

**IEEE Standard for
Local and metropolitan area networks—**

**Part 15.4: Low-Rate Wireless Personal Area
Networks (LR-WPANs)**

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

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

IEEE Std 802.15.4™-2011
(Revision of
IEEE Std 802.15.4-2006)

5 September 2011

**IEEE Standard for
Local and metropolitan area networks—**

**Part 15.4: Low-Rate Wireless Personal Area
Networks (LR-WPANs)**

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 16 June 2011

IEEE-SA Standards Board

Approved 14 August 2012

American National Standards Institute

Abstract: The protocol and compatible interconnection for data communication devices using low-data-rate, low-power, and low-complexity short-range radio frequency (RF) transmissions in a wireless personal area network (WPAN) were defined in IEEE Std 802.15.4-2006. In this revision, the market applicability of IEEE Std 802.15.4 is extended, the ambiguities in the standard are removed, and the improvements learned from implementations of IEEE Std 802.15.4-2006 are included.

Keywords: ad hoc network, IEEE 802.15.4, low data rate, low power, LR-WPAN, mobility, PAN, personal area network, radio frequency, RF, short range, wireless, wireless personal area network, WPAN

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

Copyright © 2011 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 5 September 2011. Printed in the United States of America.

IEEE is a registered trademark in the U.S. Patent & Trademark Office, owned by the Institute of Electrical and Electronics Engineers, Incorporated.

Print: ISBN 978-0-7381-6683-4 STD97126
PDF: ISBN 978-0-7381-6684-1 STDPD97126

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

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

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

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

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

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

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

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position, explanation, or interpretation of the IEEE. Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments.

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

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854-4141
USA

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

Introduction

This introduction is not part of IEEE Std 802.15.4-2011, IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs).

This is the second revision of IEEE Std 802.15.4. From the beginning, the goal of 802.15.4 was to produce a standard that enabled very low-cost, low-power communications. The initial standard, IEEE Std 802.15.4-2003, defined two optional PHYs, operating in different frequency bands with a very simple, but effective, MAC.

In 2006, the standard was revised, adding two more PHY options. The MAC was backward-compatible, but it added MAC frames with an increased version number, new security features, and a variety of MAC enhancements, including:

- Support for a shared time base with a data time stamping mechanism
- Support for beacon scheduling
- Synchronization of broadcast messages in beacon-enabled PANs

In 2007, two new PHYs were added as an amendment, one of which supported accurate ranging. As a part of this amendment, MAC capability to support ranging was added.

In 2009, two new PHY amendments were approved, one to provide operation in frequency bands specific in China and the other for operation in frequency bands specific to Japan.

The current revision of the standard was created to roll in the previous three amendments into a single document. However, IEEE Std 802.15.4 had become very popular, and there were three additional amendments, 2 PHY and 1 MAC, in process at that time. It was clear that the original organization of the standard was inadequate for the variety of applications, optional PHYs and optional MAC features to which the 802.15.4 base standard would be applied.

Thus, the major changes in the current revision are not technical but editorial. The organization of the standard was changed so that each PHY now has a separate clause. The MAC clause was split into functional description, interface specification, and security specification. In addition, a great deal of informative text, including the coexisting annex and regulatory annex, were deleted so that the document would focus on only those technical requirements needed for interoperability. The revised organization is the consensus decision of a broad group of 802.15 members, including people who were part of the original standard as well as individuals developing amendments to the standard for new applications.

The PAR for IEEE Std 802.15.4-2011 was first proposed in July 2010 and was approved in September 2010 by NesCom. After a total of 10 drafts, 3 working group ballots, and 4 sponsor ballots, the final standard was approved in June 2011, less than one year from start to finish.

Notice to users

Laws and regulations

Users of these documents should consult all applicable laws and regulations. Compliance with the provisions of this standard does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does

not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

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

Updating of IEEE documents

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

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

Errata

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

Interpretations

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

Patents

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

Participants

At the time the draft of this standard was sent to sponsor ballot, the IEEE P802.15 Working Group had the following voting members:

Robert F. Heile, *Chair*
Rick Alfvén, *Co-Vice Chair*
Patrick W. Kinney, *Co-Vice Chair*
James P. K. Gilb, *Working Group Technical Editor*
Patrick W. Kinney, *Secretary*

James P. K. Gilb, *Task Group 4i Chair and Technical Editor*

Emad Afifi	Charles Farlow	John Lampe
Gahng-Seop Ahn	John Farserotu	Zhou Lan
Roberto Aiello	Jeffrey Fischbeck	Khanh Le
Arthur Astrin	Mike Fischer	Cheolhyo Lee
Taehan Bae	George Flammer	Hyungsoo Lee
Michael Bahr	Ryosuke Fujiwara	Myung Lee
John Barr	Noriyasu Fukatsu	Daniel Lewis
Anuj Batra	Kiyoshi Fukui	Huan-Bang Li
Tuncer Baykas	John Geiger	Liang Li
Philip E. Beecher	Gregory Gillooly	Sang-Kyu Lim
Ashutosh Bhatia	Tim Godfrey	Jeremy Link
Ghulam Bhatti	Paul Gorday	Mike Lynch
Gary Birk	Elad Gottlib	Robert Mason
Mathew Boytim	Robert Hall	Tomokuni Matsumura
Peter David Bradley	Shinsuke Hara	Jeff McCullough
Nancy Bravin	Hiroshi Harada	Michael McGillan
David Britz	Timothy Harrington	Michael D. McInnis
Monique B. Brown	Rodney Hemminger	Michael McLaughlin
Sverre Brubk	Marco Hernandez	Charles Millet
Brian Buchanan	Garth Hillman	Siamak Mirnezami
John Buffington	Jin-Meng Ho	Rishi Mohindra
Kiran Bynam	Wei Hong	Emmanuel Monnerie
Brent Cain	Srinath Hosur	Rajendra Moorti
Edgar H. Callaway	David Howard	Robert Moskowitz
Chris Calvert	Jung-Hwan Hwang	Hamilton Moy
Ruben Cardozo	Taeho Hwang	Peter Murray
Douglas Castor	Ichirou Ida	Theodore Myers
Jaesang Cha	Tetsushi Ikegami	Chiu Ngo
Russell Chandler	Akio Iso	Paul Nikolich
Kuor-Hsin Chang	Yeong Min Jang	Hirohito Nishiyama
Soo-Young Chang	Adrian Jennings	David Olson
Clint Chaplin	Wuncheol Jeong	Okundu Omeni
Hind Chebbo	Steven Jillings	Ryoji Ono
Chang-Soon Choi	Noh-Gyoung Kang	Laurent Ouvry
Sangsung Choi	Tae-Gyu Kang	James Pace
Ciaran Connell	Shuzo Kato	Hyung-Il Park
David Cypher	Tatsuya Kato	Jahng Park
Matthew Dahl	Jeritt Kent	Seung-Hoon Park
David Davenport	Prithpal Khakuria	Taejoon Park
Mark Dawkins	Dae Ho Kim	Ranjeet Patro
Hendricus De Ruijter	Dong-Sun Kim	Al Petrick
Upkar Dhaliwal	Dukhyun Kim	Dalibor Pokrajac
Gang Ding	Jaehwan Kim	Daniel Popa
Paul Dixon	Jeffrey King	Stephen Pope
Guido Dolmans	Ryuji Kohno	Clinton C. Powell
Igor Dotlic	Fumihide Kojima	Richard Powell
Michael Dow	Bruce Kraemer	Chang-Woo Pyo
Dietmar Eggert	Raymond Krasinski	Mohammad Rahman
David Evans	Masahiro Kuroda	Sridhar Rajagopal

Jayaram Ramasastry
 Marc Reed
 Ivan Reede
 Richard Roberts
 Craig Rodine
 June Chul Roh
 Benjamin Rolfe
 Seung-Moon Ryu
 Didier Sagan
 Kentaro Sakamoto
 Will San Filippo
 H. Sanderford
 Kamran Sayrafian
 Timothy Schmidl
 Michael Schmidt
 Jean Schwoerer
 Cristina Seibert
 Neal Seidl
 Kunal Shah
 Steve Shearer
 Stephen Shellhammer
 Shusaku Shimada

Chang Sub Shin
 Cheol Ho Shin
 Michael Sim
 Jonathan Simon
 Jaeseung Son
 Paul Stadnik
 René Struik
 Chin-Sean Sum
 Hui-Hsia Sung
 Gu Sungi
 Kenichi Takizawa
 Hirokazu Tanaka
 Larry Taylor
 Mark Thompson
 James Tomcik
 Ichihiko Toyoda
 David Tracey
 Khanh Tran
 Jerry Upton
 Jana van Greunen
 Hartman van Wyk
 Michel Veillette

Billy Verso
 Bhupender Virk
 Joachim Walewski
 Junyi Wang
 Quan Wang
 Xiang Wang
 Andy Ward
 Scott Weikel
 Nicholas West
 Mark Wilbur
 Ludwig Winkel
 Eun Tae Won
 Alan Chi Wai Wong
 Tao Xing
 Wen-Bin Yang
 Yang Yang
 Kazuyuki Yasukawa
 Kamyaz Yazdandoost
 Kaoru Yokoo
 Mu Zhao
 Bin Zhen

Major contributions were received from the following individuals:

Philip E. Beecher
 Vern Brethour
 Monique B. Brown
 Edgar H. Callaway
 Kuor-Hsin Chang

Clint Chaplin
 James P. K. Gilb
 Patrick W. Kinney
 Michael D. McInnis

Clinton C. Powell
 Benjamin Rolfe
 Timothy Schmidl
 René Struik
 Billy Verso

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

Jon Adams
Emad Afifi
Roberto Aiello
Rick Alfvin
Nobumitsu Amachi
Mark Anderson
Tuncer Baykas
Philip E. Beecher
H. Stephen Berger
Maciej Borowka
Nancy Bravin
Vern Brethour
John Buffington
William Byrd
Edgar H. Callaway
Radhakrishna Canchi
Ruben Cardozo
Juan Carreon
Dave Cavalcanti
Kuor-Hsin Chang
Clint Chaplin
Keith Chow
Charles Cook
Todor Cooklev
Joseph Decuir
Patrick Diamond
Russell Dietz
Thomas Dineen
Sourav Dutta
Richard Eckard
Devon Gayle
John Geiger
James P. K. Gilb
Gregory Gillooly
Randall Groves
Jose Gutierrez
C. Guy
Rainer Hach
Robert F. Heile

Marco Hernandez
Oliver Hoffmann
Wei Hong
Heqing Huang
David Hunter
Noriyuki Ikeuchi
Akio Iso
Atsushi Ito
Raj Jain
Oyvind Janbu
Tal Kaitz
Shinkyo Kaku
Stuart J. Kerry
Brian Kiernan
Yongbum Kim
Youngsoo Kim
Jeffrey King
Bruce Kraemer
Geoff Ladwig
Paul Lambert
Jeremy Landt
Zhongding Lei
Jan-Ray Liao
Arthur Light
Lu Liru
HaiTao Liu
William Lumpkins
Greg Luri
Mike Lynch
Elvis Maculuba
Wayne W. Manges
Michael D. McInnis
Gary Michel
Apurva Mody
Emmanuel Monnerie
Jose Morales
Peter Murray
Michael S. Newman
Nick S. A. Nikjoo
Paul Nikolich

