



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

**Part 15.4: Wireless Medium Access Control (MAC)
and Physical Layer (PHY) Specifications for
Low-Rate Wireless Personal Area Networks
(WPANs)**

Amendment 1: Add Alternate PHYs

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

IEEE
3 Park Avenue
New York, NY 10016-5997, USA
31 August 2007

IEEE Std 802.15.4a™-2007
(Amendment to
IEEE Std 802.15.4™-2006)

*Recognized as an
American National Standard (ANSI)*

IEEE Std 802.15.4a™-2007
(Amendment to
IEEE Std 802.15.4™-2006)

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

**Part 15.4: Wireless Medium Access Control (MAC)
and Physical Layer (PHY) Specifications for
Low-Rate Wireless Personal Area Networks
(WPANs)**

Amendment 1: Add Alternate PHYs

Sponsor

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

Approved 28 August 2007

American National Standards Institute

Approved 22 March 2007

IEEE-SA Standards Board

Abstract: This standard defines 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).

Keywords: ad hoc network, 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 © 2007 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 31 August 2007. Printed in the United States of America.

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

Print: ISBN 0-7381-5583-7 SH95677
PDF: ISBN 0-7381-5584-5 SS95677

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
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 802.15.4a-2007, IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements—Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANS)—Amendment 1: Addition of Alternate PHYs.

This amendment of IEEE Std 802.15.4-2006 specifies alternate physical layers (PHYs) in addition to the PHYs specified in the base standard. These alternative PHYs are as follows:

- Ultra-wide band (UWB) PHY at frequencies of 3 GHz to 5 GHz, 6 GHz to 10 GHz, and less than 1 GHz
- Chirp spread spectrum (CSS) PHY at 2450 MHz

The UWB PHY supports an over-the-air mandatory data rate of 851 kb/s with optional data rates of 110kb/s, 6.81 Mb/s, and 27.24 Mb/s. The CSS PHY supports an over-the-air data rate of 1000 kb/s and optionally 250kb/s. The PHY chosen depends on local regulations, application, and user preference.

Notice to users

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 standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents for which a license may be required by an IEEE standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention. A patent holder has filed a statement of assurance that it will grant licenses under these rights without compensation or under reasonable rates and nondiscriminatory, reasonable terms and conditions to all applications desiring to obtain such licenses. The IEEE makes no representation as to the reasonableness of rates and/or terms and conditions of the license agreements offered by patent holders. Further information may be obtained from the IEEE Standard Department.

Participants

At the time this amendment was submitted to the IEEE-SA Standards Board for approval, the IEEE P802.15 Working Group had the following voting members:

Robert F. Heile, *Chair*
Richard Alfvén, *Vice Chair*
James D. Allen, *Vice Chair*
Patrick W. Kinney, *Vice Chair*
Patrick W. Kinney, *Secretary*
Michael D. McInnis, *Assistant Secretary and Editor*

John R. Barr, *Task Group 3b Chair*
Reed Fisher, *Task Group 3c Chair*
Robert Poor, *Task Group 4b Chair*
Myung Lee, *Task Group 5 Chair*
Erik Schylander, *WNG Chair*

Patrick W. Kinney, *Task Group 4a Chair*
Vern Brethour, *Task Group 4a Co-Editor-in-Chief*
Jay Bain, *Task Group 4a Co-Editor-in-Chief*
Patrick Houghton, *Task Group 4a Secretary*
John Lampe, *Task Group 4a Editor at Large*
Vern Brethour, *Ranging Technical Editor*
Zafer Sahinoglu, *Ranging Technical vice-Editor*
Phil Orlik, *UWB co-Technical and clause 6 editor*
Ismail Lakkis, *UWB co-Technical editor*
Rainer Hach, *CSS Technical editor*
Kyung-Kuk Lee, *CSS Technical editor*
Matt Welborn, *General Description/Document Structure Technical editor*
Jay Bain, *MAC Technical editor*
Benjamin A. Rolfe, *MAC Assistant Technical editor*
Matt Welborn, *Coexistence & Regulatory Technical editor*
Camillo Gentile, *Ranging Annex Technical editor*
Michael McLaughlin, *Interoperability Technical editor*

Jon Adams
Roberto Aiello
Richard Alfvén
James Allen
Kyu Hwan An
Mikio Aoki
Yasuyuki Arai
Takashi Arita
Larry Arnett
Naiel Askar
Arthur Astrin
Yasaman Bahreini
Jay Bain
feng Bao
John Barr
Phil Beecher
Alan Berkema
Bruce Bosco
Monique Bourgeois Brown
Mark Bowles
Charles Brabenac

David Brenner
Vern Brethour
Ronald Brown
Ed Callaway
Bill Carney
Pat Carson
Kuor-Hsin Chang
Soo-Young Chang
Jonathon Cheah
Francois Chin
Kwan-Wu Chin
Sarm-Goo Cho
Sangsung Choi
Yun Choi
Chia-Chin Chong
Chun-Ting Chou
Manoj Choudhary
Celestino Corral
Robert Charles Cragie
Joe Decuir
Javier Del Prado Pavon

Kai Dombrowski
Stefan Drude
Eryk Dutkiewicz
Michael Dydyk
Amal Ekbal
Jason Ellis
Shahriar Emami
Yew Soo Eng
Paul Everest
Mark W. Fidler
Reed Fisher
Kristoffer Fleming
Amir Freund
Ricardo Gandia Sanchez
Ian Gifford
James Gilb
Eric Gnoske
Sung-Wook Goh
Sorin Goldenberg
Paul Gorday
Bernd Grohmann

Rainer Hach
 Julian Hall
 Robert C. Hall, Jr.
 Shinsuke Hara
 Jeff Harris
 Vann Hasty
 Allen Heberling
 Robert Heile
 Eric Heinze
 Barry Herold
 Jin-Meng Ho
 Patrick Houghton
 Robert Huang
 Xiaojing Huang
 Akira Ikeda
 Hideto Ikeda
 Tetsushi Ikegami
 Oyvind Janbu
 Yeong Min Jang
 Adrian Jennings
 Ho-In Jeon
 Tzyy Hong Jiang (Chiang)
 Jeyhan Karaguz
 Michael Kelly
 Stuart Kerry
 Haksun Kim
 Jae-Hyon Kim
 Jaeyoung Kim
 JinKyeong Kim
 Yongsuk Kim
 Young Hwan Kim
 Kursat Kimyacioglu
 Matthias Kindler
 Patrick Kinney
 Guenter Kleindl
 Ryuji Kohno
 Mike Krell
 Yasushi Kudo
 Haim Kupersmidt
 Yuzo Kuramochi
 Kyung Sup Kwak
 Jiun-You Lai
 Ismail Lakkis
 John Lampe
 Jim Lansford
 Colin Lanzl
 Kyung Kuk Lee
 Myung Lee
 Wooyong Lee
 David Leeper
 Henry Li
 Huan-Bang Li
 Liang Li
 Haixiang Liang
 Darryn Lowe
 Ian Macnamara
 Tadahiko Maeda
 Akira Maeki

Patricia Martigne
 Frederick Martin
 Abbie Mathew
 Taisuke Matsumoto
 Masafumi Matsumura
 Michael Mc Laughlin
 John McCorkle
 Michael McInnis
 Charlie Mellone
 Jim Meyer
 Klaus Meyer
 Akira Miura
 Samuel Mo
 Andreas Molisch
 Mark Moore
 Marco Naeve
 Ken Naganuma
 Yves Paul Nakache
 Hiroyuki Nakase
 Saishankar Nandagopalan
 Chiu Ngo
 Erwin Noble
 Mizukoshi Nobuyuki
 Masaki Noda
 John (Jay) O'Connor
 Knut Odman
 Hiroyo Ogawa
 Yasuyuki Okuma
 Philip Orlik
 Laurent Ouvry
 John Pardee
 Young Jin Park
 Nirmalendu Patra
 Dave Patton
 Xiaoming Peng
 Tony Pollock
 Robert Poor
 Clinton Powell
 Vidyasagar Premkumar
 Yihong Qi
 Raad Raad
 Ajay Rajkumar
 Pekka Ranta
 Dani Raphaeli
 Gregg Rasor
 Charles Razzell
 Joseph Reddy
 Ivan Reede
 Yuko Rikuta
 Terry Robar
 Glyn Roberts
 Richard Roberts
 Martin Rofheart
 Benjamin Rolfe
 Philippe Rouzet
 Chandos Rypinski
 Saeid Safavi
 Zafer Sahinoglu

Tomoki Saito
 Syed Saleem
 John Sarallo
 Sidney Schrum
 Erik Schylander
 Alireza Seyedi
 Sanjeev Sharma
 Stephen Shellhammer
 Siddharth Shetty
 John (Chih-Chung) Shi
 Shusaku Shimada
 Yuichi Shiraki
 Gadi Shor
 William Shvodian
 Thomas Siep
 Michael Sim
 Kazimierz Siwiak
 Zachary Smith
 V. Somayazulu
 Carl Stevenson
 Marinus Struik
 Kazuaki Takahashi
 Kenichi Takizawa
 Teik-Kheong Tan
 Mike Tanahashi
 James Taylor
 John Terry
 Arnaud Tonnerre
 Jarvis Tou
 Jerry Upton
 Robin Vaitonis
 Bart Van Poucke
 Nanci Vogtli
 Jerry Wang
 Jing Wang
 Chris Weber
 Matthew Welborn
 Richard Wilson
 Gerald Wineinger
 Andreas Wolf
 Marcus Wong
 Stephen Wood
 Patrick Worfolk
 Tracy Wright
 Xiaodong Wu
 Yu-Ming Wu
 Hirohisa Yamaguchi
 Kamya Yekeh Yazdandoost
 Su-Khiong Yong
 Zhan Yu
 Serdar Yurdakul
 Honggang Zhang
 Bin Zhen
 Frank Xiaojun Zheng
 Zheng Zhou
 Chunhui Zhu

Major contributions were received from the following individuals:

Rick Alfvin	John Lampe	Richard Roberts
Jay Bain	Colin Lanzl	Benjamin A. Rolfe
Vern Brethour	Kyung-Kuk Lee	Phillippe Rouzet
Francois Chin	Huan-Bang Li	Saeid Safavi
Chia-Chin Chong	Gian Mario Maggio	Zafer Sahinoglu
Celestino Corral	Patricia Martigne	Ulrich Schuster
Jason Ellis	Fred Martin	Jean Schwoerer
Camillo Gentile	Lars Menzer	Enami Shahriar
Ian Gifford	Michael McLaughlin	Kai Siwiak
Rainer Hach	Andy Molisch	Arnaud Tonnerre
Patrick Houghton	Philip Orlik	Matthew Welborn
Ryuji Kohno	Laurent Ouvry	Si-Khiong Yong
Ismail Lakkis	Yihong Qi	Bin Zhen

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

Richard L. Alfvin	Atsushi Ito	Donald M. Parker
Butch Anton	Raj Jain	Gregory D. Peterson
Mikio Aoki	Oyvind Janbu	Clinton C. Powell
Lee R. Armstrong	Bobby Jose	Chuck Powers
Brian J. Baas	Shinkyo Kaku	Vikram Punj
Jay C. Bain	Stuart J. Kerry	Jose Puthenkulam
John R. Barr	Yongbum Kim	Robert A. Robinson
Gennaro Boggia	Patrick W. Kinney	Randall M. Safier
Monique J. Bourgeois	Joseph J. Kubler	Osman Sakr
Vern Brethour	Yasushi Kudoh	John H. Santhoff
Sean S. Cai	John V. Lampe	Bartien Sayogo
James T. Carlo	Jeremy A. Landt	Nicoll B. Shepherd
Juan C. Carreon	Kyung-Kuk Lee	William M. Shvodian
Yawgeng A. Chau	Daniel G. Levesque	Rajnesh D. Singh
Naftali Chayat	Jan-Ray Liao	Matthew L. Smith
Elizabeth Chesnutt	Chiwoo Lim	Jung Je Son
Aik Chindapol	Daniel M. Lubar	Amjad A. Soomro
Keith Chow	G. L. Luri	Thomas E. Starai
Bryan P. Cook	Nicolai V. Malykh	Rene Struik
Todor V. Cooklev	Jon S. Martens	Walter Struppler
Tommy P. Cooper	Gustaf S. Max	Gerald J. Stueve
Robert C. Cragie	Michael D. Mcinnis	Mark A. Sturza
Joseph Decuir	Francisco J. Melendez	Mark A. Tillinghast
Javier Del-Prado-Pavon	Gary L. Michel	Mark-Rene Uchida
Thomas J. Dineen	Apurva N. Mody	Dmitri Varsanofiev
Paul S. Eastman	Andreas F. Molisch	Stanley S. Wang
C. J. Fitzgerald	Ronald G. Murias	Mattias Wennstrom
Avraham Freedman	Juichi Nakada	Martin S. Wilcox
Devon L. Gayle	Michael S. Newman	Andreas C. Wolf
Theodore Georgantas	Richard H. Noens	Yunsong Yang
Randall C. Groves	Satoshi Obara	Oren Yuen
C. G. Guy	Knut T. Odman	Surong Zeng
Rainer Hach	Philip V. Orlik	

