



**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)

802.15.4a™

*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 Alfvín, *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 Alfvín
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	Patricia Martigne	Tomoki Saito
Julian Hall	Frederick Martin	Syed Saleem
Robert C. Hall, Jr.	Abbie Mathew	John Sarallo
Shinsuke Hara	Taisuke Matsumoto	Sidney Schrum
Jeff Harris	Masafumi Matsumura	Erik Schylander
Vann Hasty	Michael Mc Laughlin	Alireza Seyed
Allen Heberling	John McCorkle	Sanjeev Sharma
Robert Heile	Michael McInnis	Stephen Shellhammer
Eric Heinze	Charlie Mellone	Siddharth Shetty
Barry Herold	Jim Meyer	John (Chih-Chung) Shi
Jin-Meng Ho	Klaus Meyer	Shusaku Shimada
Patrick Houghton	Akira Miura	Yuichi Shiraki
Robert Huang	Samuel Mo	Gadi Shor
Xiaojing Huang	Andreas Molisch	William Shvodian
Akira Ikeda	Mark Moore	Thomas Siep
Hideto Ikeda	Marco Naeve	Michael Sim
Tetsushi Ikegami	Ken Naganuma	Kazimierz Siwiak
Oyvind Janbu	Yves Paul Nakache	Zachary Smith
Yeong Min Jang	Hiroyuki Nakase	V. Somayazulu
Adrian Jennings	Saishankar Nandagopalan	Carl Stevenson
Ho-In Jeon	Chiu Ngo	Marinus Struik
Zzyy Hong Jiang (Chiang)	Erwin Noble	Kazuaki Takahashi
Jeyhan Karaoguz	Mizukoshi Nobuyuki	Kenichi Takizawa
Michael Kelly	Masaki Noda	Teik-Kheong Tan
Stuart Kerry	John (Jay) O'Conor	Mike Tanahashi
Haksun Kim	Knut Odman	James Taylor
Jae-Hyon Kim	Hiroyo Ogawa	John Terry
Jaeyoung Kim	Yasuyuki Okuma	Arnaud Tonnerre
JinKyeong Kim	Philip Orlik	Jarvis Tou
Yongsuk Kim	Laurent Ouvre	Jerry Upton
Young Hwan Kim	John Pardee	Robin Vaitonis
Kursat Kimyacioglu	Young Jin Park	Bart Van Poucke
Matthias Kindler	Nirmalendu Patra	Nanci Vogtli
Patrick Kinney	Dave Patton	Jerry Wang
Guenter Kleindl	Xiaoming Peng	Jing Wang
Ryuji Kohno	Tony Pollock	Chris Weber
Mike Krell	Robert Poor	Matthew Welborn
Yasushi Kudo	Clinton Powell	Richard Wilson
Haim Kupersmidt	Vidyasagar Premkumar	Gerald Wineinger
Yuzo Kuramochi	Yihong Qi	Andreas Wolf
Kyung Sup Kwak	Raad Raad	Marcus Wong
Jiun-You Lai	Ajay Rajkumar	Stephen Wood
Ismail Lakkis	Pekka Ranta	Patrick Worfolk
John Lampe	Dani Raphaeli	Tracy Wright
Jim Lansford	Gregg Raser	Xiaodong Wu
Colin Lanzl	Charles Razzell	Yu-Ming Wu
Kyung Kuk Lee	Joseph Reddy	Hirohisa Yamaguchi
Myung Lee	Ivan Reede	Kamya Yekeh Yazdandoost
Wooyong Lee	Yuko Rikuta	Su-Khiong Yong
David Leeper	Terry Robar	Zhan Yu
Henry Li	Glyn Roberts	Serdar Yurdakul
Huan-Bang Li	Richard Roberts	Honggang Zhang
Liang Li	Martin Rofheart	Bin Zhen
Haixiang Liang	Benjamin Rolfe	Frank Xiaojun Zheng
Darryn Lowe	Philippe Rouzet	Zheng Zhou
Ian Macnamara	Chandos Rypinski	Chunhui Zhu
Tadahiko Maeda	Saeid Safavi	
Akira Maeki	Zafer Sahinoglu	

Major contributions were received from the following individuals:

Rick Alfvén	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. Alfvén	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	Shinkyō 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	Rajesh 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 ~~striketrough~~ (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.

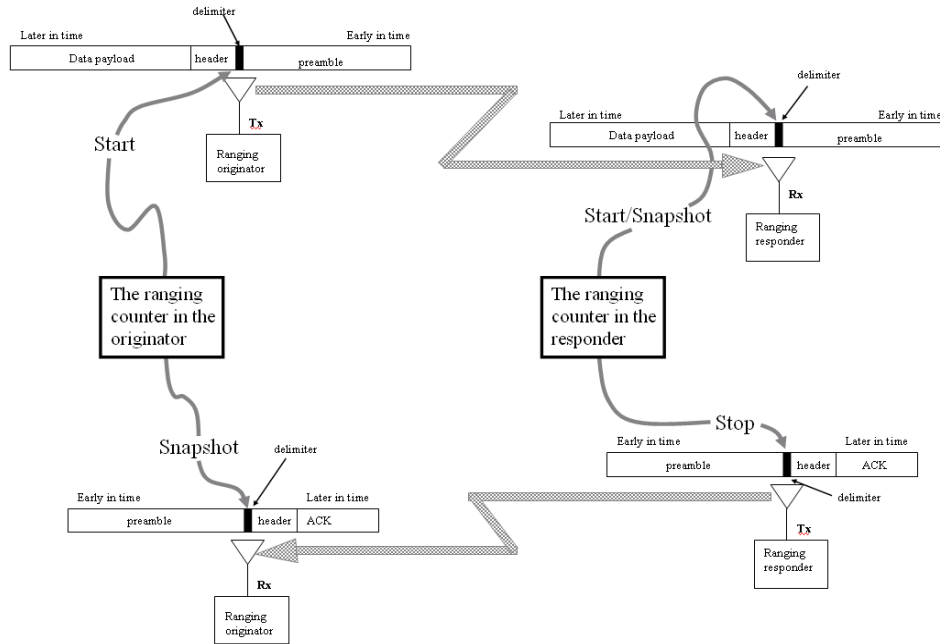


Figure 13a—The complete two-way ranging exchange

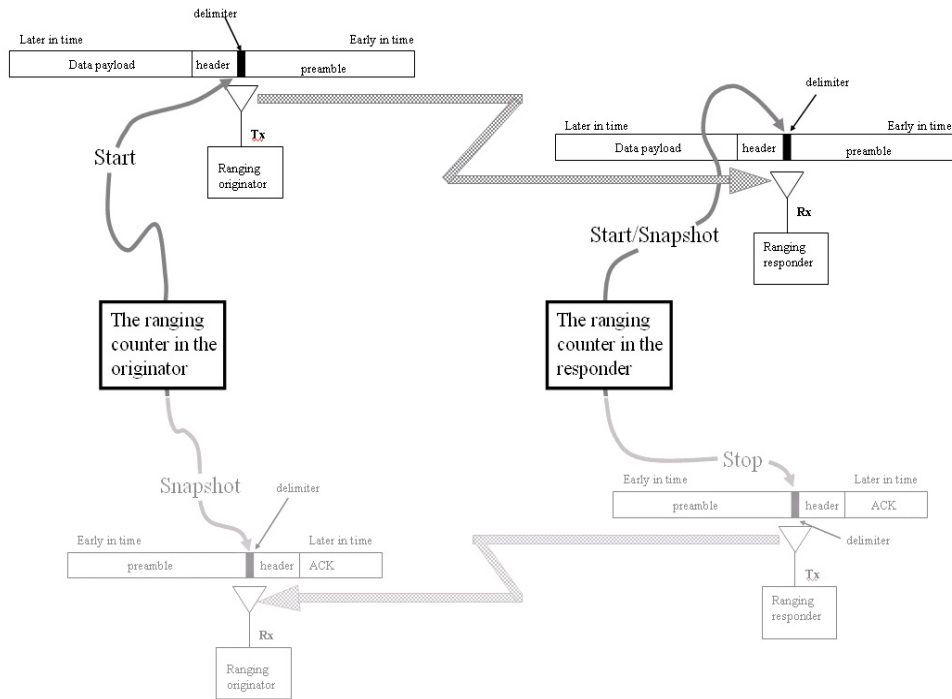


Figure 13b—The first frame of the two-way exchange

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.

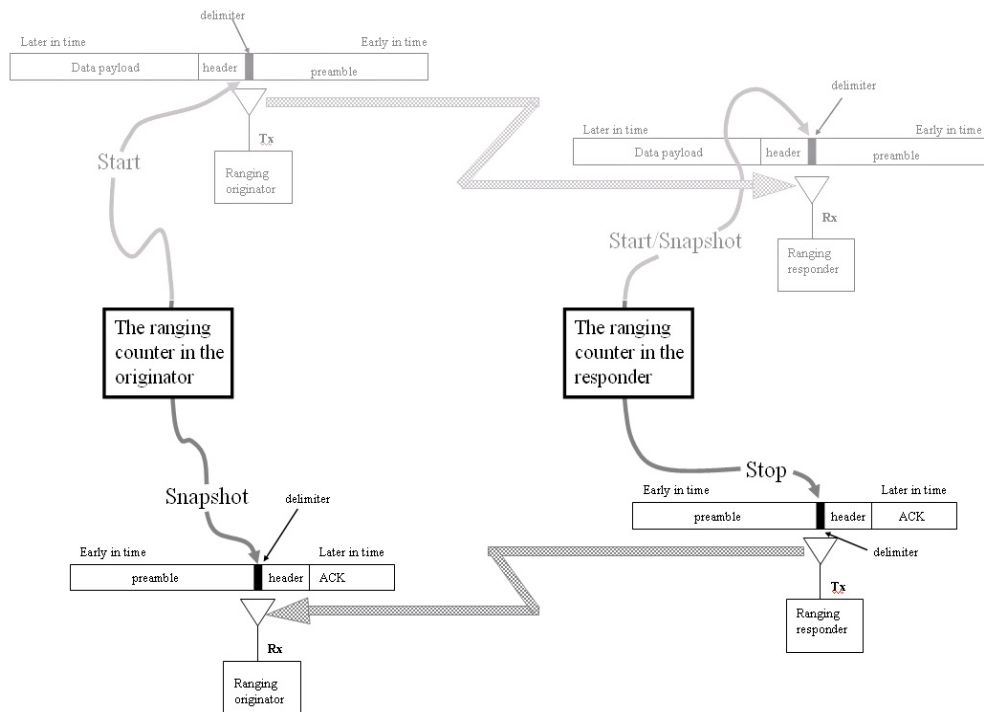


Figure 13c—The second frame of the two-way exchange

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.

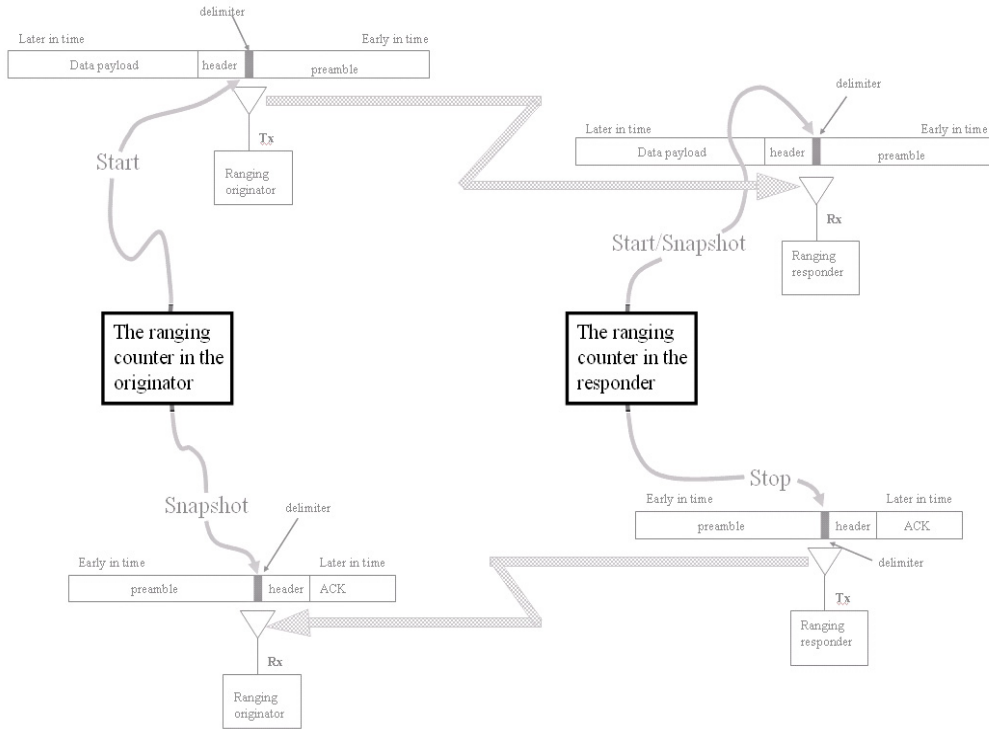


Figure 13d—The devices after the two-way exchange

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 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,
UWBPreAmbleSymbolRepetitions,
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.
UWBPreAmbleSymbolRepetitions	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. <u>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.</u>

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.
RangingTracking-Interval	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, 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.

Once every 2^{32} times, the counter will be wrapping through a value that would cause a final counter value of zero (after all corrections are applied) to be the proper and correct counter value to present in a timestamp report. There is nothing in this standard to preclude the ranging counter from wrapping through zero. However, this standard does give zero special meaning associated with devices that have no counter or have counters that are not running. Likewise, counter values of 0x00000001 have special meaning as they are presented for counter startup events. An RDEV with a running counter presenting a counter value of zero (or 0x00000001 when it is not just starting) will be algorithmically disruptive. If an RDEV with a running counter would ever normally present a counter value of zero or one, that RDEV shall instead present a value of 0x00000002. This occurrence will lead to a worst-case additional half centimeter ranging error.

The PHY does not detect or report the event if the ranging counter wraps through zero during a timed interval. Because the timestamp reports only provide the application with start times and stop times (as opposed to attempting to keep track of actual elapsed time), the application has all the information necessary to detect the instances when the counter wrapped through zero; and the application is responsible for making the corrections.

If the RDEV is not performing ranging operations, a significant power savings may be exhibited if the ranging counter is not enabled.

Insert after 6.2.2.10.3 the following new subclauses (6.2.2.11 through 6.2.2.16.3):

6.2.2.11 PLME-DPS.request (UWB PHYs only)

The PLME-DPS.request primitive attempts to set the DPS parameters. The PLME-DPS.request primitive is optional except for implementations providing ranging.

6.2.2.11.1 Semantics of the service primitive

The semantics of the PLME-DPS.request primitive is as follows:

```

PLME-DPS.request      (
                        TxDPSIndex,
                        RxDPSIndex,
                        )
    
```

Table 17a specifies the parameters for the PLME-DPS.request primitive.

Table 17a—PLME-DPS.request parameters

Name	Type	Valid range	Description
TxDPSIndex	Integer	0x00 0x0D–0x10 and 0x15–0x18	The index value for the transmitter. 0x00 disables the index and indicates that the <i>phyCurrentCode</i> value is to be used. See 6.8a.6.1 and Table 39e. 0x0D = index 13; 0x0E = index 14; 0x0F = index 15; ... 0x18 = index 24.
RxDPSIndex	Integer	0x00 0x0D–0x10 and 0x15–0x18	The index value for the transmitter. 0x00 disables the index and indicates that the <i>phyCurrentCode</i> value is to be used. See 6.8a.6.1 and Table 39e. 0x0D = index 13; 0x0E = index 14; 0x0F = index 15; ... 0x18 = index 24.

6.2.2.11.2 Appropriate usage

The PLME-DPS.request primitive is generated by the MLME and issued to its PLME whenever a change in dynamic preamble is required and whenever moving between dynamic preamble and current channel code (*phyCurrentCode*).

6.2.2.11.3 Effect on receipt

On receipt of the PLME-DPS.request primitive, the PLME will attempt to set the specified values of TxDPSIndex and RxDPSIndex. If the range of these parameters is invalid, the PLME-DPS.confirm primitive is used to report a status of INVALID_PARAMETER.

If the feature is not supported in the PHY, the PLME-DPS.confirm primitive is used to report a status of UNSUPPORTED_ATTRIBUTE.

If the requested operations are successfully completed, the PLME-DPS.confirm primitive is used to report a status of SUCCESS.

6.2.2.12 PLME-DPS.confirm (UWB PHYs only)

The PLME-DPS.confirm primitive reports the result of a request to change the settings for DPS. The PLME-DPS.confirm primitive is optional except for implementations providing ranging.

6.2.2.12.1 Semantics of the service primitive

The semantics of the PLME-DPS.confirm primitive is as follows:

```

PLME-DPS.confirm      (
                        Status
                        )
    
```

Table 17b specifies the parameter for the PLME-DPS.confirm primitive.

Table 17b—PLME-DPS.confirm parameter

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, UNSUPPORTED_ATTRIBUTE, INVALID_PARAMETER	The status of the attempt to set the DPS parameters.

6.2.2.12.2 When generated

The PLME-DPS.confirm primitive is generated by the PLME and issued to its MLME in response to a PLME-DPS.request primitive.

6.2.2.12.3 Appropriate usage

On receipt of the PLME-DPS.confirm primitive, the MLME is notified of the results of its request to set the TxDPSIndex and RxDPSIndex.

6.2.2.13 PLME-SOUNDING.request (UWB PHYs only)

The PLME-SOUNDING.request primitive attempts to have the PHY respond with channel sounding information. The PLME-SOUNDING.request primitive is optional except for implementations providing ranging. Although the PLME-SOUNDING.request primitive shall be supported by all RDEVs, the underlying sounding capability is optional in all cases.

6.2.2.13.1 Semantics of the service primitive

The semantics of the PLME-SOUNDING.request primitive is as follows:

```
PLME-SOUNDING.request      (
                             )
```

6.2.2.13.2 Appropriate usage

The PLME-SOUNDING.request primitive is generated by the MLME and issued to its PLME to request a PLME-SOUNDING.confirm primitive.

6.2.2.13.3 Effect on receipt

If the feature is supported in the UWB PHY, the PLME will issue the PLME-SOUNDING.confirm primitive with a status of SUCCESS and a list of SoundingPoints of SoundingSize in length.

If the PLME-SOUNDING.request primitive is generated by the MLME when there is no information present (e.g., when the PHY is in the process of performing a measurement), the PLME-SOUNDING.confirm primitive is used to report a status of NO_DATA.

If the PLME-SOUNDING.request primitive is generated by the MLME and the channel sounding capability is not present in the PHY, the PLME-SOUNDING.confirm primitive is used to report a status of UNSUPPORTED_ATTRIBUTE.

6.2.2.14 PLME-SOUNDING.confirm (UWB PHYs only)

The PLME-SOUNDING.confirm primitive reports the result of a request to the PHY to provide channel sounding information. The PLME-SOUNDING.confirm primitive is optional except for implementations providing ranging. Although the PLME-SOUNDING.confirm primitive shall be supported by all RDEVs, the underlying sounding capability is optional in all cases.

6.2.2.14.1 Semantics of the service primitive

The semantics of the PLME-SOUNDING.confirm primitive is as follows:

```
PLME-SOUNDING.confirm      (
                             Status,
                             SoundingSize,
                             SoundingList
                             )
```

Table 17c specifies the parameters for the PLME-SOUNDING.confirm primitive.

Table 17c—PLME-SOUNDING.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, NO_DATA, UNSUPPORTED_ATTRIBUTE	The status of the attempt to return sounding data.
SoundingSize	Unsigned Integer	0x0000–0xFFFF	Number of SoundingPoints to be returned. Each SoundingPoint is 4 octets.
SoundingList	List of Pairs of Signed Integers	0x00000000–0xFFFFFFFF for each element in the list. Each element in the list is a SoundingPoint. See Table 17d.	The list of sounding measurements. See 5.5.7.4.5.

Table 17d lists the parameters in the SoundingList. Each element of the SoundingList contains a SoundingTime and a SoundingAmplitude. The SoundingTime is a signed integer and the LSB represents a nominal 16 ps (1/128 of a chip time). A time of zero shall designate an amplitude value taken at the point indicated by RangingCounterStart. Positive time values shall indicate amplitudes that occurred earlier in time than the zero point. The SoundingAmplitude is a signed integer representing a relative linear measurement. The SoundingAmplitudes have no absolute meaning, only a relative meaning.

Table 17d—SoundingPoint subfields

Octets 3 and 2	Octets 1 and 0
SoundingTime	SoundingAmplitude

6.2.2.14.2 When generated

The PLME-SOUNDING.confirm primitive is generated by the PLME and issued to its MLME in response to a PLME-SOUNDING.request primitive. The PLME-SOUNDING.confirm primitive will return a status of SUCCESS to indicate channel sounding information is available and part of the PLME-SOUNDING.confirm parameters or return an error code of NO_DATA or UNSUPPORTED_ATTRIBUTE.

6.2.2.14.3 Appropriate usage

On receipt of the PLME-SOUNDING.confirm primitive, the MLME is notified of the results of the channel sounding information request. If the channel sounding information was available, the status parameter is set to SUCCESS. Otherwise, the status parameter will indicate an error.

6.2.2.15 PLME-CALIBRATE.request (UWB PHYs only)

The PLME-CALIBRATE.request primitive attempts to have the PHY respond with RMARKER offset information. The PLME-CALIBRATE.request primitive is optional except for implementations providing ranging.

6.2.2.15.1 Semantics of the service primitive

The semantics of the PLME-CALIBRATE.request primitive is as follows:

```
PLME-CALIBRATE.request      (
                               )
```

6.2.2.15.2 Appropriate usage

The PLME-CALIBRATE.request primitive is generated by the MLME and issued to its PLME to request a PLME-CALIBRATE.confirm primitive.

6.2.2.15.3 Effect on receipt

If the feature is supported in the UWB PHY, the PLME will issue the PLME-CALIBRATE.confirm primitive with a status of SUCCESS and a pair of integers CalTx_RMARKER_Offset and CalRx_RMARKER_Offset.

If the PLME-CALIBRATE.request primitive is generated by the MLME when there is no information present (e.g., when the PHY is in the process of performing a measurement), the PLME will issue the PLME-CALIBRATE.confirm primitive with a value of NO_DATA.

If the PLME-CALIBRATE.request primitive is generated by the MLME and the PHY does not support autonomous self-calibration, the PLME will issue the PLME-CALIBRATE.confirm primitive with a value of COMPUTATION_NEEDED. The COMPUTATION_NEEDED signals the higher layer that it should use the sounding primitives to finish the calibration (see 5.5.7.6.3).

If the PLME-CALIBRATE.request primitive is generated by the MLME and the channel sounding capability is not present in the PHY, the PLME will issue the PLME-CALIBRATE.confirm primitive with a value of UNSUPPORTED_ATTRIBUTE.

6.2.2.16 PLME-CALIBRATE.confirm (UWB PHYs only)

The PLME-CALIBRATE.confirm primitive reports the result of a request to the PHY to provide internal propagation path information. The PLME-CALIBRATE.confirm primitive is optional except for implementations providing ranging.

6.2.2.16.1 Semantics of the service primitive

The semantics of the PLME-CALIBRATE.confirm primitive is as follows:

```

PLME-CALIBRATE.confirm      (
                               Status,
                               CalTx_RMARKER_Offset,
                               CalRx_RMARKER_Offset,
                               )
    
```

Table 17e specifies the parameter for the PLME-CALIBRATE.confirm primitive.

Table 17e—PLME-CALIBRATE.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, COMPUTATION_NEEDED, NO_DATA, UNSUPPORTED_ATTRIBUTE	The status of the attempt to return sounding data.

Table 17e—PLME-CALIBRATE.confirm parameters (continued)

Name	Type	Valid range	Description
CalTx_RMARKER_Offset	Unsigned Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the propagation time from the ranging counter to the transmit antenna. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.
CalRx_RMARKER_Offset	Unsigned Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the propagation time from the receive antenna to the ranging counter. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.