John Notor
Satoshi Obara
Chris Osterloh
Clinton C. Powell
Venkatesha Prasad
Mohammad Rahman
Jayaram Ramasastry
Robert Robinson
Benjamin Rolfe
Randall Safier
Kazuyuki Sakoda
Naotaka Sato
Bartien Sayogo
Timothy Schmidl
Cristina Seibert
Jie Shen
Suresh Shrimavle
Gil Shultz
Kapil Sood
Robert Soranno
Thomas Starai
René Struik
Walter Struppler
Mark Sturza
Chin-Sean Sum
Larry Taylor
Dmitri Varsanofiev
Prabodh Varshney
Bhupender Virk
George Vlantis
Xiang Wang
Scott Weikel
Stephen Weinstein
Andreas Wolf
Tao Xing
Yang Yang
Tan Pek Yew
Oren Yuen

When the IEEE-SA Standards Board approved this standard on 16 June 2011, it had the following membership:

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

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

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

Jon Rosdahl
Sam Sciacca
Mike Seavey
Curtis Siller
Phil Winston
Howard Wolfman
Don Wright

*Member Emeritus

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

Satish K. Aggarwal, *NRC Representative*
Richard DeBlasio, *DOE Representative*
Alan H. Cookson, *NIST Representative*

Catherine Berger
IEEE Standards Program Manager, Document Development

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

Contents

1.	Overview	1
1.1	General	1
1.2	Scope	1
1.3	Purpose	2
2.	Normative references	3
3.	Definitions, acronyms, and abbreviations	4
3.1	Definitions	4
3.2	Acronyms and abbreviations	5
4.	General description	8
4.1	General	8
4.2	Components of the IEEE 802.15.4 WPAN	8
4.3	Network topologies	8
4.3.1	Star network formation	9
4.3.2	Peer-to-peer network formation	9
4.4	Architecture	10
4.4.1	Physical layer (PHY)	11
4.4.2	MAC sublayer	11
4.5	Functional overview	12
4.5.1	Superframe structure	12
4.5.2	Data transfer model	13
4.5.3	Frame structure	14
4.5.4	Improving probability of successful delivery	14
4.5.5	Power consumption considerations	15
4.5.6	Security	15
4.6	Concept of primitives	16
5.	MAC protocol	18
5.1	MAC functional description	18
5.1.1	Channel access	18
5.1.2	Starting and maintaining PANs	24
5.1.3	Association and disassociation	32
5.1.4	Synchronization	36
5.1.5	Transaction handling	39
5.1.6	Transmission, reception, and acknowledgment	40
5.1.7	GTS allocation and management	48
5.1.8	Ranging	54
5.2	MAC frame formats	57
5.2.1	General MAC frame format	57
5.2.2	Format of individual frame types	61
5.2.3	Frame compatibility	67
5.3	MAC command frames	67
5.3.1	Association request command	67

5.3.2	Association response command	69
5.3.3	Disassociation notification command	70
5.3.4	Data request command	71
5.3.5	PAN ID conflict notification command	72
5.3.6	Orphan notification command	73
5.3.7	Beacon request command	73
5.3.8	Coordinator realignment command	74
5.3.9	GTS request command	75
6.	MAC services	77
6.1	Overview	77
6.2	MAC management service	77
6.2.1	Common requirements for MLME primitives	78
6.2.2	Association primitives	79
6.2.3	Disassociation primitives	83
6.2.4	Communications notification primitives	86
6.2.5	Primitives for reading PIB attributes	90
6.2.6	GTS management primitives	91
6.2.7	Primitives for orphan notification	94
6.2.8	Primitives for resetting the MAC sublayer	96
6.2.9	Primitives for specifying the receiver enable time	97
6.2.10	Primitives for channel scanning	99
6.2.11	Primitives for writing PIB attributes	102
6.2.12	Primitives for updating the superframe configuration	103
6.2.13	Primitives for synchronizing with a coordinator	107
6.2.14	Primitives for requesting data from a coordinator	110
6.2.15	Primitives for specifying dynamic preamble	112
6.2.16	Primitives for channel sounding	113
6.2.17	Primitives for ranging calibration (for UWB PHYs)	114
6.3	MAC data service	116
6.3.1	MCPS-DATA.request	116
6.3.2	MCPS-DATA.confirm	118
6.3.3	MCPS-DATA.indication	119
6.3.4	MCPS-PURGE.request	123
6.3.5	MCPS-PURGE.confirm	124
6.4	MAC constants and PIB attributes	124
6.4.1	MAC constants	124
6.4.2	MAC PIB attributes	124
6.4.3	Calculating PHY dependent MAC PIB values	129
7.	Security	131
7.1	Overview	131
7.2	Functional description	131
7.2.1	Outgoing frame security procedure	131
7.2.2	KeyDescriptor lookup procedure	132
7.2.3	Incoming frame security procedure	133
7.2.4	DeviceDescriptor lookup procedure	135
7.2.5	SecurityLevelDescriptor lookup procedure	135
7.2.6	Incoming security level checking procedure	135
7.2.7	Incoming key usage policy checking procedure	136
7.3	Security operations	136
7.3.1	Integer and octet representation	136

7.3.2	CCM* Nonce	136
7.3.3	CCM* prerequisites	137
7.3.4	CCM* transformation data representation	137
7.3.5	CCM* inverse transformation data representation	138
7.4	Auxiliary security header	139
7.4.1	Security Control field	140
7.4.2	Frame Counter field	141
7.4.3	Key Identifier field	141
7.5	Security-related MAC PIB attributes	142
8.	General PHY requirements	146
8.1	General requirements and definitions	146
8.1.1	Operating frequency range	146
8.1.2	Channel assignments	147
8.1.3	Minimum LIFS and SIFS periods	150
8.1.4	RF power measurement	151
8.1.5	Transmit power	151
8.1.6	Out-of-band spurious emission	151
8.1.7	Receiver sensitivity definitions	151
8.2	General radio specifications	151
8.2.1	TX-to-RX turnaround time	151
8.2.2	RX-to-TX turnaround time	151
8.2.3	Error-vector magnitude (EVM) definition	152
8.2.4	Receiver maximum input level of desired signal	153
8.2.5	Receiver ED	153
8.2.6	Link quality indicator (LQI)	153
8.2.7	Clear channel assessment (CCA)	153
9.	PHY services	155
9.1	Overview	155
9.2	PHY constants	155
9.3	PHY PIB attributes	156
9.4	PHY PIB attribute values for phyMaxFrameDuration and phySHRDuration	158
10.	O-QPSK PHY	160
10.1	PPDU format	160
10.1.1	Preamble field	160
10.1.2	SFD field	160
10.1.3	Frame Length field	160
10.1.4	PSDU field	160
10.2	Modulation and spreading	161
10.2.1	Data rate	161
10.2.2	Reference modulator diagram	161
10.2.3	Bit-to-symbol mapping	161
10.2.4	Symbol-to-chip mapping	162
10.2.5	O-QPSK modulation	163
10.2.6	Pulse shape	163
10.2.7	Chip transmission order	164
10.3	O-QPSK PHY RF requirements	164
10.3.1	Operating frequency range	164
10.3.2	Transmit power spectral density (PSD) mask	164

10.3.3	Symbol rate	165
10.3.4	Receiver sensitivity	165
10.3.5	Receiver interference rejection	165
10.3.6	TX-to-RX turnaround time	166
10.3.7	RX-to-TX turnaround time	166
10.3.8	Error vector magnitude (EVM).....	166
10.3.9	Transmit center frequency tolerance.....	166
10.3.10	Transmit power	166
10.3.11	Receiver maximum input level of desired signal.....	166
10.3.12	Receiver ED	166
10.3.13	Link quality indicator (LQI)	166
10.3.14	Clear channel assessment (CCA).....	166
11.	Binary phase-shift keying (BPSK) PHY	167
11.1	PPDU format.....	167
11.2	Modulation and spreading	167
11.2.1	BPSK PHY data rates	167
11.2.2	Reference modulator diagram.....	167
11.2.3	Differential encoding	167
11.2.4	Bit-to-chip mapping	168
11.2.5	BPSK modulation	168
11.3	BPSK PHY RF requirements.....	168
11.3.1	Operating frequency range.....	168
11.3.2	915/950 MHz band transmit PSD mask	168
11.3.3	Symbol rate	169
11.3.4	Receiver sensitivity.....	169
11.3.5	Receiver interference rejection	169
11.3.6	TX-to-RX turnaround time	170
11.3.7	RX-to-TX turnaround time	170
11.3.8	Error vector magnitude (EVM).....	170
11.3.9	Transmit center frequency tolerance.....	170
11.3.10	Transmit power	170
11.3.11	Receiver maximum input level of desired signal.....	170
11.3.12	Receiver ED	170
11.3.13	Link quality indicator (LQI)	170
11.3.14	Clear channel assessment (CCA).....	170
12.	Amplitude shift keying (ASK) PHY.....	171
12.1	PPDU format.....	171
12.1.1	Preamble for ASK PHY	171
12.1.2	SFD for ASK PHY	171
12.2	Modulation and spreading	171
12.2.1	ASK PHY data rates	171
12.2.2	Reference modulator diagram.....	172
12.2.3	Bit-to-symbol mapping	172
12.2.4	Symbol-to-chip mapping	172
12.2.5	ASK modulation	173
12.2.6	Pulse shape.....	173
12.2.7	Chip transmission order	175
12.3	ASK PHY RF requirements.....	175
12.3.1	Operating frequency range.....	175
12.3.2	915 MHz band transmit PSD mask.....	175