When the IEEE-SA Standards Board approved this revision on 22 March 2007, it had the following membership:

Steve M. Mills, *Chair*
Robert M. Grow, *Vice Chair*
Don Wright, *Past Chair*
Judith Gorman, *Secretary*

Richard DeBlasio
Julian Forster*
Alex Gelman
William R. Goldbach
Arnold M. Greenspan
Joanna N. Guenin
Kenneth S. Hanus
William B. Hopf

Richard H. Hulett
Hermann Koch
Joseph L. Koepfinger*
John Kulick
David J. Law
Glenn Parsons
Ronald C. Petersen
Tom A. Prevost

Narayanan Ramachandran
Greg Ratta
Robby Robson
Anne-Marie Sahazizian
Virginia C. Sulzberger*
Malcolm V. Thaden
Richard L. Townsend
Howard L. Wolfman

*Member Emeritus

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

Satish K. Aggarwal, *NRC Representative*
Alan H. Cookson, *NIST Representative*
Virginia C. Sulzberger, *Member/TAB Representative*

Don Messina
IEEE Standards Program Manager, Document Development

Michael Kipness
IEEE Standards Program Manager, Technical Program Development

Contents

1.	Overview.....	1
1.2	Scope.....	1
1.3	Purpose.....	1
3.	Definitions	2
4.	Acronyms and abbreviations	2
5.	General description	3
5.1	Introduction.....	3
5.4	Architecture	4
5.4.1	Physical layer (PHY)	4
5.4.1.1	Advantages of the UWB PHY for LR-WPAN	5
5.4.1.2	Advantages of the CSS (2450 MHz) PHY for LR-WPAN	5
5.4.1.3	UWB band coexistence.....	5
5.5	Functional overview	6
5.5.1	Superframe structure.....	6
5.5.2	Data transfer model.....	6
5.5.2.1	Data transfer to a coordinator	6
5.5.2.2	Data transfer from a coordinator.....	6
5.5.2.3	Peer-to-peer data transfers	7
5.5.4	Improving probability of successful delivery	7
5.5.4.1a	ALOHA mechanism for the UWB device.....	7
5.5.4.4	Enhanced robustness features for the UWB PHY	7
5.5.5	Power consumption considerations	7
5.5.5.1	Additional power saving features provided by the UWB PHY	7
5.5.6	Security	8
5.5.7	General overview of ranging	8
5.5.7.1	Two-way ranging	8
5.5.7.2	Position awareness through one-way transmissions	12
5.5.7.3	The ranging counter	12
5.5.7.4	Accounting for signal arrival time	13
5.5.7.5	Management of crystal offsets	16
5.5.7.6	Accounting for internal propagation paths	18
5.5.7.7	Timestamp reports	19
5.5.7.8	Private ranging	20
5.5.8	Management of UWB options	20
5.5.8.1	Overview of UWB options	21
5.5.8.2	Modes and options for low-cost UWB devices	21
5.5.8.3	Rules for use of UWB modes and options.....	22
5.5.8.4	The optional low data rate	23
6.	PHY specification	25
6.1	General requirements and definitions	25
6.1.1	Operating frequency range.....	25
6.1.2	Channel assignments.....	26
6.1.2.1	Channel numbering.....	26
6.1.2.1a	Channel numbering for CSS PHY	26
6.1.2.1b	Channel numbering for UWB PHY	27

6.1.2.2	Channel pages	28
6.1.3	Minimum long interframe spacing (LIFS) and short interframe spacing (SIFS) periods	28
6.1.5	Transmit power	28
6.2	PHY service specifications	29
6.2.1	PHY data service	29
6.2.1.1	PD-DATA.request	29
6.2.1.2	PD-DATA.confirm	30
6.2.1.3	PD-DATA.indication	32
6.2.2	PHY management service.....	33
6.2.2.4	PLME-ED.confirm	34
6.2.2.7	PLME-SET-TRX-STATE.request.....	34
6.2.2.11	PLME-DPS.request (UWB PHYs only)	35
6.2.2.12	PLME-DPS.confirm (UWB PHYs only).....	36
6.2.2.13	PLME-SOUNDING.request (UWB PHYs only)	37
6.2.2.14	PLME-SOUNDING.confirm (UWB PHYs only)	37
6.2.2.15	PLME-CALIBRATE.request (UWB PHYs only).....	38
6.2.2.16	PLME-CALIBRATE.confirm (UWB PHYs only).....	39
6.2.3	PHY enumerations description	40
6.3	PPDU format.....	41
6.3.1	Preamble field	42
6.3.2	SFD field.....	42
6.4	PHY constants and PIB attributes.....	43
6.4.2	PIB Attributes	43
6.4.2.1	PIB values phyMaxFrameDuration, phySHRDuration for UWB	47
6.4.2.2	PIB values phyMaxFrameDuration for CSS	48
6.4.2.3	PIB values for internal propagation times UWB	48
6.5	2450 MHz PHY specifications	49
6.5a	2450 MHz PHY chirp spread spectrum (CSS) PHY	49
6.5a.1	Data rates	49
6.5a.2	Modulation and spreading	49
6.5a.2.1	Reference modulator diagram.....	49
6.5a.2.2	De-multiplexer (DEMUX).....	49
6.5a.2.3	Serial-to-parallel mapping (S/P)	50
6.5a.2.4	Data-symbol-to-bi-orthogonal-codeword mapping	50
6.5a.2.5	Parallel-to-serial converter (P/S) and QPSK symbol mapping.....	53
6.5a.2.6	DQPSK coding	54
6.5a.2.7	DQPSK-to-DQCSK modulation.....	54
6.5a.2.8	CSK generator.....	55
6.5a.2.9	Bit interleaver	55
6.5a.3	CSS frame format	55
6.5a.3.1	Preamble	55
6.5a.3.2	SFD field.....	56
6.5a.3.3	PHY header (PHR)	56
6.5a.4	Waveform and subchirp sequences.....	56
6.5a.4.1	Graphical presentation of chirp symbols (subchirp sequences).....	56
6.5a.4.2	Active usage of time gaps	56
6.5a.4.3	Mathematical representation of the continuous time CSS base-band signal.....	58
6.5a.4.4	Raised cosine window for chirp pulse shaping.....	59
6.5a.4.5	Subchirp transmission order	60
6.5a.4.6	Example of CSK signal generation.....	60
6.5a.5	2450 MHz band CSS radio specification.....	61
6.5a.5.1	Transmit power spectral density (PSD) mask and signal tolerance.....	61

	6.5a.5.2	Symbol rate	61
	6.5a.5.3	Receiver sensitivity	61
	6.5a.5.4	Receiver jamming resistance	61
6.8		868/915 MHz band (optional) O-QPSK PHY specification	62
6.8a		UWB PHY specification	62
	6.8a.1	UWB frame format	63
	6.8a.2	PPDU encoding process	64
	6.8a.3	UWB PHY symbol structure	65
	6.8a.4	PSDU timing parameters	65
	6.8a.4.1	Channel Number parameter	68
	6.8a.4.2	Peak PRF MHz parameter	68
	6.8a.4.3	Bandwidth MHz parameter	68
	6.8a.4.4	Preamble Code Length parameter	68
	6.8a.4.5	Viterbi Rate parameter	68
	6.8a.4.6	RS Rate parameter	68
	6.8a.4.7	Overall FEC Rate parameter	68
	6.8a.4.8	Burst Positions per Symbol parameter	68
	6.8a.4.9	Hop Bursts parameter	68
	6.8a.4.10	Chips per Burst parameter	69
	6.8a.4.11	Burst Duration parameter	69
	6.8a.4.12	Symbol Duration parameter	69
	6.8a.4.13	Symbol Rate parameter	69
	6.8a.4.14	Bit Rate parameter	69
	6.8a.4.15	Mean PRF parameter	69
	6.8a.5	Preamble timing parameters	69
	6.8a.6	SHR preamble	71
	6.8a.6.1	SHR SYNC field	72
	6.8a.6.2	SHR SFD	74
	6.8a.7	PHY header (PHR)	74
	6.8a.7.1	PHR rate, length, ranging, extension, preamble duration fields	75
	6.8a.7.2	PHR SECDED check bits	76
	6.8a.8	Data field	76
	6.8a.9	UWB PHY modulation	76
	6.8a.9.1	UWB PHY modulation mathematical framework	76
	6.8a.9.2	UWB PHY spreading	77
	6.8a.10	UWB PHY forward error correction (FEC)	78
	6.8a.10.1	Reed-Solomon encoding	79
	6.8a.10.2	Systematic convolutional encoding	80
	6.8a.11	PMD operating specifications	80
	6.8a.11.1	Operating frequency bands	80
	6.8a.11.2	Channel assignments	81
	6.8a.11.3	Regulatory compliance	81
	6.8a.11.4	Operating temperature range	81
	6.8a.12	Transmitter specification	82
	6.8a.12.1	Baseband impulse response	82
	6.8a.12.2	Transmit PSD mask	84
	6.8a.12.3	Chip rate clock and chip carrier alignment	84
	6.8a.13	UWB PHY optional pulse shapes	84
	6.8a.13.1	UWB PHY optional chirp on UWB (CoU) pulses	84
	6.8a.13.2	UWB PHY optional continuous spectrum (CS) pulses	86
	6.8a.13.3	UWB PHY linear combination of pulses (LCP)	87
	6.8a.14	Extended preamble for optional UWB CCA mode	87
	6.8a.15	Ranging	88
	6.8a.15.1	Ranging counter	88

6.8a.15.2	Crystal characterization	89
6.8a.15.3	Ranging FoM	89
6.9	General radio specifications	91
6.9.1	TX-to-RX turnaround time	91
6.9.3	Error-vector magnitude (EVM) definition	91
6.9.4	Transmit center frequency tolerance	91
6.9.5	Transmit power	91
6.9.6	Receiver maximum input level of desired signal	92
6.9.7	Receiver ED	92
6.9.9	CCA	92
7.	MAC sublayer specification	95
7.1	MAC sublayer service specification	95
7.1.1	MAC data service	95
7.1.1.1	MCPS-DATA.request	95
7.1.1.2	MCPS-DATA.confirm	96
7.1.1.3	MCPS-DATA.indication	98
7.1.2	MAC management service	99
7.1.5	Beacon notification primitive	99
7.1.5.1	MLME-BEACON-NOTIFY.indication	99
7.1.10	Primitives for specifying the receiver enable time	100
7.1.10.1	MLME-RX-ENABLE.request	100
7.1.10.2	MLME-RX-ENABLE.confirm	100
7.1.11	Primitives for channel scanning	101
7.1.11.2	MLME-SCAN.confirm	101
7.1.16	Primitives for requesting data from a coordinator	101
7.1.16a	Primitives for specifying dynamic preamble (for UWB PHYs)	101
7.1.16a.1	MLME-DPS.request	102
7.1.16a.2	MLME-DPS.confirm	103
7.1.16a.3	MLME-DPS.indication	104
7.1.16b	Primitives for channel sounding (for UWB PHYs)	104
7.1.16b.1	MLME-SOUNDING.request (UWB PHYs only)	104
7.1.16b.2	MLME-SOUNDING.confirm (UWB PHYs only)	105
7.1.16c	Primitives for ranging calibration (for UWB PHYs)	106
7.1.16c.1	MLME-CALIBRATE.request (UWB PHYs only)	106
7.1.16c.2	MLME-CALIBRATE.confirm (UWB PHYs only)	106
7.1.17	MAC enumeration description	108
7.4	MAC constants and PIB attributes	110
7.4.2	MAC PIB attributes	110
7.5	MAC functional description	111
7.5.2	Starting and maintaining PANs	111
7.5.2.1	Scanning through channels	111
7.5.7	GTS allocation and management	113
7.5.7a	Ranging	113
7.5.7a.1	Set-up activities before a ranging exchange	113
7.5.7a.2	Finish-up activities after a ranging exchange	114
7.5.7a.3	Managing DPS	114
7.5.7a.4	The ranging exchange	115
Annex D (normative)	Protocol implementation conformance statement (PICS) proforma	117
D.1	Introduction	117
D.1.2	Scope	117