6.2.2.16.2 When generated

The PLME-CALIBRATE.confirm primitive is generated by the PLME and issued to its MLME in response to a PLME-CALIBRATE.request primitive. The PLME-CALIBRATE.confirm primitive will return a status of SUCCESS to indicate channel propagation time information is available and part of the PLME-CALIBRATE.confirm parameters, return a status of COMPUTATION_NEEDED if the PHY lacks the computational resources to determine the offsets, or return an error code of NO_DATA or UNSUPPORTED_ATTRIBUTE.

6.2.2.16.3 Appropriate usage

On receipt of the PLME-CALIBRATE.confirm primitive, the MLME is notified of the results of the self-calibrate information request. If the RMARKER offset information was available, the status parameter is set to SUCCESS. If the PHY performed a sounding of a loopback path but lacks the computational resources to complete the processing of the sounding data, the status parameter is set to COMPUTATION_NEEDED. Otherwise, the status parameter will indicate an error.

6.2.3 PHY enumerations description

Insert the following new rows at the end of Table 18:

Table 18—PHY enumerations description

Enumeration	Value	Description
COMPUTATION_NEEDED	0x0c	PHY performed a sounding of a loopback path but lacks the computational resources to complete the processing of the sounding data. The next higher layer should use the sounding primitives to finish the calibration.
NO_DATA	0x0d	Indicates no calibration information is present, e.g., when the PHY is in the process of performing a measurement.
RX_WITH_RANGING_ON	0x0e	Indicates that the ranging counter is active.

Table 18—PHY enumerations description (continued)

Enumeration	Value	Description
PRF_OFF	0x0f	Non-UWB PHYs do not use PRF parameter; therefore, this value is used.
NOMINAL_4M	0x10	PRF is a nominal 4 MHz.
NOMINAL_16M	0x11	PRF is a nominal 16 MHz.
NOMINAL_64M	0x12	PRF is a nominal 64 MHz.
NON_RANGING	0x13	Either ranging is not supported or is not selected for the current PSDU.
ALL_RANGING	0x14	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.
PHY_HEADER_ONLY	0x15	Denotes ranging operations for this PSDU using only the ranging bit in the PHR set to one.
PSR_0	0x16	Non-UWB PHYs do not use preamble symbol repetitions parameter; therefore, this value is used.
PSR_16	0x17	Preamble symbol repetitions are 16 in number.
PSR_64	0x18	Preamble symbol repetitions are 64 in number.
PSR_1024	0x19	Preamble symbol repetitions are 1024 in number.
PSR_4096	0x1a	Preamble symbol repetitions are 4096 in number.
DATA_RATE_0	0x1b	PHYs that are neither UWB or CSS do not use the data rate parameter; therefore, this value is used.
DATA_RATE_1	0x1c	Data Rate parameter is 1.
DATA_RATE_2	0x1d	Data Rate parameter is 2.
DATA_RATE_3	0x1e	Data Rate parameter is 3.
DATA_RATE_4	0x1f	Data Rate parameter is 4.
UNSUPPORTED_PRF	0x20	Data.confirm error status returned when a corresponding Data.request command is issued with an unsupported PRF value.
UNSUPPORTED_RANGING	0x21	Data.confirm error status returned when a corresponding Data.request command is issued with Ranging = ALL_RANGING but the PHY does not support a ranging counter.

6.3 PPDU format

Change second item in the dashed list in the second paragraph of 6.3 as shown:

- A PHY header (PHR), which contains frame length information and, for UWB PHYs, rate, ranging, and preamble information

Change the fourth paragraph in 6.3 as shown:

The PPDU packet structure shall be formatted as illustrated in Figure 16, Figure 16a, or Figure 16b.

Change the title of Figure 16 as shown:

Figure 16—Format of the PPDU (except UWB and CSS)

Insert after Figure 16 the following new figures (Figure 16a and Figure 16b):

		Bits	Octets
		19	variable
Preamble	SFD	PHR (see 6.8a.7)	PSDU
SHR		PHR	PHY payload

Figure 16a—Format of the UWB PPDU

Data rate	SHR		PHR	PSDU
	Preamble	SFD		
1 Mb/s	8 symbols	4 symbols	2 symbols	variable
250 kb/s (optional)	20 symbols	4 symbols	8 symbols	variable

Figure 16b—Format of the CSS PPDU

NOTE—The preamble sequence includes the starting reference symbol, which is required for differential transmission.

6.3.1 Preamble field

Insert the following new paragraphs after the first paragraph in 6.3.1:

The Preamble field for the CSS PHY is defined in 6.5a.3.1.

The Preamble field for the UWB PHY is defined in 6.8a.6.

Change the second sentence in the paragraph starting “Preamble lengths for ASK” in 6.3.1 as shown:

For all PHYs except the ASK, CSS, and UWB PHYs, the bits

6.3.2 SFD field

Change the title of Table 20 as shown:

Table 20—SFD field length (except for ASK, CSS, and UWB PHYs)

Change the first sentence in the second paragraph of 6.3.2 as shown:

For all PHYs, except for the ASK, CSS, and UWB PHYs, the SFD

Change the title of Figure 17 as shown:

Figure 17—Format of the SFD field (except for ASK, UWB, and CSS PHYs)

Insert the following new paragraphs and table (Table 20a) after Figure 17:

Start-of-frame delimiter (SFD) bit sequences for the CSS PHY type are defined in Table 20a. Different SFD sequences are defined for the two different data rates. A SFD sequence from Table 20a shall be applied directly to both inputs (I and Q) of the QPSK mapper. A SFD sequence starts with bit 0.

Table 20a—CSS SFD sequence

Data rate	Bit (0:15)
1 Mb/s	-1 1 1 1 -1 1 -1 -1 1 -1 -1 1 1 1 -1 -1
250 kb/s (optional)	-1 1 1 1 1 -1 1 -1 -1 -1 1 -1 -1 -1 1 1

The SFD field for the UWB PHY is defined in 6.8a.6.2.

6.4 PHY constants and PIB attributes

6.4.2 PIB Attributes

Change the rows regarding *phyCCAMode*, *phyMaxFrameDuration*, *phySHRDuration*, and *phySymbolsPerOctet* in Table 23 (the entire table is not shown) as indicated and then insert the new rows at the end of the table:

Table 23—PHY PIB attributes

Attribute	Identifier	Type	Range	Description
<i>phyCCAMode</i>	0x03	Integer	1–36	The CCA mode (see 6.9.9).
<i>phyMaxFrameDuration</i> [†]	0x05	Integer	55, 212, 266, 1064 <u>except UWB and CSS PHYs</u>	The maximum number of symbols in a frame, <u>except for UWB and CSS PHYs</u> : = <i>phySHRDuration</i> + <u>ceiling</u> ([<i>aMaxPHYPacketSize</i> + 1] × <i>phySymbolsPerOctet</i>) For UWB PHYs, see 6.4.2.1. For CSS PHYs, one of two values depending on data rate. See 6.4.2.2.
<i>phySHRDuration</i> [†]	0x06	Integer	3, 7, 10, 40 <u>except UWB and CSS PHYs</u> . For UWB PHYs see 6.4.2.1 For CSS PHY, 12, 24.	The duration of the synchronization header (SHR) in symbols for the current PHY. For CSS PHY, a value of 12 corresponds to 1 Mb/s and 24 corresponds to 250 kb/s.

Table 23—PHY PIB attributes (continued)

Attribute	Identifier	Type	Range	Description
<i>phySymbolsPerOctet</i>	0x07	Float	0.4, <u>1.3</u> , 1.6, 2, <u>5.3</u> , 8	The number of symbols per octet for the current PHY. <u>For UWB PHYs, see 6.4.2.1.</u> <u>For CSS PHYs, 4/3 corresponds to 1 Mb/s and 32/6 corresponds to 250 kb/s.</u>
<i>phyPreambleSymbolLength</i>	0x08	Integer	1 or 0	0 indicates preamble symbol length is 31, 1 indicates that length 127 symbol is used. Present for UWB PHY.
<i>phyUWBDataRatesSupported</i> [†]	0x09	Bitmap	0x00–0x0f	A bit string that indicates the status (1= available, 0= unavailable) for each of the 4 valid data rates. The LSB of the bitmap refers to the lowest bit rate available in the operating channel as defined in Table 39g while the MSB indicates support for the highest bit rate in Table 39g. For example, a value of 0x0f indicates that all rates are supported while an implementation that supports only the mandatory rate shall have a value of 0x02.
<i>phyCSSLowDataRateSupported</i> [†]	0x0A	Boolean	TRUE or FALSE	A value of TRUE indicates that 250 kb/s is supported. Present for CSS PHY.
<i>phyUWBCCAModesSupported</i> [†]	0x0B	Bitmap	0x00–0x07	Representation of the three CCA modes for UWB PHYs. The LSB is for CCA mode 4, while the MSB is for CCA mode 6. UWB CCA modes are described in 6.9.9.
<i>phyUWB Pulse Shapes Supported</i> [†]	0x0C	Bitmap	0x00–0x0F	A bit string that indicates which of the optional pulse shapes the UWB PHY supports. There are a total of 4 pulse shape options: CoU (see 6.8a.13.1), CS (see 6.8a.13.2), LCP (see 6.8a.13.3), and chaotic (see Annex H). The most significant nibble (4 bits) is reserved while the lower nibble, b3,b2,b1,b0, is used to indicate CoU, CS, LCP, or chaotic pulses are supported. For example, a value of 1010 implies that CoU and LCP are supported while CS and chaotic are not.
<i>phyUWB Current Pulse Shape</i>	0x0D	Enumeration	MANDATORY, COU, CS, LCP, CHAOTIC	Indicates the current pulse shape setting of the UWB PHY. The mandatory pulse is described in 6.8a.12.1. Optional pulse shapes include CoU (see 6.8a.13.1), CS (see 6.8a.13.2), LCP (see 6.8a.13.3), and chaotic (see Annex H).

Table 23—PHY PIB attributes (continued)

Attribute	Identifier	Type	Range	Description
<i>phyUWBCoUpulse</i>	0x0E	Enumeration	CCh.1, CCh.2, ...CCh.6	Defines the slope of the frequency chirp and bandwidth of pulse. (Note that CCh.3–CCh.6 are valid only for wideband UWB channels, e.g., 4, 7, 11, or 15; see 6.8a.13.1 and Table 39k.)
<i>phyUWBSPulse</i>	0x0F	Enumeration	No.1, No.2, ..., No.6	Defines the group delay of the continuous spectrum filter. (Note that No.3–No.6 are valid only for wideband UWB channels, e.g., 4, 7, 11, or 15; see 6.8a.13.2 and Table 39l.)
<i>phyUWBChaoticPulse</i>	0x10	Boolean	TRUE or FALSE	Defines if device uses the chaotic pulse option (see Annex H).
<i>phyUWBLCWeight1</i>	0x11	Signed integer	0x00–0xFF	The weights are represented as signed 8-bit number in twos-complement form. A value of 0x80 represents –1 while a value of 0x7F represents 1. An additional constraint on the total energy in the taps is required so that the overall combined pulse has the same energy as the mandatory pulse (see 6.8a.13.3).
<i>phyUWBLCWeight2</i>	0x12	Signed integer	0x00–0xFF	The weights are represented as signed 8-bit number in twos-complement form. A value of 0x80 represents –1 while a value of 0x7F represents 1. An additional constraint on the total energy in the taps is required so that the overall combined pulse has the same energy as the mandatory pulse (see 6.8a.13.3).
<i>phyUWBLCWeight3</i>	0x13	Signed integer	0x00–0xFF	The weights are represented as signed 8-bit number in twos-complement form. A value of 0x80 represents –1 while a value of 0x7F represents 1. An additional constraint on the total energy in the taps is required so that the overall combined pulse has the same energy as the mandatory pulse (see 6.8a.13.3).
<i>phyUWBLCWeight4</i>	0x14	Signed integer	0x00–0xFF	The weights are represented as signed 8-bit number in twos-complement form. A value of 0x80 represents –1 while a value of 0x7F represents 1. An additional constraint on the total energy in the taps is required so that the overall combined pulse has the same energy as the mandatory pulse (see 6.8a.13.3).
<i>phyUWBLCDelay2</i>	0x16	integer	0x00–0xff	The delays are represented as 8-bit numbers and range from 0 to 4 ns. Thus, the resolution is $4/255 = 15.625$ ps. For example, a value of 0x00 represents 0 while 0x02 represents 31.25 ps (see 6.8a.13.3).

Table 23—PHY PIB attributes (continued)

Attribute	Identifier	Type	Range	Description
<i>phyUWBLCPDelay3</i>	0x17	integer	0x00–0xff	The delays are represented as 8-bit numbers and range from 0 to 4 ns. Thus, the resolution is $4/255 = 15.625$ ps. For example, a value of 0x00 represents 0 while 0x02 represents 31.25 ps (see 6.8a.13.3).
<i>phyUWBLCPDelay4</i>	0x18	integer	0x00–0xff	The delays are represented as 8-bit numbers and range from 0 to 4 ns. Thus, the resolution is $4/255 = 15.625$ ps. For example, a value of 0x00 represents 0 while 0x02 represents 31.25 ps (see 6.8a.13.3).
<i>phyRanging-Capabilities[†]</i>	0x19	Bitmap	0x00–0x07	The capabilities of the UWB PHY to support ranging options. The upper 5 bits are reserved. 0x01 is ranging support; 0x02 is crystal offset characterization support; 0x04 is DPS support.
<i>phyCurrentCode</i>	0x1A	Integer	0–24	For UWB and CSS PHYs, the current CDMA subchannel. 0 is for non-CDMA PHYs; for UWB PHYs, this represents the current preamble code in use by the transmitter and may be any value from 1–24 as these are the preamble code indices shown in Table 39d and Table 39e. Values 1–4 are for CSS PHY subchirps 1–4.
<i>phyNativePRF</i>	0x1B	Enumeration	0–3	For UWB PHYs, the native PRF. 0 is for non-UWB PHYs; 1 is for PRF of 4; 2 is for a PRF of 16; and 3 is for PHYs that have no preference.
<i>phyUWBScanBinsPer-Channel</i>	0x1C	Integer	0–255	Number of frequency intervals used to scan each UWB channel (scan resolution). Set to zero for non-UWB PHYs.
<i>phyUWBInserted-PreambleInterval</i>	0x1D	Enumeration	0, 4	For UWB PHYs operating with CCA mode 6, the time interval between two neighboring inserted preamble symbols in the data portion. The resolution is a data symbol duration at a nominal data rate of 850 kb/s for all channels (see 6.8a.14). Set to 4 for UWB PHY in CCA mode 6; otherwise, set to 0. See Table 39a.
<i>phyTx_RMARKER_Offset</i>	0x1E	Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the propagation time from the ranging counter to the transmit antenna. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.

Table 23—PHY PIB attributes (continued)

Attribute	Identifier	Type	Range	Description
<i>phyRx_RMARKER_Offset</i>	0x1F	Integer	0x00000000–0xFFFFFFFF	A 4-octet count of the propagation time from the receive antenna to the ranging counter. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.
<i>phyRFRAME-ProcessingTime</i>	0x20	Integer	0x00–0xFF	A 1-octet count of the processing time required by the PHY to handle an arriving RFRAME. The LSB represents 2 ms. The meaning of the value is that if a sequence of RFRAMEs arrive separated by <i>phyRFRAMEProcessingTime</i> , then the PHY can keep up with the processing indefinitely.

Insert after 6.4.2 the following new subclauses (6.4.2.1 through 6.4.2.3):

6.4.2.1 PIB values *phyMaxFrameDuration*, *phySHRDuration* for UWB

For the UWB PHY types, the values for the PIB attributes *phyMaxFrameDuration* and *phySHRDuration* vary depending upon the UWB PHY operating mode. The symbol duration varies by data rate and is different for preamble symbols and body symbols. Also note that the preamble and PHR are sent at a different data rate from the body.

$$phyMaxFrameDuration_{\eta s} = T_{SHR} + T_{PHR} + T_{PSDU} + T_{CCApreamble}$$

$$T_{SHR} = T_{psym} \times (N_{preambleSymbols} + N_{SFD})$$

$$T_{PHR} = N_{PHR} \times T_{dsym1M}$$

$$T_{PSDU} = \left[T_{dsym} \times \frac{N_{PSDUoctets} \times N_{symPerOctet}}{R_{FEC}} \right]$$

$$T_{CCApreamble} = \left[T_{psym} \times \frac{T_{PHR} + T_{PSDU}}{4 \times T_{dsym1M}} \right]$$

where

T_{psym} base symbol time for preamble symbols for the selected channel (see Table 39c)

T_{dsym} PSDU data symbol duration (see Table 39a)

T_{dsym1M} nominal 1 Mb/s data symbol duration for the selected channel (see Table 39a)

$N_{preambleSymbols}$ {16, 64, 1024, 4096} symbols
(UWBpreambleSymbolRepetitions in PD-DATA.request primitive; see 6.2.1.1)

N_{SFD} number of delimiter symbols = $\begin{pmatrix} 64 & \text{for 110 kb/s data rate} \\ 8 & \text{for other data rates} \end{pmatrix}$

N_{PHR} 16 = number of bits in the PHR

$N_{PSDUoctets}$	$aMaxPHYPacketSize$
$N_{symPerOctet}$	symbols per octet, uncoded = 8
R_{FEC}	Reed-Solomon FEC rate = 0.87 (T_{dysm} in Table 39a includes coded bits/symbol when convolutional code is used)

The values for T_{psym} and T_{dysm} are given in Table 39a and Table 39c.

The actual time for the SHR duration is

$$phySHRDuration_{\eta s} = T_{preamble} + T_{PHR}$$

The PHY PIB attribute $phySHRDuration$ should be expressed in number of symbols. We can use the value T_{psym} from Table 39b appropriate for the channel, thus

$$phySHRDuration = (T_{preamble} + T_{PHR})/T_{psym}$$

6.4.2.2 PIB values $phyMaxFrameDuration$ for CSS

For the CSS PHY type, the values of the attribute $phyMaxFrameDuration$ depend on the selected data rate of the PSDU.

For the mandatory data rate (1 Mb/s), $phyMaxFrameDuration$ is calculated as follows:

$$phyMaxFrameDuration_{1M} =$$

$$phySHRDuration_{1M} + [1.5 + 3/4 \times ceiling(4/3 \times aMaxPHYPacketSize)] \times phySymbolsPerOctet_{1M}$$

For the optional data rate (250 kb/s), $phyMaxFrameDuration$ is calculated as follows:

$$phyMaxFrameDuration_{250k} =$$

$$phySHRDuration_{250k} + 3 \times ceiling(1/3 \times [1.5 + aMaxPHYPacketSize]) \times phySymbolsPerOctet_{250k}$$

6.4.2.3 PIB values for internal propagation times UWB

For the UWB PHY types, the values for these PIB attributes represents the internal propagation times.

Attribute $phyTx_RMARKER_Offset$ is a 4-octet count of the propagation time from the ranging counter to the transmit antenna. The LSB represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.

Attribute $phyRx_RMARKER_Offset$ is a 4-octet count of the propagation time from the receive antenna to the ranging counter. The LSB represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.

6.5 2450 MHz PHY specifications

Insert after 6.5.3.4 the following new subclauses (6.5a through 6.5a.5.4):

6.5a 2450 MHz PHY chirp spread spectrum (CSS) PHY

The requirements for the 2450 MHz CSS PHY are specified in 6.5a.1 through 6.5a.5.

6.5a.1 Data rates

The data rate of the CSS (2450 MHz) PHY shall be 1 Mb/s. An additional data rate of 250 kb/s shall be optional.

6.5a.2 Modulation and spreading

This PHY uses CSS techniques in combination with differential quadrature phase-shift keying (DQPSK) and 8-ary or 64-ary bi-orthogonal coding for 1 Mb/s data rate or 250 kb/s data rate, respectively. By using alternating time gaps in conjunction with sequences of chirp signals (subchirps) in different frequency subbands with different chirp directions, this CSS PHY provides subchirp sequence division as well as frequency division.

6.5a.2.1 Reference modulator diagram

The functional block diagram in Figure 20a is provided as a reference for specifying the 2450 MHz CSS PHY modulation for both 1 Mb/s and optional 250 kb/s. The number in each block refers to the subclause that describes that function. All binary data contained in the PHR and PSDU shall be encoded using the modulation shown in Figure 20a.

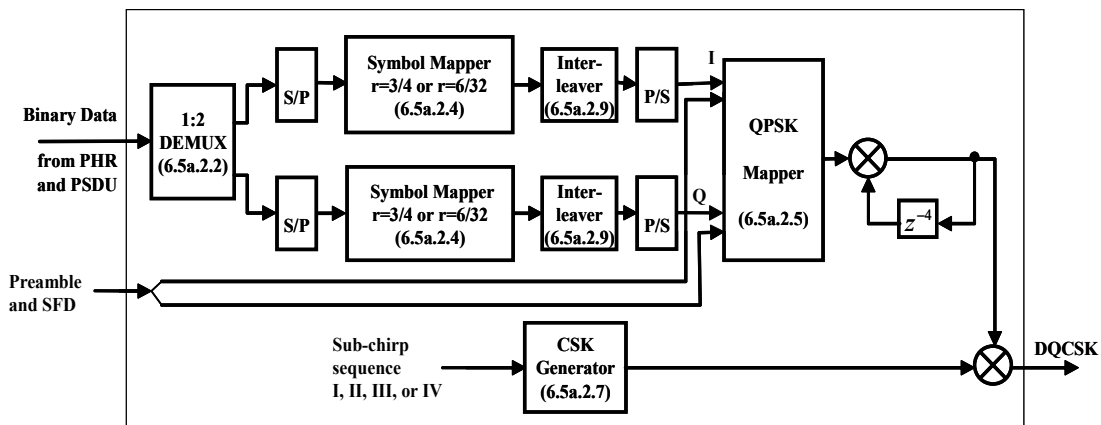


Figure 20a—Differential bi-orthogonal quaternary-chirp-shift-keying modulator and spreading ($r = 3/4$ for 8-ary 1 Mb/s, $r = 3/16$ for 64-ary 250 kb/s)

6.5a.2.2 De-multiplexer (DEMUX)

For each packet, the initial position of the DEMUX shown in Figure 20a shall be set to serve the I path (upper path). Thus the first bit of the incoming stream of information bits of a packet shall be switched to the I path, and the second bit shall be switched to the Q path.

6.5a.2.3 Serial-to-parallel mapping (S/P)

By using two serial-to-parallel converters, the substreams are independently partitioned into sets of bits to form data symbols. For the mandatory data rate of 1 Mb/s, a data symbol shall consist of three bits. Within the binary data symbol (b0,b1,b2), the first input data bit for each of I and Q is assigned b0, and the third input data bit is assigned b2. For the optional data rate of 250 kb/s, a data symbol shall consist of 6 bits. Within the binary data symbol (b0,b1,b2,b3,b4,b5), the first input data bit for each of I and Q is assigned b0, and the sixth input data bit is assigned b5.

6.5a.2.4 Data-symbol-to-bi-orthogonal-codeword mapping

Each 3-bit data symbol shall be mapped onto a 4-chip bi-orthogonal codeword (c0, c1, c2, c3) for the 1 Mb/s data rate as specified in Table 26a. Each 6-bit data symbol shall be mapped onto a 32-chip bi-orthogonal codeword (c0, c1, c2, ..., c31) for the optional 250 kb/s data rate as specified in Table 26b.