12.3.3	Symbol rate	175
12.3.4	Receiver sensitivity	175
12.3.5	Receiver interference rejection	176
12.3.6	TX-to-RX turnaround time	176
12.3.7	RX-to-TX turnaround time	176
12.3.8	Error vector magnitude (EVM).....	176
12.3.9	Transmit center frequency tolerance.....	176
12.3.10	Transmit power	176
12.3.11	Receiver maximum input level of desired signal.....	176
12.3.12	Receiver ED	177
12.3.13	Link quality indicator (LQI)	177
12.3.14	Clear channel assessment (CCA).....	177
12.3.15	Example of PSSS encoding	177
13.	Chirp spread spectrum (CSS) PHY	179
13.1	CSS PPDU format	179
13.1.1	Preamble	179
13.1.2	SFD field.....	179
13.1.3	PHY header (PHR)	180
13.2	Modulation and spreading	180
13.2.1	Data rates	180
13.2.2	Reference modulator diagram.....	180
13.2.3	De-multiplexer (DEMUX).....	180
13.2.4	Serial-to-parallel mapping	181
13.2.5	Data-symbol-to-bi-orthogonal-codeword mapping	181
13.2.6	Parallel-to-serial converter and QPSK symbol mapping.....	184
13.2.7	DQPSK coding	184
13.2.8	DQPSK-to-DQCSK modulation.....	185
13.2.9	CSK generator.....	185
13.2.10	Bit interleaver	186
13.3	Waveform and subchirp sequences.....	186
13.3.1	Graphical presentation of chirp symbols (subchirp sequences).....	186
13.3.2	Active usage of time gaps.....	186
13.3.3	Mathematical representation of the continuous time CSS base-band signal	188
13.3.4	Raised cosine window for chirp pulse shaping.....	189
13.3.5	Subchirp transmission order	190
13.3.6	Example of CSK signal generation.....	190
13.4	CSS RF requirements.....	191
13.4.1	Transmit power spectral density (PSD) mask and signal tolerance.....	191
13.4.2	Symbol rate	191
13.4.3	Receiver sensitivity.....	191
13.4.4	Receiver interference rejection	192
13.4.5	TX-to-RX turnaround time	192
13.4.6	RX-to-TX turnaround time	193
13.4.7	Transmit center frequency tolerance.....	193
13.4.8	Transmit power	193
13.4.9	Receiver maximum input level of desired signal.....	193
13.4.10	Receiver ED	193
13.4.11	Link quality indicator (LQI)	193
13.4.12	Clear channel assessment (CCA).....	193

14.	UWB PHY	194
14.1	General.....	194
14.2	UWB PPDU format	194
14.2.1	PPDU encoding process.....	195
14.2.2	UWB PHY symbol structure	196
14.2.3	PSDU timing parameters	197
14.2.4	Preamble timing parameters	200
14.2.5	SHR preamble.....	202
14.2.6	PHY header (PHR)	205
14.2.7	Data field.....	207
14.3	UWB PHY modulation	207
14.3.1	UWB PHY modulation mathematical framework	207
14.3.2	UWB PHY spreading.....	208
14.3.3	UWB PHY forward error correction (FEC)	209
14.4	UWB PHY RF requirements	211
14.4.1	Operating frequency bands	211
14.4.2	Channel assignments.....	213
14.4.3	Regulatory compliance	213
14.4.4	Operating temperature range	213
14.4.5	Baseband impulse response	213
14.4.6	Transmit PSD mask	214
14.4.7	Chip rate clock and chip carrier alignment.....	215
14.4.8	TX-to-RX turnaround time	215
14.4.9	RX-to-TX turnaround time	215
14.4.10	Transmit center frequency tolerance.....	215
14.4.11	Transmit power	215
14.4.12	Receiver maximum input level of desired signal.....	215
14.4.13	Receiver ED	215
14.4.14	Link quality indicator (LQI)	216
14.4.15	Clear channel assessment (CCA).....	216
14.5	UWB PHY optional pulse shapes.....	216
14.5.1	UWB PHY optional chirp on UWB (CoU) pulses	216
14.5.2	UWB PHY optional continuous spectrum (CS) pulses	217
14.5.3	UWB PHY linear combination of pulses (LCP).....	219
14.6	Extended preamble for optional UWB CCA mode	219
14.7	Ranging.....	220
14.7.1	Ranging counter.....	220
14.7.2	Crystal characterization	220
14.7.3	Ranging FoM.....	221
15.	GFSK PHY	223
15.1	PPDU formats	223
15.2	Modulation.....	223
15.2.1	GFSK PHY data rates	223
15.2.2	Reference modulator diagram.....	223
15.2.3	Data whitening.....	223
15.2.4	GFSK modulation.....	224
15.3	GFSK PHY RF requirements	224
15.3.1	Operating frequency range.....	224
15.3.2	Transmit PSD mask	224
15.3.3	Symbol rate	224
15.3.4	Receiver sensitivity.....	225

15.3.5	Receiver interference rejection	225
15.3.6	TX-to-RX turnaround time	225
15.3.7	RX-to-TX turnaround time	225
15.3.8	Transmit center frequency tolerance.....	225
15.3.9	Transmit power	225
15.3.10	Receiver maximum input level of desired signal.....	225
15.3.11	Receiver ED	225
15.3.12	Link quality indicator (LQI)	226
15.3.13	Clear channel assessment (CCA).....	226
Annex A (informative) Bibliography		227
Annex B (normative) CCM* mode of operation		229
B.1	Introduction.....	229
B.2	Notation and representation	229
B.2.1	Strings and string operations.....	229
B.2.2	Integers, octets, and their representation	229
B.3	Symmetric-key cryptographic building blocks.....	229
B.3.1	Block cipher	230
B.3.2	Mode of operation.....	230
B.4	Specification of generic CCM* mode of operation	230
B.4.1	CCM* mode encryption and authentication transformation.....	230
B.4.1.1	Input transformation	231
B.4.1.2	Authentication transformation	231
B.4.1.3	Encryption transformation	232
B.4.2	CCM* mode decryption and authentication checking transformation	232
B.4.2.1	Decryption transformation	233
B.4.2.2	Authentication checking transformation.....	233
B.4.3	Restrictions	233
Annex C (informative) Test vectors for cryptographic building blocks.....		235
C.1	AES block cipher	235
C.2	Mode of operation.....	235
C.2.1	MAC beacon frame.....	235
C.2.1.1	Description.....	235
C.2.1.2	CCM* mode encryption and authentication transformation.....	235
C.2.1.3	CCM* mode decryption and authentication checking transformation	237
C.2.2	MAC data frame	239
C.2.2.1	Description.....	239
C.2.2.2	CCM* mode encryption and authentication transformation.....	239
C.2.2.3	CCM* mode decryption and authentication checking transformation	241
C.2.3	MAC command frame	243
C.2.3.1	Description.....	243
C.2.3.2	CCM* mode encryption and authentication transformation.....	243
C.2.3.3	CCM* mode decryption and authentication checking transformation	245
Annex D (informative) Protocol implementation conformance statement (PICS) proforma.....		248
D.1	Introduction.....	248
D.1.1	Scope.....	248
D.1.2	Purpose.....	248
D.2	Abbreviations and special symbols.....	248

D.3	Instructions for completing the PICS proforma.....	249
D.4	Identification of the implementation.....	249
D.5	Identification of the protocol	250
D.6	Global statement of conformance	250
D.7	PICS proforma tables.....	251
D.7.1	Functional device types	251
D.7.2	Major capabilities for the PHY	251
D.7.2.1	PHY functions.....	251
D.7.2.2	Radio frequency (RF)	252
D.7.2.3	Channel capabilities for UWB PHY	252
D.7.3	Major capabilities for the MAC sublayer	255
D.7.3.1	MAC sublayer functions.....	255
D.7.3.2	MAC frames	256
Annex E (informative)	Location topics	258
E.1	Overview.....	258
E.1.1	Two-way ranging.....	258
E.1.2	Position awareness through one-way transmissions.....	260
E.1.3	The ranging counter	260
E.1.4	Accounting for signal arrival time	261
E.1.4.1	Leading edge search during the acquisition preamble.....	261
E.1.4.2	FoM for bad times.....	262
E.1.4.3	Other opportunities for leading edge search refinement.....	262
E.1.4.4	Managing the preamble length for leading edge search	262
E.1.4.5	PHY deferral of the computations for leading edge search	262
E.1.4.6	PHY deferral of the computations for self-calibration	263
E.1.5	Management of crystal offsets.....	263
E.1.5.1	Characterizing crystal offsets with digital tracking loops	264
E.1.5.2	Characterizing crystal offsets with analog tracking loops	264
E.1.5.3	Characterizing crystal offsets with different tracking loops	264
E.1.5.4	Size of units	265
E.1.6	Accounting for internal propagation paths	265
E.1.6.1	PIB attributes for internal propagation paths.....	266
E.1.6.2	Support for self-calibration and one-way ranging	266
E.1.6.3	Use of the calibrate primitives	266
E.1.6.4	Use of the COMPUTATION_NEEDED status.....	266
E.1.7	Timestamp reports	267
E.1.7.1	Presentation of timestamp reports.....	267
E.1.7.2	Start and stop times in the timestamp report.....	267
E.1.8	Private ranging.....	267
E.1.8.1	Simple encryption of the timestamp reports	267
E.1.8.2	Dynamic preamble selection (DPS).....	267
E.2	Time-of-arrival estimation from channel sounding	268
E.3	Time-of-arrival estimation in non-line-of-sight (NLOS) conditions	270
E.4	Asynchronous ranging	271
E.4.1	Two-way ranging (TWR)	272
E.4.2	Symmetric double-sided two-way ranging (SDS-TWR).....	273
E.5	Location estimation from range data	274
E.5.1	Time of arrival	275
E.5.2	Time difference of arrival	276
E.5.2.1	Mode 1	276
E.5.2.2	Mode 2	276

E.6	Network location algorithms	276
E.6.1	Ad hoc algorithms	277
E.6.2	Centralized algorithms	279
E.6.3	Convex optimization algorithms	280
E.6.4	Location estimation using multipath delays	282
Annex F (informative) Example UWB PHY transmit data frame encoding		284
F.1	Channel used in the example	284
F.2	Encoding progression	284
F.2.1	Transmit PSDU	284
F.2.2	PSDU bits	284
F.2.3	Reed-Solomon encoded bits	284
F.2.4	Convolutional encoder input bits	285
F.2.5	Convolutional encoder output bits	285
F.2.6	Scrambler output bits	285
F.2.7	Ternary output symbols	286
Annex G (informative) MPSK PHY		289
G.1	General	289
G.2	780 MHz band data rates	289
G.3	Modulation and spreading	289
G.3.1	Reference modulator diagram	289
G.3.2	Bit-to-symbol mapping	289
G.3.3	Symbol-to-chip mapping	290
G.3.4	Pre-processing	291
G.3.5	PSK modulation	291
G.3.6	Pulse shape	291
G.3.7	Chip transmission order	292
G.4	MPSK PHY RF requirements	292
G.4.1	Transmit power	292
G.4.2	Operating frequency range	292
G.4.3	Transmit PSD mask	292
G.4.4	Symbol rate	292
G.4.5	Receiver sensitivity	292
G.4.6	Receiver interference rejection	293
Annex H (informative) Considerations for the 950 MHz band		294
H.1	General	294
H.2	Listen before talk (LBT) considerations	294

IEEE Standard for Local and metropolitan area networks—

Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)

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

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

1. Overview

1.1 General

Wireless personal area networks (WPANs) are used to convey information over relatively short distances. Unlike wireless local area networks (WLANs), connections effected via WPANs involve little or no infrastructure. This feature allows small, power-efficient, inexpensive solutions to be implemented for a wide range of devices.

1.2 Scope

This standard defines the physical layer (PHY) and medium access control (MAC) sublayer specifications for low-data-rate wireless connectivity with fixed, portable, and moving devices with no battery or very limited battery consumption requirements typically operating in the personal operating space (POS) of 10 m.