D.1.2	Purpose	117
D.5	Identification of the protocol	117
D.7	PICS proforma tables	117
D.7.2	Major capabilities for the PHY	117
D.7.3	Major capabilities for the MAC sublayer	120
Annex D1 (informative)	Location topics	121
D1.1	Overview	121
D1.2	Time-of-arrival estimation from channel sounding	121
D1.2.1	Time-of-arrival estimation in non-line-of-sight (NLOS) conditions	123
D1.3	Asynchronous ranging	124
D1.3.1	Two-way ranging (TWR)	124
D1.3.2	Symmetric double-sided two-way ranging (SDS-TWR)	126
D1.4	Location estimation from range data	127
D1.4.1	Time of arrival	127
D1.4.2	Time difference of arrival	128
D1.5	Network location algorithms	129
D1.5.1	Ad hoc algorithms	130
D1.5.2	Centralized algorithms	132
D1.5.3	Convex optimization algorithms	133
D1.5.4	Location estimation using multipath delays	134
D1.5.5	Reduced dimension approach to bad geometric dilution of precision (GDOP) problem	135
Annex E (informative)	Coexistence with other IEEE standards and proposed standards	139
E.1	Introduction	139
E.2	Standards and proposed standards characterized for coexistence	139
E.3	General coexistence issues	140
E.3.2	Modulation	140
E.3.4	Low duty cycle	140
E.3.5	Low transmit power	142
E.3.6	Channel alignment	143
E.3.7	Dynamic channel selection	143
E.4	2400 MHz band coexistence performance (except for CSS PHYs)	143
E.5	800/900 MHz bands coexistence performance	144
E.6	2400 MHz band coexistence performance for CSS PHYs	144
E.6.1	Assumptions for coexistence performance	144
E.6.2	Coexistence simulation results	146
E.6.3	Low-duty-cycle assumption	148
E.6.4	Impact of increased data rate	148
E.6.5	Co-channel scenario	148
E.7	UWB coexistence performance	159
E.7.1	Specific regulatory requirements for UWB coexistence	159
E.7.2	Mitigation of interference from UWB devices using low PAN duty cycles	159
E.7.3	Coexistence assurance: methodology and assumptions	159
E.7.4	UWB PHY coexistence	160
E.7.5	Path loss model	161
E.7.6	BER as a function of SIR	162
E.7.7	Temporal model	163
E.7.8	Coexistence analysis	163
E.7.9	Impact of IEEE 802.15.4a devices on ECMA 368 networks	169
E.7.10	Impact of IEEE 802.15.4a devices on IEEE P802.22 networks	171

E.7.11	Conclusions	171
E.8	~~E.6~~ Notes on the calculations.....	171
Annex F (informative) ~~IEEE 802.15.4~~ Regulatory requirements.....		173
F.1	IEEE Std 802.15.4	173
F.1.1	~~F.1~~ Introduction	173
F.1.2	~~F.2~~ Applicable U.S. (FCC) rules	173
F.1.3	~~F.3~~ Applicable European rules	173
F.1.4	~~F.4~~ Applicable Japanese rules	173
F.1.5	~~F.5~~ Emissions specification analysis with respect to known worldwide regulations	173
F.1.6	~~F.6~~ Summary of out-of-band spurious emissions limits	173
F.1.7	~~F.7~~ Phase noise requirements inferred from regulatory limits	173
F.1.8	~~F.8~~ Summary of transmission power levels	174
F.2	IEEE 802.15.4a UWB.....	174
F.2.1	Introduction	174
F.2.2	Applicable U.S. (FCC) rules	174
F.2.3	Applicable European rules	175
F.2.4	Applicable Japanese rules	177
Annex H (informative) UWB PHY optional chaotic pulses.....		181
Annex I (informative) Example UWB PHY transmit data frame encoding.....		183
I.1	Channel used in the example	183
I.2	Encoding progression	183
I.2.1	Transmit PSDU	183
I.2.2	PSDU bits.....	183
I.2.3	Reed-Solomon encoded bits.....	184
I.2.4	Convolutional encoder input bits	184
I.2.5	Convolutional encoder output bits	184
I.2.6	Scrambler output bits	184
I.2.7	Ternary output symbols.....	185

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

**Part 15.4: Wireless Medium Access Control (MAC)
and Physical Layer (PHY) Specifications for
Low-Rate Wireless Personal Area Networks
(WPANs)**

Amendment 1: Add Alternate PHYs

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

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

1. Overview

1.2 Scope

Insert the following new paragraph at the end of 1.2:

In addition, alternative physical layers (PHYs) for data communication devices with precision ranging, extended range, and enhanced robustness and mobility are specified.

1.3 Purpose

Insert the following new paragraph at the end of 1.3:

This standard also provides an international standard for an ultra-low complexity, ultra-low cost, ultra-low power consumption alternate PHY for IEEE Std 802.15.4™ (comparable to the goals for IEEE Std 802.15.4-2003). To satisfy an evolutionary set of industrial and consumer requirements for wireless personal area network (WPAN) communications, the precision ranging capability will be accurate to one meter or better,

and the communication range, robustness, and mobility improved over IEEE Std 802.15.4-2003. The requirements to support coexisting networks of sensors, controllers, and logistic and peripheral devices in multiple compliant co-located systems are addressed.

3. Definitions

Insert the following new definition alphabetically into Clause 3:

3.6a burst: A group of ultra-wide band (UWB) pulses occurring at consecutive chip periods.

3.6b chirp: Linear frequency sweep (frequency may either increase or decrease).

3.6c chirp symbol: One subchirp sequence followed by a time gap.

3.6d complex channel: A combination of a channel [radio frequency (RF) center frequency] and a code that applies to ultra-wide band (UWB) and chirp spread spectrum (CSS) physical layer (PHY) types. For UWB, code is a ternary code sequence, and for CSS, a subchirp sequence.

3.16a hybrid modulation: The modulation used in the ultra-wide band (UWB) physical layer (PHY) that combines both binary phase-shift keying (BPSK) and pulse position modulation (PPM) so that both coherent and noncoherent receivers can be used to demodulate the signal.

3.24a mean pulse repetition frequency (PRF): The total number of pulses within a symbol divided by the symbol duration.

3.32a peak pulse repetition frequency (PRF): The maximum rate at which a ultra-wide band (UWB) physical layer (PHY) emits pulses.

3.37a ranging-capable device (RDEV): A device containing an implementation capable of supporting ranging. As a practical matter, it means that a ultra-wide band (UWB) device supports the ranging counter.

3.37b ranging counter: An abstraction used to characterize the behavior of a ranging-capable device (RDEV) as it produces ranging counter values.

3.37c ranging figure of merit (FoM): A single octet that expresses the quality of a ranging measurement.

3.37d ranging frame (RFRAME): A ultra-wide band (UWB) frame having the ranging bit set in the physical layer (PHY) header (PHR).

3.37e ranging marker (RMARKER): The first ultra-wide band (UWB) pulse of the first bit of the physical layer (PHY) header (PHR) of a ranging frame (RFRAME).

3.41a solver: The node in a ranging network that computes relative positions from timestamp reports.

3.41b subband: The frequency band that can be either lower or upper half of the total occupied band.

3.41c subchirp: The chirp signal with amplitude shaping that occupies one of the two subbands.

3.41d subchirp sequence: A sequence of four subchirps.

4. Acronyms and abbreviations

Insert the following abbreviations alphabetically into Clause 4:

ADC	analog-to-digital converter
AF	activity factor
AGC	automatic gain control

AM	amplitude modulation
AWN	affected wireless network
BPM	burst position modulation
COOK	chaotic on-off keying
CoU	chirp on UWB
CS	continuous spectrum
CSK	chirp-shift keying (The phases of subsequent chirps are modulated.)
CSS	chirp spread spectrum
DAA	detect and avoid
DEMUX	de-multiplexer
DPS	dynamic preamble selection
DQCSK	differential quadrature chirp-shift keying (The phases of subsequent chirps are modulated with DQPSK values.)
DQPSK	differential quadrature phase-shift keying
ERP	extended rate PHY conforming to Clause 19 of IEEE Std 802.11™-2007
FBWA	fixed broadband wireless access
FEC	forward error correction
FoM	figure of merit
GDOP	geometric dilution of precision
HR/DSSS	High Rate direct sequence spread spectrum (see Clause 18 of IEEE Std 802.11-2007)
IWN	interfering wireless network
LCP	linear combination of pulses
LFSR	linear feedback shift register
LOS	line-of-sight
NLOS	non-line-of-sight
PLL	phase-locked loop
PPM	pulse position modulation
PRBS	pseudo-random binary sequence
PRF	pulse repetition frequency
RDEV	ranging-capable device
RFRAME	ranging frame
RMARKER	ranging marker
RSSI	receive signal strength indicator
SDS-TWR	symmetric double-sided two-way ranging
SOP	simultaneously operating piconets
U-NII	unlicensed national information infrastructure
UWB	ultra-wide band

5. General description

5.1 Introduction

Change the following items in the dashed list of 5.1 as shown:

- Over-the-air data rates of 851 kb/s, 250 kb/s, 110 kb/s, 40 kb/s, and 20 kb/s
- Carrier sense multiple access with collision avoidance (CSMA-CA) or ALOHA [ultra-wide band (UWB)] channel access
- 16 channels in the 2450 MHz band, 30 channels in the 915 MHz band, ~~and~~ 3 channels in the 868 MHz band, 14 overlapping chirp spread spectrum (CSS) channels in the 2450 MHz band, and 16 channels in three UWB bands (500 MHz and 3.1 GHz to 10.6 GHz).

Insert the following new paragraph at the end of 5.1:

In addition, two optional PHYs are specified. A UWB PHY with optional ranging is one option while a CSS PHY operating in the 2450 MHz band is the second.

5.4 Architecture

5.4.1 Physical layer (PHY)

Change the third paragraph in 5.4.1 as shown:

The features of the PHY are activation and deactivation of the radio transceiver, energy detection (ED), link quality indication (LQI), channel selection, clear channel assessment (CCA), and transmitting as well as receiving packets across the physical medium. The optional UWB PHY also has the optional feature of precision ranging. The radio operates at one or more of the following unlicensed bands:

- 868–868.6 MHz (e.g., Europe)
- 902–928 MHz (e.g., North America)
- 2400–2483.5 MHz (worldwide)
- 3100–10 600 MHz (UWB varies by region)

Insert the following new text at the end of 5.4.1:

Low-rate WPANs (LR-WPANs) can operate in multiple independent license-free bands. A LR-WPAN device can be implemented in a single band or multiple bands, but in each band implemented it supports the mandatory channel set for that band to ensure interoperability for devices that share a common band. For UWB devices, there are three independent bands: the sub-gigahertz band (250–750 MHz), the low band (3.1–5 GHz), and the high band (6–10.6 GHz). Each UWB band has a single mandatory channel, and devices in each band operate independently of the other band. Devices in the three different UWB bands use the same bandwidths and chipping rates to simplify design and implementation, with each band having different performance and regulatory constraints in different regions of the world. The three different UWB bands provide flexibility to allow LR-WPAN devices to operate in different regions as the UWB regulations are defined and updated over time.

The specification for UWB LR-WPAN devices also incorporates a number of optional enhancements to potentially improve performance, reduce power consumption, or enhance coexistence characteristics. These optional enhancements

- Do not compromise the existing LR-WPAN model so that all devices operating in a common band will always be able to interoperate with a single default mandatory mode.
- Do not raise the baseline complexity for compliant devices, but recognize that some LR-WPAN applications or implementations may need enhanced performance or coexistence capabilities while still maintaining full interoperability.
- Provide the capability to UWB LR-WPAN devices to operate under a wider range of radio frequency (RF) channel conditions while still providing robust performance and precision ranging.

In general, the mechanisms for managing enhancement options and modes are out of scope for this standard to ensure maximal flexibility while still supporting interoperability.

CSS is a spread spectrum system that is similar to both direct sequence spread spectrum (DSSS) and UWB and offers some significant properties in addition to these systems due to the different modulation methods.

A chirp is a linear frequency modulated pulse. It could be thought of as sweeping the band at a very high speed. The type of CSS system defined for this standard uses patterns of smaller chirps, or “subchirps,” to build one larger chirp symbol. This allows multiple networks to use the same frequency channel simultaneously and also offers more robust performance. This CSS system also uses differential quadrature phase-shift keying (DQPSK) modulation to further enhance performance.

The channel plan for the 2450 MHz CSS was chosen to be identical to that of IEEE 802.11 High Rate (HR)/DSSS and IEEE 802.11 extended rate PHY (ERP) systems in order to enhance coexistence. This channel plan offers primary nonoverlapping channels and secondary partially overlapping channels, which vary by regulatory region. Since CSS can benefit more fully from more spectrum, CSS did not use the narrower channels used by 2450 MHz DSSS.

Insert after 5.4.1 the following new subclauses (5.4.1.1 through 5.4.1.3):

5.4.1.1 Advantages of the UWB PHY for LR-WPAN

The UWB LR-WPAN specification is designed to provide robust performance for LR-WPAN applications while leveraging the unique capability of UWB waveforms to support precision ranging between devices. The UWB PHY design is intended to make use of the wide bands of spectrum being made available for UWB operation around the world. This spectrum, combined with advances in low-cost and low-power process technology, enables the implementation of LR-WPAN devices that can provide enhanced resistance to multipath fading for robust performance with very low transmit power.

5.4.1.2 Advantages of the CSS (2450 MHz) PHY for LR-WPAN

The CSS LR-WPAN specification is designed to provide robust performance for LR-WPAN applications while leveraging the unique capability of CSS waveforms to support long-range links or to support links to mobile devices moving at higher speeds. The CSS PHY is intended to take advantage of the global deployability of the 2450 MHz band due to favorable regulations, both indoors and out, while offering enhanced robustness, range, and mobility. In addition to the robustness mechanisms described in 5.5, the properties of CSS give LR-WPAN devices enhanced immunity to multipath fading and extended range for robust performance with very low transmit power.

