k} \right\rceil$$

where $N'_{PSDU} = 8(N_{MACheader} + N_{MACframebody} + N_{FCS})$ is the number of bits in the PSDU, $N_{MACheader}$ is the number of octets in the MAC header, $N_{MACframebody}$ is the number of octets in the MAC frame body, and N_{FCS} is the number of octets of the FCS.

If the $rem(N'_{PSDU}, k) \neq 0$, the last codeword requires $N_{bs} = N_{CW}k - N'_{PSDU}$ bits stuffing. Hence, the total number of bits before encoding is given by $N_{PSDU} = N'_{PSDU} + N_{bs}$.

9.6.2.1 BCH(63,51) encoder

The generator polynomial for a BCH ($n = 63, k = 51$) encoder is given by Equation (80):

$$g(x) = 1 + x^3 + x^4 + x^5 + x^8 + x^{10} + x^{12} \quad (80)$$

The parity bits are determined by computing the remainder polynomial $r(x)$ as shown in Equation (81):

$$r(x) = \sum_{i=0}^{11} r_i x^i = x^{12} m(x) \bmod g(x) \quad (81)$$

where $m(x)$ is the message polynomial given by Equation (82):

$$m(x) = \sum_{i=0}^{50} m_i x^i \quad (82)$$

and $r_i, i = 0, \dots, 11$ and $m_i, i = 0, \dots, 50$ are elements of GF(2).

The message polynomial is created as follows: m_{50} is the first bit of the message and m_0 is the last bit of the message. The order of the parity bits is as follows: r_{11} is the first parity bit transmitted and r_0 is the last parity bit transmitted.

9.6.2.2 BCH(126,63) encoder

In case of high QoS mode operation, the shortened BCH($n = 126, k = 63$) encoder shall be used according to the HARQ mechanism described in 9.15. Such shortened BCH(126,63) code is derived from the mother code BCH(127,64), whose generator polynomial is given by Equation (83):

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

$$g(x) = 1 + x^2 + x^5 + x^{15} + x^{18} + x^{19} + x^{21} + x^{22} + x^{23} + x^{24} + x^{25} + x^{26} + x^{30} + x^{31} + x^{32} + x^{33} + x^{35} + x^{36} + x^{38} + x^{40} + x^{47} + x^{48} + x^{49} + x^{51} + x^{53} + x^{55} + x^{56} + x^{61} + x^{63} \quad (83)$$

The parity bits are determined similarly to 9.6.2.1.

9.6.3 Pad bits

Pad bits shall be appended to the input bit stream to align on a symbol boundary. The number of pad bits is given by Equation (84):

$$N_{pad} = \log_2(M) \left\lceil \frac{N_{PSDU} + (n-k)N_{CW}}{\log_2(M)} \right\rceil - [N_{PSDU} + (n-k)N_{CW}] \quad (84)$$

where M is the cardinality of the constellation of a given modulation scheme.

All appended pad bits shall be set to zero. In the case of uncoded transmission, N_{CW} shall be set to zero.

The total number of bits on the air is given by Equation (85):

$$N_T = N_{PSDU} + (n-k)N_{CW} + N_{pad} \quad (85)$$

9.6.4 Bit interleaving

Bit interleaving shall be applied prior to modulation to provide robustness against error propagation. The algebraic interleaver shall be defined as shown in Equation (86):

$$\Pi(n) = n b_s \text{ Mod } N_I \quad (86)$$

where N_I is the interleaver's length, $\Pi(n) \in [0, N_I - 1]$ denotes the new position to which index n is permuted or interleaved, and $\text{Mod } N_I$ represents modulo N_I arithmetic.

The interleaver's length shall be set to $N_I = 192$, and seeding parameter shall be set to $b = 37$.

The interleaver is applied in blocks of N_I bits over the total number of bits N_T .

If $N_{rem} = \text{rem}(N_T, N_I) \neq 0$, in the last interleaved block, N_I shall be set to N_{rem} .

9.7 PHR construction

The PHR contains information about the data rate of the PSDU, length of the MAC frame body, pulse shape, burst mode, HARQ, and scrambler seed. The PHR construction is illustrated in Figure 133.




Figure 133 —PHR construction

9.7.1 PHR frame

The PHR information frame shall be formed by 24 bits as illustrated in Table 55.

Table 55—PHR frame structure

Bit 0	1	2	3	4	5	6	7	8	9	10	11	12	13
R ₀	R ₁	R ₂	r	L ₀	L ₁	L ₂	L ₃	L ₄	L ₅	L ₆	L ₇	r	R
Data rate				MAC frame body length									

14	15	16	17	18	19	20	21	22	23
B	W ₀	W ₁	H ₀	H ₁	SS	K _m	r	r	r
	Pulse type		HARQ						

r = reserved, B = burst mode, SS = scrambler seed, K_m = on-off constellation mapper.

The description of the different fields of the PHR is as follows.

9.7.1.1 Data rate

Data rates (R₀, R₁, R₂), where R₂ is the most significant bit (MSB) and R₀ is the least significant bit (LSB), for IR-UWB are defined in Table 67 and Table 68. Data rates for FM-UWB are defined in Table 71.

9.7.1.2 MAC frame body length

A variable frame length is indicated with eight bits (L₀, L₁, L₂, L₃, L₄, L₅, L₆, L₇), where L₇ is the MSB and L₀ is the LSB.

9.7.1.3 Burst mode (BM)

The MAC shall set the burst mode bit (B) as defined in Table 56. The burst mode supports higher throughput by allowing the transmission of consecutive frames without ACK. In the burst mode (B = 1), the interframe spacing shall be equal to pMIFS (see 9.18.1 and Table 74).

Table 56—Burst mode

B	Next packet status
0	Next packet is not part of burst
1	Next packet is part of burst

9.7.1.4 Pulse type

The employed pulse shape for transmission is indicated by (W₀, W₁) and defined in Table 57.

IEEE Std 802.15.6-2012
IEEE Standard for Local and metropolitan area networks—
Part 15.6: Wireless Body Area Networks

Table 57—Pulse type

W_0	W_1	Pulse type
0	0	Chirp pulse
0	1	Chaotic pulse
1	0	Short pulse shape
1	1	Reserved

9.7.1.5 HARQ

The HARQ retransmission flow is controlled by H_0 and H_1 (see 9.15) and defined in Table 58.

Table 58—HARQ

H_0	H_1	HARQ state
0	0	Disable
1	0	BCH encoding: $D + P \leftarrow \text{BCH}$ & Send D
0	1	Send D
1	1	Send P

D = systematic bits, P = parity bits.

In the default mode, H_0 and H_1 shall be set to (0,0) and optionally to (1,0).

9.7.1.6 Scrambler seed

The MAC shall set the scrambler seed bit as defined in Table 54 (see 9.6.1).

9.7.1.7 Constellation mapper for on-off modulation

The constellation mapper used for on-off modulation is indicated in Table 59.

Table 59—Constellation mapper

K_m	Symbol mapper
0	Table 63 ($K = 4, M = 16$)
1	Table 64 ($K = 1, M = 2$)

The mandatory mapping shall be $K_m = 1$.