Physical layers (PHYs) are defined for

- Devices operating in the license-free 868–868.6 MHz, 902–928 MHz, and 2400–2483.5 MHz bands
- Devices with precision ranging, extended range, and enhanced robustness and mobility
- Devices operating according the Chinese regulations, Radio Management of P. R. of China doc. #6326360786867187500 or current document, for one or more of the 314–316 MHz, 430–434 MHz, and 779–787 MHz frequency bands
- Devices operating in the 950–956 MHz allocation in Japan and coexisting with passive tag systems in the band

1.3 Purpose

The standard provides for ultra low complexity, ultra low cost, ultra low power consumption, and low data rate wireless connectivity among inexpensive devices. The raw data rate is high enough (250 kb/s) to satisfy a set of applications but is also scaleable down to the needs of sensor and automation needs (20 kb/s or below) for wireless communications.

In addition, one of the alternate PHYs provides precision ranging capability that is accurate to one meter.

Multiple PHYs are defined to support a variety of frequency bands including

- 868–868.6 MHz
- 902–928 MHz
- 2400–2483.5 MHz
- 314–316 MHz, 430–434 MHz, and 779–787 MHz band for LR-WPAN systems in China
- 950–956 MHz in Japan

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.

FIPS Pub 197, Advanced Encryption Standard (AES).¹

NITS/CWPAN Part 15.4, Chinese standard for the Wireless Medium Access Control (MAC) and the Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks, GB/T 15629.15-2010.²

¹FIPS publications are available from the National Technical Information Service (NTIS), U.S. Dept. of Commerce, 5285 Port Royal Rd., Springfield, VA 22161, USA (<http://www.ntis.org/>).

²CWPAN publications are available from either the China Electronics Standardization Institute (CESI), No.1 Andingmen East Street, Dongcheng District, Beijing, China (<http://www.cssn.net.cn>) or the China National Institute of Standardization (CNIS), No. 4 Zhichun Road, Haidian District, Beijing, China.

3. Definitions, acronyms, and abbreviations

3.1 Definitions

For the purposes of this document, the following terms and definitions apply. *The IEEE Standards Dictionary: Glossary of Terms & Definitions* should be consulted for terms not defined in this clause.³

association: The service used to establish membership for a device in a network.

authentication tag: Information that allows the verification of authenticity of a message.

beacon-enabled personal area network (PAN): A PAN in which all coordinators emit regular beacons.

chirp: Linear frequency sweep (frequency may either increase or decrease).

contention access period: The period of time immediately following a beacon frame during which devices wishing to transmit will compete for channel access using a slotted carrier sense multiple access with collision avoidance mechanism.

coordinator: A device in an low rate wireless personal area network (LR WPAN) that provides synchronization services to other devices in the LR WPAN.

data authentication: The process whereby an entity receiving a message corroborates evidence about the true source of the information in the message and, thereby, evidence that the message has not been modified in transit.

data authenticity: Assurance about the source of information.

device: Any entity containing an implementation of the IEEE 802.15.4 medium access control and physical interface to the wireless medium. A device may be a reduced-function device or a full-function device.

encryption: The transformation of a message into a new representation so that privileged information is required to recover the original representation.

frame: The format of aggregated bits from a medium access control sublayer entity that are transmitted together in time.

full-function device: A device capable of operating as a coordinator.

group key: A key that is known only to a set of devices.

key: Privileged information that may be used, for example, to protect information from disclosure to, and/or undetectable modification by, parties that do not have access to this privileged information.

key establishment: A process whereby two or more parties establish a key.

keying material: The combination of a key and associated security information (e.g., a nonce value).

key management: The collection of processes for the establishment and maintenance of keying relationships over a system's lifetime.

³The *IEEE Standards Dictionary: Glossary of Terms & Definitions* is available at <http://shop.ieee.org/>.

link key: A secret key that is shared between precisely two devices.

mobile device: A device whose location in the network may change during use.

nonbeacon-enabled personal area network (PAN): A PAN in which coordinators do not emit regular beacons.

nonce: A nonrepeating value, such as an increasing counter, a sufficiently long random string, or a time stamp.

orphaned device: A device that has lost contact with its associated coordinator.

packet: The formatted, aggregated bits that are transmitted together in time across the physical medium.

payload data: The contents of a data message that is being transmitted.

personal area network (PAN) coordinator: A coordinator that is the principal controller of a PAN. An IEEE 802.15.4 network has exactly one PAN coordinator.

plain text: A string of unscrambled information.

ranging-capable device: A device containing an implementation capable of supporting ranging.

reduced-function device (RFD): A device that is not capable of operating as a coordinator.

symmetric key: A secret key that is shared between two or more parties that may be used for encryption/decryption or integrity protection/integrity verification, depending on its intended use.

transaction: The exchange of related, consecutive frames between two peer medium access control (MAC) entities, required for a successful transmission of a MAC command or data frame.

3.2 Acronyms and abbreviations

ADC	analog-to-digital converter
AES	advanced encryption standard
AGC	automatic gain control
AR	acknowledgment request
ASK	amplitude shift keying
AWGN	additive white Gaussian noise
BER	bit error rate
BLE	battery life extension
BPM	burst position modulation
BPSK	binary phase-shift keying
BSN	beacon sequence number
CAP	contention access period
CBC-MAC	cipher block chaining message authentication code
CCA	clear channel assessment
CCM	counter mode encryption and cipher block chaining message authentication code
CCM*	extension of counter mode encryption and cipher block chaining message authentication code
CFP	contention-free period

CRC	cyclic redundancy check
CoU	chirp on ultra-wide band
CS	continuous spectrum
CSK	chirp-shift keying
CSMA-CA	carrier sense multiple access with collision avoidance
CSS	chirp spread spectrum
CTR	counter mode
CWPAN	Chinese wireless personal area network
DAA	detect and avoid
DEMUX	de-multiplexer
DPS	dynamic preamble selection
DQCSK	differential quadrature chirp-shift keying
DQPSK	differential quadrature phase-shift keying
DSN	data sequence number
DSSS	direct sequence spread spectrum
ED	energy detection
EIRP	effective isotropic radiated power
ERP	effective radiated power
EVM	error-vector magnitude
FCS	frame check sequence
FEC	forward error correction
FFD	full-function device
FH	frequency hopping
FHSS	frequency hopping spread spectrum
FoM	figure of merit
GDOP	geometric dilution of precision
GFSK	Gaussian frequency-shift keying
GTS	guaranteed time slot
IFS	interframe space or spacing
ISM	industrial, scientific, and medical
LBT	listen before talk
LCP	linear combination of pulses
LFSR	linear feedback shift register
LIFS	long interframe spacing
LOS	line-of-sight
LQI	link quality indication
LPDU	logical link control protocol data unit
LR-WPAN	low-rate wireless personal area network
LSB	least significant bit
MAC	medium access control
MCPS	MAC common part sublayer
MCPS-SAP	MAC common part sublayer service access point
MFR	MAC footer
MHR	MAC header
MIC	message integrity code
MLME	MAC sublayer management entity

MLME-SAP	MAC sublayer management entity service access point
MSB	most significant bit
MPDU	MAC protocol data unit
MPSK	m-ary phase shift keying
MSDU	MAC service data unit
NLOS	non-line-of-sight
OCDM	orthogonal code division multiplexing
O-QPSK	offset quadrature phase-shift keying
PAN	personal area network
PC	personal computer
PD	PHY data
PD-SAP	PHY data service access point
PER	packet error rate
PHR	PHY header
PHY	physical layer
PIB	personal area network information base
PICS	protocol implementation conformance statement
PLL	phase-locked loop
PLME	physical layer management entity
PLME-SAP	physical layer management entity service access point
PN	pseudo-random noise
PPDU	PHY protocol data unit
PPM	pulse position modulation
PRBS	pseudo-random binary sequence
PRF	pulse repetition frequency
PSD	power spectral density
PSDU	PHY service data unit
PSSS	parallel sequence spread spectrum
RDEV	ranging-capable device
RF	radio frequency
RFD	reduced-function device
RFRAME	ranging frame
RMARKER	ranging marker
RSSI	receive signal strength indicator
RX	receive or receiver
SDS-TWR	symmetric double-sided two-way ranging
SFD	start-of-frame delimiter
SHR	synchronization header
SIFS	short interframe spacing
SNR	signal-to-noise ratio
TRX	transceiver
TX	transmit or transmitter
UWB	ultra-wide band
WPAN	wireless personal area network

4. General description

4.1 General

An LR-WPAN is a simple, low-cost communication network that allows wireless connectivity in applications with limited power and relaxed throughput requirements. The main objectives of an LR-WPAN are ease of installation, reliable data transfer, extremely low cost, and a reasonable battery life, while maintaining a simple and flexible protocol.

Some of the capabilities provided by this standard are as follows:

- Star or peer-to-peer operation
- Unique 64-bit extended address or allocated 16-bit short address
- Optional allocation of guaranteed time slots (GTSs)
- Carrier sense multiple access with collision avoidance (CSMA-CA) or ALOHA channel access
- Fully acknowledged protocol for transfer reliability
- Low power consumption
- Energy detection (ED)
- Link quality indication (LQI)

This standard defines multiple PHYs operating in a variety of frequency bands, as described in 8.1.1.

Two different device types can participate in an IEEE 802.15.4 network: a full-function device (FFD) and a reduced-function device (RFD). An FFD is a device that is capable of serving as a personal area network (PAN) coordinator or a coordinator. An RFD is a device that is not capable of serving as either a PAN coordinator or a coordinator. An RFD is intended for applications that are extremely simple, such as a light switch or a passive infrared sensor; it does not have the need to send large amounts of data and only associates with a single FFD at a time. Consequently, the RFD can be implemented using minimal resources and memory capacity.

4.2 Components of the IEEE 802.15.4 WPAN

A system conforming to this standard consists of several components. The most basic is the device. Two or more devices communicating on the same physical channel constitute a WPAN. However, this WPAN includes at least one FFD, which operates as the PAN coordinator.

A well-defined coverage area does not exist for wireless media because propagation characteristics are dynamic and uncertain. Small changes in position or direction often result in drastic differences in the signal strength or quality of the communication link. These effects occur whether a device is stationary or mobile, as moving objects affect station-to-station propagation.

4.3 Network topologies

Depending on the application requirements, an IEEE 802.15.4 LR-WPAN operates in either of two topologies: the star topology or the peer-to-peer topology. Both are shown in Figure 1. In the star topology, the communication is established between devices and a single central controller, called the PAN coordinator. A device typically has some associated application and is either the initiation point or the termination point for network communications. A PAN coordinator can also have a specific application, but it can be used to initiate, terminate, or route communication around the network. The PAN coordinator is the primary controller of the PAN. All devices operating on a network of either topology have unique addresses,

referred to as extended addresses. A device will use either the extended address for direct communication within the PAN or the short address that was allocated by the PAN coordinator when the device associated. The PAN coordinator will often be mains powered, while the devices will most likely be battery powered. Applications that benefit from a star topology include home automation, personal computer (PC) peripherals, games, and personal health care.