Two data rates are specified for CSS in order to offer implementers the flexibility to select the rate and properties best suited for their applications, as the following guidelines illustrate:

- The lower/coded rate would be appropriate in quiet additive white Gaussian noise (AWGN) and high multipath environments.
- The higher rate would be appropriate for low-energy consumption and burst interference environments.

5.4.1.3 UWB band coexistence

LR-WPAN devices using UWB bands operate in spectrum different from other PHYs that use unlicensed spectrum. For this reason, it is important that UWB LR-WPAN devices provide good coexistence performance with respect to other systems using spectrum overlaid by the UWB bands. Due to these special considerations, a number of extra features have been included with the UWB PHY design to support coexistence with other spectrum users as well as with other UWB systems.

The UWB PHY provides the following coexistence features:

- Low power spectral density (PSD) in accordance with regulations for UWB in different parts of the world, including unprecedented low out-of-band emission requirements.
- Multiple bands and operating frequencies within each band to allow devices the option to avoid bands that might be in use or otherwise unavailable.
- Optional modes to operate with shorter symbol timing to minimize emissions and channel occupancy when applications and channel conditions allow.
- Specific commands that provide a basic framework to allow higher layers to control the radio for coexistence functions and possible interference mitigation.
- Optional spectral control features based on pulse shaping to allow enhanced coexistence with other spectrum users.

Further details of these coexistence features and the expected potential impact of UWB LR-WPAN devices is included in Annex E.

5.5 Functional overview

Change the text of 5.5 as shown:

A brief overview of the general functions of a LR-WPAN is given in 5.5.1 through ~~5.5.6~~ 5.5.8 and includes information on the superframe structure, the data transfer model, the frame structure, ~~improving probability of successful delivery, robustness,~~ power consumption considerations, precision ranging, and security.

5.5.1 Superframe structure

Change the next to last sentence of the first paragraph in 5.5.1 as shown:

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.

5.5.2 Data transfer model

5.5.2.1 Data transfer to a coordinator

Change the text of 5.5.2.1 as shown (the figures remain unchanged):

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, using slotted CSMA-CA or ALOHA, as appropriate, to the coordinator. The coordinator may acknowledge the successful reception of the data by transmitting an optional acknowledgment frame. This sequence is summarized in Figure 6.

When a device wishes to transfer data in a nonbeacon-enabled PAN, it simply transmits its data frame, using unslotted CSMA-CA or ALOHA, as appropriate, to the coordinator. The coordinator acknowledges the successful reception of the data by transmitting an optional acknowledgment frame. The transaction is now complete. This sequence is summarized in Figure 7.

5.5.2.2 Data transfer from a coordinator

Change the text of 5.5.2.2 as shown (the figures remain unchanged):

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, using slotted CSMA-CA or ALOHA, as appropriate. The coordinator acknowledges the successful reception of the data request by transmitting an acknowledgment frame. The pending data frame is then sent using slotted CSMA-CA or ALOHA, as appropriate, or, if possible, immediately after the acknowledgment (see 7.5.6.3). The device may acknowledge the successful reception of the data by transmitting an optional acknowledgment frame. 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. This sequence is summarized in Figure 8.

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 may make contact by transmitting a MAC command requesting the data, using unslotted CSMA-CA or ALOHA, as appropriate, to its coordinator at an application-defined rate. 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, using unslotted CSMA-CA or ALOHA, as appropriate, to the device. 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 (see 7.5.6.3). If requested, the device acknowledges the successful reception of the data frame by transmitting an acknowledgment frame. This sequence is summarized in Figure 9.

5.5.2.3 Peer-to-peer data transfers

Change the text of 5.5.2.3 as shown:

In a peer-to-peer personal area network (PAN), every device may communicate with every other device in its radio sphere of influence. 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 using unslotted CSMA-CA or ALOHA, as appropriate. In the latter case, other measures need to be taken in order to achieve synchronization. Such measures are beyond the scope of this standard.

5.5.4 Improving probability of successful delivery

Insert between 5.5.4.1 and 5.5.4.2 the following new subclause (5.5.4.1a):

5.5.4.1a ALOHA mechanism for the UWB device

In the ALOHA protocol, a device transmits when it desires to transmit 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. In addition to the benefits of sparse node distribution diminishing the mutual interference, the UWB PHY provides a large amount of processing gain so that even two simultaneous transmissions that do collide may both result in a successful packet transfer.

Insert after 5.5.4.3 the following new subclause (5.5.4.4):

5.5.4.4 Enhanced robustness features for the UWB PHY

The UWB PHY was specifically designed to provide enhanced robustness for LR-WPAN applications. This enhanced robustness is a result of several PHY features:

- Extremely wide bandwidth characteristics (UWB) that can provide very robust performance under harsh multipath and interference conditions
- Concatenated forward error correction (FEC) system to provide flexible and robust performance under harsh multipath conditions
- Optional UWB pulse control features to provide improved performance under some channel conditions while supporting reliable communications and precision ranging capabilities

5.5.5 Power consumption considerations

Insert after 5.5.5 the following new subclause (5.5.5.1):

5.5.5.1 Additional power saving features provided by the UWB PHY

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.

5.5.6 Security

Insert after 5.5.6 the following new subclauses (5.5.7 through 5.5.8.4):

5.5.7 General overview of ranging

Ranging is an optional capability that has additional options within it. Ranging capability is achieved through support of a number of specific PHY capabilities as well as defined MAC behaviors and protocols. The key protocol in ranging is a two-way frame exchange that is presented as a sequence of individual steps in 5.5.7.1. Another ranging approach, which the standard does not preclude, is one-way position awareness, and this is noted briefly in 5.5.7.2. The single most critical device capability for ranging is the ranging counter, which is described in 5.5.7.3. Another essential capability for ranging is the determination of the arrival time of a signal at the device receiver, and this is discussed in 5.5.7.4. When performing ranging computations, there are several significant sources of error to be managed. One of the sources of error is the relative frequency offset of the crystals in the devices performing ranging, and management of this error source is described in 5.5.7.5. Another significant potential source of error in ranging is internal propagation paths, and the management of that error source is discussed in 5.5.7.6. To manage ranging at the application solver layer, it is necessary for the PHY to make an estimate of the quality of each individual ranging measurement. The quality information is carried from the PHY to the application with the ranging figure of merit (FoM) discussed in 5.5.7.7. This standard supports private ranging, described in 5.5.7.8, which is an optional mode for enhancing the integrity of ranging traffic in the face of a disruptive hostile device.

5.5.7.1 Two-way ranging

UWB devices that have implemented optional ranging support are called *ranging-capable devices* (RDEVs). UWB PHYs have a bit in the PHY header (PHR) called the *ranging bit*, which is set by the transmitting PHY for frames used in ranging, and the bit serves to signal the receiver that this particular frame is intended for ranging. A UWB frame with the ranging bit set in the PHR is called a *ranging frame* (RFRAME). There is nothing else (beyond the ranging bit set in the PHR) that makes an RFRAME unique. RFRAMEs can carry data, RFRAMEs can be acknowledgments, and RFRAMEs do not even (for the case of one-way ranging) necessarily require an acknowledgment. As far as ranging is concerned, the critical instant in a frame is the first pulse of the PHR. The first pulse of the PHR is the *ranging marker* (RMARKER). This standard is primarily structured to support the two-way time-of-flight computation of distance between two RDEVs. The overview of two-way ranging is shown in Figure D1.4 (in Annex D1). The figures of this subclause describe the two-way ranging sequence step by step. Figure 13a shows the complete sequence for two-way ranging. This sequence is disassembled and presented step by step in Figure 13b, Figure 13c, and Figure 13d.

In Figure 13b, the bottom half of the sequence is light gray so the reader can focus on the first frame of the two-way exchange. This RFRAME is sent from the originating device to the responding device. A ranging counter start value is captured in the originator device upon the RMARKER departure from the originator, and a ranging counter start value is captured in the responding device upon RMARKER arrival at the responder. The RFRAME has the acknowledge request bit set in the MAC header. In the most general case, the counter in the responder PHY may have already started running when a previous RFRAME arrived, but the previous RFRAME was not intended for this device and thus did not get an acknowledge from this device. In the figures, the counter activity is labeled “start/snapshot” from the PHY perspective. For the PHY, the counter function is “start” for the first arriving frame and “snapshot” for subsequent frames. Snapshot means that the value of the counter is captured and stored at the instant of the snapshot, but the counter continues to count as if snapshot had not happened. The responder PHY initiates PD-DATA.indication primitives with counter snapshot values for all arriving RFRAMEs. The responder MAC discards the snapshot values that are for RFRAMEs not intended for the responder device. At the end of the first frame transmission in Figure 13b, the counters are running in both devices.




Figure 13a—The complete two-way ranging exchange




Figure 13b—The first frame of the two-way exchange

In Figure 13c, the top half of the sequence is light gray so the reader can focus on the second RFRAME of the two-way exchange. This RFRAME is an acknowledgment sent from the responding device to the originating device. The ranging counter stop value is snapshot in the responding device upon RMARKER departure from the responder. The responder PHY is now in transmit mode, and the counter is still running. Because the PHY is in transmit mode, it will not be receiving any frames or taking any counter snapshots. Leaving the counter running in the responder at this point in the algorithmic flow only serves to deplete the battery of a mobile device. For overview purposes, in Figure 13a, Figure 13b, Figure 13c, and Figure 13d, the counter action is labeled stop, not because it really is stopped (it is not), but because the algorithmic flow is done with it and because it will appear to the application as if it has stopped because it will generate no more snapshots. The counter stop value in the originator device is snapshot and saved upon RMARKER arrival at the originator. The originator MAC verifies that the frame was from the responder and ultimately the application will then stop the counter with a MLME-RX-ENABLE.request/PLME-SET-TRX-STATE.request primitive pair. The originator PHY is in receive mode; therefore, until the counter is stopped, that PHY will continue to generate PD-DATA.indication primitives for all future arriving RFRAMES.




Figure 13c—The second frame of the two-way exchange

In Figure 13d, both frame transmissions are light gray to call the reader's attention to the ranging counter in each of the devices after the exchange is complete. When the application in the responding device learned that the acknowledgment had been sent, it stopped the ranging counter in the PHY. When the application at the originator device discovered that the acknowledgment frame was for the originator device, it stopped the counter in the PHY. Thus in Figure 13d, all the counters have stopped, and the values are located in the respective devices.




Figure 13d—The devices after the two-way exchange

The discussion above in this subclause risks confusion because it includes the general case of arriving frames not intended for the devices in the figures. That behavior is important for algorithmic robustness, but for understanding basic ranging, it is a distraction. In Figure 13d, the ranging pair holds two different sets of counter values, with a start value and a stop value in each set. Along the way, the application may have discarded counter snapshots due to frames destined for other devices; but in any case, what remains at the end of the exchange are pairs of counter values that when subtracted represent the elapsed times between the arrival and departure of the intended frames.

At the system state represented by Figure 13d, the necessary information required to compute the range between the devices is known. However, the information is still distributed in the system; and before the ranging computation can be accomplished, the data are brought to a common compute node. The difference of the counter start and stop values in the originator device represents the total elapsed time from the departure of the first message to the arrival of the acknowledgment. The difference of the counter start and stop values in the responder represents the total elapsed time from arrival of the data message to the departure of the acknowledgment. After these values are all brought together at a common compute node, they are subtracted, the difference is divided by 2, and the time of flight (and thus the inferred range) is known.

While management of the ultimate disposition of the counter values is outside the scope of this standard, the most immediately obvious resolution is for the application at the responder device to send the counter value

to the originator device in a data frame. The originator device started the exchange; therefore, it might be assumed that it is the point at which the application desires to have the range information. The responder device was just in radio contact with the originator device; therefore, a communication channel is very likely to be available.

The obvious solution suggested in the previous paragraph is often the wrong solution. This standard supports applications that bring a large number of ranging measurements together at a single computation device, and there the application solves not just for the ranges between individual devices but also the two- or three-dimensional relative location of the devices. A discussion of typical activities at a central compute node is included in Annex D1.

5.5.7.2 Position awareness through one-way transmissions

The primary intent of the ranging support in this standard is to support ranging through two-way time-of-flight measurements. To establish that capability, this standard defines a whole suite of capabilities and behaviors by which two-way ranging is enabled. The capabilities required to accomplish one-way ranging are sufficiently similar that this standard allows operation in that mode as well. The form of one-way ranging allowed by this standard is described as “Mode 2” in D1.4.2. One-way ranging requires an infrastructure of RDEVs and some means to establish a common notion of time across those devices. The protocol to establish the common notion of time is outside the scope of this standard. What this standard does provide is a bit in the PHR (which any UWB device can optionally set), and this bit serves to signal an RDEV (in this example, it will be an infrastructure RDEV) that location awareness is desired. To operate in a one-way environment, any UWB device (not necessarily an RDEV) simply sends a frame having an appropriate preamble length, the ranging bit set in the PHR (making the frame an RFRAME), and the acknowledge bit off in the MAC header. As described in 5.5.7.1, after the PHY counter is running in an infrastructure RDEV, the PHY initiates a PD-DATA.indication for all arriving RFRAMES. The MAC in the infrastructure RDEV does not send an acknowledgment when the acknowledge request bit is off, but rather initiates an MCPS-DATA.indication primitive to the next higher layer. By this mechanism, the application acquires a list of counter values representing arrival times of all one-way ranging messages received at all infrastructure devices. As with the case of two-way ranging, nothing useful can happen with the lists of counter values until they are brought together at a common compute node; and as with two-way ranging, that bringing together activity is beyond the scope of this standard. After the counter values have been brought together, the computation of the relative locations proceeds as shown in D1.4.2.