Table 26a—8-ary bi-orthogonal mapping (r = 3/4, 1 Mb/s)

Data symbol (decimal)	Data symbol (binary) (b0 b1 b2)	Codeword (co c1 c2 c3)
0	000	1 1 1 1
1	001	1 -1 1 -1
2	010	1 1 -1 -1
3	011	1 -1 -1 1
4	100	-1 -1 -1 -1
5	101	-1 1 -1 1
6	110	-1 -1 1 1
7	111	-1 1 1 -1

Table 26b—64-ary bi-orthogonal mapping (r = 3/16, 250 kb/s)

Data symbol (decimal)	Data symbol (binary) (b0 b1 b2 b3 b4 b5)	Codeword (co c1 c2 ... c31)
0	000000	1 1
1	000001	1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1
2	000010	1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1
3	000011	1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1
4	000100	1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1
5	000101	1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1

Table 26b—64-ary bi-orthogonal mapping ($r = 3/16$, 250 kb/s) (continued)

Data symbol (decimal)	Data symbol (binary) (b0 b1 b2 b3 b4 b5)	Codeword (c0 c1 c2 ... c31)
6	000110	1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1
7	000111	1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1
8	001000	1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1
9	001001	1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1
10	001010	1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1
11	001011	1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1
12	001100	1 1 1 1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1
13	001101	1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1
14	001110	1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1
15	001111	1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1
16	010000	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1
17	010001	1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1
18	010010	1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1
19	010011	1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1
20	010100	1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1
21	010101	1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1
22	010110	1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1
23	010111	1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1
24	011000	1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1
25	011001	1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1

Table 26b—64-ary bi-orthogonal mapping ($r = 3/16$, 250 kb/s) (continued)

Data symbol (decimal)	Data symbol (binary) (b0 b1 b2 b3 b4 b5)	Codeword (c0 c1 c2 ... c31)
26	011010	1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1
27	011011	1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1
28	011100	1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 -1 -1 -1 -1
29	011101	1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1
30	011110	1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1
31	011111	1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1
32	100000	-1 -1
33	100001	-1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1
34	100010	-1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1
35	100011	-1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1
36	100100	-1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1
37	100101	-1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1
38	100110	-1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1
39	100111	-1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1
40	101000	-1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1
41	101001	-1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1
42	101010	-1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1
43	101011	-1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1
44	101100	-1 -1 -1 -1 1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 1 -1 -1 -1
45	101101	-1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1

Table 26b—64-ary bi-orthogonal mapping ($r = 3/16$, 250 kb/s) (continued)

Data symbol (decimal)	Data symbol (binary) (b0 b1 b2 b3 b4 b5)	Codeword (c0 c1 c2 ... c31)
46	101110	-1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1
47	101111	-1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1
48	110000	-1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
49	110001	-1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1
50	110010	-1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1
51	110011	-1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1
52	110100	-1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1
53	110101	-1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1
54	110110	-1 -1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1
55	110111	-1 1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1
56	111000	-1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1
57	111001	-1 1 -1 1 -1 1 -1 1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1
58	111010	-1 -1 1 1 -1 -1 1 1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1
59	111011	-1 1 1 -1 -1 1 1 -1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1
60	111100	-1 -1 -1 -1 1 1 1 1 1 1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 1
61	111101	-1 1 -1 1 1 -1 1 -1 1 -1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1
62	111110	-1 -1 1 1 1 1 -1 -1 1 1 -1 -1 -1 -1 1 1 1 1 -1 -1 -1 -1 1 1 -1 -1 1 1 1 1 -1 -1
63	111111	-1 1 1 -1 1 -1 -1 1 1 -1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 -1 -1 1 1 -1 1 -1 -1 1

6.5a.2.5 Parallel-to-serial converter (P/S) and QPSK symbol mapping

Each bi-orthogonal codeword shall be converted to a serial chip sequence. Within each 4-chip codeword (c0, c1, c2, c3) for the 1 Mb/s data rate, the least significant chip c0 is processed first, and the most significant chip c3 is processed last for I and Q, respectively. Within each 32-chip codeword (c0, c1,

c2, ... , c31) for the 250 kb/s data rate, the least significant chip c0 is processed first, and the most significant chip c31 is processed last for I and Q, respectively. Each pair of I and Q chips shall be mapped onto a QPSK symbol as specified in Table 26c.

Table 26c—QPSK symbol mapping

Input chips ($I_{n,k}$ $Q_{n,k}$)	Magnitude	Output phase (rad)
1, 1	1	0
-1, 1	1	$\pi/2$
1, -1	1	$-\pi/2$
-1, -1	1	π

6.5a.2.6 DQPSK coding

The stream of QPSK symbols shall be differentially encoded by using a differential encoder with a QPSK symbol feedback memory of length 4. (In other words, the phase differences between QPSK symbol 1 and 5, 2 and 6, 3 and 7, 4 and 8, and so on are computed.) For a detailed explanation of the index variables n and k , see 6.5a.4.3.

DQPSK output

$$e^{j\theta_{n,k}} = e^{j\theta_{n-1,k}} \times e^{j\phi_{n,k}}$$

where

$e^{j\phi_{n,k}}$ is DQPSK input

$e^{j\theta_{n-1,k}}$ is stored in feedback memory

For every packet, the initial values of all four feedback memory stages of the differential encoder shall be set

$$e^{j\pi/4}$$

or equivalently

$$\theta_{0,k} = \frac{\pi}{4} [rad]$$

6.5a.2.7 DQPSK-to-DQCSK modulation

The stream of DQPSK symbols shall be modulated onto the stream of subchirps that is generated by the chirp-shift keying (CSK) generator. The effect of the differential quadrature chirp-shift keying (DQCSK) modulation shall be that each subchirp is multiplied with a DQPSK value that has unit magnitude and has constant phase for the duration of the subchirp. An example of this operation can be found in 6.5a.4.6.

6.5a.2.8 CSK generator

The CSK generator shall periodically generate one of the four defined subchirp sequences (chirp symbols) as specified in 6.5a.4.3. Since each chirp symbol consists of four subchirps, the subchirp rate is four times higher than the chirp symbol rate.

6.5a.2.9 Bit interleaver

The bit interleaver is applied only for the optional data rate of 250 kb/s. The 32 chip bi-orthogonal codewords for the optional 250 kb/s data rate are interleaved prior to the parallel to serial converter. Bit interleaving provides robustness against double intra-symbol errors caused by the differential detector. The interleaver permutes the chips across two consecutive codewords for each of I and Q, independently.

The memory of the interleaver shall be initialized with zeros before the reception of a packet.

The data stream going into the interleaver shall be padded with zeros if the number of octets to be transmitted does not align with the bounds of the interleaver blocks.

The input-output relationship of this interleaver shall be given as follows:

Input

even-symbol (c0, c1, c2, c3, c4, c5, c6, c7, c8, c9, c10, c11, c12, c13, c14, c15,
c16, c17, c18, c19, c20, c21, c22, c23, c24, c25, c26, c27, c28, c29, c30, c31)
odd-symbol (d0, d1, d2, d3, d4, d5, d6, d7, d8, d9, d10, d11, d12, d13, d14, d15,
d16, d17, d18, d19, d20, d21, d22, d23, d24, d25, d26, d27, d28, d29, d30, d31)

Output

even-symbol (c0, c1, c2, c3, d20, d21, d22, d23, c8, c9, c10, c11, d28, d29, d30, d31,
c16, c17, c18, c19, d4, d5, d6, d7, c24, c25, c26, c27, d12, d13, d14, d15)
odd-symbol (d0, d1, d2, d3, c20, c21, c22, c23, d8, d9, d10, d11, c28, c29, c30, c31,
d16, d17, d18, d19, c4, c5, c6, c7, d24, d25, d26, d27, c12, c13, c14, c15)

NOTE—As shown in Figure 20a, coding is applied to every bit following the SFD. The first codeword generated shall be counted as zero and thus is even.

6.5a.3 CSS frame format

6.5a.3.1 Preamble

The preamble for 1 Mb/s consists of 8 chirp symbols, and the preamble for optional 250 kb/s consists of 20 chirp symbols as specified in Table 26d. The preamble sequence from Table 26d should be applied directly to both I input and the Q input of QPSK.

Table 26d—Preamble sequence

Data rate	Preamble sequence
1 Mb/s	ones(0:31)
250 kb/s	ones(0:79)

where

$\text{ones}(0:N)$ for integer number N is defined as a 1-by- N matrix of ones

6.5a.3.2 SFD field

The SFD field for CSS is defined in 6.3.2.

6.5a.3.3 PHY header (PHR)

The format of the PHR is shown in Figure 20b.

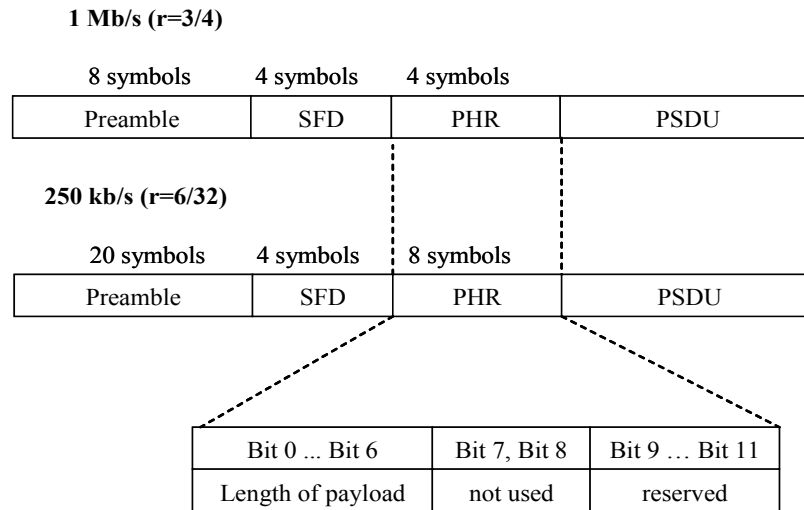


Figure 20b—Format of the CSS PHR

6.5a.4 Waveform and subchirp sequences

Four individual chirp signals, here called subchirps, shall be concatenated to form a full chirp symbol (subchirp sequence), which occupies two adjacent frequency subbands. Four different subchirp sequences are defined. Each subchirp is weighted with a raised cosine window in the time domain.

6.5a.4.1 Graphical presentation of chirp symbols (subchirp sequences)

Four different sequences of subchirp signals are available for use. Figure 20c shows the four different chirp symbols (subchirp sequences) as time frequency diagrams. It can be seen that four subchirps, which have either a linear down-chirp characteristic or a linear up-chirp characteristic, and a center frequency, which has either a positive or a negative frequency offset, are concatenated. The frequency discontinuities between subsequent chirps will not impact the spectrum because the signal amplitude will be zero at these points.

6.5a.4.2 Active usage of time gaps

In conjunction with the subchirp sequence, different pairs of time gaps are defined. The time gaps are chosen to make the four sequences even closer to being orthogonal. The time gaps shall be applied alternatively between subsequent chirp symbols as shown in Figure 20d. The values of the time gaps are calculated from the timing parameters specified in Table 26g (in 6.5a.4.3).




Figure 20c—Four different combinations of subchirps

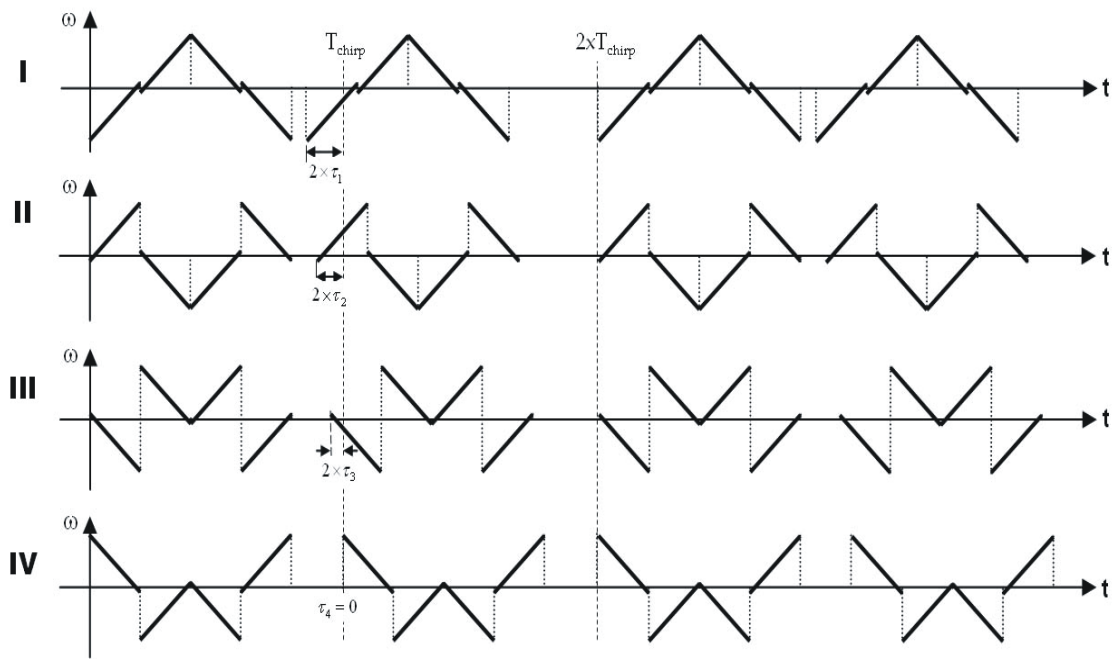


Figure 20d—Four different time-gap pairs for the four different subchirp sequences

6.5a.4.3 Mathematical representation of the continuous time CSS base-band signal

The mathematical representation of the continuous time-domain base-band signal $\tilde{s}^m(t)$ built of chirp symbols (subchirp sequences) as shown in Figure 20c with alternating time gaps as shown in Figure 20d is given by Equation (1a). The subchirp sequence with its associated time gap is defined to be a chirp symbol.

$$\begin{aligned} \tilde{s}^m(t) &= \sum_{n=0}^{\infty} \tilde{s}^m(t, n) \\ &= \sum_{n=0}^{\infty} \sum_{k=1}^4 \tilde{c}_{n,k} \exp \left[j \left(\hat{\omega}_{k,m} + \frac{\mu}{2} \xi_{k,m} (t - T_{n,k,m}) \right) (t - T_{n,k,m}) \right] \times P_{RC}(t - T_{n,k,m}) \end{aligned} \quad (1a)$$

Where $m = 1, 2, 3, 4$ (I, II, III, and IV in Figure 20c) defines which of the four different possible chirp symbols (subchirp sequences) is used, $n = 0, 1, 2 \dots$ is the sequence number of the chirp symbols.

The $\tilde{c}_{n,k}$ of $s(t)$ in the Equation (1a) is the sequence of the complex data that consists of in-phase data $a_{n,k}$ and quadrature-phase data $b_{n,k}$ as the output of DQPSK coding.

The possible value of $a_{n,k}$ and $b_{n,k}$ are +1 or -1.

$$\tilde{c}_{n,k} = a_{n,k} + jb_{n,k}$$

where

- n is the sequence number of chirp symbols
- $k = 0, 1, 2, \text{ and } 3$ is the subchirp index
- j is $\sqrt{-1}$

$\hat{\omega}_{k,m} = 2\pi \times f_{k,n}$ are the center frequencies of the subchirp signals. This value depends on m and $k=1, 2, 3, 4$, which defines the subchirp number in the subchirp sequence.

$T_{n,k,m}$ as expressed in Equation (1b) defines the starting time of the actual subchirp signal to be generated. It is determined by T_{chirp} , which is the average duration of a chirp symbol, and by T_{sub} , which is the duration of a subchirp signal.

$$T_{n,k,m} = \left(k + \frac{1}{2} \right) T_{sub} + n T_{chirp} - (1 - (-1)^n) \tau_m \quad (1b)$$

The constant μ defines the characteristics of the subchirp signal. A value of $\mu = 2\pi \times 7.3158 \times 10^{12} [\text{rad}/\text{sec}^2]$ shall be used.

The function P_{RC} , which is defined in 6.5a.4.4, is a windowing function that is equal to zero at the edges and outside of the subchirp centered at time zero.

The constant τ_m is either not added or added twice and thus determines (but is not identical to) the time gap that was applied between two subsequent subchirp sequences as shown in Figure 20d.

Table 26e shows the values for the subband center frequencies, Table 26f the subchirp directions, and Table 26g the timing parameters in Equation (1a). It should be noted that these time and frequency parameters are assumed to be derived from a reference crystal in a locked manner. In other words, any relative errors in chirp subband center frequencies, chirp rate, and time gaps are equal.

Table 26e—Equation (1a) numerical parameters subband center frequencies, $f_{k,m}$ (MHz)

$m \backslash k$	1	2	3	4
1	$f_c - 3.15$	$f_c + 3.15$	$f_c + 3.15$	$f_c - 3.15$
2	$f_c + 3.15$	$f_c - 3.15$	$f_c - 3.15$	$f_c + 3.15$
3	$f_c - 3.15$	$f_c + 3.15$	$f_c + 3.15$	$f_c - 3.15$
4	$f_c + 3.15$	$f_c - 3.15$	$f_c - 3.15$	$f_c + 3.15$

Table 26f—Equation (1a) numerical parameters subchirp directions, $\zeta_{k,m}$

$m \backslash k$	1	2	3	4
1	+1	+1	-1	-1
2	+1	-1	+1	-1
3	-1	-1	+1	+1
4	-1	+1	-1	+1

Table 26g—Equation (1a) numerical parameters—timing parameters

Symbol	Value	Multiple of 1/32MHz
T_{chirp}	6 μ s	192
T_{sub}	1.1875 μ s	38
τ_1	468.75 μ s	15
τ_2	312.5 ns	10
τ_4	156.25 ns	5
τ_4	0 ns	0

6.5a.4.4 Raised cosine window for chirp pulse shaping

The raised-cosine time-window described by Equation (1c) shall be used to shape the subchirp. The raised cosine window $P_{RC}(t)$ is applied to every subchirp signal in the time domain. See Figure 20e.

$$P_{RC}(t) = \begin{cases} 1 & |t| \leq \frac{(1-\alpha)T_{sub}}{(1+\alpha)2} \\ \frac{1}{2} \left[1 + \cos \left(\frac{(1+\alpha)\pi}{\alpha T_{sub}} \left(|t| - \frac{(1-\alpha)T_{sub}}{(1+\alpha)2} \right) \right) \right] & \frac{(1-\alpha)T_{sub}}{(1+\alpha)2} < |t| \leq \frac{T_{sub}}{2} \\ 0 & |t| > \frac{T_{sub}}{2} \end{cases} \quad (1c)$$

where

$$\alpha = 0.25$$




Figure 20e—Subchirp time-domain pulse shaping

6.5a.4.5 Subchirp transmission order

During each chirp symbol period, subchirp 1 ($k = 1$) is transmitted first, and subchirp 4 ($k = 4$) is transmitted last.

6.5a.4.6 Example of CSK signal generation

An example for the modulation of one chirp symbol is provided in this subclause to illustrate each step from DEMUX to the output of the reference modulator as shown in Figure 20a. The scenario parameters are as follows:

- The initial values of all four feedback memory stages of the differential encoder is set $e^{(j\pi)/4}$.
- The data bit rate is 1 Mb/s.

Input binary data

0 1 0 1 1 0

Demux

I-path: 0 0 1
Q-path: 1 1 0

Serial-to-parallel mapping

I-path: {1 0 0}
Q-path: {1 1 0}

Bi-orthogonal mapping ($r = 3/4$)

I-path: 1 -1 1 -1
Q-path: -1 -1 1 1

Parallel-to-serial and QPSK symbol mapping

Mapper input: $(1 - j)$, $(-1 - j)$, $(1 + j)$, $(-1 + j)$
QPSK output phase: $-\pi/2$, π , 0 , $\pi/2$

D-QPSK coding

Initial phase of four feedback memory for D-QPSK: all $\pi/4$
D-QPSK coder output phase: $-\pi/4$, $-3\pi/4$, $\pi/4$, $3\pi/4$

D-QPSK-to-D-QCSK modulation output and subchirp sequence of D-QCSK output

$[\exp(-j\pi/4) \times \text{subchirp}(k=1), \exp(-j3\pi/4) \times \text{subchirp}(k=2), \exp(j\pi/4) \times \text{subchirp}(k=3), \exp(j3\pi/4) \times \text{subchirp}(k=4)]$

6.5a.5 2450 MHz band CSS radio specification

In addition to meeting regional regulatory requirements, CSS devices operating in the 2450 MHz band shall also meet the radio requirements in 6.5a.5.1 through 6.5a.5.4.

6.5a.5.1 Transmit power spectral density (PSD) mask and signal tolerance

The transmitted spectral power density of a CSS signal $s(t)$ shall be within the relative limits specified in the template shown in Figure 20f. The average spectral power shall be made using 100 kHz resolution bandwidth and a 1 kHz video bandwidth. For the relative limit, the reference level shall be the highest average spectral power measured within ± 11 MHz of the carrier frequency. Specifically, the normalized frequency spectrum to the peak value in the signal bandwidth $|f - f_c| \leq 7$ MHz shall be less than or equal to -30 dB in the stop band $11 \text{ MHz} \leq |f - f_c| \leq 22 \text{ MHz}$ and shall be less than or equal to -50 dB in the stop band $|f - f_c| > 22 \text{ MHz}$. For testing the transmitted spectral power density, a $2^{15} - 1$ pseudo-random binary sequence (PRBS) shall be used as input data.

As additional criteria for the compliance of a CSS signal, the mean square error shall be used. Let $s(t)$ be the baseband CSS signal that is given in Equation (1a). Then the implemented signal, $s_{impl}(t)$, shall satisfy the following equation:

$$mmse = \min_{A, \tau_d, \phi} \left(\frac{\int_0^T |s^m(t) - A \times s_{impl}^m(t - \tau_d) e^{j\phi}|^2 dt}{\int_0^T |s^m(t)|^2 dt} \right) \leq 0.005$$

where the constants A , τ_d , and ϕ are used to minimize the mean squared error. The constant T_{chirp} is the period of the CSS symbol. The $c_{n,k}$ of $s(t)$ in Equation (1a) is the constant data $(1 + j1)$ for the measurement for all n and k .

6.5a.5.2 Symbol rate

The 2450 MHz PHY DQCSK symbol rate shall be 166.667 ksymbol/s (1/6 Msymbol/s) ± 40 ppm.

6.5a.5.3 Receiver sensitivity

Under the conditions specified in 6.1.6, a compliant device shall be capable of achieving a sensitivity of -85 dBm or better for 1 Mb/s and -91 dBm or better for 250 kb/s.

6.5a.5.4 Receiver jamming resistance

Table 26h gives minimum jamming resistance levels. A nonoverlapping adjacent channel is defined to have a center frequency offset of 25 MHz. A nonoverlapping alternate channel is defined to have a center frequency offset of 50 MHz. The adjacent channel rejection shall be measured as follows: The desired signal shall be a compliant 2450 MHz CSS signal of pseudo-random data. The desired signal is input to the receiver at a level 3 dB above the maximum allowed receiver sensitivity given in 6.5a.5.3. In the adjacent or the alternate channel, a CSS signal of the same or a different subchirp sequence as the victim device is input at the relative level specified in Table 26h. The test shall be performed for only one interfering signal at a time. The receiver shall meet the error rate criteria defined in 6.1.6 under these conditions.




Figure 20f—Transmit PSD mask

Table 26h—Minimum receiver jamming resistance levels for 2450 MHz CSS PHY

Data rate	Nonoverlapping adjacent channel rejection (25 MHz offset) (dB)	Nonoverlapping alternate channel rejection (50 MHz offset) (dB)
1 Mb/s	34	48
250 kb/s (optional)	38	52

6.8 868/915 MHz band (optional) O-QPSK PHY specification

Insert after 6.8.3.5 the following new subclauses (6.8a through 6.8a.15.3):

6.8a UWB PHY specification

The UWB PHY waveform is based upon an impulse radio signaling scheme using band-limited data pulses. The UWB PHY supports three independent bands of operation:

- The sub-gigahertz band, which consists of a single channel and occupies the spectrum from 249.6 MHz to 749.6 MHz
- The low band, which consists of four channels and occupies the spectrum from 3.1 GHz to 4.8 GHz
- The high band, which consists of eleven channels and occupies the spectrum from 6.0 GHz to 10.6 GHz

Within each channel, there is support for at least two complex channels that have unique length 31 SHR preamble codes. The combination of a channel and a preamble code is termed a *complex channel*. A compliant device shall implement support for at least one of the channels (0,3 or 9) in Table 39i (in 6.8a.11.1). In addition, each device shall support the two unique length 31 preamble codes for the implemented channels as defined in Table 39d (in 6.8a.6.1). Support for the other channels listed in Table 39i is optional.