Figure 1—Star and peer-to-peer topology examples

The peer-to-peer topology also has a PAN coordinator; however, it differs from the star topology in that any device is able to communicate with any other device as long as they are in range of one another. Peer-to-peer topology allows more complex network formations to be implemented, such as mesh networking topology. Applications such as industrial control and monitoring, wireless sensor networks, asset and inventory tracking, intelligent agriculture, and security would benefit from such a network topology. A peer-to-peer network allows multiple hops to route messages from any device to any other device on the network. Such functions can be added at the higher layer, but they are not part of this standard.

Each independent PAN selects a unique identifier. This PAN identifier allows communication between devices within a network using short addresses and enables transmissions between devices across independent networks. The mechanism by which identifiers are chosen is outside the scope of this standard.

The network formation is performed by the higher layer, which is not part of this standard. However, 4.3.1 and 4.3.2 provide a brief overview on how each supported topology can be formed.

4.3.1 Star network formation

The basic structure of a star network is illustrated in Figure 1. After an FFD is activated, it can establish its own network and become the PAN coordinator. All star networks operate independently from all other star networks currently in operation. This is achieved by choosing a PAN identifier that is not currently used by any other network within the radio communications range. Once the PAN identifier is chosen, the PAN coordinator allows other devices, potentially both FFDs and RFDs, to join its network. The higher layer can use the procedures described in 5.1.2 and 5.1.3 to form a star network.

4.3.2 Peer-to-peer network formation

In a peer-to-peer topology, each device is capable of communicating with any other device within its radio communications range. One device is nominated as the PAN coordinator, for instance, by virtue of being the first device to communicate on the channel. Further network structures are constructed out of the peer-to-peer topology, and it is possible to impose topological restrictions on the formation of the network.

An example of the use of the peer-to-peer communications topology is the cluster tree. The cluster tree network is a special case of a peer-to-peer network in which most devices are FFDs. An RFD connects to a cluster tree network as a leaf device at the end of a branch because RFDs do not allow other devices to associate. Any FFD is able to act as a coordinator and provide synchronization services to other devices or other coordinators. Only one of these coordinators is the overall PAN coordinator, potentially because it has greater computational resources than any other device in the PAN. The PAN coordinator forms the first cluster by choosing an unused PAN identifier and broadcasting beacon frames to neighboring devices. A contention resolution mechanism is required if two or more FFDs simultaneously attempt to establish themselves as PAN coordinators; however, such a mechanism is outside the scope of this standard. A candidate device receiving a beacon frame is able to request to join the network at the PAN coordinator. If the PAN coordinator permits the device to join, it adds the new device as a child device in its neighbor list. Then the newly joined device adds the PAN coordinator as its parent in its neighbor list and begins transmitting periodic beacons; other candidate devices are able to then join the network at that device. If the original candidate device is not able to join the network at the PAN coordinator, it will search for another parent device. The detailed procedures describing how a PAN is started and how devices join a PAN are found in 5.1.2 and 5.1.3.

The simplest form of a cluster tree network is a single cluster network, but larger networks are possible by forming a mesh of multiple neighboring clusters. Once predetermined application or network requirements are met, the first PAN coordinator instructs a device to become the PAN coordinator of a new cluster adjacent to the first one. Other devices gradually connect and form a multicluster network structure, such as the one seen in Figure 2. The lines in Figure 2 represent the parent-child relationships of the devices and not the communication flow. The advantage of a multicluster structure is increased coverage area, while the disadvantage is an increase in message latency.




Figure 2—Cluster tree network

4.4 Architecture

The IEEE 802.15.4 architecture is defined in terms of a number of blocks in order to simplify the standard. These blocks are called layers. Each layer is responsible for one part of the standard and offers services to the higher layers.

The interfaces between the layers serve to define the logical links that are described in this standard.

An LR-WPAN device comprises at least one PHY, which contains the radio frequency (RF) transceiver along with its low-level control mechanism, and a MAC sublayer that provides access to the physical channel for all types of transfer. Figure 3 shows these blocks in a graphical representation, which are described in more detail in 4.4.1 and 4.4.2.




Figure 3—LR-WPAN device architecture

The upper layers, shown in Figure 3, consist of a network layer, which provides network configuration, manipulation, and message routing, and an application layer, which provides the intended function of the device. The definition of these upper layers is outside the scope of this standard.

4.4.1 Physical layer (PHY)

The PHY provides two services: the PHY data service and the PHY management service. The PHY data service enables the transmission and reception of PHY protocol data units (PPDUs) across the physical radio channel. The general PHY requirements are described in Clause 8.

The features of the PHY are activation and deactivation of the radio transceiver, ED, LQI, channel selection, clear channel assessment (CCA), and transmitting as well as receiving packets across the physical medium. The UWB PHY also has the feature of precision ranging.

A discussion of the coexistence of the various IEEE 802.15.4 PHYs with other wireless systems is given in “Coexistence analysis of IEEE Std 802.15.4 with other IEEE standards and proposed standards” [B4].⁴

4.4.2 MAC sublayer

The MAC sublayer provides two services: the MAC data service and the MAC management service interfacing to the MAC sublayer management entity (MLME) service access point (SAP) (known as MLME-SAP). The MAC data service enables the transmission and reception of MAC protocol data units (MPDUs) across the PHY data service.

The features of the MAC sublayer are beacon management, channel access, GTS management, frame validation, acknowledged frame delivery, association, and disassociation. In addition, the MAC sublayer provides hooks for implementing application-appropriate security mechanisms.

Clause 5 and Clause 6 contains the specifications for the MAC sublayer.

⁴The numbers in brackets correspond to the numbers of the bibliography in Annex A.

4.5 Functional overview

A brief overview of the general functions of a LR-WPAN is given in this subclause.

4.5.1 Superframe structure

This standard allows the optional use of a superframe structure. The format of the superframe is defined by the coordinator. The superframe is bounded by network beacons sent by the coordinator, as illustrated in Figure 4a), and is divided into 16 slots of equal duration. Optionally, the superframe can have an active and an inactive portion, as illustrated in Figure 4b). During the inactive portion, the coordinator is able to enter a low-power mode. The beacon frame transmission starts at the beginning of the first slot of each superframe. If a coordinator does not wish to use a superframe structure, it will turn off the beacon transmissions. The beacons are used to synchronize the attached devices, to identify the PAN, and to describe the structure of the superframes.




Figure 4—Superframe structure

Any device wishing to communicate during the contention access period (CAP) between two beacons competes with other devices using a slotted CSMA-CA or ALOHA mechanism, as appropriate. For low-latency applications or applications requiring specific data bandwidth, the PAN coordinator dedicates portions of the active superframe to that application. These portions are called guaranteed time slots (GTSs). The GTSs form the contention-free period (CFP), which always appears at the end of the active superframe starting at a slot boundary immediately following the CAP, as shown in Figure 5. The PAN coordinator allocates up to seven of these GTSs, and a GTS is allowed to occupy more than one slot period. However, a sufficient portion of the CAP remains for contention-based access of other networked devices or new devices wishing to join the network. All contention-based transactions are completed before the CFP begins. Also each device transmitting in a GTS ensures that its transaction is complete before the time of the next GTS or the end of the CFP. More information on the superframe structure can be found in 5.1.1.1.




Figure 5—Structure of the active periods with GTSs

4.5.2 Data transfer model

Three types of data transfer transactions exist. The first one is the data transfer to a coordinator in which a device transmits the data. The second transaction is the data transfer from a coordinator in which the device receives the data. The third transaction is the data transfer between two peer devices. In star topology, only two of these transactions are used because data is exchanged only between the coordinator and a device. In a peer-to-peer topology, data is exchanged between any two devices on the network; consequently all three transactions are used in this topology.

The mechanisms for each transfer type depend on whether the network supports the transmission of periodic beacons. A beacon-enabled PAN is used in networks that either require synchronization or support for low-latency devices, such as PC peripherals. If the network does not need synchronization or support for low-latency devices, it can elect not to use the beacon for normal transfers. However, the beacon is still required for network discovery. The structure of the frames used for the data transfer is specified in 5.2.

4.5.2.1 Data transfer to a coordinator

When a device wishes to transfer data to a coordinator in a beacon-enabled PAN, it first listens for the network beacon. When the beacon is found, the device synchronizes to the superframe structure. At the appropriate time, the device transmits its data frame to the coordinator. The coordinator will acknowledge the successful reception of the data by transmitting an acknowledgment frame, if requested.

When a device wishes to transfer data in a nonbeacon-enabled PAN, it simply transmits its data frame to the coordinator. The coordinator acknowledges the successful reception of the data by transmitting an optional acknowledgment frame, completing the transaction.

4.5.2.2 Data transfer from a coordinator

When the coordinator wishes to transfer data to a device in a beacon-enabled PAN, it indicates in the network beacon that the data message is pending. The device periodically listens to the network beacon and, if a message is pending, transmits a MAC command requesting the data. The coordinator acknowledges the successful reception of the data request by transmitting an acknowledgment frame. The pending data frame is then sent by the coordinator. The device acknowledges the successful reception of the data by transmitting an acknowledgment frame, if requested. The transaction is now complete. Upon successful completion of the data transaction, the message is removed from the list of pending messages in the beacon.

When a coordinator wishes to transfer data to a device in a nonbeacon-enabled PAN, it stores the data for the appropriate device to make contact and request the data. A device requests data by transmitting a MAC command requesting the data to its coordinator. The coordinator acknowledges the successful reception of the data request by transmitting an acknowledgment frame. If a data frame is pending, the coordinator transmits the data frame. If a data frame is not pending, the coordinator indicates this fact either in the acknowledgment frame following the data request or in a data frame with a zero-length payload, as defined in 5.1.6.3. If requested, the device acknowledges the successful reception of the data frame by transmitting an acknowledgment frame.

4.5.2.3 Peer-to-peer data transfers

In a peer-to-peer PAN, every device communicates directly with every other device in its radio communications range. In order to do this effectively, the devices wishing to communicate will need to either receive constantly or synchronize with each other. In the former case, the device can simply transmit its data. In the latter case, other measures need to be taken in order to achieve synchronization. Such measures are beyond the scope of this standard.

4.5.3 Frame structure

The frame structures have been designed to keep the complexity to a minimum while at the same time making them sufficiently robust for transmission on a noisy channel. Each successive protocol layer adds to the structure with layer-specific headers and footers. This standard defines four MAC frame structures:

- A beacon frame, used by a coordinator to transmit beacons
- A data frame, used for all transfers of data
- An acknowledgment frame, used for confirming successful frame reception
- A MAC command frame, used for handling all MAC peer entity control transfers

The MAC frames are passed to the PHY as the PHY service data unit (PSDU), which becomes the PHY payload. The PPDU is illustrated in Figure 6.