5.5.7.2.1 Processing sequential arriving RFRAMES

Subclause 5.5.7.2 says that the PHY will initiate a PD-DATA.indication for all arriving RFRAMES. That can be thought of as a great goal; and if the RFRAMES arrive infrequently, it is a very achievable goal. However, the PHY does not get to choose how quickly a sequence of RFRAMES might arrive; and in real world applications, the RFRAMES may arrive more quickly than the PHY can deal with them. The processing time for an RFRAME is very dependent on the implementation of the PHY. While the PHY has no control over the inter-arrival time of RFRAMES, the application very well might. The application can discover the RFRAME processing time by reading the PHY PAN information base (PIB) attribute *phyRFRAMEProcessingTime*. The intended use of this attribute is that if an application discovers the processing capabilities of the devices in the network it can structure the traffic so that devices are not overrun. The PIB value is a single octet and the least significant bit (LSB) nominally represents 4 ms. In other words, the maximum value is a nominal second. The only purpose is to help prevent overruns; therefore, there is no need for high precision in expressing this value.

5.5.7.3 The ranging counter

At the most fundamental level, ranging capability as described in this standard is enabled by the ranging counter shown in the center boxes of Figure 13a. The ranging counter has the capability of assigning values to the precise instant that RMARKERS are transmitted and received from the device. Once that counter and

the ability for it to precisely snapshot a timestamp are in place, then conceptually, the computation of the time of flight is simple.

The ranging counter, which is a key to an RDEV, is a set of behavioral properties and capabilities of the RDEV that produce ranging counter values in response to PHY primitives and events at the device antenna.

An actual physical counter that exhibits the behavior attributed to the “ranging counter” does not need to exist in any PHY implementation. The ranging counter described is simply an abstraction that is used to specify the required PHY behavior. For example, in the case of a freshly asserted PLME-SET-TRX-STATE with parameter RX_WITH_RANGING_ON, the ranging counter is specified in 6.2.2.7.3 to start counting from 0x00000001 upon the arrival at the antenna of the first pulse of the header of the first arriving PHY protocol data unit (PPDU) for which the ranging bit is asserted in the header. A literal implementation of that counter would require a prediction machine that could somehow pre-know the contents of a bit in the PHR that has not yet arrived.

The LSB of the ranging counter represents a time interval so small that an actual physical counter would have to run at a nominal 64 GHz to produce values with the required resolution. It is unlikely that a device intended for low-cost battery-powered operation would implement a counter running at 64 GHz.

The lack of an actual physical realization does not in any way preclude the use of the ranging counter as an abstraction in this standard to visualize and specify the behavior of an RDEV. From an algorithmic and computational viewpoint, the RDEV will appear to an application as if it possesses a 64 GHz counter with the ability to start and stop based on the state of bits *before* they arrive.

The implementation of the ranging counter is beyond the scope of this standard; however, the following ideas are suggested:

- Take snapshots of the counter at every event when counter value might be required and then discard unneeded snapshots after the future becomes known.
- Do not build a 64 GHz counter, but rather generate the less significant bits of the ranging counter values using computational techniques like those described in D1.2.

5.5.7.4 Accounting for signal arrival time

This standard specifies that the start counter values represent the time of arrival of the first pulse of the first symbol of the header of a PPDU. That necessary task is not trivial when the signal arrives in a channel with significant multipath. The first pulse of the header arrives at the antenna once for every multipath reflection in the environment. In addition, the first pulse also arrives off the reflections of the reflections in the environment. The result is that in an indoor environment, the “first pulse” can seem to arrive a multitude of different times. Accurate ranging requires discriminating the leading edge of the cluster of signal arrivals that accounts for the first pulse of the header. The technique for achieving this discrimination is beyond the scope of this standard, but it is helpful to discuss a typical approach. In a typical approach, the counter value is snapshot relative to a position on the arriving waveform where the PHY-tracking loop has achieved and is (hopefully) maintaining a “signal lock spot.” Then the offset from the lock spot to the leading edge of the pulse energy is determined. After a time offset from the UWB signal lock spot to the leading edge of the energy cluster is found, the rest is very straight-forward: that offset is just subtracted from the counter value exactly as the other correction factors are. The time offset to the leading edge is discovered by sampling the energy ahead of the acquisition lock spot over multiple different offsets to discover the earliest point with discernable energy. To achieve the required precision, the sample values in the vicinity of the leading edge could be further refined using techniques like those shown in D1.2. For those computations (and other techniques, generally known as *up-sampling*), it is critical that the noise in the samples be well suppressed. This standard supports this leading edge activity by allowing a long as well as a very long acquisition preamble keeping the signal steady and data free for a protracted time.

The necessary characterization of the channel multipath response is generically called a channel sounding. The techniques that accomplish that task can be numerically intense. This standard does not preclude the system designer accomplishing that task in the PHY. However, this standard provides a sounding mechanism [involving both MAC sublayer management entity (MLME) and physical layer management entity (PLME) primitives]. This allows the PHY to present raw data to a higher layer should the PHY lack significant computational capability and the system designer wishes to employ numerically intensive channel sounding algorithms. The raw data can then be moved to whatever device has sufficient computational resources to support the desired algorithms. The primitives supplied are very simple and assume that both the MAC and the application are well behaved and give priority attention to sounding activities. For example, the primitives do not include tags to associate a particular sounding with a particular packet. The application is responsible for making sure that sounding activities are conducted in a timely way so that the sounding information is associated with the last packet received. The actual handling of the raw data after the sounding operation facilitates uploading to a next higher layer is beyond the scope of this standard.

5.5.7.4.1 Leading edge search during the acquisition preamble

Upon acquisition of the signal, the PHY is not aware of how much time remains in the preamble before the delimiter. A reasonable goal is to do the best possible job of bracketing the leading edge with whatever time is available and then reporting how well the leading edge was bracketed by way of the ranging FoM. In a typical implementation, if the delimiter arrives very quickly after the acquisition threshold was satisfied, then the leading edge equipment will still be using coarse steps to characterize the energy. The PHY will make the best judgment it can about the leading edge based on the coarse steps and then report a FoM value appropriate for coarse steps. If the leading edge search engine has ample time before the delimiter arrives, then not only can it have progressed to using a very fine search step, but it can also have integrated many samples to drive down the noise in the computation. In this case, the correction representing the leading edge offset is applied to the counter value (and it might have been the very same correction value as was applied in the previous case when the search time was short), but this time the FoM value is reported for a very small characterization bin and very high confidence that the leading edge truly was in that bin. Again note that the counter value returned with a good FoM can have the same value as the counter value returned with a bad FoM, i.e., the counter value is independent of the FoM.

5.5.7.4.2 FoM for bad times

If the PHY performs a short leading edge search (as will happen after recovering from an acquisition false alarm, for example), it still makes its best guess for a leading edge correction and goes on with the ranging algorithm. Even when the final counter value represents a known error-prone measurement, the PHY should not return a FoM of zero. Zero means “no FoM,” which is neither correct nor useful. An appropriate FoM to report for the most error-prone cases is 0x79. That value decodes to tell the application that even if the other RDEV calculated its half of the measurement perfectly, given the expected error just due to this RDEV’s measurement alone, the PHY is 80% confident that the computed range will be wrong by more than 2 m.

5.5.7.4.3 Other opportunities for leading edge search refinement

The previous discussion was framed as if all channel characterization had to stop upon the arrival of the delimiter. In fact, PHYs can do additional things after the delimiter to further refine the estimate of the leading edge offset. The optional UWB CCA pulses described in 6.8a.14 (if used) offer additional opportunities to look at the signal in a known state after the delimiter. A very sophisticated PHY may perform additional characterization during the time that data are on the air if the algorithm designer is willing to “back out” the effects of the data after demodulation (and thus after the data are “known”). From the application’s point of view, all the application sees for the difference between a very capable PHY and a sloppy, mediocre PHY is a difference in the FoM values being reported.

5.5.7.4.4 Managing the preamble length for leading edge search

One of the most distinguishing traits of ranging UWB radio transmissions is the long preambles. This standard allows the application to specify preambles that are either 1024 or 4096 symbol repetitions long. The selection is a function of the channel multipath, the signal-to-noise ratio (SNR) in the link, and the capability of the receiving PHY. It is theoretically possible that a very capable PHY that does leading edge refinement using the data could do ranging accurately with a preamble that is 16 symbols long. It is also possible (likely, in fact) that a PHY with a poorly designed search engine will not do a good job in heavy multipath even with a 4096 symbol preamble. The upper layers are responsible for picking the preamble length. It is suggested that the application start ranging operations using the 1024 symbol preamble and keep a history of how the FoMs are reported. The FoMs are the critical feedback information that tells the application how the various PHYs are doing, and the application can make future adjustments to the preamble length based on that history.

5.5.7.4.5 PHY deferral of the computations for leading edge search

As discussed in 5.5.7.4, this standard provides a mechanism to optionally allow the PHY to pass the computational burden of leading edge processing to a higher layer. If the computations are not done in the PHY, then the value in the timestamp report for RangingCounterStart is not corrected for the leading edge. The RangingFOM is used to signal this condition to the higher layer, which will have to compute a correction based on data acquired using the sounding primitives. See 6.8a.15.3. The higher layer issues a MLME-SOUNDING.request primitive, which in turn causes a PLME-SOUNDING.request primitive. The associated MLME- and PLME-SOUNDING.confirm responses return a list SoundingPoints where each SoundingPoint is a pair of integers representing data taken by the PHY at time offsets from the point on the waveform represented by the uncorrected value in RangingCounterStart. A time of zero in the list designates an amplitude value taken at the point indicated by RangingCounterStart. Positive time values indicate amplitudes that occurred earlier in time than the zero point. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz. The amplitude values do not represent any particular voltage. They are only meaningful in a relative sense and in the context of each other. The values are linear (not logarithmic). The amplitude values are all consistent with each other. For example, it is acceptable for an automatic gain control (AGC) circuit to change the gain during the measurement of the amplitudes, if the numbers are corrected so that the effect of the gain change is removed and the numbers returned in a SOUNDING.confirm primitive are the values that would have been measured had the gain been perfectly stable and unchanged for all measurements. The list of measurements returned by the SOUNDING.confirm primitive begins with the size of the list. The maximum size that can be represented is 65 K value pairs. That large value is only because a single octet would not be adequate to represent lists larger than 255 pairs. In practical systems, lists larger than 255 pairs can occur. Two octets would be the next choice to represent the list size, but that does not mean that lists approaching 65K pairs would be appropriate. There is no particular acceptable or unacceptable list size. Generally, a larger list is superior: see 5.5.7.4.1. In the case where the PHY is deferring the leading edge computation to an upper layer, the PHY does not assign a FoM to the timestamp report. That does not mean that a FoM is unneeded by algorithms at the higher layers; it just means that the PHY will not be the source. In cases where the PHY defers computation, the upper layer will typically compute a FoM for itself based on the size and quality of the list returned with the SOUNDING.confirm primitive.

5.5.7.4.6 PHY deferral of the computations for self-calibration

As discussed in 5.5.7.4.4, the sounding primitives provide a mechanism to offload from the PHY the computational burden of analyzing a channel sounding for the leading edge of an arriving signal. A very similar problem arises in the self-calibration of a ranging UWB PHY. One excellent technique for self-calibration is the “sounding” of a loopback path in the radio. In this technique, the PHY actually transmits to itself through reflections associated with the transmission path (like a transmit/receive switch or the antenna itself). When using this technique, the issue comes up again that it can be computationally intense to discover the moment of arrival of the (often small) amplitude disturbances associated with elements in the

path to the antenna. Implementers may choose to achieve this computational effort in the PHY. In an alternate implementation, the sounding mechanism can be used to offload this computational burden to a higher layer. When performing a sounding for leading edge computation, the instant associated with time zero is the point associated with signal tracking. Sounding for calibration is slightly different in that the time associated with zero is the launching of the pulse event from logic at the level of the ranging counter.

5.5.7.5 Management of crystal offsets

The numbers that will be subtracted in the range computation will typically represent times on the order of 5 ms. The time of flight for a 10 m link is about 30 ns. The expected difference of the counter values will be twice the time of flight, or something like 60 ns, in this example.

When a subtraction of values representing 5 ms is supposed to yield a meaningful answer on the order of 60 ns, even small percentage errors in the relative measurement of the 5 ms numbers yield large errors in the difference. The root cause of these errors is the fact that each of the individual 5 ms measurements was made with different crystals in different devices.

The crystals' oscillators in different devices may generate frequency errors of 20 ppm. A 20 ppm error in a 5 ms number can account for 100 ns. This 100 ns error is disconcerting considering the 60 ns time-of-flight result.