A combination of burst position modulation (BPM) and binary phase-shift keying (BPSK) is used to support both coherent and noncoherent receivers using a common signaling scheme. The combined BPM-BPSK is used to modulate the symbols, with each symbol being composed of an active burst of UWB pulses. The various data rates are supported through the use of variable-length bursts. Figure 27a shows the sequence of processing steps used to create and modulate a UWB PHY packet. The sequence of steps indicated here for the transmitter is used as a basis for explaining the creation of the UWB PHY waveform specified in the PHY of this standard. Note that the receiver portion of Figure 27a is informative and meant only as a guide to the essential steps that any compliant UWB receiver needs to implement in order to successfully decode the transmitted signal.

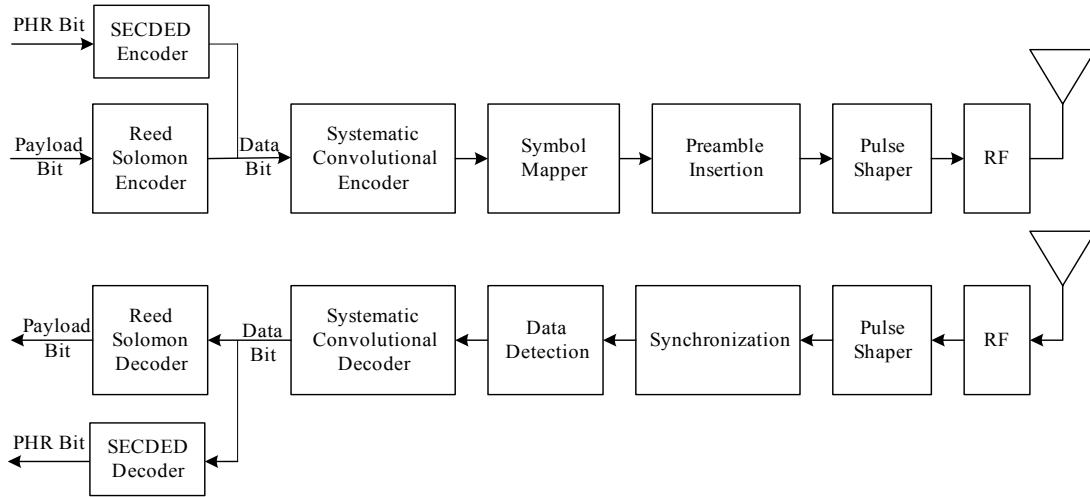


Figure 27a—PHY signal flow

6.8a.1 UWB frame format

Figure 27b shows the format for the UWB frame, which is composed of three major components: the SHR preamble, the PHR, and the PSDU. For convenience, the PPDU packet structure is presented so that the leftmost field as written in this standard shall be transmitted or received first. All multiple octet fields shall be transmitted or received least significant octet first, and each octet shall be transmitted or received LSB first. The same transmission order should apply to data fields transferred between the PHY and MAC sublayer.

The SHR preamble is first, followed by the PHR, and finally the PSDU. As shown in Figure 27b, the SHR preamble is always sent at the base rate for the preamble code defined in Table 39b (in 6.8a.5). Note that each UWB-compliant device shall support the length 31 preamble codes specified in Table 39d and that two base rates corresponding to the two mandatory PRFs result for this code length. The mandatory SHR preamble base rates are, therefore, 1.01 Msymbol/s and 0.25 Msymbol/s as indicated in Table 39b.

The PHR is sent at a nominal rate of 850 kb/s for all data rates above 850 kb/s and at a nominal of 110 kb/s for the nominal data rate of 110 kb/s. The PSDU is sent at the desired information data rate as defined in Table 39a.

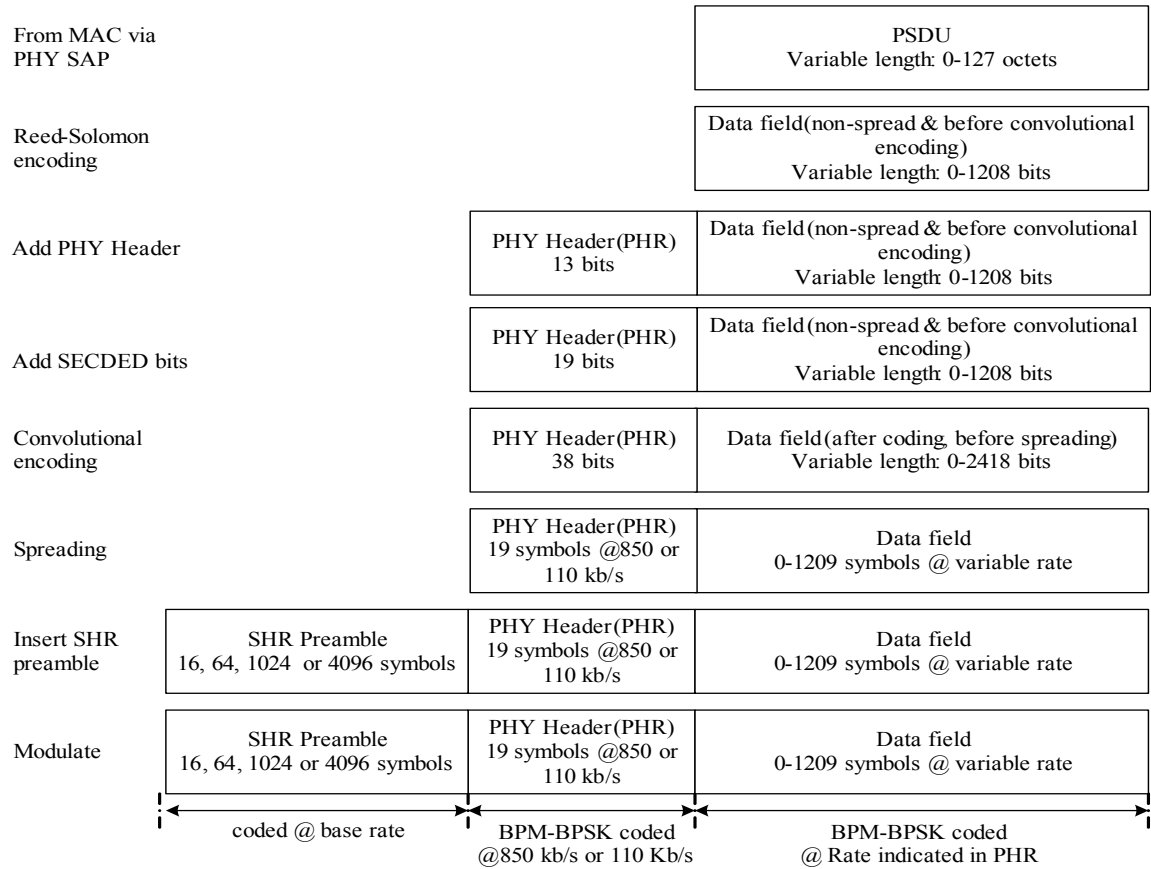


Figure 27b—PPDU encoding process

6.8a.2 PPDU encoding process

The encoding process is composed of many steps as illustrated in Figure 27b. The details of these steps are fully described in later subclauses, as noted in the following list, which is intended to facilitate an understanding of those details:

- Perform Reed-Solomon encoding on PSDU as described in 6.8a.10.1.
- Produce the PHR as described in 6.8a.7.1.
- Add SECDED check bits to PHR as described in 6.8a.7.2 and prepend to the PSDU.
- Perform further convolutional coding as described in 6.8a.10.2. Note that in some instances at the 27 Mb/s data rate, the convolutional encoding of the data field is effectively bypassed and two data bits are encoded per BPM-BPSK symbol.
- Modulate and spread PSDU according to the method described in 6.8a.9.1 and 6.8a.9.2 The PHR is modulated using BPM-BPSK at either 850 kb/s or 110 kb/s and the data field is modulated at the rate specified in the PHR.
- Produce the SHR preamble field from the SYNC field (used for AGC convergence, diversity selection, timing acquisition, and coarse frequency acquisition) and the SFD field (used to indicate the start of frame). The SYNC and SFD fields are described in 6.8a.6.1 and 6.8a.6.2, respectively.

6.8a.3 UWB PHY symbol structure

In the BPM-BPSK modulation scheme, a UWB PHY symbol is capable of carrying two bits of information: one bit is used to determine the position of a burst of pulses while an additional bit is used to modulate the phase (polarity) of this same burst.

The structure and timing of a UWB PHY symbol is illustrated in Figure 27c. Each symbol shall consist of an integer number of possible chip positions, N_c , each with duration T_c . The overall symbol duration denoted by T_{dsym} is given by $T_{dsym} = N_c T_c$. Furthermore, each symbol is divided into two BPM intervals each with duration $T_{BPM} = T_{dsym}/2$, which enables binary position modulation.

A burst is formed by grouping N_{cpb} consecutive chips and has duration $T_{burst} = N_{cpb} T_c$. The location of the burst in either the first half or second half of the symbol indicates one bit of information. Additionally, the phase of the burst (either -1 or $+1$) is used to indicate a second bit of information.

In each UWB PHY symbol interval, a single burst event shall be transmitted. The fact that burst duration is typically much shorter than the BPM duration, i.e., $T_{burst} \ll T_{BPM}$, provides for some multi-user access interference rejection in the form of time hopping. The total number of burst durations per symbol, N_{burst} , is given by $N_{burst} = T_{dsym}/T_{burst}$. In order to limit the amount of inter-symbol interference caused by multipath, only the first half of each T_{BPM} period shall contain a burst. Therefore, only the first $N_{hop} = N_{burst}/4$ possible burst positions are candidate hopping burst positions within each BPM interval. Each burst position can be varied on a symbol-to-symbol basis according to a time hopping code as described in 6.8a.9.

6.8a.4 PSDU timing parameters

The PSDU rate-dependent parameters and timing-related parameters are summarized in Table 39a. Within each UWB channel $\{0:15\}$, the peak PRF shall be 499.2 MHz. This rate corresponds to the highest frequency at which a compliant transmitter shall emit pulses. Additionally, the mean PRF is defined as the total number of pulses emitted during a symbol period divided by the length of the symbol duration. During the SHR preamble portion of a UWB frame, the peak and mean PRFs are essentially the same since pulses are emitted uniformly during each preamble symbol. During the data portion of a PPDU, however, the peak and mean PRFs differ due to the grouping of pulses into consecutive chip durations.

There are two possible preamble code lengths (31 or 127) and three possible mean PRFs (15.6 MHz, 3.90 MHz, and 62.4 MHz). A compliant device shall implement support for the preamble code length of 31 and shall also support both the 15.6 MHz and 3.90 MHz mean PRFs for the PSDU as depicted in Table 39a. The use of the length 127 code is optional; when implemented, the mean PRF of the PSDU shall be 62.4 MHz.

UWB channels $\{4, 7, 11, 15\}$ are all optional channels and are differentiated from other UWB channels by the larger bandwidth (> 500 MHz) of the transmitted signals. These channels overlap the existing lower bandwidth channels. The larger bandwidth enables devices operating in these channels to transmit at a higher power (for fixed PSD constraints), and thus they may achieve longer communication range. The larger bandwidth pulses offer enhanced multipath resistance. Additionally, larger bandwidth leads to more accurate range estimates. The admissible data rates, preamble code lengths, PRFs, and modulation timing parameters are listed in Table 39a.

Each UWB channel allows for several data rates (table column “Bit Rate”) that are obtained by modifying the number of chips within a burst (“# Chips Per Burst”), while the total number of possible burst positions (“#Burst Positions Per Symbol”) remains constant. Therefore, the symbol duration, T_{dsym} , changes to obtain the stated symbol rate and bit rates.




Figure 27c—UWB PHY symbol structure

Table 39a—UWB PHY rate-dependent and timing-related parameters

Channel Number	Peak PRF MHz	Bandwidth MHz	Preamble Code Length	Modulation & Coding			Data Symbol Structure						Data		
				Viterbi Rate	RS Rate	Overall FEC Rate	#Burst Positions per Symbol N_{burst}	# Hop Bursts N_{hop}	# Chips Per Burst N_{cpb}	#Chips Per Symbol	Burst Duration T_{burst} (ns)	Symbol Duration T_{sym} (ns)	Symbol Rate (MHz)	Bit Rate Mb/s	Mean PRF (MHz)
{0,3, 5;6, 8;10, 12;14}	499.2	499.2	31	0.5	0.87	0.44	32	8	128	4096	256.41	8205.13	0.12	0.11	15.60
	499.2	499.2	31	0.5	0.87	0.44	32	8	16	512	32.05	1025.64	0.98	0.85	15.60
	499.2	499.2	31	0.5	0.87	0.44	32	8	2	64	4.01	128.21	7.80	6.81	15.60
	499.2	499.2	31	1	0.87	0.87	32	8	1	32	2.00	64.10	15.60	27.24	15.60
{0,3, 5;6, 8;10, 12;14}	499.2	499.2	31	0.5	0.87	0.44	128	32	32	4096	64.10	8205.13	0.12	0.11	3.90
	499.2	499.2	31	0.5	0.87	0.44	128	32	4	512	8.01	1025.64	0.98	0.85	3.90
	499.2	499.2	31	0.5	0.87	0.44	128	32	2	256	4.01	512.82	1.95	1.70	3.90
	499.2	499.2	31	1	0.87	0.87	128	32	1	128	2.00	256.41	3.90	6.81	3.90
{0,3, 5;6, 8;10, 12;14}	499.2	499.2	127	0.5	0.87	0.44	8	2	512	4096	1025.64	8205.13	0.12	0.11	62.40
	499.2	499.2	127	0.5	0.87	0.44	8	2	64	512	128.21	1025.64	0.98	0.85	62.40
	499.2	499.2	127	0.5	0.87	0.44	8	2	8	64	16.03	128.21	7.80	6.81	62.40
	499.2	499.2	127	0.5	0.87	0.44	8	2	2	16	4.01	32.05	31.20	27.24	62.40
{4, 11}	499.2	1331.2	31	0.5	0.87	0.44	32	8	128	4096	256.41	8205.13	0.12	0.11	15.60
	499.2	1331.2	31	0.5	0.87	0.44	32	8	16	512	32.05	1025.64	0.98	0.85	15.60
	499.2	1331.2	31	0.5	0.87	0.44	32	8	2	64	4.01	128.21	7.80	6.81	15.60
	499.2	1331.2	31	1	0.87	0.87	32	8	1	32	2.00	64.10	15.60	27.24	15.60
7	499.2	1331.2	127	0.5	0.87	0.44	8	2	512	4096	1025.64	8205.13	0.12	0.11	62.40
	499.2	1331.2	127	0.5	0.87	0.44	8	2	64	512	128.21	1025.64	0.98	0.85	62.40
	499.2	1331.2	127	0.5	0.87	0.44	8	2	8	64	16.03	128.21	7.80	6.81	62.40
	499.2	1331.2	127	0.5	0.87	0.44	8	2	2	16	4.01	32.05	31.20	27.24	62.40
7	499.2	1081.6	31	0.5	0.87	0.44	32	8	128	4096	256.41	8205.13	0.12	0.11	15.60
	499.2	1081.6	31	0.5	0.87	0.44	32	8	16	512	32.05	1025.64	0.98	0.85	15.60
	499.2	1081.6	31	0.5	0.87	0.44	32	8	2	64	4.01	128.21	7.80	6.81	15.60
	499.2	1081.6	31	1	0.87	0.87	32	8	1	32	2.00	64.10	15.60	27.24	15.60
7	499.2	1081.6	127	0.5	0.87	0.44	8	2	512	4096	1025.64	8205.13	0.12	0.11	62.40
	499.2	1081.6	127	0.5	0.87	0.44	8	2	64	512	128.21	1025.64	0.98	0.85	62.40
	499.2	1081.6	127	0.5	0.87	0.44	8	2	8	64	16.03	128.21	7.80	6.81	62.40
	499.2	1081.6	127	0.5	0.87	0.44	8	2	2	16	4.01	32.05	31.20	27.24	62.40
15	499.2	1354.97	31	0.5	0.87	0.44	32	8	128	4096	256.41	8205.13	0.12	0.11	15.60
	499.2	1354.97	31	0.5	0.87	0.44	32	8	16	512	32.05	1025.64	0.98	0.85	15.60
	499.2	1354.97	31	0.5	0.87	0.44	32	8	2	64	4.01	128.21	7.80	6.81	15.60
	499.2	1354.97	31	1	0.87	0.87	32	8	1	32	2.00	64.10	15.60	27.24	15.60
15	499.2	1354.97	127	0.5	0.87	0.44	8	2	512	4096	1025.64	8205.13	0.12	0.11	62.40
	499.2	1354.97	127	0.5	0.87	0.44	8	2	64	512	128.21	1025.64	0.98	0.85	62.40
	499.2	1354.97	127	0.5	0.87	0.44	8	2	8	64	16.03	128.21	7.80	6.81	62.40
	499.2	1354.97	127	0.5	0.87	0.44	8	2	2	16	4.01	32.05	31.20	27.24	62.40

As stated above, the UWB PHY contains several optional data rates, preamble code lengths, and PRF. Each row in Table 39a completely describes all timing parameters shown in Figure 27c for each permitted combination of channel number, preamble code length, and PRF. Subclauses 6.8a.4.1 through 6.8a.4.15 describe in more detail how each field in Table 39a is computed.

6.8a.4.1 Channel Number parameter

This column identifies the UWB PHY channel numbers where the remaining PSDU timing parameters in the current row are valid. Association between channel number and center frequency is given in Table 39i.

6.8a.4.2 Peak PRF MHz parameter

The peak PRF states the highest frequency in megahertz at which a compliant transmitter shall emit pulses. The peak PRF is also used to derive the chip duration T_c by the formula $T_c = 1/(peakPRF)$. The value of T_c is approximately 2 ns.

6.8a.4.3 Bandwidth MHz parameter

The bandwidth denotes the 3 dB bandwidth of the UWB pulses. Note that the bandwidth is not necessarily the inverse of the chip duration T_c . Pulse shape and bandwidth are further defined in 6.8a.12.1.

6.8a.4.4 Preamble Code Length parameter

The value denotes the length of the preamble code length to be used during the SHR portion of a data frame. The code length together with the channel number defines a complex channel. Individual codes to be used on each channel are given in Table 39d (length 31) and Table 39e (length 127).

6.8a.4.5 Viterbi Rate parameter

This value determines the rate of the convolutional code applied to the PSDU data bits. A value of 1 indicates that no convolutional coding is applied while a value of 0.5 indicates that a rate 1/2 code as described in 6.8a.10.2 is applied to the PSDU data bits.

6.8a.4.6 RS Rate parameter

This is the (63,55) Reed-Solomon code rate, which is approximately 0.87. The Reed-Solomon code is applied to all the PSDU data bits that are transmitted by the UWB PHY. Reed-Solomon encoding is further described in 6.8a.10.1.

6.8a.4.7 Overall FEC Rate parameter

The overall FEC rate is determined by the product of the Viterbi rate and the Reed-Solomon rate and has either a value of 0.44 or 0.87.

6.8a.4.8 Burst Positions per Symbol parameter

This is the total number of possible burst positions in a data symbol duration. N_{burst} has been chosen so that for each mean PRF a data symbol consists of a fixed number of burst durations.

6.8a.4.9 Hop Bursts parameter

This is the number of burst positions that may contain an active burst, that is, a burst containing UWB pulses. The value is computed as $N_{hop} = N_{burst}/4$.

6.8a.4.10 Chips per Burst parameter

This is the number of chip T_c durations within each burst period T_{burst} . Each burst consists of a multiple number of consecutive chips (see Figure 27c). Depending on the data rate to be used in the transmission of the PSDU, the number of chips in a burst varies, e.g., for low data rates, the burst consists of more chip periods than for high data rates. Particular, values of N_{cpb} have been selected so that the following is a valid data rate: $(2 \times \text{Overall FEC rate}) / (N_{cpb} \times N_{burst} \times T_c)$.

6.8a.4.11 Burst Duration parameter

This is simply the duration of a burst and is computed as $T_{burst} = N_{cpb} \times T_c$.

6.8a.4.12 Symbol Duration parameter

This is a the duration of a modulated and coded PSDU symbol on the air and is computed as follows:
 $T_{dsym} = N_{burst} \times T_{burst}$.

6.8a.4.13 Symbol Rate parameter

This is the inverse of the PSDU symbol duration $1/T_{dsym}$.

6.8a.4.14 Bit Rate parameter

This is the user information rate considering FEC and is computed as follows:

$$\text{Bit Rate} = 2 \times (\text{Overall FEC Rate}) / T_{dsym}$$

6.8a.4.15 Mean PRF parameter

This is the average PRF during the PSDU portion of a PHY frame and is computed as follows:

$$\text{Mean PRF} = N_{cpb} / T_{dsym}$$

6.8a.5 Preamble timing parameters

Due to the variability in the preamble code length and the PRF, there are several admissible values for the timing parameters of a preamble symbol. These values are summarized in Table 39b. In this subclause, a preamble symbol is defined as the waveform consisting of one whole repetition of the modulated preamble code (either length 31 or 127). Details on the construction of the preamble symbol for various code lengths and PRFs are given in 6.8a.6. For each target PRF, the preamble is constructed from a preamble code, C_p , by inserting a number of chip durations between code symbols. The number of chip durations to insert is denoted by δ_L , values for each code length and PRF are given in Table 39b, and the chip insertion is detailed in Equation (9a).

Table 39b presents the timing parameters during the SHR portion of a UWB PHY frame while Table 39a presents the timing parameters for the PSDU portion of the frame. First, note that the preamble is sent at a slightly higher mean PRF than the data (see Table 39a). This is due to the fact that length 31 or 127 ternary codes are being used within the SHR, and the number of chips within the SHR is no longer a power of 2. For example, for the two mandatory PRFs in channels {0:3, 5:6, 8:10, 12:14}, the peak PRFs during the preamble are 31.2 MHz and 7.8 MHz, respectively, and the corresponding mean PRFs during the preamble are 16.10 MHz and 4.03 MHz, respectively. The corresponding mean PRFs during the data (PSDU) are 15.60 MHz and 3.90 MHz, respectively. The remaining peak and mean PRF values for other optional UWB channels and the optional length 127 code are listed in Table 39b.

Table 39b—UWB PHY preamble parameters

Bands	Preamble						
Channel Number	C_i Code Length	Peak PRF (MHz)	Mean PRF (MHz)	Delta Length δ_L	#Chips Per Symbol	Symbol Duration T_{psym} (ns)	Base Rate Msymbol/s
{0:15}	31	31.20	16.10	16	496	993.59	1.01
{0:3, 5:6, 8:10, 12:14}	31	7.80	4.03	64	1984	3974.36	0.25
{0:15}	127	124.80	62.89	4	508	1017.63	0.98

The base symbol rate is defined as the rate at which the preamble symbols are sent. The base rates corresponding to the two mandatory mean PRFs of 16.10 MHz and 4.03 MHz are 1 Msymbol/s and 0.25 Msymbol/s, respectively, and are listed in the column with the heading “Base Rate” in Table 39b. These symbol rates correspond to a preamble symbol duration, T_{psym} , of 993.59 ns and 3974.36 ns for the two mandatory PRFs.

Finally, for each UWB frame consisting of the SHR, SFD, PHR, and a data field, there are four possible durations of the SHR. This is due to the four possible lengths of SYNC field in the SHR (see 6.8a.6). The SYNC field consists of repetitions of the preamble symbol. The number of preamble symbol repetitions are 16, 64, 1024, and 4096. These different SYNC field lengths yield different time durations of the UWB frame. The relationship between SYNC field length and frame duration is shown in Table 39c. For each UWB channel, the number of chips in an individual preamble symbol is shown in the row titled “ N_c .” N_c is a function of the PRF used within the channel and, therefore, has either two or three values. For each value of N_c , the admissible preamble symbol durations T_{psym} are defined, and the duration of the SYNC portion of the SHR for each length (16, 64, 1024, or 4096) is denoted as T_{sync} . After the insertion of the SFD (the SFD may be either 8 or 64 preamble symbols long), the total length (in preamble symbols) of the SHR may any of the N_{pre} values shown in Table 39c, and this in turn leads to the possible SHR durations denoted as T_{pre} . After creation of the SHR, the frame is appended with the PHR whose length, N_{hdr} , is 16 symbols and duration is denoted as T_{hdr} . The values of the frame duration parameters are shown in Table 39c for each of the UWB channels.