Figure 6—Schematic view of the PPDU

The format of the SHR and PHR is defined for each of the PHYs in their respective clause.

The MHR and MFR are defined in 5.2.1.

4.5.4 Improving probability of successful delivery

The IEEE 802.15.4 LR-WPAN employs various mechanisms to improve the probability of successful data transmission. These mechanisms are discussed in 4.5.4.1 through 4.5.4.4.

4.5.4.1 CSMA-CA mechanism

The IEEE 802.15.4 LR-WPAN uses two types of channel access mechanism, depending on the network configuration. Nonbeacon-enabled PANs use an unslotted CSMA-CA channel access mechanism, as described in 5.1.1.4. Each time a device wishes to transmit data frames or MAC commands, it waits for a random period. If the channel is found to be idle, following the random backoff, the device transmits its data. If the channel is found to be busy following the random backoff, the device waits for another random period before trying to access the channel again. Acknowledgment frames are sent without using a CSMA-CA mechanism.

Beacon-enabled PANs use a slotted CSMA-CA channel access mechanism, where the backoff periods are aligned with the start of the beacon transmission. The backoff periods of all devices within one PAN are aligned to the PAN coordinator. Each time a device wishes to transmit data frames during the CAP, it locates the boundary of the next backoff period and then waits for a random number of backoff periods. If the channel is busy, following this random backoff, the device waits for another random number of backoff periods before trying to access the channel again. If the channel is idle, the device begins transmitting on the next available backoff period boundary. Acknowledgment and beacon frames are sent without using a CSMA-CA mechanism.

4.5.4.2 ALOHA mechanism

In the ALOHA protocol, a device transmits without sensing the medium or waiting for a specific time slot. The ALOHA mechanism is appropriate for lightly loaded networks since the probability of collision is reasonably small if the probability of clear channel is sufficiently large.

4.5.4.3 Frame acknowledgment

A successful reception and validation of a data or MAC command frame is optionally confirmed with an acknowledgment, as described in 5.1.6.4. If the receiving device is unable to handle the received data frame for any reason, the message is not acknowledged.

If the originator does not receive an acknowledgment after some period, it assumes that the transmission was unsuccessful and retries the frame transmission. If an acknowledgment is still not received after several retries, the originator can choose either to terminate the transaction or to try again. When the acknowledgment is not required, the originator assumes the transmission was successful.

4.5.4.4 Data verification

A cyclic redundancy check (CRC) is used to detect errors in every PSDU, as defined in 5.2.1.9.

4.5.5 Power consumption considerations

In many applications that use this standard, devices will be battery powered, and battery replacement or recharging in relatively short intervals is impractical. Therefore, the power consumption is of significant concern. This standard was developed with limited power supply availability in mind. However, the physical implementation of this standard will require additional power management considerations that are beyond the scope of this standard.

The protocol has been developed to favor battery-powered devices. However, in certain applications, some of these devices could potentially be mains powered. Battery-powered devices will require duty-cycling to reduce power consumption. These devices will spend most of their operational life in a sleep state; however, each device periodically listens to the RF channel in order to determine whether a message is pending. This mechanism allows the application designer to decide on the balance between battery consumption and message latency. Higher powered devices have the option of listening to the RF channel continuously.

In addition to the power saving features of the LR-WPAN system, the UWB PHY also provides a hybrid modulation that enables very simple, noncoherent receiver architectures to further minimize power consumption and implementation complexity.

4.5.6 Security

From a security perspective, wireless ad hoc networks are no different from any other wireless network. They are vulnerable to passive eavesdropping attacks and active tampering because physical access to the wire is not required to participate in communications. The very nature of ad hoc networks and their cost objectives impose additional security constraints, which perhaps make these networks the most difficult environments to secure. Devices are low-cost and have limited capabilities in terms of computing power, available storage, and power drain; and it cannot always be assumed they have a trusted computing base nor a high-quality random number generator aboard. Communications cannot rely on the online availability of a fixed infrastructure and might involve short-term relationships between devices that never have communicated before. These constraints severely limit the choice of cryptographic algorithms and protocols and influence the design of the security architecture because the establishment and maintenance of trust relationships between devices need to be addressed with care. In addition, battery lifetime and cost constraints put severe limits on the security overhead these networks can tolerate, something that is of far

less concern with higher bandwidth networks. Most of these security architectural elements can be implemented at higher layers and, therefore, are outside the scope of this standard.

The cryptographic mechanism in this standard is based on symmetric-key cryptography and uses keys that are provided by higher layer processes. The establishment and maintenance of these keys are outside the scope of this standard. The mechanism assumes a secure implementation of cryptographic operations and secure and authentic storage of keying material.

The cryptographic mechanism provides particular combinations of the following security services:

- *Data confidentiality*: Assurance that transmitted information is only disclosed to parties for which it is intended.
- *Data authenticity*: Assurance of the source of transmitted information (and, hereby, that information was not modified in transit).
- *Replay protection*: Assurance that duplicate information is detected.

The actual frame protection provided can be adapted on a frame-by-frame basis and allows for varying levels of data authenticity (to minimize security overhead in transmitted frames where required) and for optional data confidentiality. When nontrivial protection is required, replay protection is always provided.

Cryptographic frame protection uses either a key shared between two peer devices (link key) or a key shared among a group of devices (group key), thus allowing some flexibility and application-specific tradeoffs between key storage and key maintenance costs versus the cryptographic protection provided. If a group key is used for peer-to-peer communication, protection is provided only against outsider devices and not against potential malicious devices in the key-sharing group.

For more detailed information on the cryptographic security mechanisms used for protected MAC frames following this standard, refer to Clause 7.

4.6 Concept of primitives

The services of a layer are the capabilities it offers to the user in the next higher layer or sublayer by building its functions on the services of the next lower layer. This concept is illustrated in Figure 7, showing the service hierarchy and the relationship of the two correspondent users and their associated layer (or sublayer) peer protocol entities.




Figure 7—Service primitives

The services are specified by describing the information flow between the N-user and the N-layer. This information flow is modeled by discrete, instantaneous events, which characterize the provision of a service. Each event consists of passing a service primitive from one layer to the other through a layer SAP associated with an N-user. Service primitives convey the required information by providing a particular service. These

service primitives are an abstraction because they specify only the provided service rather than the means by which it is provided. This definition is independent of any other interface implementation.

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

A primitive can be one of four generic types:

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

5. MAC protocol

5.1 MAC functional description

The MAC sublayer handles all access to the physical radio channel and is responsible for the following tasks:

- Generating network beacons if the device is a coordinator
- Synchronizing to network beacons
- Supporting PAN association and disassociation
- Supporting device security
- Employing the CSMA-CA mechanism for channel access
- Handling and maintaining the GTS mechanism
- Providing a reliable link between two peer MAC entities

There are two device types: a full-function device (FFD) and a reduced-function device (RFD). The FFD may operate in three modes serving as a personal area network (PAN) coordinator, a coordinator, or a device. An RFD shall only operate as a device.

Throughout this subclause, the receipt of a frame is defined as the successful receipt of the frame by the PHY and the successful verification of the frame check sequence (FCS) by the MAC sublayer, as described in 5.2.1.9.

Constants and PAN information base (PIB) attributes that are specified and maintained by the MAC sublayer or PHY layer are written in the text in italics. Constants have a general prefix of “a”, e.g., *aBaseSlotDuration*, and the MAC constants are listed in Table 51, while the PHY constants are listed in Table 70. MAC PIB attributes have a general prefix of “mac”, e.g., *macAckWaitDuration*, and are listed in Table 52, while the security attributes are listed in Table 60. PHY PIB attributes have a general prefix of “phy”, e.g., *phyCurrentChannel*, and are listed in Table 71.

The next higher layer accesses the services provided by the MAC through the MAC sublayer management entity (MLME) service access point (SAP), as described in 6.2, and the MAC common part sublayer (MCPS) SAP, as described in 6.3. The primitives for the MLME SAP are written in all capital letters prefixed with MLME, e.g., MLME-SCAN.confirm. The primitives for the MCPS SAP are written in all capital letters prefixed with MCPS, e.g., MCPS-DATA.request.

5.1.1 Channel access

This subclause describes the mechanisms for accessing the physical radio channel.

5.1.1.1 Superframe structure

A coordinator on a PAN can optionally bound its channel time using a superframe structure. A superframe is bounded by the transmission of a beacon frame and can have an active portion and an inactive portion. The coordinator may enter a low-power (sleep) mode during the inactive portion.

The structure of this superframe is described by the values of *macBeaconOrder* and *macSuperframeOrder*. The MAC PIB attribute *macBeaconOrder* describes the interval at which the coordinator shall transmit its beacon frames. The value of *macBeaconOrder* and the beacon interval, *BI*, are related as follows:

$$BI = aBaseSuperframeDuration \times 2^{macBeaconOrder}$$

for

$$0 \leq \text{macBeaconOrder} \leq 14$$

If $\text{macBeaconOrder} = 15$, the coordinator shall not transmit beacon frames except when requested to do so, such as on receipt of a beacon request command. The value of $\text{macSuperframeOrder}$ shall be ignored if $\text{macBeaconOrder} = 15$.

The MAC PIB attribute $\text{macSuperframeOrder}$ describes the length of the active portion of the superframe, which includes the beacon frame. The value of $\text{macSuperframeOrder}$, and the superframe duration, SD , are related as follows:

$$SD = aBaseSuperframeDuration \times 2^{\text{macSuperframeOrder}}$$

for

$$0 \leq \text{macSuperframeOrder} \leq \text{macBeaconOrder} \leq 14$$

If $\text{macSuperframeOrder} = 15$, the superframe shall not remain active after the beacon. If $\text{macBeaconOrder} = 15$, the superframe shall not exist (the value of $\text{macSuperframeOrder}$ shall be ignored), and macRxOnWhenIdle shall define whether the receiver is enabled during periods of transceiver inactivity.

The active portion of each superframe shall be divided into $aNumSuperframeSlots$ equally spaced slots of duration $2^{\text{macSuperframeOrder}} \times aBaseSlotDuration$ and is composed of three parts: a beacon, a CAP, and a CFP. The beacon shall be transmitted, without the use of CSMA, at the start of slot 0, and the CAP shall commence immediately following the beacon. The start of slot 0 is defined as the point at which the first symbol of the beacon PPDU is transmitted. The CFP, if present, follows immediately after the CAP and extends to the end of the active portion of the superframe. Any allocated GTSSs shall be located within the CFP.

The MAC sublayer shall ensure that the integrity of the superframe timing is maintained, e.g., compensating for clock drift error.

PANs that wish to use the superframe structure (referred to as beacon-enabled PANs) shall set macBeaconOrder to a value between 0 and 14, both inclusive, and $\text{macSuperframeOrder}$ to a value between 0 and the value of macBeaconOrder , both inclusive.