Management of the errors due to crystal differences is essential to ranging. Correcting for a crystal difference algebraically at the time of the subtraction computation is straight-forward if the difference is precisely known. The crux of the problem is to determine the crystal difference.

The mechanisms to characterize the crystal difference will be present and functioning in typical UWB PHY implementations. This crystal characterization equipment is the signal tracking loop in the receiver. A UWB pulse occupying 500 MHz has a nominal envelope width of 2 ns. The receiver tracking loop in a UWB PHY will stay "locked on" to this envelope for the duration of a packet. In the case of ranging packets, this will typically amount to milliseconds. In actual practice, the tracking loop will hold the sampling point on the envelope much tighter than 2 ns; therefore, by the end of the transmission the tracking loop has the information to hold its sample point steady (with respect to its local crystal) on the received signal (which is sourced by the other devices crystal). In other words, by the end of the packet, the tracking loop has exactly measured the crystal difference. This crystal difference is the very thing necessary to correct the values in the ranging computation.

5.5.7.5.1 Characterizing crystal offsets with digital tracking loops

For a digital tracking loop, the most convenient way to express the crystal difference is with two numbers. A tracking interval number is the total number of units during which the signal was tracked, and a tracking offset number is a count of the number of times the tracking loop had to add or drop a unit to hold the sample point steady on the incoming signal. If the other oscillator frequency was lower than the tracking loops local oscillator, then the tracking loop would be adding units to hold the sample point steady. The tracking offset is simply a count and a sign bit. All numbers are expressed from the local device's point of view; therefore, positive counts are characterizing crystals in the other device that are running slower (so the receiver was adding time units to match it) and negative numbers characterize faster crystals in the other device (so the receiver was dropping local units to keep up). The offset is thus a signed magnitude integer (not the twos-complement number that might be expected). The actual units (generally called "parts") that are called out in the count are whatever units happen to be convenient for a given PHY implementation. Since the numbers are used only as a ratio, the type of unit need not be specified as long as the numbers express the same unit.

5.5.7.5.2 Characterizing crystal offsets with analog tracking loops

PHYs that use analog phase-locked loops (PLLs) to track the received signal do not lend themselves as directly to the expression of the tracking offsets as counts. However, the PLL steady-state error signal is still a direct measure of the crystal offset. The analog PLL-based PHY can convert the PLL error signal to a number [for example, with an analog-to-digital converter (ADC)], put that result in the offset count field (taking care to get the sign bit correct), and put a convenient scaling number (like a million) into the total tracking interval field so that the ratio again expresses the difference of the crystals from the local oscillator's point of view.

5.5.7.5.3 Characterizing crystal offsets with different tracking loops

The use of the receiver's tracking loop to characterize the crystal offset is convenient for some PHY implementations, but it is not required for compliance. In fact, RDEVs are not required to support crystal characterization at all. If two RDEVs are involved in a ranging exchange and only one of them is supporting crystal characterization, all the information needed for a good ranging computation is available. (If both RDEVs support crystal characterization, they will get the same ratio with opposite sign; therefore, there is little new information.)

If neither RDEV supports crystal characterization, the application puts more traffic on the air to support ranging. For this situation, the application does the measurement twice. The first time is simply the normal exchange. On the second measurement, the roles are reversed. The device that was the originator on the first measurement is the responder for the second measurement, and likewise the responder on the first measurement becomes the originator for the second measurement. Then the application does the range computation twice. Because neither measurement provided for correcting for crystal offsets, the answers for both measurements are likely to be totally wrong. But, the computations did involve the same crystals so the error in the measurements is the same. Because the application reversed the measurement sequence between the two measurements, the answers have errors with opposite sign. The bottom line is that while the two individual answers are both hopelessly inaccurate by themselves, the average of the two individual answers will be exactly correct. A further refinement of this technique is called *symmetric double-sided two-way ranging* (SDS-TWR) and is discussed in D1.3.2. The refinement shown in Annex D1 seeks greater efficiency by combining the two independent measurements into a single stream with the originator sending an acknowledgment for the responder's acknowledgment. While the "acknowledge for an acknowledge" approach is absolutely sound mathematically for ranging and the additional efficiency is tempting, the "acknowledge for an acknowledge" message sequence construct is beyond the scope of this standard.

5.5.7.5.4 Size of units

As introduced in 5.5.7.5.1, the units of measurement in the crystal characterization ratio are not rigidly defined in this standard to allow vendors the freedom to choose a value that works well with their PHY implementation. Design freedom is good, but to ensure at least a minimum level of ranging accuracy, this standard insists on a value that allows the ratio to express individual parts per million of oscillator difference. When the ranging computation is done for typical packet sizes and turnaround times, the desired answer will typically only be tens of parts per million compared to the numbers being subtracted. It is in this context that this standard calls for using the nominal 500 MHz chip time (nominally 2 ns) as the largest acceptable unit for the crystal characterization numbers. While an implementation using this value will be compliant, it would typically yield ranging errors on the order of a meter due to poor crystal characterization. To achieve reduced ranging errors, it is recommended that smaller units for crystal characterization be used since this translates directly to reduced errors in the ranging computation. When both RDEVs of a two-way ranging exchange are supporting crystal characterization, the application has a choice of which set of numbers will be used to make the correction. The ranging application would be wise to manage the reality that different devices might present different quality results for the crystal characterization.

The LSB of the ranging counter caps the highest ranging accuracy that can be achieved by compliant devices (which use the straight-forward two-way ranging techniques described here). The LSB represents a nominal 16 ps, which corresponds to about half a centimeter of flight for energy in free space.

5.5.7.6 Accounting for internal propagation paths

Subclause 5.5.7.3 introduced the ranging counter for measuring events at the device antenna very precisely. It is understood that an actual implementation will not be trivial. Typically the device's PHY will have a counter somewhere in the digital section, multiple correction values stored in registers, and some arithmetic hardware to apply the correction values. The end result of all this is that it appears (from the numbers reported) that the PHY has a counter that is somehow magically positioned right at the antenna of the device and is taking snapshots of the counter values for events as they happen right at the antenna. That is important because the computation is supposed to be for the time of flight through the air, not through some impedance matching network feeding an antenna. Subtracting the correcting values for internal propagation times is not hard. What is hard is actually knowing the values of the internal propagation times. This standard provides a CALIBRATE mechanism (involving both MLME and PLME primitives) that allows an application to cause a PHY (at a time that the application chooses) to invoke whatever capability that the PHY might have to characterize the internal propagation times of the PHY. Inclusion of a device capability for antenna loopback with an associated self-calibration algorithm is encouraged, but beyond the scope of this standard. A defensively written ranging application can maintain tables of correction factors at the computation nodes where the table entries are individually associated with the unique devices it is using for ranging. In this way, the application can compensate (after the fact) for devices in the ranging environment that may have done a poor job of self-calibration.

5.5.7.6.1 PIB attributes for internal propagation paths

This standard provides a defined place to go to for the correction factors characterizing the delays of the internal propagation paths. There are two separate PHY PIB attributes that separately cover the transmit and receive paths. The LSB of these values represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz. This time interval represents about half a centimeter of travel at the speed of light and is the same size as used throughout the ranging computations. The intended use of these PIB attributes is for them to be written by the application at a time of the application's choosing and for the values to stay with the device until rewritten. One possible way for the application to learn what values to write to the PIB attributes is to invoke the CALIBRATE primitives. This standard does not mandate that the CALIBRATE primitives must be used. The standard simply makes them available for use, if desired.

5.5.7.6.2 Support for self-calibration and one-way ranging

As discussed in 5.5.7.2, this standard does not preclude position awareness through one-way ranging. Successful one-way ranging requires that the internal propagation paths to the transmit antenna and from the receive antenna be accounted for separately. The PHY PIB attribute *phyTx_RMARKER_Offset* represents the time from the internal ranging counter reference to the transmit antenna. Likewise, the PHY PIB attribute *phyRx_RMARKER_Offset* represents the time from the receive antenna to the internal ranging counter reference.

5.5.7.6.3 Use of the calibrate primitives

The CALIBRATE.confirm primitives return either the values that are correct for the RMARKER offsets, if the PHY takes care of all computations itself, or the status COMPUTATION_NEEDED. The actual implementation of the self-calibration could be as simple as returning hardwired values for the RMARKER offsets. In this situation, the hardwired values would be selected by the vendor at the time of device manufacture to represent a best guess of what the offsets might ultimately be. Alternatively, the self-calibrate implementation might involve a full channel sounding of a loopback path and a sophisticated pattern-matching algorithm to determine from the sounding waveform when an internally generated pulse

reflected back from an antenna. In either of the two scenarios above, the status `COMPUTATION_NEEDED` would not be used because either (in the first case) the calibrate implementation was so crude that there was nothing to do, or (in the second case) the calibrate implementation was so sophisticated that the PHY took care of all of the computations without assistance.

It is a property of the algorithms that a node which only does ranging transmissions within a one-way infrastructure-based ranging application need not support calibration in any form.

5.5.7.6.4 Use of the `COMPUTATION_NEEDED` status

Subclause 5.5.7.6.2 described two extreme implementations of PHY self-calibration. This standard supports a reasonable middle ground implementation where the PHY does a channel sounding of a loopback path using the same hardware resources as normally used for leading edge scanning, but then the PHY defers the actual computations associated with the channel sounding to a higher layer. As discussed in 5.5.7.4.5, the computations associated with processing a channel sounding can be numerically intense and may well be beyond the resources of a particular PHY implementation.

When the higher layer receives a status of `COMPUTATION_NEEDED` in response to a `CALIBRATE.request` primitive, the higher layer will then use the sounding primitives (see 5.5.7.4.6) to get a list of `SoundingPoints` from the PHY. The higher layer will then process the `SoundingPoints` to determine the values of the `RMARKER` offsets.

5.5.7.7 Timestamp reports

Subclause 5.5.7.1 introduced the two ranging counter values (start and stop). Subclause 5.5.7.4 introduced the ranging FoM value. Subclause 5.5.7.5 introduced the two values that (as a ratio) characterize the crystal offsets. All of these values (five in all) characterize a single ranging measurement. The five individual numbers that characterize a measurement are referred to in a group as a *timestamp report*. It then takes (at least) two timestamp reports to do a time-of-flight computation. There are a total of 16 octets in a timestamp report. The numbers in a single timestamp report have meaning in the context of each other. As such, they are generated by the PHY as a set and not split apart during subsequent data handling.

5.5.7.7.1 Presentation of timestamp reports

Subclause 5.5.7.7 described the timestamp report. It should be noted that these reports will occur at seemingly nonintuitive times in the actual primitives and the message sequence charts. For example, a timestamp report is included in the `PD-DATA.confirm` primitive. When the `PD-DATA.confirm` primitive is used following an initial transmission, the elements of the timestamp report are not all known. However, when the `PD-DATA.confirm` primitive is used following an acknowledgment in a ranging message sequence, all of the elements are known. Likewise, the `PD-DATA.indication` primitive includes a timestamp report, but when the `PD-DATA.indication` primitive is used in response to the initial reception of the first message of a ranging message sequence, not all of the elements of the timestamp report are known. However, when the `PD-DATA.indication` primitive is used following reception of the acknowledgment message, all of the elements are known.

5.5.7.7.2 Start and stop times in the timestamp report

The timestamp report as both a start time and a stop time with 4 octets for each. This may appear to be counter intuitive since either start or stop number by itself is useless and that the only real utility for the numbers is in their difference. A different strategy would be to have the PHY do arithmetic on the pair of numbers and present only the difference in the timestamp report. In this standard, the numbers are handled separately by the PHY to allow ranging by PHY implementations having few arithmetic or logic resources. Another reason is to allow an infrastructure node in a one-way ranging environment to issue a new timestamp report for each arriving `RFRAME` without being concerned about when the “start time” was.

5.5.7.8 Private ranging

It is important to note that for some applications of this standard, the range information will be the critical deliverable information for the entire system. As such, it is reasonable to protect this information as well as safeguard the integrity of the ranging traffic itself.

5.5.7.8.1 Simple encryption of the timestamp reports

As discussed in 5.5.7.1, at the end of the two-way exchange, half of the information necessary for the range computation is in each of two devices. Either half, by itself, is useless to the desired ranging application as well as to any undesired hostile device. When the two halves of the information are brought together, the range is computed by simple arithmetic. The single most critical and effective thing that an application can do to keep hostile devices from learning range information is encrypting the time reports whenever they are being transmitted. There is no problem doing this, as the reports are moved after the time-critical ranging exchange is complete and there is nothing time critical about movement of the timestamp reports.

5.5.7.8.2 Dynamic preamble selection (DPS)

It is anticipated that typical ranging traffic will take place using the normal channel codes and preambles in regular network use. Therefore, even if the time reports are encrypted and a hostile device is denied knowledge of the ranges, a hostile device can monitor traffic and listen for long preambles. It can then turn on its transmitter, spoof the acknowledgment transmission, and generally disrupt the ranging traffic. This hostile behavior creates a race between the hostile device and the legitimate responder, but the hostile device can expect to win the race because the legitimate responder will be parsing the MAC header to discover whether an acknowledgment is appropriate before it starts transmitting. To defeat this spoof attack, this standard offers the DPS option, where RDEVs are allowed to move their long preamble RFRAMES to codes that are altogether different from the codes in normal use. Furthermore, the different preambles are precoordinated using encrypted messages so that the hostile device is denied knowledge of the preambles that will be used. And finally, there are no retries allowed with these preambles so that a “jam and spoof the retry” attack will also be defeated.