The motivation for two different PRFs stems from the fact that the devices will operate in environments with widely varying delay spreads. The low PRF is mainly intended for operation in environments with high delay spreads, particularly if the receiver uses energy detection. For coherent reception in high-delay-spread environments and for any receivers in low-delay-spread environments, operation with the higher PRF is preferable.

The higher PRF allows the use of lower peak voltages for a fixed output power; note, however, that there is no requirement for the transmitter to transmit with maximum admissible (by the frequency regulators) output power. Thus, a transmitter can—if so desired by the designer—operate the low PRF with the same peak voltage as for the high-PRF case. Note that this case is still beneficial for noncoherent receivers as it reduces the intersymbol interference (compared to the high-PRF case).

Finally, it is noteworthy that the implementation of a dual PRF does not lead to a significant increase in the complexity of either transmitter or receiver since the PRFs are integer multiples of each other. The transmit signal corresponding to a low PRF can thus be generated with the same pulse generator as for the high PRF; the generator is simply excited less frequently. Similarly, the receive signal corresponding to the low PRF can be obtained by subsampling of the signal corresponding to the high PRF. Thus, even the synchronization procedure (when the PRF is unknown) can use the same sampler and ADC and just search both the fast-sampled (i.e., for high PRF) signal and (possibly in parallel) a subsampled version that is obtained from the fast-sampled signal by retaining only every fourth sample.

Table 39c—UWB PHY frame-dependent parameters

Parameter	Description	Value			
<i>Channel</i>	UWB PHY Channel Number	{0:15}		{0:3, 5:6, 8:10, 12:14}	
<i>PRF_{mean}</i>	Mean PRF (MHz)	16.10	62.89	4.03	
<i>N_c</i>	Number of chips per preamble symbol	496	508	1984	
<i>T_{psym}</i>	Preamble Symbol Duration (ns)	993.6	1017.6	3974.4	
<i>N_{sync}</i>	Number of symbols in the packet sync sequence	Short	16		
		Default	64		
		Medium	1024		
		Long	4096		
<i>T_{sync}</i>	Duration of the packet sync sequence (μs)	Short	15.9	16.3	63.6
		Default	63.6	65.1	254.4
		Medium	1017.4	1042.1	4069.7
		Long	4069.7	4168.2	NA*
<i>N_{sfd}</i>	Number of symbols in the SFD	8 or 64			
<i>T_{sfd}</i>	Duration of the frame sequence (μs)	7.9 or 63.6	8.1 or 65.1	31.8 or 254.4	
<i>N_{pre}</i>	Number of symbols in the SHR Preamble	Short	24 or 80		
		Default	72 or 128		
		Medium	1032 or 1088		
		Long	4104 or 4160		
<i>T_{pre}</i>	Duration of the SHR Preamble (μs)	Short	23.8 or 79.5	24.4 or 270.6	95.4 or 128.7
		Default	71.5 or 127.2	73.3 or 319.5	286.2 or 319.5
		Medium	1025.4 or 1081.0	1050.2 or 1296.4	4101.5 or 4134.9
		Long	4077.7 or 4133.3	4176.3 or 4422.6	NA*
<i>N_{hdr}</i>	Number of symbols in the PHY Header	16			
<i>T_{hdr}</i>	Duration of the PHY Header field (μs)	16.4	16.8	65.6	
<i>N_{data}</i>	Number of symbols in the data field	16 x LENGTH + 96			
<i>T_{data}</i>	Duration of the Data Field (μs)	$N_{data} \times T_{dsym}$			
<i>N_{CCA_PHR}</i>	Number of multiplexed preamble symbols in PHR	4 or 32			
<i>N_{CCA_data}</i>	Number of multiplexed preamble symbols in data field	$T_{data} / (4 \times T_{dsymIM})$			

6.8a.6 SHR preamble

A SHR preamble shall be added prior to the PHR to aid receiver algorithms related to AGC setting, antenna diversity selection, timing acquisition, coarse and fine frequency recovery, packet and frame synchronization, channel estimation, and leading edge signal tracking for ranging.

In this subclause, four different mandatory preambles are defined: a default preamble, a short preamble, a medium preamble, and a long preamble. The preamble to be used in the transmission of the current frame is determined by the value of *UWB Preamble Symbols Repetitions* in the PD-DATA.request primitive.

Figure 27d shows the structure of the SHR preamble. The preamble can be subdivided into two distinct portions: SYNC (packet synchronization, channel estimation, and ranging sequence) and SFD (frame delimiter sequence). The duration of these portions are provided in Table 39c. Subclauses 6.8a.6.1 and 6.8a.6.2 detail the different portions of the preamble.

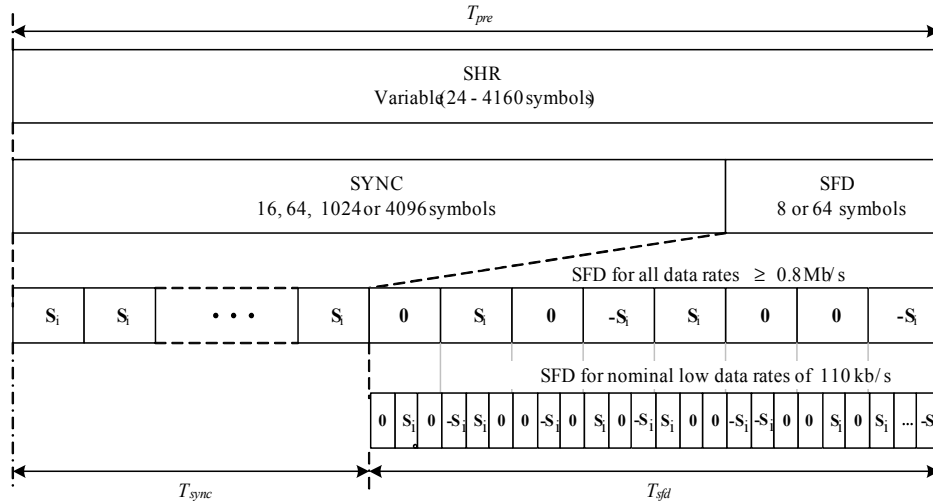


Figure 27d—SHR preamble structure

6.8a.6.1 SHR SYNC field

Each PAN operating on one of the UWB PHY channels {0–15} is also identified by a preamble code. The preamble code is used to construct symbols that constitute the SYNC portion of the SHR preamble as shown in Figure 27d. The UWB PHY supports two lengths of preamble code: a length 31 code and an optional length 127 code. Each preamble code is a sequence of code symbols drawn from a ternary alphabet $\{-1,0,1\}$ and selected for use in the UWB PHY because of their perfect periodic autocorrelation properties. The length 31 code sequences are shown in Table 39d while the length 127 code sequences are shown in Table 39e where they are indexed from 1–24 ($C_i, i = 1,2,\dots,24$). The first 8 codes (index 1–8) are length 31 while the remaining 16 (index 9–24) are length 127. Which codes may be used in each of the UWB PHY channels is restricted, and the particular code assignments are made in Table 39d and Table 39e. Specifically, the last column in each table indicates the set of UWB channel numbers that permit use of the code. This restriction of codes is to ensure that codes with the lowest cross-correlation are used in the same UWB PHY channel. Additionally, 8 of the length 127 codes are reserved for use with the private ranging protocol only and are not used during normal WPAN operation. This restriction is indicated in the third column of Table 39e as well.

Table 39d—Length 31 ternary codes

Code index	Code sequence	Channel number ^a
1	-0000+0-0+++0+000+-+++00+0-00	0, 1, 8, 12
2	0+0+-0+0+000-++0-+---00+00++000	0, 1, 8, 12
3	-+0++000-+--+00++0+00-0000-0+0-	2, 5, 9, 13
4	0000+-00-00-++++0+-+000+0-0++0-	2, 5, 9, 13
5	-0+-00+++0000-0+++0-0+0000-00	3, 6, 10, 14
6	++00+00---+0++-000+0+0-+0+0000	3, 6, 10, 14
7	+0000+-0+0+00+000+0+---0-+00-+	4, 7, 11, 15
8	0+00-0-0++0000--+00-+0+-++0+00	4, 7, 11, 15

^aNote that codes indexed 1 through 6 may also be used for UWB channels 4, 7, 11, and 15 (i.e., channels whose bandwidth is wider than 500 MHz) if interchannel communication is desired.

Table 39e—Optional length 127 ternary codes

Code index	Code sequence	Channel number ^a
9	+00+000-0--00--+0+0+00-+--+0+0000++-000+00-00--0+0+0--0-+++0++000+-0+00-0++-0+++00-+00+0+0-0++-+--000000+00000-+0000-0-000--+	0-4, 5, 6, 8-10, 12-14
10	++00+0-+00+00+000000-000-00--000-0+--+0-0+-0+00000+-00++0-0+00--+00++-+0+-0+0000-0-0-0-++-+0+00+0+000-+0+++000----+++0000++0--	0-4, 5, 6, 8-10, 12-14
11	-+0000+00--00000-0+0+0+-0+00+00+0-00-+++0+000-+0+0-0000++++++-+0--0+-0++-0-000+0+00+0+---000-000000-+00+-0++000++-00++-0-0	0-4, 5, 6, 8-10, 12-14
12	-+0++000000-0+0+-0---+--+00-+0++0+0+0+000-00-00-+00+-+000+-0-++0-0+++++0-00-0++0+0+00++-00+000+-000-0--0000-0000--0+0000+--	0-4, 5, 6, 8-10, 12-14
13	+000--0000--++0-++++0-0++0+0-00-+0+00++-0+0+-+0-00+00-0--000-+-0+0+0000-0++-00000+-0-000000-0-+--+000-0+0+0++-00--00+0+000	0-15; DPS only
14	+000++0-0+0-00+-0+0-00+0+0000+0+-0000++0+0+++++-+0-0+-0--0+---000---0+000+0+0+-000000+--+0-00++000-00+00++-00--++00-00000	0-15; DPS only
15	0+-00+0-000-++0000---++000+0+-0-+00-+000--0-00--0-++++-+0-++00+-++0+00000+0-0+++0-0+00+000-0000+00--0+0+0+0-00-0+-0+0+00000	0-15; DPS only
16	++0000+000+00+-0-+--+0-000--00+-0+00++000++0+0+0-0-0+00+00+0++----+00++-+0+-0--000000-0-0000-+0--00+00000+-++000-0+0+0	0-15; DPS only
17	+--000-0-0000+-00000+000000+-+--+0-0+0+00+-00+++0-++0-00+0-+000++0+++0-0-0+0+0-0-00-00+000-++0000+0++-+00+0+0+-00-0-000+00+	4, 7, 11, 15
18	--0+++0000++++---000+++0+-000+0+00+0+--+0-0-0-0000+0-0+--+00+-00+0-0++00-+00000+-0-0+0+-0-000--00-000-000000+00+00+-0+00++	4, 7, 11, 15
19	-0-++00-++000++0-00+-000000-000----+0+00+-0+000-0--++0-0--0+0-+++++0000-0+0+-000+00+++0-0-0+00+00+0-+0+0+0-00000--00+0000+-0	4, 7, 11, 15
20	--+00000+0--0000-0000+-0-000-+000+00-++00+0+00++0-00-0++++0-0++0+-0-000++-+00+-00-00-000+0+0+0+0+0+-00++-+---0+0+0-000000++0+-	4, 7, 11, 15
21	+0+00--00-+++0+0+0-000+-+--+00-000000-0-+00000-++0-0000+00-+-000--0-00+00-0+-0+0+-+00++0+-00-0+0+0+0+++0+-+0000--000+000	0-15; DPS only
22	0-00-+-+00-++00+00-000++00--0-+0+00000+-0+0+000+0---000--++0+-0-0+0+-+++++0+00++0000-+0+0000+0+00-0+-0+00-0+0-0+000+0000	0-15; DPS only
23	000++0+0+-0-00-0+0+0+0+-00+0000-000+00+00-+++0-0+00000+0++00+++0+-++++-0-00-0-000+-00+-0+0+000++-0000++-000-0+00+000	0-15; DPS only
24	+0+-0-000++-+00000+00--0+-0000-0-000000+-0-0+0-++00+----++0+00+00+0-0+-0-0+0+00+++000++00+0+00--000-0++-0-++00+000+0000++0	0-15; DPS only

^aNote that codes indexed 9 through 13 may also be used for UWB channels 4, 7, 11, and 15 (i.e., channels whose bandwidth is wider than 500 MHz) if interchannel communication is desired.

Note that the assignment of preamble codes to channels has been done to enable interchannel communication. In other words, it is possible that a device operating on a wideband channel {4,7,11,15} may communicate with a device on a channel with which it overlaps.

For a WPAN using the ternary code indexed by i , the SYNC field shall consist of N_{sync} repetitions of the symbol S_i , where S_i is the code C_i spread by the delta function δ_L of length L as shown in Table 39b. The spreading operation, where code C_i is extended to the preamble symbol duration indicated in Table 39b, is described mathematically in Equation (9a). In Equation (9a), the operator \otimes indicates a Kronecker product.

After the Kronecker operation, a preamble symbol is formed as depicted in Figure 27e, where $L - 1$ zeros have been inserted between each ternary element of C_i .

$$S_i = C_i \otimes \delta_L(n)$$

$$\delta_L(n) = \begin{cases} 1 & n = 0 \\ 0 & n = 1, 2, \dots, L-1 \end{cases} \quad (9a)$$

The spreading factor L , number of chips per symbol, preamble symbol duration T_{psym} , and base symbol rate for different channels are given in Table 39b.




Figure 27e—Construction of symbol S_i from code C_i

6.8a.6.2 SHR SFD

A SFD shall be added to establish frame timing. The UWB PHY supports a mandatory short SFD for default and medium data rates and an optional long SFD for the nominal low data rate of 110 kb/s as shown in Figure 27d. The mandatory short SFD shall be $[0 +1 0 -1 +1 0 0 -1]$ spread by the preamble symbol S_i , where the leftmost bit shall be transmitted first in time. The optional long SFD shall be obtained by spreading the sequence $[0 +1 0 -1 +1 0 0 -1 0 +1 0 -1 +1 0 0 -1 -1 0 0 +1 0 -1 0 +1 0 +1 0 0 0 -1 0 -1 0 -1 0 0 +1 0 -1 -1 0 -1 +1 0 0 0 0 +1 +1 0 0 -1 -1 -1 +1 -1 +1 +1 0 0 0 0 +1 +1]$ by the preamble sequence S_i . Note that the long SFD is eight times longer than the short SFD and consists of 64 preamble symbols, only 32 of which are active, and the other 32 are zeros. The structure of the SHR preamble and the two possible SFDs are shown in Figure 27d.

6.8a.7 PHY header (PHR)

A PHR, as shown in Figure 27f, shall be added after the SHR preamble. The PHR consists of 19 bits and conveys information necessary for a successful decoding of the packet to the receiver. The PHR contains information about the data rate used to transmit the PSDU, the duration of the current frame's preamble, and the length of the frame payload. Additionally, six parity check bits are used to further protect the PHR against channel errors.

Bit 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
R1	R0	L6	L5	L4	L3	L2	L1	L0	RNG	EXT	P1	P0	C5	C4	C3	C2	C1	C0
Data Rate	Frame Length								Ranging Packet Extension Header	Preamble Duration	SECDED Check Bits							

Figure 27f—PHR bit assignment

The PHR shall be transmitted using the BPM-BPSK modulation outlined in 6.8a.9. The PHR shall be transmitted at the nominal rate of 850 kb/s for all data rates above 850 kb/s and at the nominal rate of 110 kb/s for the nominal low data rates of 110 kb/s.

6.8a.7.1 PHR rate, length, ranging, extension, preamble duration fields

The Data Rate field shall consist of two bits (R1, R0) that indicate the data rate of the received PSDU. The bits R1–R0 shall be set, dependent on the mean PRF, according to Table 39g. The default value of the bits R1–R0 shall be set to 01 as this is the only mandatory data rate that is supported by a UWB-compliant PHY implementation. Support for other data rates listed in Table 39g is optional.

The Frame Length field, L6–L0, shall be an unsigned 7-bit integer number that indicates the number of octets in the PSDU that the MAC sublayer is currently requesting the PHY to transmit.

The Ranging Packet bit, RNG, indicates that the current frame is an RFRAME if it is set to 1; otherwise, it is set to 0.

The Header Extension bit, EXT, is reserved for future extension of the PHR. This bit shall be set to 0.

The Preamble Duration field, P1–P0, represents the length (in preamble symbols) of the SYNC portion of the SHR. P1–P0 shall be set according to Table 39f. The default setting Preamble Duration setting is 01, which corresponds to a SYNC field of length 64 preamble symbols.

Table 39f—Preamble Duration field values

P1–P0	SYNC length (symbols) (Si)
00	16
01	64
10	1024
11	4096

The Preamble Duration field is intended for use during ranging operations and is used by a receiver of the PHY frame to help determine at which preamble symbol the UWB PHY acquired and began tracking the preamble. A receiver may use the Preamble Duration field to set the value of its own preamble duration based upon the received value when communicating a ranging ACK packet.

Table 39g—Nominal data rates

R1–R0	Mean PRF 15.60 or 62.40 MHz	Mean PRF 3.90 MHz
00	0.11	0.11
01	0.85	0.85
10	6.81	1.70
11	27.24	6.81

6.8a.7.2 PHR SECDED check bits

The SECDED (single error correct, double error detect) field, C5–C0, is a set of six parity check bits that are used to protect the PHR from errors caused by noise and channel impairments. The SECDED bits are a simple Hamming block code that enables the correction of a single error and the detection of two errors at the receiver. The SECDED bit values depend on PHR bits 0–12 and are computed as follows:

$$\begin{aligned} C0 &= XOR(R0, R1, L0, L2, L4, L5, EXT, P1) \\ C1 &= XOR(R1, L2, L3, L5, L6, RNG, EXT, P0) \\ C2 &= XOR(R0, L0, L1, L5, L6, RNG, EXT) \\ C3 &= XOR(L0, L1, L2, L3, L4, RNG, EXT) \\ C4 &= XOR(P0, P1) \\ C5 &= XOR(R0, R1, L5, L6, C3, C4) \end{aligned}$$

6.8a.8 Data field

The Data field is the last component of the PPDU and is encoded as shown in Figure 27g.




Figure 27g—Data field encoding process

The data field shall be formed as follows:

- Encode the PSDU using systematic Reed-Solomon block code, which adds 48 parity bits as described in 6.8a.10.1.
- Encode the output of the Reed-Solomon block code using a systematic convolutional encoder as described in 6.8a.10.2.
- Spread and modulate the encoded block using BPM-BPSK modulation as described in 6.8a.9.

6.8a.9 UWB PHY modulation

6.8a.9.1 UWB PHY modulation mathematical framework

The UWB PHY transmit waveform during the k^{th} symbol interval may be expressed as shown in Equation (9b):

$$x^{(k)}(t) = [1 - 2g_1^{(k)}] \sum_{n=1}^{N_{cpb}} [1 - 2s_{n+kN_{cpb}}] \times p(t - g_0^{(k)}T_{BPM} - h^{(k)}T_{burst} - nT_c) \quad (9b)$$

Equation (9b) describes the time hopping with polarity scrambling, which improves interference rejection capabilities of the UWB PHY. The k^{th} symbol interval carries two information bits $g_0^{(k)}$ and $g_1^{(k)} \in \{0, 1\}$. Bit $g_0^{(k)}$ is encoded into the burst position whereas bit $g_1^{(k)}$ is encoded into the burst polarity. The sequence $s_{n+kN_{cpb}} \in \{0, 1\}, n = 0, 1, \dots, N_{cpb} - 1$ is the scrambling code used during the k^{th} symbol interval, $h^{(k)} \in \{0, 1 - N_{hop} - 1\}$ is the k^{th} burst hopping position, and $p(t)$ is the transmitted pulse shape at the

antenna input. The burst hopping sequence $h^{(k)}$ provides for multiuser interference rejection. The chip scrambling sequence $s_{n+kN_{cpb}}$ provides additional interference suppression among coherent receivers as well as spectral smoothing of the transmitted waveform. Note that Equation (9b) defines the transmitted signal during the valid burst interval; at all other possible burst positions, no signal shall be transmitted. A reference modulator illustrating the BPM-BPSK modulation is shown in Figure 27h.

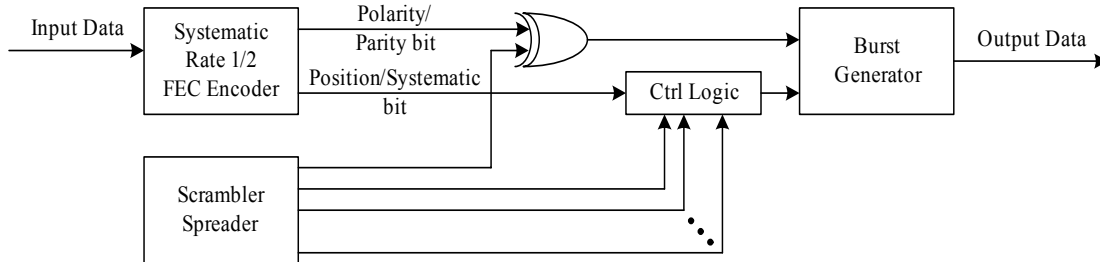


Figure 27h—Reference symbol modulator

6.8a.9.2 UWB PHY spreading

The time-varying spreader sequence $s_{n+kN_{cpb}}$ and the time-varying burst hopping sequence $h^{(k)}$ shall be generated from a common PRBS scrambler.

The polynomial for the scrambler generator shall be $g(D) = 1 + D^{14} + D^{15}$

where D is a single chip delay, T_c , element. This polynomial forms not only a maximal length sequence, but also is a primitive polynomial. By the given generator polynomial, the corresponding scrambler output is generated as

$$s_n = s_{n-14} \oplus s_{n-15} \quad n = 0, 1, 2, \dots$$

where \oplus denotes modulo-2 addition.

A linear feedback shift register (LFSR) realization of the scrambler is shown in Figure 27i. The LFSR shall be initialized upon the transmission of bit 0 of the PHR. Note that N_{cpb} may change depending on the data rate and PRF in use during the PSDU. The LFSR shall not be reset after transmission of the PHR.

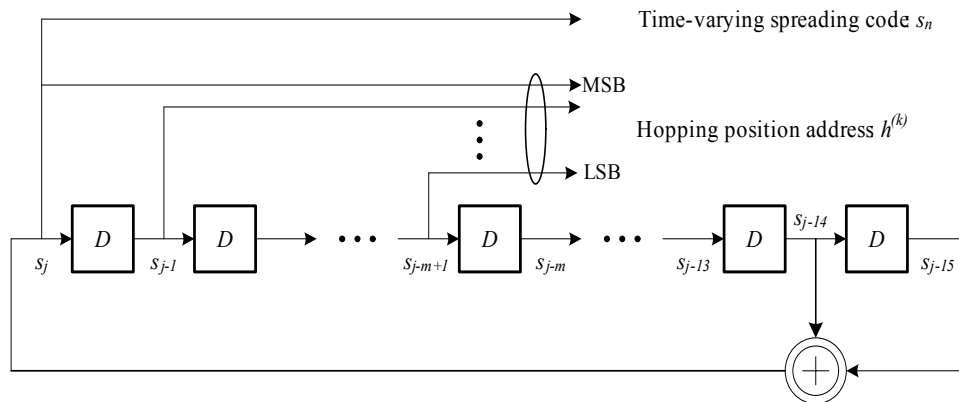


Figure 27i—LFSR implementation of the scrambler

The initial state of the LFSR shall be determined from the preamble code by first removing all the 0s in the ternary code and then replacing all the -1s with a zero. The first 15 bits of the resulting binary state shall be loaded into the LFSR. Table 39h shows an example of the above procedure for preamble code, C_6 (length 31, preamble code index 6, see Table 39d). The table shows the initial state as well as the first 16 output bits from the scrambler.