PANs that do not wish to use the superframe structure (referred to as nonbeacon-enabled PANs) shall set both macBeaconOrder and $\text{macSuperframeOrder}$ to 15. In this case, a coordinator shall not transmit beacons, except upon receipt of a beacon request command; all transmissions, with the exception of acknowledgment frames and any data frame that quickly follows the acknowledgment of a data request command, as described in 5.1.6.3, shall use an unslotted CSMA-CA mechanism to access the channel. In addition, GTSSs shall not be permitted.

An example of a superframe structure is shown in Figure 8. In this case, the beacon interval, BI , is twice as long as the active superframe duration, SD , and the CFP contains two GTSSs.

5.1.1.1.1 Contention access period (CAP)

The CAP shall start immediately following the beacon and complete before the beginning of the CFP on a superframe slot boundary. If the CFP is zero length, the CAP shall complete at the end of the active portion of the superframe. The CAP shall be at least $aMinCAPLength$, unless additional space is needed to




Figure 8—An example of the superframe structure

temporarily accommodate the increase in the beacon frame length needed to perform GTS maintenance, as described in 5.2.2.1.3, and shall shrink or grow dynamically to accommodate the size of the CFP.

All frames, except acknowledgment frames and any data frame that quickly follows the acknowledgment of a data request command, as described in 5.1.6.3, transmitted in the CAP shall use a slotted CSMA-CA mechanism to access the channel. A device transmitting within the CAP shall ensure that its transaction is complete (i.e., including the reception of any acknowledgment) one interframe spacing (IFS) period, as described in 5.1.1.3, before the end of the CAP. If this is not possible, the device shall defer its transmission until the CAP of the following superframe.

MAC command frames shall always be transmitted in the CAP.

5.1.1.1.2 Contention-free period (CFP)

The CFP shall start on a slot boundary immediately following the CAP, and it shall complete before the end of the active portion of the superframe. If any GTSs have been allocated by the PAN coordinator, they shall be located within the CFP and occupy contiguous slots. The CFP shall therefore grow or shrink depending on the total length of all of the combined GTSs.

No transmissions within the CFP shall use a CSMA-CA mechanism to access the channel. A device transmitting in the CFP shall ensure that its transmissions are complete one IFS period, as described in 5.1.1.3, before the end of its GTS.

5.1.1.2 Incoming and outgoing superframe timing

On a beacon-enabled PAN, a coordinator that is not the PAN coordinator shall maintain the timing of both the superframe in which its coordinator transmits a beacon (the incoming superframe) and the superframe in which it transmits its own beacon (the outgoing superframe). The relative timing of these superframes is defined by the StartTime parameter of the MLME-START.request primitive, as described in 6.2.12.1 and 5.1.2.4. The relationship between incoming and outgoing superframes is illustrated in Figure 9.

If a device receives a coordinator realignment command frame from its coordinator indicating that the coordinator will begin using a new superframe configuration, the device shall ensure that its own beacons do not overlap with the beacons transmitted by the coordinator. If the new superframe configuration causes the incoming and outgoing superframes to overlap, the device shall stop transmitting its beacons immediately and notify the next higher layer via the MLME-SYNC-LOSS.indication primitive, as described in 6.2.13.2, with the LossReason parameter set to SUPERFRAME_OVERLAP.

The beacon order and superframe order shall be equal for all superframes on a PAN. All devices shall interact with the PAN only during the active portion of a superframe.




Figure 9—The relationship between incoming and outgoing beacons

5.1.1.3 Interframe spacing (IFS)

The MAC sublayer needs a finite amount of time to process data received by the PHY. To allow for this, two successive frames transmitted from a device shall be separated by at least an IFS period; if the first transmission requires an acknowledgment, the separation between the acknowledgment frame and the second transmission shall be at least an IFS period. The length of the IFS period is dependent on the size of the frame that has just been transmitted. Frames (i.e., MPDUs) of up to *aMaxSIFSFrameSize* shall be followed by a short interframe space (SIFS) period of a duration of at least *macSIFSPeriod*. Frames (i.e., MPDUs) with lengths greater than *aMaxSIFSFrameSize* shall be followed by a long interframe spacing (LIFS) period of a duration of at least *macLIFSPeriod*. These concepts are illustrated in Figure 10.




Figure 10—IFS

The CSMA-CA algorithm shall take this requirement into account for transmissions in the CAP. The timing of the transmission of an acknowledgment frame, t_{ack} , is defined in 5.1.6.3 and 5.1.6.4.2.

5.1.1.4 CSMA-CA algorithm

The CSMA-CA algorithm shall be used before the transmission of data or MAC command frames transmitted within the CAP, unless the frame can be quickly transmitted following the acknowledgment of a data request command, as defined in 5.1.6.3. The CSMA-CA algorithm shall not be used for the transmission of beacon frames in a beacon-enabled PAN, acknowledgment frames, or data frames transmitted in the CFP.

If periodic beacons are being used in the PAN, the MAC sublayer shall employ the slotted version of the CSMA-CA algorithm for transmissions in the CAP of the superframe. Conversely, if periodic beacons are not being used in the PAN or if a beacon could not be located in a beacon-enabled PAN, the MAC sublayer shall transmit using the unslotted version of the CSMA-CA algorithm. In both cases, the algorithm is

implemented using units of time called backoff periods, where one backoff period shall be equal to $aUnitBackoffPeriod$.

In slotted CSMA-CA, the backoff period boundaries of every device in the PAN shall be aligned with the superframe slot boundaries of the PAN coordinator; i.e., the start of the first backoff period of each device is aligned with the start of the beacon transmission. In slotted CSMA-CA, the MAC sublayer shall ensure that the PHY commences all of its transmissions on the boundary of a backoff period. In unslotted CSMA-CA, the backoff periods of one device are not related in time to the backoff periods of any other device in the PAN.

Each device shall maintain three variables for each transmission attempt: NB , CW , and BE . NB is the number of times the CSMA-CA algorithm was required to back off while attempting the current transmission; this value shall be initialized to zero before each new transmission attempt. CW is the contention window length, defining the number of backoff periods that need to be clear of channel activity before the transmission can commence. This value shall be initialized to CW_0 before each transmission attempt and reset to CW_0 each time the channel is assessed to be busy. For operation in the Japanese 950 MHz band, CW_0 shall be set to one; otherwise, CW_0 shall be set to two. The CW variable is only used for slotted CSMA-CA. BE is the backoff exponent, which is related to how many backoff periods a device shall wait before attempting to assess a channel. In unslotted systems, or slotted systems with the received battery life extension (BLE) field, as defined in Figure 41, set to zero, BE shall be initialized to the value of $macMinBE$. In slotted systems with the received BLE field set to one, this value shall be initialized to the lesser of two and the value of $macMinBE$. Note that if $macMinBE$ is set to zero, collision avoidance will be disabled during the first iteration of this algorithm.

Although the receiver of the device is enabled during the CCA analysis portion of this algorithm, the device may discard any frames received during this time.

Figure 11 illustrates the steps of the CSMA-CA algorithm. If the algorithm ends in “Success,” the MAC is allowed to begin transmission of the frame. Otherwise, the algorithm terminates with a channel access failure.

In a slotted CSMA-CA system with the BLE field set to zero, the MAC sublayer shall ensure that, after the random backoff, the remaining CSMA-CA operations can be undertaken and the entire transaction can be transmitted before the end of the CAP. If the number of backoff periods is greater than the remaining number of backoff periods in the CAP, the MAC sublayer shall pause the backoff countdown at the end of the CAP and resume it at the start of the CAP in the next superframe. If the number of backoff periods is less than or equal to the remaining number of backoff periods in the CAP, the MAC sublayer shall apply its backoff delay and then evaluate whether it can proceed. The MAC sublayer shall proceed if the remaining CSMA-CA algorithm steps [i.e., two CCA analyses, or a single continuous CCA analysis of at least $phyCCADuration$ for the regulatory domains that require listen before talk (LBT) such as the 950 MHz band in Japan, as described in Annex H], the frame transmission, and any acknowledgment can be completed before the end of the CAP. If the MAC sublayer can proceed, it shall request that the PHY perform the CCA in the current superframe. If the MAC sublayer cannot proceed, it shall wait until the start of the CAP in the next superframe and apply a further random backoff delay before evaluating whether it can proceed again.

In a slotted CSMA-CA system with the BLE field set to one, the MAC sublayer shall ensure that, after the random backoff, the remaining CSMA-CA operations can be undertaken and the entire transaction can be transmitted before the end of the CAP. The backoff countdown shall only occur during the first $macBattLifeExtPeriods$ full backoff periods after the end of the IFS period following the beacon. The MAC sublayer shall proceed if the remaining CSMA-CA algorithm steps [two CCA analyses, or a single continuous CCA analysis of at least $phyCCADuration$ for the regulatory domains that require listen before talk (LBT) such as the 950 MHz band in Japan, as described in Annex H], the frame transmission, and any acknowledgment can be completed before the end of the CAP, and the frame transmission will start in one of the first $macBattLifeExtPeriods$ full backoff periods after the IFS period following the beacon. If the




Figure 11—CSMA-CA algorithm

MAC sublayer can proceed, it shall request that the PHY perform the CCA in the current superframe. If the MAC sublayer cannot proceed, it shall wait until the start of the CAP in the next superframe and apply a further random backoff delay before evaluating whether it can proceed again.

5.1.2 Starting and maintaining PANs

This subclause specifies the procedures for scanning through channels, identifying PAN identifier conflicts, and starting PANs.

5.1.2.1 Scanning through channels

All devices shall be capable of performing passive and orphan scans across a specified list of channels. In addition, an FFD shall be able to perform energy detection (ED) and active scans. The next higher layer should submit a scan request for a particular channel page containing a list of channels chosen only from the channels specified by *phyChannelsSupported* for that particular channel page.

A device is instructed to begin a channel scan through the MLME-SCAN.request primitive, as described in 6.2.10.1. Channels are scanned in order from the lowest channel number to the highest. For the duration of the scan, the device shall suspend beacon transmissions, if applicable, and shall only accept frames received over the PHY data service that are relevant to the scan being performed. For UWB and CSS PHYs, each preamble code appropriate to the specified channel is scanned. Upon the conclusion of the scan, the coordinator device in a beacon-enabled PAN shall recommence beacon transmissions. The results of the scan shall be returned via the MLME-SCAN.confirm primitive as described in 6.2.10.2.

5.1.2.1.1 ED channel scan

An ED scan allows a device to obtain a measure of the peak energy in each requested channel. This could be used by a prospective PAN coordinator to select a channel on which to operate prior to starting a new PAN. During an ED scan, the MAC sublayer shall discard all frames received over the PHY data service.

An ED scan over a specified set of channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate an ED scan. For each channel, the MLME shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly, and then repeatedly perform an ED measurement for $[aBaseSuperframeDuration \times (2^n + 1)]$, where n is the value of the ScanDuration parameter in the MLME-SCAN.request primitive. The maximum ED measurement obtained during this period shall be noted before moving on to the next channel in the channel list. A device shall be able to store at least one channel ED measurement.