When the DPS option is invoked by the devices that support it, the hazard is created where if the two-way ranging packet is not received as expected, the devices waiting and listening for special unique length 127 preambles will have become lost. To render this hazard safe, this standard provides for an additional timeout, `DPSIndexDuration`, which is used whenever DPS is used to ensure that devices at risk of becoming lost are always returned to an interoperable state. See 7.5.7a.3.

DPS provides no additional ranging capability beyond resistance to attacks by hostile nodes. PHYs that do not implement DPS do not be give up any one-way or two-way ranging capabilities.

5.5.8 Management of UWB options

The UWB PHY is designed to address a broad spectrum of applications; and as a result, the specification of the UWB PHY includes a rich set of optional modes and operational configurations. An overview of the optional modes and behaviors is given in 5.5.8.1. Having a rich set of options does not preclude the cost-effective devices discussed in 5.5.8.2. A very rigid framework of option management rules (see Table 23 and 6.4.2) govern the use of optional modes in a way that ensures interoperability of UWB devices conforming to this standard. An overview of the rules is given in 5.5.8.3. The optional low data rate is handled differently from the other rates and is discussed in 5.5.8.4.

5.5.8.1 Overview of UWB options

The UWB PHY allows for operation using possibilities selected from lists of the following variables:

- Center frequencies
- Bandwidths occupied
- Pulse repetition frequencies (PRFs)
- Chipping rates
- Data rates
- Preamble codes
- FEC options (no FEC, or Reed-Solomon only, or convolutional only, or Reed-Solomon with convolutional)
- Waveforms
- CCA mode (on or off)
- Preamble symbol lengths
- Ranging
- Private ranging

Clearly the richness of the menu above enables PHYs that support an exceptionally broad set of service conditions. However, a fundamental goal of this standard is low cost; and if a standard is to mean anything, interoperation of compliant devices is crucial. The rich menu might seem a contradiction to both these important goals. The enablement of low-cost devices is discussed in 5.5.8.2. The interoperability of compliant devices is assured through the rules in 5.5.8.3.

5.5.8.2 Modes and options for low-cost UWB devices

The low-cost imperative of this standard is met through the fact that only a small subset of the combinations of capabilities represent mandatory modes.

A UWB PHY is required to be able to turn on and off a bit in the PHR in response to a single, specific primitive attribute. That single bit in the header is the only required ranging “support” for a UWB PHY. What that bit achieves is the ability of that simple device to have its location determined (if some higher level application turns the bit on) when it is in the operational range of an infrastructure of UWB PHYs that have implemented optional support for ranging and are running an application supporting one-way ranging (see 5.5.7.2).

A compliant UWB PHY need support only the following:

- One single band (see 6.8a.11.1)
- One mandatory center frequency (see Table 39i in 6.8a.11.1)
- One mandatory data rate (see 6.8a.7.1)
- One mandatory bandwidth (see 6.8a.11.1)
- One mandatory pulse shape (see 6.8a.12.1)
- One chipping rate (see Table 39a in 6.8a.4)

The compliant PHY does have to support two preamble codes, but needs to use only one mandatory preamble symbol length. No special support for CCA (see 6.9.9) is required, nor is support for FEC required when receiving (see 6.8a.10).

For interoperability, this standard does require all UWB PHYs to both transmit and receive using both a nominal 16 Mpulse/s as well as a 4 Mpulse/s PRF, but a PHY is not required to transmit with the same

power level at 4 Mpulse/s as it does at 16 Mpulse/s. Therefore, it is not necessary to build any large voltage swing output drivers. Low-cost devices are additionally enabled by signal modulation, which is designed to be received by devices that do not sample coherently.

A compliant UWB device need support FEC only when transmitting frames. Encoding for transmission is straight-forward (see 6.8a.10). While it is noted that decoding FEC in the receiver may be costly and power consuming, it should be further noted that this behavior does not need to be supported. The FEC codes are designed to be “systematic”; in other words, the receiver is free to toss aside the redundant bits and do no error correction at all.

5.5.8.3 Rules for use of UWB modes and options

The UWB PHY specification allows operation in any of three bands:

- A sub-gigahertz band
- A low band, which is roughly between the 2.45 GHz industrial, scientific, and medical (ISM) band and the 5 GHz unlicensed national information infrastructure (U-NII) band
- A high band, which is above the U-NII band

The implementer is free to choose either one or several of the bands to be supported by an implementation. Within a band, there is only one mandatory channel.

There are five UWB waveforms supported by this standard. However, all beacon frames are transmitted using the mandatory waveform, and a PAN is allowed to use an optional waveform (for nonbeacon traffic) only after it is determined by a coordinator that all devices associated to the PAN are capable of supporting the optional waveform. Even after a PAN has transitioned its traffic to an optional waveform, new devices can learn about the PAN’s existence from the beacon frames. If a new device is allowed to join a PAN that is using a nonmandatory waveform and the new device is not capable of supporting that nonmandatory waveform, the entire PAN is returned by the controller to the mandatory waveform. The capabilities of an individual PHY are determined by reading the PHY PIB. The mechanism for communicating PHY capabilities between devices is accomplished by layers above the MAC sublayer and is beyond the scope of this standard. The rationale for the decision to allow or disallow a device to join a network (based on the device’s capabilities or any other reason) is out of scope of this standard.

The optional UWB CCA signaling mode described in 6.8a.14 is strictly a signaling mode. It is not associated with the use of any particular waveform. The UWB CCA mode can be optionally used (or not) with either the mandatory waveform or any of the optional waveforms. However, the impact on interoperability of the UWB CCA signaling mode is the same as if it were an optional waveform. In other words, devices using the CCA mode cannot communicate with devices that are not using the CCA mode. The guiding philosophy of the UWB PHY in this standard is that while signaling options are provided and allowed in homogenous networks, interoperability concerns always take precedence. To maintain interoperability, CCA signaling is treated as if it were an optional waveform. Use of optional CCA signaling is restricted by the same “homogenous network” rules as optional waveforms in the preceding paragraph.

Devices in a PAN are allowed to use optional data rates when communicating with each other, but the network beacon broadcasts at the mandatory data rate.

UWB PHYs need not implement the low data rate; however, a well-designed device that does not implement the low data rate can still survive gracefully if it finds itself in the operational vicinity of devices that are using the low data rate. Graceful survival is easily accomplished because each of the first two symbols of the low-data-rate header are identical to the single symbol used in the mandatory-data-rate header. A second delimiter detection indicates to a PHY that low-data-rate devices are operating nearby. A device not supporting the low data rate can keep its delimiter detector running while it is attempting to demodulate the PHR. If the delimiter detector triggers during what was expected to be the PHR, the mandatory-only receiver

can know that the packet is coming in at the (unsupported) low data rate and just give up and save power. Less graceful behavior involves demodulation of the extended low-data-rate delimiter as if it were the PHR and having checksums fail and wasting power. Other than wasting some power, no harm is done: The frame was not intended for the device not supporting low data rate.

5.5.8.4 The optional low data rate

This standard supports an optional low data rate of 110 kb/s to enable long links or provide high processing gain to cost-effective PHYs that might need the extra help. For the mandatory as well as the optional higher data rates, the PHR is transmitted at the mandatory data rate. The definition of the data rate for the rest of the PPDU is carried in the data rate subfield of the PHR (see 6.8a.7). That cannot work for the low data rate. The low-data-rate frame announces itself to the receiver by use of the extended delimiter (see Figure 27d in 6.8a.6). In frames transmitted using the low data rate, the PHR is also transmitted at the low data rate, and the PHY has the opportunity to demodulate the data rate subfield in the header to verify what it already learned by detecting the extended preamble.

6. PHY specification

6.1 General requirements and definitions

Insert the following item at the end of the dashed list in 6.1:

- Precision ranging for UWB PHYs

Insert the following paragraph at the end of 6.1:

In further additions to the rates supported in IEEE Std 802.15.4-2006, two high-data-rate PHYs have been added. They are CSS operating in the 2.4 GHz band and UWB operating both in the sub-gigahertz band and the 3–10 GHz band.

6.1.1 Operating frequency range

Change Table 1 (the entire table is not shown) as indicated:¹

Table 1—Frequency bands and data rates

PHY (MHz)	Frequency band (MHz)	Spreading parameters		Data parameters		
		Chip rate (kchip/s)	Modulation	Bit rate (kb/s)	Symbol rate (ksymbol/s)	Symbols
2450 <u>DSSS</u>	2400–2483.5	2000	O-QPSK	250	62.5	16-ary Orthogonal
<u>UWB sub-gigahertz (optional) (see note)</u>	<u>250–750</u>	<u>See 6.8a.11.1</u>				
<u>2450 CSS (optional) (see note)</u>	<u>2400–2483.5</u>	<u>See 6.5a.2</u>		<u>250</u>	<u>166.667 (see 6.5a.5.2)</u>	
		<u>See 6.5a.2</u>		<u>1000</u>	<u>166.667 (see 6.5a.5.2)</u>	
<u>UWB low band (optional) (see note)</u>	<u>3244–4742</u>	<u>See 6.8a.11.1</u>				
<u>UWB high band (optional) (see note)</u>	<u>5944–10 234</u>	<u>See 6.8a.11.1</u>				
<u>NOTE—UWB PHYs may operate in one or more of three distinct bands: sub-gigahertz, low band, and high band. IEEE 802.15.4a UWB device types and CSS PHY types are options to IEEE Std 802.15.4-2006.</u>						

Change the third sentence in the third paragraph in 6.1.1 as shown:

IEEE 802.15.4 and IEEE 802.15.4a devices conforming to this standard shall also comply with specific regional legislation.

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

6.1.2 Channel assignments

Change the first sentence in 6.1.2 as shown:

The introduction of the ~~“868/915 MHz band (optional) amplitude shift keying (ASK) PHY specifications”~~ and ~~“868/915 MHz band (optional) O-QPSK PHY specifications”~~ several optional PHY types operating in several frequency bands results in

6.1.2.1 Channel numbering

Insert the following new sentence at the end of the last paragraph in 6.1.2.1:

An exception to this is the UWB PHY where specific mandatory and optional behaviors are as defined in 6.8a.11.1.

Insert after 6.1.2.1 the following new subclauses (6.1.2.1a and 6.1.2.1b):

6.1.2.1a Channel numbering for CSS PHY

A total of 14 frequency channels, numbered 0 to 13, are available across the 2.4 GHz band (see Table 1a). Different subsets of these frequency channels are available in different regions of the world. In North America and Europe, three frequency channels can be selected so that the nonoverlapping frequency channels are used.

Table 1a—Center frequencies of CSS

Frequency channel number	Frequency (MHz)
0	2412
1	2417
2	2422
3	2427
4	2432
5	2437
6	2442
7	2447
8	2452
9	2457
10	2462
11	2467
12	2472
13	2484

A channel frequency defines the center frequency of each band for CSS.

$$F_c = 2412 + 5(k - 1) \text{ in megahertz, for } k = 1, 2, \dots, 13$$

$$F_c = 2484 \text{ in megahertz, for } k = 14$$

where

k is the band number

Fourteen different frequency bands in combination with four different subchirp sequences form a set of $14 \times 4 = 56$ complex channels.

6.1.2.1b Channel numbering for UWB PHY

Sixteen channels, divided into three bands, are defined for the UWB PHY (see Table 1b). A compliant UWB device shall be capable of transmitting in at least one of three specified bands. A UWB device that implements the sub-gigahertz band shall implement channel 0. A UWB device that implements the low band shall support channel 3. The remaining low-band channels are optional. A UWB device that implements the high band shall support channel 9. The remaining high-band channels are optional.

Table 1b—UWB PHY channel frequencies

Channel number	Center frequency (MHz)	UWB band/mandatory
0	499.2	Sub-gigahertz
1	3494.4	Low band
2	3993.6	Low band
3	4492.8	Low band mandatory
4	3993.6	Low band
5	6489.6	High band
6	6988.8	High band
7	6489.6	High band
8	7488.0	High band
9	7987.2	High band mandatory
10	8486.4	High band
11	7987.2	High band
12	8985.6	High band
13	9484.8	High band
14	9984.0	High band
15	9484.8	High band

6.1.2.2 Channel pages

Change Table 2 (the entire table is not shown) as indicated:

Table 2—Channel page and channel number

Channel page (decimal)	Channel page (binary) (b31,b30,b29,b28,b27)	Channel number(s) (decimal)	Channel number description
<u>3</u>	<u>00011</u>	<u>0–13</u>	<u>Channels for CSS PHY</u>
<u>4</u>	<u>00100</u>	<u>0</u>	<u>Channel 0 is sub-gigahertz band for UWB PHY</u>
		<u>1–4</u>	<u>Channels 1 to 4 are low band for UWB PHY</u>
		<u>5–15</u>	<u>Channels 5 to 15 are high band for UWB PHY</u>
<u>35–31</u>	00011–1111	Reserved	Reserved

6.1.3 Minimum long interframe spacing (LIFS) and short interframe spacing (SIFS) periods

Insert the following new rows at the end of Table 3:

Table 3—Minimum LIFS and SIFS period

PHY	<i>macMinLIFSPeriod</i>	<i>macMinSIFSPeriod</i>	Units
2400–2483.5 MHz CSS	40	12	Symbols
UWB	40	12	Preamble symbols (see note)
NOTE—For the UWB PHY only, the IFS and SIFS periods are measured in units of preamble symbols for a preamble code length of 31. The actual time that this value represents depends on the PRF in use and the channel. See Table 39b (in 6.8a.5). UWB PHY preamble parameters in 6.8a.4 for these values.			