Table 39h—Example LFSR initial state for preamble code 6

Initial state ($s_{-15}, s_{-14}, \dots, s_{-1}$)	LFSR output: First 16 bits s_0, s_1, \dots, s_{15} (s_0 first in time)
111000101101101	0010011101101110

Note that even though each device within a PAN uses the same initial LFSR setting, the communication in WPAN is asynchronous so that the hopping and scrambling provides interference rejection.

The LFSR shall be clocked at the peak PRF of 499.2 MHz as specified in Table 39a. During the k^{th} symbol interval, the LFSR shall be clocked N_{cpb} times, and the scrambler output shall be the k^{th} scrambling code $s_{n+kN_{cpb}}$, $n = 0, 1, \dots, N_{cpb} - 1$. Furthermore, the k^{th} burst hopping position, shall be computed as follows:

$$h^{(k)} = 2^0 s_{kN_{cpb}} + 2^1 s_{1+kN_{cpb}} + \dots + 2^{m-1} s_{m-1+kN_{cpb}}$$

where

$$m = \log_2(N_{hop})$$

As shown in Table 39a, the number of hopping burst N_{hop} is always a power of two, and consequently m is always an integer. Note that for $N_{cpb} < m$, the LFSR is clocked N_{cpb} times, not m times.

For the mandatory modes with mean data PRFs of 15.60 MHz and 3.90 MHz, the numbers of hopping bursts are 8 and 32, respectively, as indicated in Table 39a, and consequently m takes on the values 3 and 5, respectively. The corresponding hopping sequences are as follows:

$$h^{(k)} = s_{kN_{cpb}} + 2s_{1+kN_{cpb}} + 4s_{2+kN_{cpb}} \quad \text{Mean PRF} = 15.60 \text{ MHz}$$

$$h^{(k)} = s_{kN_{cpb}} + 2s_{1+kN_{cpb}} + 4s_{2+kN_{cpb}} + 8s_{3+kN_{cpb}} + 16s_{4+kN_{cpb}} \quad \text{Mean PRF} = 3.90 \text{ MHz}$$

6.8a.10 UWB PHY forward error correction (FEC)

The FEC used by the UWB PHY is a concatenated code consisting of an outer Reed-Solomon systematic block code and an inner half-rate systematic convolutional code. The inner convolutional code is not necessarily enabled at all data rates; the rows of Table 39a that have a Viterbi rate of 1 indicate that the inner convolutional code is disabled for the PSDU part of the PHY frame.

The FEC encoding of a block of M PSDU bits, b_0, b_1, \dots, b_{M-1} , is shown in Figure 27j. The Reed-Solomon encoder shall append 48 parity bits, p_0, p_1, \dots, p_{47} , to the original block. This results in a Reed-Solomon encoded block of length $M + 48$. A half-rate systematic convolutional encoder shall encode the Reed-Solomon encoded block into a systematic coded block of length $2M + 96$ bits. The convolutional systematic bits shall be used to encode the position of the burst whereas the convolutional parity bits shall be used to

encode the polarity of the pulses within a burst. A noncoherent receiver cannot see the convolutional parity bits (parity bits), and consequently a noncoherent receiver may use only a Reed-Solomon decoder to improve its performance. A coherent receiver may use either or both Reed-Solomon and convolutional decoding algorithms. Note here that since both the Reed-Solomon and the convolutional codes are both systematic, a receiver (either coherent or noncoherent) may be implemented without an FEC decoder. In this case, the information bits are simply recovered by demodulating the position of the burst. There will be additional parity check bits as a result of the Reed-Solomon encoding, but these may be simply ignored.




Figure 27j—FEC encoding process

Subclauses 6.8a.10.1 and 6.8a.10.2 provide details of Reed-Solomon and convolutional encoding.

6.8a.10.1 Reed-Solomon encoding

The systematic Reed-Solomon code is over Galois field, $GF(2^6)$, which is built as an extension of $GF(2)$. The systematic Reed-Solomon code shall use the generator polynomial

$$g(x) = \prod_{k=1}^8 (x + \alpha^k) = x^8 + 55x^7 + 61x^6 + 37x^5 + 48x^4 + 47x^3 + 20x^2 + 6x + 22$$

where $\alpha = 010000$ is a root of the binary primitive polynomial $1 + x + x^6$ in $GF(2^6)$.

In Reed-Solomon encoding $RS_6(K + 8, K)$, a block of I bits (with $K = \lceil I/6 \rceil$) is encoded into a codeword of $I + 48$ bits. The Reed-Solomon encoding procedure is performed in the following five steps:

- Addition of dummy bits.* The block of I information bits is expanded to 330 bits by adding $330 - I$ dummy (zero) bits to the beginning of the block. The expanded block is denoted as $\{d_0, d_1, \dots, d_{329}\}$ where d_0 is the first in time.
- Bit-to-symbol conversion.* The 330 bits $\{d_0, d_1, \dots, d_{329}\}$ are converted into 55 Reed-Solomon symbols $\{D_0, D_1, \dots, D_{54}\}$ having the following polynomial representation:

$$D_k = \alpha^5 d_{6k+5} + \alpha^4 d_{6k+4} + \alpha^3 d_{6k+3} + \alpha^2 d_{6k+2} + \alpha d_{6k+1} + d_{6k}, \quad k = 0:54$$

Resulting 6-bit symbols are presented as $D_k = \{d_{6k+5}, d_{6k+4}, d_{6k+3}, d_{6k+2}, d_{6k+1}, d_{6k}\}$, where d_{6k+5} is the MSB and d_{6k} is the LSB.

- Encoding.* The information symbols $\{D_0, D_1, \dots, D_{54}\}$ are encoded by systematic $RS_6(63,55)$ code with output symbols $\{U_0, U_1, \dots, U_{62}\}$ ordered as follows:

$$U_k = \begin{cases} D_k & (k = 0, 1, \dots, 54) \\ P_k & (k = 55, 56, \dots, 62) \end{cases}$$

where P_k are parity check symbols added by $RS_6(63,55)$ encoder.

The information polynomial associated with the information symbols $\{D_0, D_1, \dots, D_{54}\}$ is denoted as $D(x) = x^{54}D_0 + x^{53}D_1 + \dots + xD_{53} + D_{54}$. The parity check polynomial associated with the parity check symbols is denoted as $P(x) = x^7P_{55} + x^6P_{56} + \dots + xP_{61} + P_{62}$. The parity check symbols are calculated as:

$$P(x) = \text{remainder}[x^8D(x)/g(x)]$$

$$U(x) = x^8D(x) + P(x)$$

- d) *Symbol-to-bit conversion.* The output symbols $\{U_0, U_1, \dots, U_{62}\}$ are converted into binary form with LSB coming out first, resulting in a block of 378 bits $\{u_0, u_1, \dots, u_{377}\}$.
- e) *Removal of dummy bits.* The $330 - I$ dummy bits added in the first step are removed. Only the last $I + 48$ bits are transmitted, i.e., $\{u_{330-I}, u_{331-I}, \dots, u_{377}\}$ with u_{330-I} being first in time.

6.8a.10.2 Systematic convolutional encoding

The inner convolutional encoder shall use the rate $R = 1/2$ code with generator polynomials $g_0 = [010]_2$ and $g_1 = [101]_2$ as shown in Figure 27k. Upon transmission of each PPDU, the encoder shall be initialized to the all zero state. Additionally, the encoder shall be returned to the all zero state by appending two zero bits to the PPDU. Note that since the generator polynomials are systematic, they are also noncatastrophic.

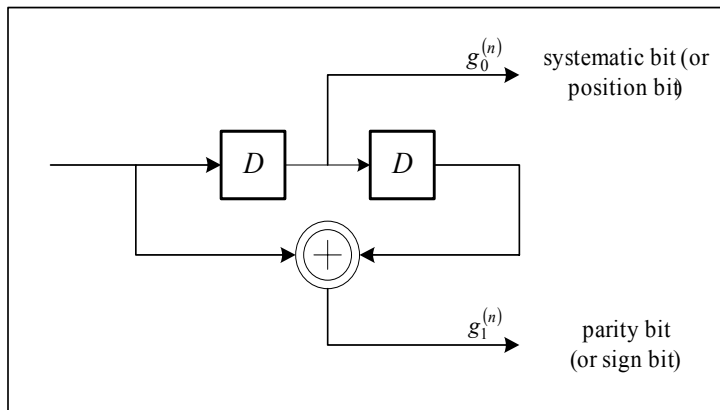


Figure 27k—Systematic convolutional encoder

6.8a.11 PMD operating specifications

6.8a.11.1 Operating frequency bands

The set of operating frequency bands is defined in Table 39i. For the sub-gigahertz operation, channel 0 is defined as the mandatory channel; for the low-band operation, channel 3 is the mandatory channel; and for the high-band operation, channel 9 is the mandatory channel.

Table 39i—UWB PHY band allocation

Band group ^a (decimal)	Channel number (decimal)	Center frequency, f_c (MHz)	Band width (MHz)	Mandatory/Optional
0	0	499.2	499.2	Mandatory below 1 GHz
1	1	3494.4	499.2	Optional
	2	3993.6	499.2	Optional
	3	4492.8	499.2	Mandatory in low band
	4	3993.6	1331.2	Optional
2	5	6489.6	499.2	Optional
	6	6988.8	499.2	Optional
	7	6489.6	1081.6	Optional
	8	7488.0	499.2	Optional
	9	7987.2	499.2	Mandatory in high band
	10	8486.4	499.2	Optional
	11	7987.2	1331.2	Optional
	12	8985.6	499.2	Optional
	13	9484.8	499.2	Optional
	14	9984.0	499.2	Optional
	15	9484.8	1354.97	Optional

^aNote bands indicate a sequence of adjacent UWB center frequencies: band 0 is the sub-gigahertz channel, band 1 has the low-band UWB channels, and band 2 has the high-band channels.

Figure 27l is a graphical representation of the data presented in Table 39i. The figure shows each UWB PHY channel as a heavy black line centered on the channel's center frequency. The length of the lines depicts the channel bandwidth. This figure is useful for visualizing the relationship among the various channels, specifically, channel overlap.

6.8a.11.2 Channel assignments

A total of 32 complex channels are assigned for operation, two channels in each of the 16 defined operating frequency bands. A compliant implementation shall support at least the two logical channels for one of the mandatory bands.

6.8a.11.3 Regulatory compliance

The maximum allowable output PSD shall be in accordance with practices specified by the appropriate regulatory bodies.

6.8a.11.4 Operating temperature range

A conformant implementation shall meet all of the specifications in this standard for ambient temperatures from 0 to 40 °C.




Figure 27I—UWB PHY band plan

6.8a.12 Transmitter specification

6.8a.12.1 Baseband impulse response

The transmitted pulse shape $p(t)$ of the UWB PHY shall be constrained by the shape of its cross-correlation function with a standard reference pulse, $r(t)$. The normalized cross-correlation between two waveforms is defined as

$$\phi(\tau) = \frac{1}{\sqrt{E_r E_p}} \operatorname{Re} \int_{-\infty}^{\infty} r(t) p^*(t + \tau) dt$$

In the above, E_r and E_p are the energies of $r(t)$ and $p(t)$, respectively. The reference $r(t)$ pulse used in the calculation of $\phi(\tau)$ is a root raised cosine pulse with roll-off factor of $\beta = 0.6$. Mathematically this is

$$r(t) = \frac{4\beta}{\pi\sqrt{T_p}} \frac{\cos\left[(1 + \beta)\pi t/T_p + \frac{\sin[(1 - \beta)\pi t/T_p]}{4\beta(t/T_p)}\right]}{(4\beta t/T_p)^2 - 1}$$

In the above equation, T_p is the pulse duration. Table 39j shows the required pulse duration for each channel.

Table 39j—Required reference pulse durations in each channel

Channel number	Pulse duration, T_p (ns)	Main lobe width, T_w (ns)
{0:3, 5:6, 8:10, 12:14}	2.00	0.5
7	0.92	0.2
{4, 11}	0.75	0.2
15	0.74	0.2

In order for a UWB PHY transmitter to be compliant with this standard, the transmitted pulse $p(t)$ shall have a magnitude of the cross-correlation function $|\phi(\tau)|$ whose main lobe is greater or equal to 0.8 for a duration of at least T_w (see Table 39j), and any sidelobe shall be no greater than 0.3. For the purposes of testing a pulse for compliance, the following are defined: Let $|\phi(\tau)|$ be the magnitude of the cross-correlation of $p(t)$ and $r(t)$, and let τ_i $i = 1, 2, \dots$ be a set of critical points, i.e., points at which $\frac{d}{d\tau}|\phi(\tau)|_{\tau=\tau_i} = 0$. The maximum of the function occurs at one of these critical points, τ_{max} where $|\phi(\tau_{max})| \geq |\phi(\tau)|$ for all values of τ . The requirement above thus states that for some continuous set of values that contain the point τ_{max} the function $|\phi(\tau)|$ is greater than 0.8. In addition, the second constraint on the value of sidelobes may be stated mathematically as $|\phi(\tau_i)| \leq 0.3$ for all τ_i .

Figure 27m shows an example UWB-compliant pulse, $p(t)$ (left plot), along with the root raised cosine reference pulse $r(t)$ (middle plot) with $T_p = 2.0$ ns and the magnitude of the cross-correlation $|\phi(\tau)|$ (right plot). The pulse $p(t)$ is an 8 order butterworth pulse with a 3 dB bandwidth of 500 MHz. The figure is intended to show that this example pulse meets the requirements for compliance. Specifically, the main lobe is above 0.8 for nearly 1 ns, and no sidelobe is greater than 0.3 (in this case, the largest sidelobe peak is 0.2). The pulse $p(t)$ is a compliant pulse for channels {0:3, 5:6, 8:10, 12:14}.




Figure 27m—Compliant pulse example

Note that it is not the intention of this standard to imply that pulse shaping shall occur at baseband, only that the measurements described here occur on the pulse envelope if shaping is done at passband.

6.8a.12.2 Transmit PSD mask

The transmitted spectrum shall be less than -10 dBr (dB relative to the maximum spectral density of the signal) for $0.65/T_p < |f-f_c| < 0.8/T_p$ and -18 dBr for $|f-f_c| > 0.8/T_p$. For example, the transmit spectrum mask for channel 4 is shown in Figure 27n. The measurements shall be made using 1 MHz resolution bandwidth and a 1 kHz video bandwidth.




Figure 27n—Transmit spectrum mask for band 4

6.8a.12.3 Chip rate clock and chip carrier alignment

A UWB transmitter shall be capable of chipping at the peak PRF given in Table 39a with an accuracy of ± 20 ppm. In addition, for each UWB PHY channel, the center of transmitted energy shall be within the values listed in Table 39i also with an accuracy of ± 20 ppm. The measurements shall be made using 1 MHz resolution bandwidth and a 1 kHz video bandwidth.

6.8a.13 UWB PHY optional pulse shapes

The UWB PHY offers the capability to transmit several optional pulse types. These are described in detail in 6.8a.13.1 through 6.8a.13.3. The use of these options is controlled by the PAN coordinator and shall be limited to the nonbeacon frames. In other words, beacon frames shall be transmitted using the mandatory pulse shape as defined in 6.8a.12.1 but all other frames may be transmitted using the optional pulse shapes if all devices in the PAN are capable of supporting the optional pulse shapes. PANs that use the optional pulse shapes shall indicate the use of a specific option via the *phyUWBCurrentPulseShape* PIB attribute. Devices choosing to join a PAN using one of the optional pulse shapes should make their decision based on the value of *phyUWBCurrentPulseShape* that is reported during the scan procedure.

6.8a.13.1 UWB PHY optional chirp on UWB (CoU) pulses

This subclause specifies an optional mode of CoU pulses. The purpose of CoU pulses is to provide an additional dimension (besides frequency and DS codes) to support simultaneously operating piconets (SOP).

Since CoU is an optional mode of pulse shapes in addition to the mandatory pulse shape, all modulation specifications shall be the same as they are for the mandatory pulse shape except those defined for the CoU pulses when a device implements the CoU option.

A mathematical representation of a CoU pulse at baseband is given by Equation (9c), and a graphical example of CoU pulse is shown in Figure 27o.

$$p_{CoU}(t) = \begin{cases} p(t) \exp\left(-j\frac{\pi\beta t^2}{2}\right) & -\frac{T}{2} \leq t \leq \frac{T}{2} \\ 0 & otherwise \end{cases} \quad (9c)$$

where

- $p(t)$ denotes a mandatory pulse shape that satisfies constraints in 6.8a.12.1
- $\beta = B/T$ is the chirping rate (chirping slope). Moreover, B and T are the bandwidth and time duration of the CoU pulse, respectively.




Figure 27o—Graphical view of a CoU pulse

It can be seen from Equation (9c) that CoU is an operation added to the mandatory pulse. When a CoU pulse is transmitted, the receiver needs to perform a matched de-chirp operation to demodulate the signal.

The optional CoU pulses are admitted with two slopes per each DS code per each 500 MHz bandwidth. The chirp slopes are denoted as CCh.1 and CCh.2. Within channels 4, 7, 11, 15, there are chirp slopes admitted per each DS code. These are denoted as CCh.3 through CCh.6. The values for each chirp slope are listed in Table 39k.

Table 39k—CoU channel slopes

CoU number	β (slopes)
CCh.1	500 MHz/2.5 ns
CCh.2	-500 MHz/2.5 ns
CCh.3	1 GHz/5 ns
CCh.4	-1 GHz/5 ns
CCh.5	1 GHz/10 ns
CCh.6	-1 GHz/10 ns

6.8a.13.2 UWB PHY optional continuous spectrum (CS) pulses

This subclause specifies optional CS pulses. A CS pulse is obtained by passing the mandatory pulse through an all-passing CS filter. The CS filter introduces controlled group delays to the input pulse. The purpose of the optional CS pulses is to reduce the interference level between different PANs to enhance SOP performance.

Since CS is an optional mode of pulse shapes in addition to the mandatory pulse shape, all modulation specifications shall be the same as they are for the mandatory pulse shape except those defined for the CS pulses when a device implements the CS option.

An optional CS pulse $p_{CS}(t)$ is defined by Equation (9d), and some examples of CS pulses are shown in Figure 27p.

$$p_{CS}(t) = \int P(f) \exp[-j2\pi f(t - (\tau \times f))] df \quad (9d)$$

where

τ represents the group delay (s/Hz)

$P(f)$ represents the Fourier transform [see Equation (9e)] of $p(t)$, where $p(t)$ is any pulse shape that meets the requirements defined in 6.8a.12.1

$$P(f) = \int p(t) \exp[-j2\pi ft] dt \quad (9e)$$

It can be seen from Equation (9d) that CS filtering is an operation added to the mandatory pulse. When a CS pulse is transmitted, the receiver needs to perform an inverse CS filtering (CS^{-1}) operation to demodulate the signal.




Figure 27p—Examples of CS pulses

Each 500 MHz band shall use No.1 or No.2 pulses, while each 1.5 GHz band shall use one of No.3 through No.6 pulses (see Table 39l).

Table 39I—CS group delays

CS pulse number	τ (Group delay)
No.1	2 ns/500 MHz
No.2	−2 ns/500 MHz
No.3	5 ns/1 GHz
No.4	−5 ns/1 GHz
No.5	10 ns/1 GHz
No.6	−10 ns/1 GHz

6.8a.13.3 UWB PHY linear combination of pulses (LCP)

This subclause specifies an optional pulse shape that consists of a weighted linear combination of the pulses. LCP can be used in regulatory regions where “detect and avoid” (DAA) schemes are required by regulators. Using LCP pulses enables a PAN to limit interference to incumbent wireless systems. This new optional pulse shape is denoted $p_{LCP}(t)$ and is the sum of N weighted and delayed pulses $p(t)$ as follows:

$$p_{LCP}(t) = \sum_{i=1}^N a_i p(t - \tau_i)$$

where

$p(t)$ is any pulse that satisfies the cross-correlation constraints outlined in 6.8a.12.1

The number of pulses N that can be combined is set to a fixed value of 4 (although smaller values can be realized by setting the amplitudes of some of the pulses to zero). The values of the pulse delays shall be limited to $0 \leq \tau_i \leq 4ns$. The value of τ_1 is assumed to be zero, and thus the remaining delays are considered as relative delay time with respect to the nominal pulse location. The values for these delays are stored as the PIB values *phyUWBLCPPDelay2*, *phyUWBLCPPDelay3*, and *phyUWBLCPPDelay4*. The values for the amplitudes a_i are stored as the PIB values *phyUWBLCPPweight1* through *phyUWBLCPPDelay4* (see Table 23). The numerical values of the delays and amplitudes of the pulses shall be transmitted following the general framework of optional pulse shapes, as defined in 6.8a.13. The method to compute the weights and delay values is outside the scope of this standard.

6.8a.14 Extended preamble for optional UWB CCA mode

The PHY may provide the capability to perform the optional UWB CCA mode 6 (see 6.9.9). This CCA mode shall be supported by the modified frame structure where preamble symbols are multiplexed with the data symbols in the PHR and the PSDU of a frame.

Figure 27q shows the modified frame structure with multiplexed preamble symbols. One preamble symbol is inserted after each PHR and PSDU segment. The inserted preamble symbol shall be the same as the symbol used in the SHR (see 6.8a.6.1) of the same frame. The time interval between two neighboring inserted preamble symbols, which is also the time duration of each PHR or PSDU segment, is independent of the current data rate. The PIB attribute *phyUWBInsertedPreambleInterval* defines a constant time interval based on the nominal data rate of 850 kb/s for all operation bands. The nominal data rate of 850 kb/s is listed in Table 39g with the data rate field index R1 – R0 = 01. The value of the PIB attribute *phyUWBInsertedPreambleInterval* is fixed to 4. Distinguished from the CCA in narrowband systems, which is used to detect

6.8a.15.2 Crystal characterization

An RDEV that implements optional crystal characterization shall produce a tracking offset value and a tracking interval value for every timestamp report that is produced. The tracking offset and the tracking interval are computed from measurements taken during an interval that includes the interval bounded by the ranging counter start value and the ranging counter stop value. Note that crystal characterization is relevant only if it is characterizing the crystal that affects the ranging counter.

6.8a.15.2.1 Ranging tracking offset

The UWB ranging tracking offset is a signed magnitude integer. The integer magnitude part of the number shall be 19 bits. The LSB of the integer represents a “part.” The sign bit of the signed magnitude integer shall be equal to zero when the oscillator at the transmitter is a higher frequency than the oscillator at the receiver, and the sign bit shall be 1 when the oscillator at the receiver is a higher frequency than the transmitter. The value of the integer shall be a number that represents the difference in frequency between the receiver’s oscillator and the transmitter’s oscillator after the tracking offset integer is divided by the ranging tracking interval integer of 6.8a.15.2.2. For example, if the difference between the oscillators is 10 ppm, then an acceptable value of the ranging tracking offset would be 10 when the ranging tracking interval is 1 million. Another acceptable value for the ranging tracking offset is 15 when the ranging tracking interval is 1.5 million.

6.8a.15.2.2 Ranging tracking interval

The UWB ranging tracking interval shall be a 32-bit unsigned integer. The LSB of the ranging tracking interval represents a “part” that shall be exactly equal to the “part” in the LSB of the ranging tracking offset of 6.8a.15.2.1. The size of the “part” is a time period that shall be smaller than or equal to a chip time at the mandatory chipping rate of 499.2 MHz. Use of smaller “parts” for the LSB is encouraged. See 5.5.7.5.4.

6.8a.15.3 Ranging FoM

An RDEV shall produce a ranging FoM for every ranging counter value that is produced. The UWB ranging FoM is an octet as shown in Figure 27p. The FoM is composed of three subfields and an extension bit. The FoM Confidence Level subfield is defined in Table 39o. The confidence level is the probability that the leading edge of the pulse will arrive during the confidence interval. The FoM Confidence Interval subfield is defined in Table 39p. The confidence interval width in Table 39p is the entire interval width, not a plus or minus number. The FoM Confidence Interval Scaling Factor subfield is defined in Table 39q. Thus the overall confidence interval is obtained according to the formula *overall confidence interval = confidence interval × confidence interval scaling field*. The MSB of the FoM octet is the extension bit. When the extension bit is set to zero, the subfields have the normal meanings given in Table 39o, Table 39p, and Table 39q. When the extension bit is 1, the FoM has meaning given in Table 39r.