The ED scan shall terminate when either the number of channel ED measurements stored equals the implementation-specified maximum or energy has been measured on each of the specified channels.

5.1.2.1.2 Active and passive channel scan

An active or passive channel scan allows a device to locate any coordinator transmitting beacon frames within its radio communications range. An active scan uses the beacon request command to extract the beacon from a coordinator. In a passive scan, the beacon request command is not transmitted. A message sequence chart for active scan is illustrated in Figure 12 and for passive scan in Figure 13.

During an active or passive scan, the MAC sublayer shall discard all frames received over the PHY data service that are not beacon frames. If a beacon frame is received that contains the address of the scanning device in its list of pending addresses, the scanning device shall not attempt to extract the pending data.

Before commencing an active or passive scan, the MAC sublayer shall store the value of *macPANId* and then set it to 0xffff for the duration of the scan. This enables the receive filter to accept all beacons rather than just the beacons from its current PAN, as described in 5.1.6.2. On completion of the scan, the MAC sublayer shall restore the value of *macPANId* to the value stored before the scan began.




Figure 12—Active scan message sequence chart

An active or passive scan over a specified set of channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate an active or passive scan. For each channel, the device shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly. For an active scan, the device shall send a beacon request command, as described in 5.3.7. For UWB and CSS PHYs, the scan process shall be repeated for each mandatory preamble code, setting the *phyCurrentCode* appropriately. Upon successful transmission of the beacon request command for an active scan or after switching to the channel for a passive scan, the device shall enable its receiver for at most $[aBaseSuperframeDuration \times (2^n + 1)]$, where n is the value of the ScanDuration parameter. During this time, the device shall reject all nonbeacon frames and record the information contained in all unique beacons in a PAN descriptor structure, as described in Table 17, including the channel information and, if required, the preamble code.

If a beacon frame is received when *macAutoRequest* is set to TRUE, the list of PAN descriptor structures shall be stored by the MAC sublayer until the scan is complete; at this time, the list shall be sent to the next higher layer in the PANDescriptorList parameter of the MLME-SCAN.confirm primitive. A device shall be able to store at least one PAN descriptor. A beacon frame shall be assumed to be unique if it contains both a PAN identifier and a source address that has not been seen before during the scan of the current channel.

If a beacon frame is received when *macAutoRequest* is set to FALSE, each recorded PAN descriptor is sent to the next higher layer in a separate MLME-BEACON-NOTIFY.indication primitive as described in 6.2.4.1. A received beacon frame containing one or more octets of payload shall also cause the PAN descriptor to be sent to the next higher layer via the MLME-BEACON-NOTIFY.indication primitive. Once




Figure 13—Passive scan message sequence chart

the scan with *macAutoRequest* set to FALSE is complete, the MLME-SCAN.confirm shall be issued to the next higher layer with a null PANDescriptorList.

For UWB and CSS PHYs, the beacon request is repeated for each preamble code.

If a protected beacon frame is received, i.e., the Security Enabled field is set to one, the device shall attempt to unsecure the beacon frame using the unsecuring process described in 7.2.3.

The security-related elements of the PAN descriptor corresponding to the beacon, as defined in Table 17, shall be set to the corresponding parameters returned by the unsecuring process. The SecurityStatus element of the PAN descriptor shall be set to SUCCESS if the status from the unsecuring process is SUCCESS and set to one of the other status codes indicating an error in the security processing otherwise.

The information from the unsecured frame shall be recorded in the PAN descriptor even if the status from the unsecuring process indicated an error.

If a coordinator of a beacon-enabled PAN receives the beacon request command, it shall ignore the command and continue transmitting its periodic beacons as usual. If a coordinator of a nonbeacon-enabled PAN receives this command, it shall transmit a single beacon frame using unslotted CSMA-CA.

If *macAutoRequest* is set to TRUE, the active scan on a particular channel shall terminate when the number of beacons found equals the implementation-specified limit or the channel has been scanned for the full time, as specified in 5.1.2.1.2. If *macAutoRequest* is set to FALSE, the active scan on a particular channel

shall terminate when the channel has been scanned for the full time. If a channel was not scanned for the full time, it shall be considered to be unscanned.

If *macAutoRequest* is set to TRUE, the entire scan procedure shall terminate when the number of PAN descriptors stored equals the implementation-specified maximum or every channel in the set of available channels has been scanned. If *macAutoRequest* is set to FALSE, the entire scan procedure shall only terminate when every channel in the set of available channels has been scanned.

5.1.2.1.3 Orphan channel scan

An orphan scan allows a device to attempt to relocate its coordinator following a loss of synchronization. During an orphan scan, the MAC sublayer shall discard all frames received over the PHY data service that are not coordinator realignment command frames.

An orphan scan over a specified set of channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate an orphan scan. For each channel, the device shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly, and for UWB and CSS PHYs, setting the preamble code *phyCurrentCode* appropriately, and then send an orphan notification command, as described in 5.3.6. Upon successful transmission of the orphan notification command, the device shall enable its receiver for at most *macResponseWaitTime*. If the device successfully receives a coordinator realignment command, as described in 5.3.8, within this time, the device shall terminate the scan. For the UWB and CSS PHYs, if the coordinator realignment command is not received, the process shall be repeated for each preamble code until a realignment command is received or all preamble codes for the PHY have been used.

The orphan scan shall terminate when the device receives a coordinator realignment command or the specified set of channels has been scanned.

A example message sequence chart for orphan scan and realignment is shown in Figure 14.

If a coordinator receives the orphan notification command, the MLME shall send the MLME-ORPHAN.indication primitive, as described in 6.2.7.1, to the next higher layer. The next higher layer should then search its device list for the device indicated by the primitive. If the next higher layer finds a record of the device, it should send a coordinator realignment command to the orphaned device using the MLME-ORPHAN.response primitive, as described in 6.2.7.2, with the AssociatedMember parameter set to TRUE and the ShortAddress parameter set to the corresponding short address allocated to the orphaned device. The process of searching for the device and sending the coordinator realignment command shall occur within *macResponseWaitTime*. The coordinator realignment command shall contain its current PAN identifier, *macPANId*, its current channel and channel page, and the short address of the orphaned device. If the next higher layer of the coordinator finds no record of the device, it should send the MLME-ORPHAN.response primitive to the MLME with the AssociatedMember parameter set to FALSE.

Figure 15 illustrates the sequence of messages necessary for a coordinator to issue a notification of an orphaned device.

5.1.2.2 PAN identifier conflict resolution

In some instances a situation could occur in which two PANs exist in the same radio communications range with the same PAN identifier. If this conflict happens, the PAN coordinator and its devices shall perform the PAN identifier conflict resolution procedure.

This procedure is optional for an RFD.




Figure 14—Orphaned device realignment message sequence chart




Figure 15—Message sequence chart for orphan notification

5.1.2.2.1 Detection

The PAN coordinator shall conclude that a PAN identifier conflict is present if either of the following apply:

- A beacon frame is received by the PAN coordinator with the PAN Coordinator field, as defined in 5.2.2.1.2, set to one and the PAN identifier equal to *macPANId*.
- A PAN ID conflict notification command, as defined in 5.3.5, is received by the PAN coordinator from a device associated with it on its PAN.

A device that is associated through the PAN coordinator (i.e., *macAssociatedPANCoord* is set to TRUE) shall conclude that a PAN identifier conflict is present if the following applies:

- A beacon frame is received by the device with the PAN Coordinator field set to one, the PAN identifier equal to *macPANId*, and an address that is equal to neither *macCoordShortAddress* nor *macCoordExtendedAddress*.

A device that is associated through a coordinator that is not the PAN coordinator shall not be capable of detecting a PAN identifier conflict.

5.1.2.2.2 Resolution

On the detection of a PAN identifier conflict by a device, it shall generate the PAN ID conflict notification command, as defined in 5.3.5, and send it to its PAN coordinator. Because the PAN ID conflict notification command has the AR field set to request an acknowledgment, the PAN coordinator shall confirm its receipt by sending an acknowledgment frame. Once the device has received the acknowledgment frame from the PAN coordinator, the MLME shall issue an MLME-SYNC-LOSS.indication primitive, as described in 6.2.13.2, with the LossReason parameter set to PAN_ID_CONFLICT. If the device does not receive an acknowledgment frame, the MLME shall not inform the next higher layer of the PAN identifier conflict.

On the detection of a PAN identifier conflict by the PAN coordinator, the MLME shall issue an MLME-SYNC-LOSS.indication to the next higher layer with the LossReason parameter set to PAN_ID_CONFLICT. The next higher layer of the PAN coordinator may then perform an active scan and, using the information from the scan, select a new PAN identifier. The algorithm for selecting a suitable PAN identifier is outside the scope of this standard. If the next higher layer does select a new PAN identifier, it may then issue an MLME-START.request with the CoordRealignment parameter set to TRUE in order to realign the PAN, as described in 5.1.2.3.

5.1.2.3 Starting and realigning a PAN

This subclause specifies procedures for the PAN coordinator starting a PAN, coordinators realigning a PAN, and devices being realigned on a PAN.

5.1.2.3.1 Starting a PAN

A PAN should be started by an FFD only after having first performed a MAC sublayer reset, by issuing the MLME-RESET.request primitive, as described in 6.2.8.1, with the SetDefaultPIB parameter set to TRUE, an active channel scan, and a suitable PAN identifier selection. The algorithm for selecting a suitable PAN identifier from the list of PAN descriptors returned from the active channel scan procedure is outside the scope of this standard. In addition, an FFD should set *macShortAddress* to a value less than 0xffff.

An FFD is instructed to begin operating a PAN through the use of the MLME-START.request primitive, as defined in 6.2.12.1, with the PANCoordinator parameter set to TRUE and the CoordRealignment parameter set to FALSE. On receipt of this primitive, the MAC sublayer shall update the superframe configuration and channel parameters as specified in 5.1.2.3.4. After completing this, the MAC sublayer shall issue the MLME-START.confirm primitive, as described in 6.2.12.2, with a status of SUCCESS and begin operating as the PAN coordinator.

A message sequence chart for starting a PAN is illustrated in Figure 16.

5.1.2.3.2 Realigning a PAN

If a coordinator receives the MLME-START.request primitive, as defined in 6.2.12.1, with the CoordRealignment parameter set to TRUE, the coordinator shall attempt to transmit a coordinator

realignment command containing the new parameters for PANId, ChannelNumber, and, if present, ChannelPage.




Figure 16—PAN start message sequence chart—PAN coordinator

When the coordinator is already transmitting beacons and the CoordRealignment parameter is set to TRUE, the next scheduled beacon shall be transmitted on the current channel using the current superframe configuration, with the Frame Pending field set to one. Immediately following the transmission of the beacon, the coordinator realignment command shall also be transmitted on the current channel using CSMA-CA.

When the coordinator is not already transmitting beacons and the CoordRealignment parameter is set to TRUE, the coordinator realignment command shall be transmitted immediately on the current channel using CSMA-CA.