6.1.5 Transmit power

Insert the following new sentence at the end of the paragraph in 6.1.5:

For UWB PHYs, the parameter *phyTransmitPower* refers to the total power transmitted across the entire occupied bandwidth (not the dBm/MHz usually found in regulations).

6.2 PHY service specifications

6.2.1 PHY data service

6.2.1.1 PD-DATA.request

6.2.1.1.1 Semantics of the service primitive

Insert the following new parameters at the end of the list in 6.2.1.1.1 (before the closing parenthesis):

UWBPRF,
Ranging,
UWBPreableSymbolRepetitions,
DataRate

Insert the following new rows at the end of Table 6:

Table 6—PD-DATA.request parameters

Name	Type	Valid range	Description
UWBPRF	Enumeration	PRF_OFF, NOMINAL_4_M, NOMINAL_16_M, NOMINAL_64_M	The pulse repetition value of the transmitted PPDU. Non-UWB PHYs use a value of PRF_OFF.
Ranging	Enumeration	NON_RANGING, ALL_RANGING, PHY_HEADER_ONLY	A value of NON_RANGING indicates that ranging is not to be used with the PHY service data unit (PSDU) to be transmitted. A value of ALL_RANGING denotes ranging operations for this PSDU using both the ranging bit set to one in the PHR and counter operation enabled. A value of PHY_HEADER_ONLY denotes ranging operations for this PSDU using only the ranging bit in the PHR set to one. A value of NON_RANGING is used for non-UWB PHYs.
UWBPreableSymbolRepetitions	Enumeration	PSR_0, PSR_16, PSR_64, PSR_1024, PSR_4096	The preamble symbol repetitions of the UWB PHY frame to be transmitted by the PHY entity. PSR_16 indicates 16 preamble symbols are transmitted, PSR_64 indicates 64 preamble symbols are transmitted, and so on. A value of PSR_0 is used for non-UWB PHYs.

Table 6—PD-DATA.request parameters (continued)

Name	Type	Valid range	Description
DataRate	Enumeration	DATA_RATE_0, DATA_RATE_1, DATA_RATE_2, DATA_RATE_3, DATA_RATE_4	The data rate of the PHY frame to be transmitted by the PHY entity. A value of DATA_RATE_0 is used with a non-UWB or non-CSS PHY. A value of DATA_RATE_1 or DATA_RATE_2 is used with CSS PHYs (DATA_RATE_1 corresponds to 250 kb/s rate and 2 corresponds to DATA_RATE_2 Mb/s rate). DATA_RATE_1 through DATA_RATE_4 is used with UWB PHYs. See 6.8a.7.1 for UWB rate definitions.

6.2.1.1.3 Effect on receipt

Insert the following new paragraph at the end of 6.2.1.1.3:

If the PD-DATA.request primitive is received by a ranging-capable PHY with Ranging parameter equal to ALL_RANGING, then the ranging counter will begin counting from 0x00000001 as RMARKER leaves the transmit antenna.

6.2.1.2 PD-DATA.confirm

6.2.1.2.1 Semantics of the service primitive

Insert the following new parameters at the end of the list in 6.2.1.2.1 (before the closing parenthesis):

RangingReceived,
RangingCounterStart,
RangingCounterStop,
RangingTrackingInterval,
RangingOffset,
RangingFOM

Change Table 7 as shown:

Table 7—PD-DATA.confirm parameters

Name	Type	Valid range	Description
status	Enumeration	SUCCESS, RX_ON, TRX_OFF, or BUSY_TX, or UNSUPPORTED PRF UNSUPPORTED RANGING	The result of the request to transmit a packet. A value of UNSUPPORTED_PRF indicates that the PHY is not capable of transmitting at the requested PRF. A value of UNSUPPORTED_RANGING is returned if the PHY does not implement a ranging counter.

Table 7—PD-DATA.confirm parameters (continued)

Name	Type	Valid range	Description
<u>RangingReceived</u>	<u>Boolean</u>	<u>TRUE or FALSE</u>	<u>A value of FALSE indicates that ranging is either not supported in a UWB PHY or not to be indicated for the PSDU received. A value of TRUE denotes ranging operations requested for this PSDU. A value of FALSE is used for non-UWB PHYs.</u>
<u>RangingCounterStart</u>	<u>Unsigned Integer</u>	<u>0x00000000–0xFFFFFFFF</u>	<u>A 4-octet count of the time units corresponding to an RMARKER at the antenna at the beginning of a ranging exchange. A value of x00000000 is used if ranging is not supported or not enabled or this is not a UWB PHY. The value x00000000 is also used if the counter is not used for this PPDU. See 6.8a.15.1.</u>
<u>RangingCounterStop</u>	<u>Unsigned Integer</u>	<u>0x00000000–0xFFFFFFFF</u>	<u>A 4-octet count of the time units corresponding to an RMARKER at the antenna at the end of a ranging exchange. A value of x00000000 is used if ranging is not supported or not enabled or this is not a UWB PHY. The value x00000000 is also used if the counter is not used for this PPDU. See 6.8a.15.1.</u>
<u>RangingTrackingInterval</u>	<u>Unsigned integer</u>	<u>0x00000000–0xFFFFFFFF</u>	<u>A 4-octet count of the time units in a message exchange over which the tracking offset was measured. If tracking-based crystal characterization is not supported or this is not a UWB PHY, a value of x00000000 is used. See 6.8a.15.2.2</u>
<u>RangingOffset</u>	<u>Signed Magnitude Integer</u>	<u>0x000000–0x0FFFFF</u>	<u>3-octet count of the time units slipped or advanced by the radio tracking system over the course of the entire tracking interval. The top 4 bits are reserved and set to zero. The most significant of the active bits is the sign bit. See 6.8a.15.2.1.</u>
<u>RangingFOM</u>	<u>Integer</u>	<u>0x00–0x7F</u>	<u>One-octet FoM characterizing the ranging measurement. The most significant bit (MSB) is reserved and is zero. The remaining 7 bits are used in three subfields: Confidence Level, Confidence Interval, and Confidence Interval Scaling Factor. See 6.8a.15.3.</u>

6.2.1.3 PD-DATA.indication

6.2.1.3.1 Semantics of the service primitive

Insert the following new parameters at the end of the list in 6.2.1.3.1 (before the closing parenthesis):

UWBPRF,
UWBPreableSymbolRepetitions,
DataRate,
RangingReceived,
RangingCounterStart,
RangingCounterStop,
RangingTrackingInterval,
RangingOffset,
RangingFOM

Insert the following new rows at the end of Table 8:

Table 8—PD-DATA.indication parameters

Name	Type	Valid range	Description
UWBPRF	Enumeration	OFF, NOMINAL_4_M, NOMINAL_16_M, NOMINAL_64_M	The pulse repetition value of the received PPDU. Non-UWB PHYs use a value of OFF.
UWBPreable-SymbolRepetitions	Enumeration	PSR_0, PSR_16, PSR_64, PSR_1024, PSR_4096	The preamble symbol repetitions of the UWB PHY frame received by the PHY entity. PSR_16 indicates 16 preamble symbols, PSR_64 indicates 64 preamble symbols, and so on. A value of PSR_0 is used for non-UWB PHYs.
DataRate	Enumeration	DATA_RATE_0, DATA_RATE_1, DATA_RATE_2, DATA_RATE_3, DATA_RATE_4	The data rate of the PHY frame to be transmitted by the PHY entity. A value of DATA_RATE_0 is used with a non-UWB or non-CSS PHY. A value of DATA_RATE_1 or DATA_RATE_2 is used with CSS PHYs (DATA_RATE_1 corresponds to 250 kb/s rate and 2 corresponds to DATA_RATE_2 Mb/s rate). DATA_RATE_1 though DATA_RATE_4 is used with UWB PHYs (DATA_RATE_1 corresponds to 110 kb/s rate and DATA_RATE_4 corresponds to highest rate allowed for the current channel and PRF). See 6.8a.7.1 for UWB rate definitions.
RangingReceived	Boolean	TRUE or FALSE	A value of FALSE indicates that ranging is either not supported in a UWB PHY or not to be used for the PSDU received. A value of TRUE denotes ranging operations requested for this PSDU. A value of FALSE is used for non-UWB PHYs.

Table 8—PD-DATA.indication parameters (continued)

Name	Type	Valid range	Description
RangingCounterStart	Unsigned Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the time units corresponding to an RMARKER at the antenna at the beginning of a ranging exchange. A value of x00000000 is used if ranging is not supported or not enabled or this is not a UWB PHY. The value x00000000 is also used if the counter is not used for this PPDU. See 6.8a.15.1.
RangingCounterStop	Unsigned Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the time units corresponding to an RMARKER at the antenna at the end of a ranging exchange. A value of x00000000 is used if ranging is not supported or not enabled or this is not a UWB PHY. The value x00000000 is also used if the counter is not used for this PPDU. See 6.8a.15.1.
RangingTrackingInterval	Unsigned Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the time units in a message exchange over which the tracking offset was measured. If tracking-based crystal characterization is not supported or this is not a UWB PHY, a value of x00000000 is used. See 6.8a.15.2.2.
RangingOffset	Signed Magnitude Integer	0x000000–0x0FFFFFFF	A 3-octet count of the time units slipped or advanced by the radio tracking system over the course of the entire tracking interval. The top 4 bits are reserved and set to zero. The most significant of the active bits is the sign bit. See 6.8a.15.2.1.
RangingFOM	Integer	0x00–0x7F	A 1-octet FoM characterizing the ranging measurement. The MSB is reserved and is zero. The remaining 7 bits are used in three subfields: Confidence Level, Confidence Interval, and Confidence Interval Scaling Factor. See 6.8a.15.3.

6.2.2 PHY management service

Insert the following new rows at the end of Table 9:

Table 9—PLME-SAP primitives

PLME-SAP primitive	Request	Confirm
PLME-DPS	6.2.2.11	6.2.2.12
PLME-SOUNDING	6.2.2.13	6.2.2.14
PLME-CALIBRATE	6.2.2.15	6.2.2.16

6.2.2.4 PLME-ED.confirm

6.2.2.4.1 Semantics of the service primitive

Change the second row of Table 11 as shown:

Table 11—PLME-ED.confirm parameters

Name	Type	Valid range	Description
EnergyLevel	Integer or List of Integers	0x00–0xFF	ED level for the current channel. If status is set to SUCCESS, this is the ED level for the ~~current channel~~ or, for UWB PHY types, a list of ED levels of size phyUWBScanBinsPerChannel. If status is not SUCCESS, this parameter is meaningless. ~~Otherwise, the value of this parameter will be ignored.~~

6.2.2.7 PLME-SET-TRX-STATE.request

Change the third item in the dashed list in the first paragraph of 6.2.2.7 as shown:

- Receiver enabled (RX_ON) and for ranging receivers, Receiver and ranging enabled (RX_WITH_RANGING_ON)

6.2.2.7.1 Semantics of the service primitive

Change Table 14 as shown:

Table 14—PLME-SET-TRX-STATE.request parameters

Name	Type	Valid range	Description
Status	Enumeration	RX_ON, TRX_OFF, FORCE_TRX_OFF, ~~or~~ TX_ON, <u>RX_WITH_RANGING_ON</u>	The new state in which to configure the transceiver. The value of <u>RX_WITH_RANGING_ON</u> is <u>present only for UWB PHYs.</u>

6.2.2.7.3 Effect on receipt

Insert the following new paragraphs at the end of 6.2.2.7.3:

For RDEV devices, RX_ON is extended with RX_WITH_RANGING_ON. If RX_ON is selected, then the receiver state is changed, but the ranging counter is not enabled. Behavior of the receiver when RX_WITH_RANGING_ON is selected is as described for RX_ON in the preceding paragraph. What is unique about RX_WITH_RANGING_ON is how it affects the ranging counter.

If RX_WITH_RANGING_ON is selected and the ranging counter is not already counting, the ranging counter begins counting from 0x00000001 upon the arrival at the receive antenna of the RMARKER of the next RFRAME. If the ranging counter is already counting when RX_WITH_RANGING_ON is asserted, then there is no effect. In this case, the counter continues counting from wherever it was. While the counter is counting, the PHY will capture the value of the counter upon the arrival at the receive antenna of the RMARKER of all RFRAMEs.