Bit 7	6	5	4	3	2	1	0
Extension	Confidence Interval Scaling Factor subfield		Confidence Interval subfield		Confidence Level subfield		

Figure 27p—Ranging FoM

Table 39o—Confidence Level subfield

Confidence level	Bit 2	Bit 1	Bit 0
No FoM	0	0	0
20%	0	0	1
55%	0	1	0
75%	0	1	1
85%	1	0	0
92%	1	0	1
97%	1	1	0
99%	1	1	1

Table 39p—Confidence Interval subfield

Confidence interval	Bit 1	Bit 0
100 ps	0	0
300 ps	0	1
1 ns	1	0
3 ns	1	1

Table 39q—Confidence Interval Scaling Factor subfield

Confidence interval scaling factor	Bit 1	Bit 0
Confidence interval \times 1/2	0	0
Confidence interval \times 1	0	1
Confidence interval \times 2	1	0
Confidence interval \times 4	1	1

Table 39r—FoM values with the extension bit set

	Bit 7	6	5	4	3	2	1	0
UWB Ranging Start is uncorrected	1	0	0	0	0	0	0	0
Reserved	1	Any nonzero value						

The FoM characterizes the accuracy of the PHY estimate of the arrival time of the leading edge of the first pulse of the header at the antenna. The FoM in a particular timestamp report shall characterize the accuracy of the first pulse of the header that corresponds to the timer counter value in the same timestamp report.

The FoM value of 0x80 is specifically used to signal the upper layer that the RangingCounterStart value is not correct and the upper layer must use the sounding primitives. The FoM value of 0x00 is special and means “no FoM.” No FoM means that there simply is no information about the quality of a ranging measurement. That is different from reporting a very low quality measurement, but it is known that the measurement cannot be trusted. The FoM value 0x00 is not used to report untrustworthy measurements. The most untrustworthy measurement reportable is 0x79.

Change the text in 6.9 as shown:

6.9 General radio specifications

The specifications in 6.9.1 through 6.9.9 apply to ~~both~~ the 2450 MHz DS PHY described in 6.5.1 through 6.5.3, the CSS PHY described in 6.5a, the UWB PHY described in 6.8a, and the 868/915 MHz PHYs described in 6.6 through 6.8 and, with the exception of 6.9.3 and 6.9.5, apply to all PHY implementations including the alternate PHYs. The specification of 6.9.3 does not apply to the CSS PHY nor the UWB PHY. The specification of 6.9.5 does not apply to the UWB PHY.

Change the text in 6.9.1 as shown:

6.9.1 TX-to-RX turnaround time

The TX-to-RX turnaround time shall be less than or equal to $aTurnaroundTime$ (see 6.4.1).

The TX-to-RX turnaround time for DS modulation is defined as the shortest time possible at the air interface from the trailing edge of the last part/chip (of the last symbol) of a transmitted PPDU to the leading edge of the first part/chip (of the first symbol) of the next received PPDU.

The TX-to-RX turnaround time for CSS modulation is defined as the shortest time possible at the air interface from the trailing edge of the last chirp (of the last symbol) of a transmitted PPDU to the leading edge of the first chirp (of the first symbol) of the next received PPDU.

The TX-to-RX turnaround time shall be less than or equal to the RX-to-TX turnaround time.

Change the first sentence in the last paragraph in 6.9.3 as shown:

6.9.3 Error-vector magnitude (EVM) definition

With the exception of the UWB PHY transmitter as described in 6.8a and the CSS PHY transmitter as described in 6.5a, a transmitter shall have EVM values of less than 35% when measured for 1000 chips.

Change the text of 6.9.4 through 6.9.7 and 6.9.9 as shown:

6.9.4 Transmit center frequency tolerance

The transmitted center frequency tolerance shall be ± 40 ppm maximum except in the case of the UWB PHY in which the tolerance on the chipping clock given in 6.8a.12.3 takes precedence and the center frequency tolerance is ± 20 ppm. It should be noted that a tighter frequency tolerance could facilitate a more precise range outcome for UWB devices.

6.9.5 Transmit power

A transmitter shall be capable of transmitting at least -3 dBm with the exception of UWB devices, which have no minimum transmit power dictated by this standard. Devices should transmit lower power when possible in order to reduce interference to other devices and systems.

The maximum transmit power is limited by local regulatory bodies.

6.9.6 Receiver maximum input level of desired signal

The receiver maximum input level is the maximum power level of the desired signal present at the input of the receiver for which the error rate criterion in 6.1.7 is met. A receiver shall have a receiver maximum input level greater than or equal to -20 dBm with the exception of a UWB receiver, which shall have a maximum input level greater than or equal to -45 dBm/MHz.

6.9.7 Receiver ED

The receiver ED measurement is intended for use by a network layer as part of a channel selection algorithm. It is an estimate of the received signal power within the bandwidth of the channel. No attempt is made to identify or decode signals on the channel. Except for the UWB PHY, the ED measurement time, to average over, shall be equal to 8 symbol periods. For the UWB PHY, the averaging period is implementation specific.

The ED result shall be reported to the MLME using PLME-ED.confirm primitive (see 6.2.2.4) as an 8-bit integer ranging from 0x00 to 0xff. The minimum ED value (zero) shall indicate received power less than 10 dB above the specified receiver sensitivity (see 6.5.3.3 and 6.6.3.4), and the range of received power spanned by the ED values shall be at least 40 dB. Within this range, the mapping from the received power in decibels to ED value shall be linear with an accuracy of ± 6 dB.

For UWB PHY types, the ED measurement for each channel may be performed as a series of measurements, each made at a fraction of the total channel bandwidth, in which case *phyUWBScanBinsPerChannel* specifies the number of frequency increments used. When this value is greater than 1, the ED result reported using PLME-ED.confirm primitive shall be a list of ED measurements, one for each frequency increment measurement. An implementation may provide multiple ED measurements, for example, to provide information to a higher layer that detects non-UWB services for the purpose of active DAA procedures as may be required in some environments.

6.9.9 CCA

The PHY shall provide the capability to perform CCA on the channel specified by *phyCurrentChannel* and *phyCurrentPage* according to at least one of the following ~~three~~ six methods (modes 4, 5, and 6 apply only to the UWB PHY):

- *CCA Mode 1: Energy above threshold.* CCA shall report a busy medium upon detecting any energy above the ED threshold.
- *CCA Mode 2: Carrier sense only.* CCA shall report a busy medium only upon the detection of an IEEE 802.15.4 signal ~~compliant with this standard~~ with the same modulation and spreading characteristics of the PHY that is currently in use by the device. This signal may be above or below the ED threshold.
- *CCA Mode 3: Carrier sense with energy above threshold.* CCA shall report a busy medium using a logical combination of
 - Detection of a signal with the modulation and spreading characteristics of ~~this standard~~ IEEE Std 802.15.4 and
 - Energy above the ED threshold, where the logical operator may be AND or OR.
- *CCA Mode 4: ALOHA.* CCA shall always report an idle medium.
- *CCA Mode 5: UWB preamble sense based on the SHR of a frame.* CCA shall report a busy medium upon detection of a preamble symbol as specified in 6.8a.6. An idle channel shall be reported if no preamble symbol is detected up to a period not shorter than the maximum packet duration plus the maximum period for acknowledgment.

- CCA mode 6: UWB preamble sense based on the packet with the multiplexed preamble as specified in 6.8a.14. CCA shall report a busy medium upon detection of a preamble symbol as specified in 6.8a.6.

For any of the CCA modes, if the PLME-CCA.request primitive (see 6.2.2.1) is received by the PHY during reception of a PPDU, CCA shall report a busy medium. PPDU reception is considered to be in progress following detection of the SFD, and it remains in progress until the number of octets specified by the decoded PHR has been received.

A busy channel shall be indicated by the PLME-CCA.confirm primitive (see 6.2.2.2) with a status of BUSY.

A clear channel shall be indicated by the PLME-CCA.confirm primitive with a status of IDLE.

The PHY PIB attribute *phyCCAMode* (see 6.4) shall indicate the appropriate operation mode. The CCA parameters are subject to the following criteria:

- a) The ED threshold shall correspond to a received signal power of at most 10 dB above the specified receiver sensitivity (see 6.5.3.3, 6.5a.5.3, 6.6.3.4, 6.7.3.4, and 6.8.3.4).
- b) The CCA detection time shall be equal to 8 symbol periods for the 868/915 MHz and the 2450 MHz bands. The UWB CCA detection time for CCA mode 6 shall be equal to 40 mandatory symbol periods, which includes at least 8 (multiplexed) preamble symbols (see 6.8a.14).

7. MAC sublayer specification

Change the fifth item in the dashed list as shown:

- Employing the CSMA-CA mechanism for channel access, except for UWB PHYs where ALOHA is used

7.1 MAC sublayer service specification

7.1.1 MAC data service

7.1.1.1 MCPS-DATA.request

7.1.1.1.1 Semantics of the service primitive

Insert the following new parameters at the end of the list in 7.1.1.1.1 (before the closing parenthesis):

UWBPRF,
Ranging,
UWBPreAmbleSymbolRepetitions,
DataRate

Insert the following new rows at the end of Table 41:

Table 41—MCPS-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	OFF, ALL_RANGING, PHY_HEADER_ONLY	A value of OFF indicates that ranging is not to be used for the 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 OFF is used for non-UWB PHYs.
UWBPreAmble-SymbolRepetitions	Enumeration	0, 16, 64, 1024, 4096	The preamble symbol repetitions of the UWB PHY frame to be transmitted by the PHY entity. A value of 0 is used for non-UWB PHYs.

Table 41—MCPS-DATA.request parameters (continued)

Name	Type	Valid range	Description
DataRate	Enumeration	0, 1, 2, 3, 4	The data rate of the PHY frame to be transmitted by the PHY entity. A value of 0 is used with a non-UWB or non-CSS PHY. A value of 1 or 2 is used with CSS PHYs (1 corresponds to 250 kb/s rate and 2 corresponds to 1 Mb/s rate) and 1 though 4 with UWB PHYs. See Table 39g (in 6.8a.7.1) for UWB rate definitions.

7.1.1.2 MCPS-DATA.confirm

7.1.1.2.1 Semantics of the service primitive

Insert the following new parameters at the end of the list in 7.1.1.2.1 (before the closing parenthesis):

RangingReceived,
RangingCounterStart,
RangingCounterStop,
RangingTrackingInterval,
RangingOffset,
RangingFOM

Change Table 42 (the entire table is not shown) as indicated:

Table 42—MCPS-DATA.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, TRANSACTION_OVERFLOW, TRANSACTION_EXPIRED, CHANNEL_ACCESS_FAILURE, INVALID_ADDRESS, INVALID_GTS, NO_ACK, COUNTER_ERROR, FRAME_TOO_LONG, UNAVAILABLE_KEY, UNSUPPORTED_SECURITY, or INVALID_PARAMETER, UNSUPPORTED_PREF, or UNSUPPORTED_RANGING	The status of the last MSDU transmission.
<u>Ranging-Received</u>	<u>Boolean</u>	<u>OFF, ON</u>	<u>A value of OFF indicates that ranging is either not supported in a UWB PHY or not to be indicated for the PSDU received. A value of ON denotes ranging operations requested for this PSDU. A value of OFF is used for non-UWB PHYs.</u>

Table 42—MCPS-DATA.confirm parameters (continued)

Name	Type	Valid range	Description
<u>Ranging-CounterStart</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>Ranging-CounterStop</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>Ranging-TrackingInterval</u>	<u>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>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.</u>
<u>RangingFOM</u>	<u>Integer</u>	<u>0x00–0x7F</u>	<u>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.</u>

7.1.1.3 MCPS-DATA.indication

7.1.1.3.1 Semantics of the service primitive

Insert the following new parameters at the end of the list in 7.1.1.3.1 (before the closing parenthesis):

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

Insert the following new rows at the end of Table 43:

Table 43—MCPS-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	0, 16, 64, 1024, 4096	The preamble symbol repetitions of the UWB PHY frame received by the PHY entity. A value of 0 is used with a non-UWB PHY.
DataRate	Enumeration	0, 1, 2, 3, 4	The data rate of the PHY frame received by the PHY entity. A value of 0 is used with a non-UWB or non-CSS PHY.
RangingReceived	Enumeration	NO_RANGING_ REQUESTED, RANGING_ACTIVE, RANGING_ REQUESTED_BUT_ NOT_SUPPORTED	A value of RANGING_REQUESTED_BUT_NOT_SUPPORTED indicates that ranging is not supported in a UWB PHY but has been requested. A value of NO_RANGING_REQUESTED indicates that no ranging is requested for the PSDU received. A value of RANGING_ACTIVE denotes ranging operations requested for this PSDU. A value of NO_RANGING_REQUESTED is used for non-UWB PHYs.
Ranging-CounterStart	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.14.1.

Table 43—MCPS-DATA.indication parameters (continued)

Name	Type	Valid range	Description
Ranging-CounterStop	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.14.1.
Ranging-TrackingInterval	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	0x000000–0x0FFFFFF	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.

7.1.2 MAC management service

Insert the following new rows at the end of Table 46:

Table 46—Summary of the primitives accessed through the MLME-SAP

Name	Request	Indication	Response	Confirm
MLME-DPS	7.1.16a.1	7.1.16a.3		7.1.16a.2
MLME-SOUNDING	7.1.16b.1			7.1.16b.2
MLME-CALIBRATE	7.1.16c.1			7.1.16c.2

7.1.5 Beacon notification primitive

7.1.5.1 MLME-BEACON-NOTIFY.indication

7.1.5.1.1 Semantics of the service primitive

Insert the following new rows at the end of Table 55:

Table 55—Elements of PANDescriptor

Name	Type	Valid range	Description
SubChannelCode	Bitmap	0x000000– 0xFFFFFFFF (UWB) 0x000000– 0x00000F (CSS)	Subchannel preamble codes detected for UWB or subchirp codes for CSS PHY. Specifies the code(s) in use when the channel was detected. The LSB corresponds to code index 1 while the MSB corresponds to code index 24 in Table 39e and Table 39d. For other PHY types, this parameter is set to zero.

7.1.10 Primitives for specifying the receiver enable time

7.1.10.1 MLME-RX-ENABLE.request

7.1.10.1.1 Semantics of the service primitive

Insert the following new parameter at the end of the list in 7.1.10.1.1 (before the closing parenthesis):

RangingRxControl

Insert the following new row at the end of Table 65:

Table 65—MLME-RX-ENABLE.request parameters

Name	Type	Valid range	Description
RangingRxControl	Enumeration	RANGING_OFF, RANGING_ON	Configure the transceiver to Rx with ranging for a value of RANGING_ON or to not enable ranging for RANGING_OFF. Ranging control is only valid for UWB PHYs.

7.1.10.2 MLME-RX-ENABLE.confirm

7.1.10.2.1 Semantics of the service primitive

Change Table 66 as shown:

Table 66—MLME-RX-ENABLE.confirm parameter

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, PAST_TIME, ON_TIME_TOO_LONG, or <u>INVALID_PARAMETER_or</u> <u>RANGING_NOT_SUPPORTED</u>	The result of the request to enable or disable the receiver.

7.1.11 Primitives for channel scanning**7.1.11.2 MLME-SCAN.confirm****7.1.11.2.1 Semantics of the service primitive**

Insert the following new parameters at the end of the list in 7.1.11.2.1 (before the closing parenthesis):

DetectedCategory
UWBEnergyDetectList

Insert the following new rows at the end of Table 68:

Table 68—MLME-SCAN.confirm parameters

Name	Type	Valid range	Description
DetectedCategory	Integer	0x00–0xFF	Categorization of energy detected in channel with the following values: 0: Category detection is not supported 1: UWB PHY detected 2: Non-UWB PHY signal source detected 3–25: Reserved for future use
UWBEnergy-DetectList	Array of Integers	0x00–0xFF for each element of array	For UWB PHYs, the list of energy measurements taken. The total number of measurement is indicated by ResultListSize. This parameter is null for active, passive, and orphan scans. It is also null for non-UWB PHYs.

7.1.11.2.2 When generated

Change the second sentence of the first paragraph of 7.1.11.2.2 as shown:

If the MLME-SCAN.request primitive requested an active, passive, or orphan scan, the EnergyDetectList and UWBEnergyDetectList parameters will be null.

Insert the following new paragraph after the first paragraph of 7.1.11.2.2:

If the MLME-SCAN.request primitive requested an ED and the PHY type is UWB (the *phyChannelPage* provides this information), then the UWBEnergyDetectList contains the results for the UWB channels scanned, and the EnergyDetectList and PANDescriptorList are null. The UWB scan is fully described in 7.5.2.1.

7.1.16 Primitives for requesting data from a coordinator

Insert after 7.1.16.3 the following new subclauses (7.1.16a through 7.1.16c.2.3):

7.1.16a Primitives for specifying dynamic preamble (for UWB PHYs)

MLME-SAP DPS primitives define how a device can enable or disable DPS as well as define the value of dynamic preamble for transmission and reception for a given time.

7.1.16a.1 MLME-DPS.request

The MLME-DPS.request primitive allows the next higher layer to request that the PHY utilize a given pair of preamble codes for a single use pending expiration of the DPSIndexDuration.

7.1.16a.1.1 Semantics of the service primitive

The semantics of the MLME-DPS.request primitive is as follows:

```
MLME-DPS.request
    (
        TxDPSIndex
        RxDPSIndex
        DPSIndexDuration
    )
```

Table 77a specifies the parameters for the MLME-DPS.request primitive.

Table 77a—MLME-DPS.request parameters

Name	Type	Valid range	Description
TxDPSIndex	Integer	0x00 0x0D–0x10 and 0x15–0x18	The index value for the transmitter. 0x00 disables the index and indicates that the <i>phyCurrentCode</i> value is to be used. See 6.8a.6.1 and Table 39e. 0x0D = index 13; 0x0E = index 14; 0x0F = index 15; ... 0x18 = index 24.
RxDPSIndex	Integer	0x00 0x0D–0x10 and 0x15–0x18	The index value for the receiver. 0x00 - disables the index and indicates that the <i>phyCurrentCode</i> value is to be used. See 6.8a.6.1 and Table 39e. 0x0D = index 13; 0x0E = index 14; 0x0F = index 15; ... 0x18 = index 24.
DPSIndexDuration	Integer	0x000000–0xfffff	The number of symbols for which the transmitter and receiver will utilize the respective DPS indices if a MCPS-DATA.request primitive is not issued.

7.1.16a.1.2 Appropriate usage

The MLME-DPS.request primitive is generated by the next higher layer and issued to the MLME to enable the transmitter and receiver dynamic preambles for a single use. The DPSIndexDuration constrains the length of time that the MLME-DPS.request primitive may be outstanding before a following MCPS-DATA.request primitive occurs. This primitive may also be generated to cancel a previously generated

request to enable the transmitter and receiver dynamic preambles. The use of the index for the transmitter and receiver is enabled or disabled exactly once per primitive request.

7.1.16a.1.3 Effect on receipt

The MLME will treat the MLME-DPS.request primitive as two parts. The first part operates in the MLME and consists of the timer that assures that the device returns to a normal operating state with default preambles if a following MCPS-DATA.request primitive does not occur. The second part of the MLME-DPS.request primitive initiates the PLME-DPS.request primitive to the PHY entity. A PLME-DPS.request primitive additionally occurs if the DPSIndexDuration expires prior to a MCPS-DATA.request primitive being received by the MLME. The content of the second PLME-DPS.request primitive is defined as a zero value for the RxDPSIndex and the TxDPSIndex.

Upon completion of initiating the timer in the MLME and receiving a PLME-DPS.confirm primitive, the MLME initiates a MLME-DPS.confirm primitive with the appropriate status parameter enumeration.

If any parameter in the MLME-DPS.request primitive is not supported or is out of range, the MAC sublayer issues the MLME-DPS.confirm primitive with a status of DPS_NOT_SUPPORTED. If the request to enable or disable the DPS was successful, the MLME issues the MLME-DPS.confirm primitive with a status of SUCCESS.

7.1.16a.2 MLME-DPS.confirm

The MLME-DPS.confirm primitive reports the results of the attempt to enable or disable the DPS.

7.1.16a.2.1 Semantics of the service primitive

The semantics of the MLME-DPS.confirm primitive is as follows:

```
MLME-DPS.confirm      (
                        Status
                        )
```

Table 77b specifies the parameter for the MLME-DPS.confirm primitive.

Table 77b—MLME-DPS.confirm parameter

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, DPS_NOT_SUPPORTED	The result of the request to enable or disable dynamic preambles

7.1.16a.2.2 When generated

The MLME-DPS.confirm primitive is generated by the MLME and issued to its next higher layer in response to an MLME-DPS.request primitive.

7.1.16a.2.3 Appropriate usage

On receipt of the MLME-DPS.confirm primitive, the next higher layer is notified of its request to enable or disable the dynamic preamble index for the transmitter and receiver. This primitive returns a status of either SUCCESS if the request to enable or disable was successful or DPS_NOT_SUPPORTED if a failure.

7.1.16a.3 MLME-DPS.indication

The MLME-DPS.indication primitive reports the results of the expiration of the DPSIndexDuration and the resetting of the DPS values in the PHY.

7.1.16a.3.1 Semantics of the service primitive

The semantics of the MLME-DPS.indication primitive is as follows:

```
MLME-DPS.indication      (
                          Status
                          )
```

Table 77c specifies the parameter for the MLME-DPS.indication primitive.

Table 77c—MLME-DPS.indication parameter

Name	Type	Valid range	Description
Status	Enumeration	RESET_OF_DPS	If a MCPS-DATA.request primitive is not received before the timer expires, this status is indicated to the next higher layer. The MLME will issue a PLME-DPS.request primitive to reset the DPS values.

7.1.16b Primitives for channel sounding (for UWB PHYs)

MLME-SAP channel sounding primitives define how a device can obtain the results of a channel sounding from an RDEV that supports the optional sounding capability.

7.1.16b.1 MLME-SOUNDING.request (UWB PHYs only)

The MLME-SOUNDING.request primitive allows the next higher layer to request that the PHY respond with channel sounding information. The MLME-SOUNDING.request primitive is optional except for implementations providing ranging. Although the MLME-SOUNDING.request primitive shall be supported by all RDEVs, the underlying sounding capability is optional in all cases.

7.1.16b.1.1 Semantics of the service primitive

The semantics of the MLME-SOUNDING.request primitive is as follows:

```
MLME-SOUNDING.request    (
                          )
```

7.1.16b.1.2 Appropriate usage

The MLME-SOUNDING.request primitive is generated by the next higher layer and issued to the MLME to request a MLME-SOUNDING.confirm primitive.

7.1.16b.1.3 Effect on receipt

If the feature is supported in the UWB PHY, the MLME issues the MLME-SOUNDING.confirm primitive with a status of SUCCESS and a list of SoundingPoints of SoundingSize in length.

If the MLME-SOUNDING.request primitive is generated by the next higher layer when there is no information present, e.g., when the PHY is in the process of performing a measurement, the MLME issues the MLME-SOUNDING.confirm primitive with a value of NO_DATA.

If the MLME-SOUNDING.request primitive is generated by the next higher layer and the channel sounding capability is not present in the PHY, the MLME issues the MLME-SOUNDING.confirm primitive with a value of UNSUPPORTED_ATTRIBUTE.

7.1.16b.2 MLME-SOUNDING.confirm (UWB PHYs only)

The MLME-CHANNEL.confirm primitive reports the result of a request to the PHY to provide channel sounding information. The MLME-SOUNDING.confirm primitive is optional except for implementations providing ranging. Although the MLME-SOUNDING.confirm primitive shall be supported by all RDEVs, the underlying sounding capability is optional in all cases.

7.1.16.2b.1 Semantics of the service primitive

The semantics of the MLME-SOUNDING.confirm primitive is as follows:

```
MLME-SOUNDING.confirm
(
    Status,
    SoundingSize,
    SoundingList
)
```

Table 77d specifies the parameters for the MLME-SOUNDING.confirm primitive

Table 77d—MLME-SOUNDING.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, NO_DATA, SOUNDING_NOT_SUPPORTED	The status of the attempt to return sounding data.
SoundingSize	Unsigned Integer	0x0000–0xFFFF	Number of SoundingPoints to be returned. Each SoundingPoint is 4 octets.
SoundingList	List of Pairs of Signed Integers	0x00000000–0xFFFFFFFF for each element in the list. Each element in the list is a SoundingPoint. See Table 17d.	The list of sounding measurements. See 5.5.7.4.5.

Table 17d lists the parameters in the SoundingList. Each element of the SoundingList contains a SoundingTime and a SoundingAmplitude. The SoundingTime is a signed integer, and the LSB represents a nominal 16 ps (1/128 of a chip time). The SoundingAmplitude is a signed integer representing a relative measurement. The SoundingAmplitudes have no absolute meaning, only a relative meaning.

7.1.16b.2.2 When generated

The MLME-SOUNDING.confirm primitive is generated by the MLME and issued to its next higher layer in response to a MLME-SOUNDING.request primitive. The MLME-SOUNDING.confirm primitive returns a status of SUCCESS to indicate channel sounding information is available or returns an error code of NO_DATA or SOUNDING_NOT_SUPPORTED. In the case of status of SUCCESS, the MLME-SOUNDING.confirm primitive also returns the primitive parameters SoundingSize and SoundingList.

7.1.16.2b.3 Appropriate usage

On receipt of the MLME-SOUNDING.confirm primitive, the next higher layer is notified of the results of the channel sounding information request. If the channel sounding information was available, the status parameter is set to SUCCESS. Otherwise, the status parameter will indicate an error.

7.1.16c Primitives for ranging calibration (for UWB PHYs)

MLME-SAP ranging calibration primitives define how a device can obtain the results of a ranging calibration request from an RDEV.

7.1.16c.1 MLME-CALIBRATE.request (UWB PHYs only)

The MLME-CALIBRATE.request primitive attempts to have the PHY respond with RMARKER offset information. The MLME-CALIBRATE.request primitive is optional except for implementations providing ranging.

7.1.16c.1.1 Semantics of the service primitive

The semantics of the MLME-CALIBRATE.request primitive is as follows:

```
MLME-CALIBRATE.request      (
                               )
```

7.1.16c.1.2 Appropriate usage

The MLME-CALIBRATE.request primitive is generated by the next higher layer and issued to its MLME to request a MLME-CALIBRATE.confirm primitive.

7.1.16c.1.3 Effect on receipt

If the feature is supported in the UWB PHY, the MLME issues the MLME-CALIBRATE.confirm primitive with a status of SUCCESS and a pair of integers CalTx_RMARKER_Offset and CalRx_RMARKER_Offset. If the MLME-CALIBRATE.request primitive is generated by the next higher layer when there is no information present, e.g., when the PHY is in the process of performing a measurement, the MLME issues the MLME-CALIBRATE.confirm primitive with a value of NO_DATA.

If the MLME-CALIBRATE.request primitive is generated by the MLME and the PHY does not support autonomous self-calibration, the MLME issues the MLME-CALIBRATE.confirm primitive with a value of COMPUTATION_NEEDED. The COMPUTATION_NEEDED signals the next higher layer that it should use the sounding primitives to finish the calibration (see 5.5.7.6.3).

If the MLME-CALIBRATE.request primitive is generated by the MLME and the channel sounding capability is not present in the PHY, the MLME issues the MLME-CALIBRATE.confirm primitive with a value of UNSUPPORTED_ATTRIBUTE.

7.1.16c.2 MLME-CALIBRATE.confirm (UWB PHYs only)

The MLME-CALIBRATE.confirm primitive reports the result of a request to the PHY to provide internal propagation path information. The MLME-CALIBRATE.confirm primitive is optional except for implementations providing ranging.

7.1.16c.2.1 Semantics of the service primitive

The semantics of the MLME-CALIBRATE.confirm primitive is as follows:

```
MLME-CALIBRATE.confirm
    (
        Status,
        CalTx_RMARKER_Offset,
        CalRx_RMARKER_Offset,
    )
```

Table 77e specifies the parameters for the MLME-CALIBRATE.confirm primitive.

Table 77e—MLME-CALIBRATE.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS, COMPUTATION_ NEEDED, NO_DATA, UNSUPPORTED_ ATTRIBUTE	The status of the attempt to return sounding data.
CalTx_RMARKER_Offset	Unsigned Integer	0x00000000– 0xFFFFFFFF	A 4-octet count of the propagation time from the ranging counter to the transmit antenna. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.
CalRx_RMARKER_Offset	Unsigned Integer	0x00000000– 0xFFFFFFFF	A 4-octet count of the propagation time from the receive antenna to the ranging counter. The LSB of a time value represents 1/128 of a chip time at the mandatory chipping rate of 499.2 MHz.

7.1.16c.2.2 When generated

The MLME-CALIBRATE.confirm primitive is generated by the MLME and issued to its next higher layer in response to a MLME-CALIBRATE.request primitive. The MLME-CALIBRATE.confirm primitive returns a status of SUCCESS to indicate channel propagation time information is available and part of the MLME-CALIBRATE.confirm parameters, or returns a status of COMPUTATION_NEEDED if the PHY lacks the computational resources to determine the offsets, or returns an error code of NO_DATA or UNSUPPORTED_ATTRIBUTE.

7.1.16c.2.3 Appropriate usage

On receipt of the MLME-CALIBRATE.confirm primitive, the next higher layer is notified of the results of the self-calibrate information request. If the RMARKER offset information was available, the status parameter is set to SUCCESS. If the PHY performed a sounding of a loopback path but lacks the computational resources to complete the processing of the sounding data, the status parameter is set to COMPUTATION_NEEDED. Otherwise, the status parameter will indicate an error.

7.1.17 MAC enumeration description

Change rows two and three in Table 78 (the entire table is not shown) as indicated and then insert the new rows in alphabetical order into the table:

Table 78—MAC enumerations description

Enumeration	Value	Description
—	0x01–0x da <u>7f</u>	Reserved for MAC command status and reason code values.
—	0x80–0x da <u>af</u> , 0xcc–0xcf, 0xfe–0xff	Reserved.
<u>ALL_RANGING</u>	<u>0xb5</u>	<u>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.</u>
<u>COMPUTATION_NEEDED</u>	<u>0xc6</u>	<u>The value returned when the next higher layer should use the sounding primitives to finish the calibration.</u>
<u>DATA_RATE_0</u>	<u>0xbf</u>	<u>PHYs that are neither UWB or CSS do not use the data rate parameter; therefore, this value is used.</u>
<u>DATA_RATE_1</u>	<u>0xc0</u>	<u>Data rate parameter is 1.</u>
<u>DATA_RATE_2</u>	<u>0xc1</u>	<u>Data rate parameter is 2.</u>
<u>DATA_RATE_3</u>	<u>0xc2</u>	<u>Data rate parameter is 3.</u>
<u>DATA_RATE_4</u>	<u>0xc3</u>	<u>Data rate parameter is 4.</u>
<u>DPS_NOT_SUPPORTED</u>	<u>0xc4</u>	<u>The value when dynamic preamble select is not supported.</u>
<u>NOMINAL_4_M</u>	<u>0xb1</u>	<u>PRF with a nominal 4 MHz.</u>
<u>NOMINAL_16_M</u>	<u>0xb2</u>	<u>PRF with a nominal 16 MHz.</u>
<u>NOMINAL_64_M</u>	<u>0xb3</u>	<u>PRF with a nominal 64 MHz.</u>
<u>NON_RANGING</u>	<u>0xb4</u>	<u>Ranging either is not supported or is not selected for the PSDU to be transmitted.</u>
<u>NO_RANGING_REQUESTED</u>	<u>0xc7</u>	<u>A value of NO_RANGING_REQUESTED indicates that no ranging is requested for the PSDU received.</u> <u>A value of NO_RANGING_REQUESTED is used for non-UWB PHYs.</u>
<u>PHY_HEADER_ONLY</u>	<u>0xb6</u>	<u>PHY_HEADER_ONLY denotes ranging operations for this PSDU using only the ranging bit in the PHR set to one.</u>

Table 78—MAC enumerations description (continued)

Enumeration	Value	Description
<u>PRF_OFF</u>	<u>0xb0</u>	<u>Non-UWB PHYs do not use PRF parameter; therefore, this value is used.</u>
<u>PSR_0</u>	<u>0xba</u>	<u>Non-UWB PHYs do not use preamble symbol repetitions parameter; therefore, this value is used.</u>
<u>PSR_16</u>	<u>0xbb</u>	<u>Preamble symbol repetitions are 16 in number.</u>
<u>PSR_64</u>	<u>0xbc</u>	<u>Preamble symbol repetitions are 64 in number.</u>
<u>PSR_1024</u>	<u>0xbd</u>	<u>Preamble symbol repetitions are 1024 in number.</u>
<u>PSR_4096</u>	<u>0xbe</u>	<u>Preamble symbol repetitions are 4096 in number.</u>
<u>RANGING_ACTIVE</u>	<u>0xc8</u>	<u>A value of RANGING_ACTIVE denotes ranging operations requested for this PSDU.</u>
<u>RANGING_NOT_SUPPORTED</u>	<u>0xb7</u>	<u>The value when the receiver does not support ranging.</u>
<u>RANGING_OFF</u>	<u>0xb8</u>	<u>Configure the transceiver to Rx with ranging set to off.</u>
<u>RANGING_ON</u>	<u>0xb9</u>	<u>Configure the transceiver to Rx with ranging set to on.</u>
<u>RANGING_REQUESTED_BUT_NOT_SUPPORTED</u>	<u>0xc9</u>	<u>A value of RANGING_REQUESTED_BUT_NOT_SUPPORTED indicates that ranging is not supported in a UWB PHY but has been requested.</u>
<u>SOUNDING_NOT_SUPPORTED</u>	<u>0xc5</u>	<u>The value when sounding is not supported.</u>
<u>UNSUPPORTED_PRF</u>	<u>0xca</u>	<u>Data.confirm error status returned when a corresponding Data.request command is issued with an unsupported PRF value.</u>
<u>UNSUPPORTED_RANGING</u>	<u>0xcb</u>	<u>Data.confirm error status returned when a corresponding Data.request command is issued with Ranging = ALL_RANGING, but the PHY does not support a ranging counter.</u>

7.4 MAC constants and PIB attributes

7.4.2 MAC PIB attributes

Insert the following text after the last paragraph in 7.4.2 (i.e., before Table 86):

For the UWB PHY, where the CCA method is ALOHA and there are two octets in the PHR, the formula for *macAckWaitDuration* reduces to Equation (14a).

$$macAckWaitDuration = aTurnaroundTime + phySHRDuration + ceil(7 \times phySymbolsPerOctet) \quad (14a)$$

where

7 represents the number of PHR octets (2) plus the number of PSDU octets in an acknowledgment frame (5)

phySHRDuration is calculated as shown in 6.4.2.1

For UWB PHYs, *macMaxBE*, *macMaxCSMABackoffs*, and *m* become zero; therefore, the formula for *macMaxFrameTotalWaitTime* reduces to Equation (14b).

$$macMaxFrameTotalWaitTime = phyMaxFrameDurations/T_{psym} \quad (14b)$$

where

T_{psym} is from Table 39b (appropriate for the channel, PRF, and preamble code)

phyMaxFrameDurations is also given in 6.4.2.1

Note that T_{psym} depends on the transmit parameters UWBPRF, UWBPreableSymbolRepetitions, and DataRate provided by the next higher layer with the MCPS-DATA.request primitive. The formula and the T_{psym} values in Table 39b are intended to aid the next higher layer implementer in determining appropriate values for turnaround and timeouts.

For the CSS PHY, the values of the read-only attribute *macAckWaitDuration* depend on the selected data rate of the PSDU.

For the mandatory data rate (1 Mb/s), *macAckWaitDuration* is calculated as shown in Equation (14c).

$$macAckWaitDuration_{1M} = aUnitBackoffPeriod + aTurnaroundTime + phySHRDuration_{1M} + [1.5 + 3/4 \times ceiling(4/3 \times 5)] \times phySymbolsPerOctet_{1M} \quad (14c)$$

For the optional data rate (250 kb/s), *macAckWaitDuration* is calculated as shown in Equation (14d).

$$macAckWaitDuration_{250k} = aUnitBackoffPeriod + aTurnaroundTime + phySHRDuration_{250k} + 3 \times ceiling(1/3 \times [1.5 + 5]) \times phySymbolsPerOctet_{250k} \quad (14d)$$

Change the `macAckWaitDuration` and `macMaxFrameTotalWaitTime` rows in Table 86 (the entire table is not shown) as indicated and then insert the new row in alphabetical order into the table:

Table 86—MAC PIB attributes

Attribute	Identifier	Type	Range	Description	Default
<code>macAckWaitDuration</code>	0x40	Integer	See Equation (13), Equation (14a), Equation (14c), and Equation (14d) for range as a function of PHY type	The maximum number of symbols to wait for an acknowledgment frame to arrive following a transmitted data frame. This value is dependent on the supported PHY, which determines both the selected logical channel and channel page. The calculated value is the time to commence transmitting the ACK plus the length of the ACK frame. The commencement time is described in 7.5.6.4.2.	Dependent on currently selected PHY, indicated by <code>phyCurrentPage</code> . Also dependent on PHY operating parameters when UWB or CSS is the selected PHY.
<code>macMaxFrameTotalWaitTime</code>	0x58	Integer	See Equation (14) and Equation (14b)	The maximum number of CAP symbols in a beacon-enabled PAN, or symbols in a nonbeacon-enabled PAN, to wait either for a frame intended as a response to a data request frame or for a broadcast frame following a beacon with the Frame Pending subfield set to one. This attribute, which shall only be set by the next higher layer, is dependent upon <code>macMinBE</code> , <code>macMaxBE</code> , <code>macMaxCSMA-Backoffs</code> and the number of symbols per octet. See 7.4.2 for the formula relating the attributes.	Dependent on currently selected PHY, indicated by <code>phyCurrentPage</code> .
<code>macRangingSupported</code> †	0x60	Boolean	TRUE or FALSE	This indicates whether the MAC sublayer supports the optional ranging features*.	Dependent on supported PHY and MAC capability.

7.5 MAC functional description

7.5.2 Starting and maintaining PANs

7.5.2.1 Scanning through channels

Insert the following new sentence after the second sentence in the second paragraph in 7.5.2.1:

For UWB and CSS PHYs, each preamble code appropriate to the specified channel is scanned.

7.5.2.1.1 ED channel scan

Insert the following new sentence after the second sentence in the first paragraph of 7.5.2.1.1:

It may also be used by a next higher layer to detect potential interferer and/or victim devices in the implementation of a DAA procedure.

Insert the following new paragraph after the second paragraph in 7.5.2.1.1:

The ED procedure implemented in the PHY is implementation-specific. For UWB PHY implementations, the UWB channel measurement may be accomplished in frequency increments, each a fraction of the full channel bandwidth. The number of increments is indicated in the *phyUWBScanBinsPerChannel* PHY PIB parameter. When a UWB ED scan is performed, the resulting *UWBEnergyDetectList* will have one entry per channel increment for each channel scanned, for a total of (*phyUWBScanBinsPerChannel* × number of channels scanned) or up to the implementation-specified maximum number of ED measurements.

7.5.2.1.2 Active channel scan

Change the fourth paragraph in 7.5.2.1.2 as shown:

An active scan over a specified set of logical channels is requested using the *MLME-SCAN.request* primitive with the *ScanType* parameter set to indicate an active scan. For each logical channel, the device shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly, and send a beacon request command (see 7.3.7). For UWB and CSS PHYs, the scan process shall be repeated for each mandatory preamble code, setting the *phyCurrentCode* appropriately. Upon successful transmission of the beacon request command, the device shall enable its receiver for at most [*aBaseSuperframeDuration* × ($2^n + 1$)] symbols, where *n* is the value of the *ScanDuration* parameter. During this time, the device shall reject all nonbeacon frames and record the information contained in all unique beacons in a PAN descriptor structure (see Table 55 in 7.1.5.1.1), including the channel information and the preamble code. If a beacon frame is received when *macAutoRequest* is set to TRUE, the list of PAN descriptor structures shall be stored by the MAC sublayer until the scan is complete; at this time, the list shall be sent to the next higher layer in the *PANDescriptorList* parameter of the *MLME-SCAN.confirm* primitive. A device shall be able to store between one and an implementation-specified maximum number of PAN descriptors. A beacon frame shall be assumed to be unique if it contains both a PAN identifier and a source address that has not been seen before during the scan of the current channel. If a beacon frame is received when *macAutoRequest* is set to FALSE, each recorded PAN descriptor is sent to the next higher layer in a separate *MLME-BEACON-NOTIFY.indication* primitive. For UWB and CSS PHYs, the beacon request is repeated for each preamble code. A received beacon frame containing one or more octets of payload shall also cause the PAN descriptor to be sent to the next higher layer via the *MLME-BEACON-NOTIFY.indication* primitive.

7.5.2.1.3 Passive channel scan

Change the fourth paragraph in 7.5.2.1.3 as shown:

A passive scan over a specified set of logical channels is requested using the *MLME-SCAN.request* primitive with the *ScanType* parameter set to indicate a passive scan. For each logical channel, the device shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly, and for UWB and CSS PHYs, setting the preamble code *phyCurrentCode* appropriately, and then enable its receiver for at most [*aBaseSuperframeDuration* × ($2^n + 1$)] symbols, where *n* is the value of the *ScanDuration* parameter. During this time, the device shall reject all nonbeacon frames and record the information contained in all unique beacons in a PAN descriptor structure (see Table 55 in 7.1.5.1.1). If a beacon frame is received when *macAutoRequest* is set to TRUE, the list of PAN descriptor structures shall be stored by the MAC sublayer until the scan is complete; at this time, the list shall be sent to the next higher layer in the *PANDescriptorList* parameter of the *MLME-SCAN.confirm* primitive. A device shall be able to store

between one and an implementation- specified maximum number of PAN descriptors. A beacon frame shall be assumed to be unique if it contains both a PAN identifier and a source address that has not been seen before during the scan of the current channel. If a beacon frame is received when *macAutoRequest* is set to FALSE, each recorded PAN descriptor is sent to the next higher layer in a separate MLME-BEACON-NOTIFY.indication primitive. For UWB and CSS PHYs, the channel scan shall be repeated for each preamble code. Once the scan is complete, the MLME-SCAN.confirm shall be issued to the next higher layer with a null PANDescriptorList. A received beacon frame containing one or more octets of payload shall also cause the PAN descriptor to be sent to the next higher layer via the MLME-BEACON-NOTIFY.indication primitive.

7.5.2.1.4 Orphan channel scan

Change the second paragraph in 7.5.2.1.4 as shown:

An orphan scan over a specified set of logical channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate an orphan scan. For each logical channel, the device shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly, and for UWB and CSS PHYs, setting the preamble code *phyCurrentCode* appropriately, and then send an orphan notification command (see 7.3.6). Upon successful transmission of the orphan notification command, the device shall enable its receiver for at most *macResponseWaitTime* symbols. If the device successfully receives a coordinator realignment command (see 7.3.8) within this time, the device shall terminate the scan. If the coordinator realignment command is not received, the process shall be repeated for each preamble code until a realignment command is received or all preamble codes for the PHY have been used.

7.5.7 GTS allocation and management

Insert after 7.5.7.6 the following new subclauses (7.5.7a through 7.5.7a.4):

7.5.7a Ranging

Ranging is an option for UWB PHYs. The fundamental measurements for ranging are achieved using a data-acknowledgment frame sequence. Ranging capabilities are enabled in an RDEV with the MCPS-DATA.request primitive. Whenever ranging is enabled in an RDEV, the RDEV delivers timestamp reports to the next higher layer as a result of events at the device antenna.

7.5.7a.1 Set-up activities before a ranging exchange

The mandatory part of ranging is limited to the generation of timestamp reports during the period that ranging is enabled in an RDEV. It is possible that an RDEV will consume more power when ranging is enabled; therefore, a natural default for an application would be to have ranging disabled. Prior to a two-way ranging exchange, both RDEVs involved in the exchange shall already have ranging enabled. Furthermore, if the optional DPS capability is to be used, there shall have been some sort of precoordination of preambles prior to the two-way ranging exchange. How this precoordination and enabling actually is accomplished is beyond the scope of this standard. It may be perfectly acceptable to accomplish the precoordination and enabling with a clock and a look-up table that says what a device should do at a particular time. Because precoordination generally involves communication and because the PHYs are designed to achieve communication, it is natural to suggest that the UWB PHY be used for precoordination.

7.5.7a.2 Finish-up activities after a ranging exchange

At the end of a two-way exchange, each device is in position of a timestamp report. To accomplish anything useful, both of those timestamp reports shall eventually come to be at the same node where computations are performed. How this movement of timestamp reports is accomplished is beyond the scope of this standard. Timestamp reports are just data. Because movement of data involves communication and because the PHYs are designed to achieve communication, it is natural to suggest that the UWB PHY be used for the final consolidation of timestamp reports.

The application is responsible for enabling the ranging mode in the RDEV before a ranging exchange. After a ranging exchange, the application is again responsible for disabling the ranging mode in the RDEV. If the application fails to disable the ranging mode in the RDEV, there will be no algorithmic harm. Ranging mode is fully compatible with other uses of the RDEV, and the only result of leaving ranging enabled when it is not really being used is that the RDEV will generate useless timestamp reports while potentially consuming significant power.

7.5.7a.3 Managing DPS

Figure 73a shows a suggested message sequence for two-way ranging. The messages represented in the two top dotted boxes are simply suggestions showing how the communications capability of the RDEV can be used to accomplish the ranging setup activities. The messages in the bottom dotted box are suggestions showing how the communications capability of the RDEV can be used to accomplish the ranging finish-up activities.

The top dotted box in Figure 73a illustrates the use of a data exchange to effect the precoordination of the preambles to be used for a two-way ranging exchange. The precoordination of preambles is needed only when using the optional DPS capability of the PHY. If optional DPS is not used, the communication sequence in the top box can be thought of as arranging for the recipient RDEV to become aware that a ranging exchange is desired and that the recipient next higher layer should enable ranging in the recipient PHY. The middle dotted box in Figure 73a illustrates the use of the MLME-DPS.request, PLME-DPS.request, MLME-DPS.confirm, and PLME-DPS.confirm primitive set. Use of this primitive set is unique to the optional DPS mode of ranging. The PLME primitives are described in 6.2.2.11 and 6.2.2.12. The MLME primitives are described in 7.1.16a.1 and 7.1.16a.2.

Upon the assertion of the PLME.DPS.confirm primitives in Figure 73a, both of the PHYs have switched from the normal length 31 preamble symbols to length 127 preamble symbols. This is desirable behavior intended to help hide the PHYs' transmissions from malicious nodes and protect the PHYs from transmissions by malicious nodes. A side effect of this mode is that neither PHYs can communicate with the rest of the network. To prevent the PHYs from becoming lost as a result of this optional behavior, the MAC sublayers on both sides of the link shall initiate timers after the receipt of a PLME-DPS.confirm. The timer shall run for DPSIndexDuration. If the timer duration is exceeded before the MAC sublayer receives the MCPS-DATA.confirm (for the originator) or the MCPS-DATA.indication primitive (for the recipient), then the MAC sublayer shall initiate a MCPS-DPS.indication to the next higher layer as described in 7.1.16a.3. Not shown in Figure 73a, one responsibility of the application, if the optional DPS capability is used, is to initiate the MLME-DPS.request primitive on both sides of the ranging link at the completion of the ranging exchange. Most typically, this MLME-DPS.request primitive would be part of the finish-up activities and would have both TxDPSIndex and RxDPSIndex set to zero to return the PHYs to using phyCurrentCode from the PIB. Also not shown in Figure 73a, another responsibility of the application is to initiate a MLME-DPS.request primitive in response to an MLME-DPS.indication. Most typically, this MLME-DPS.request primitive would have both TxDPSIndex and RxDPSIndex set to zero and return the PHY to using phyCurrentCode from the PIB.

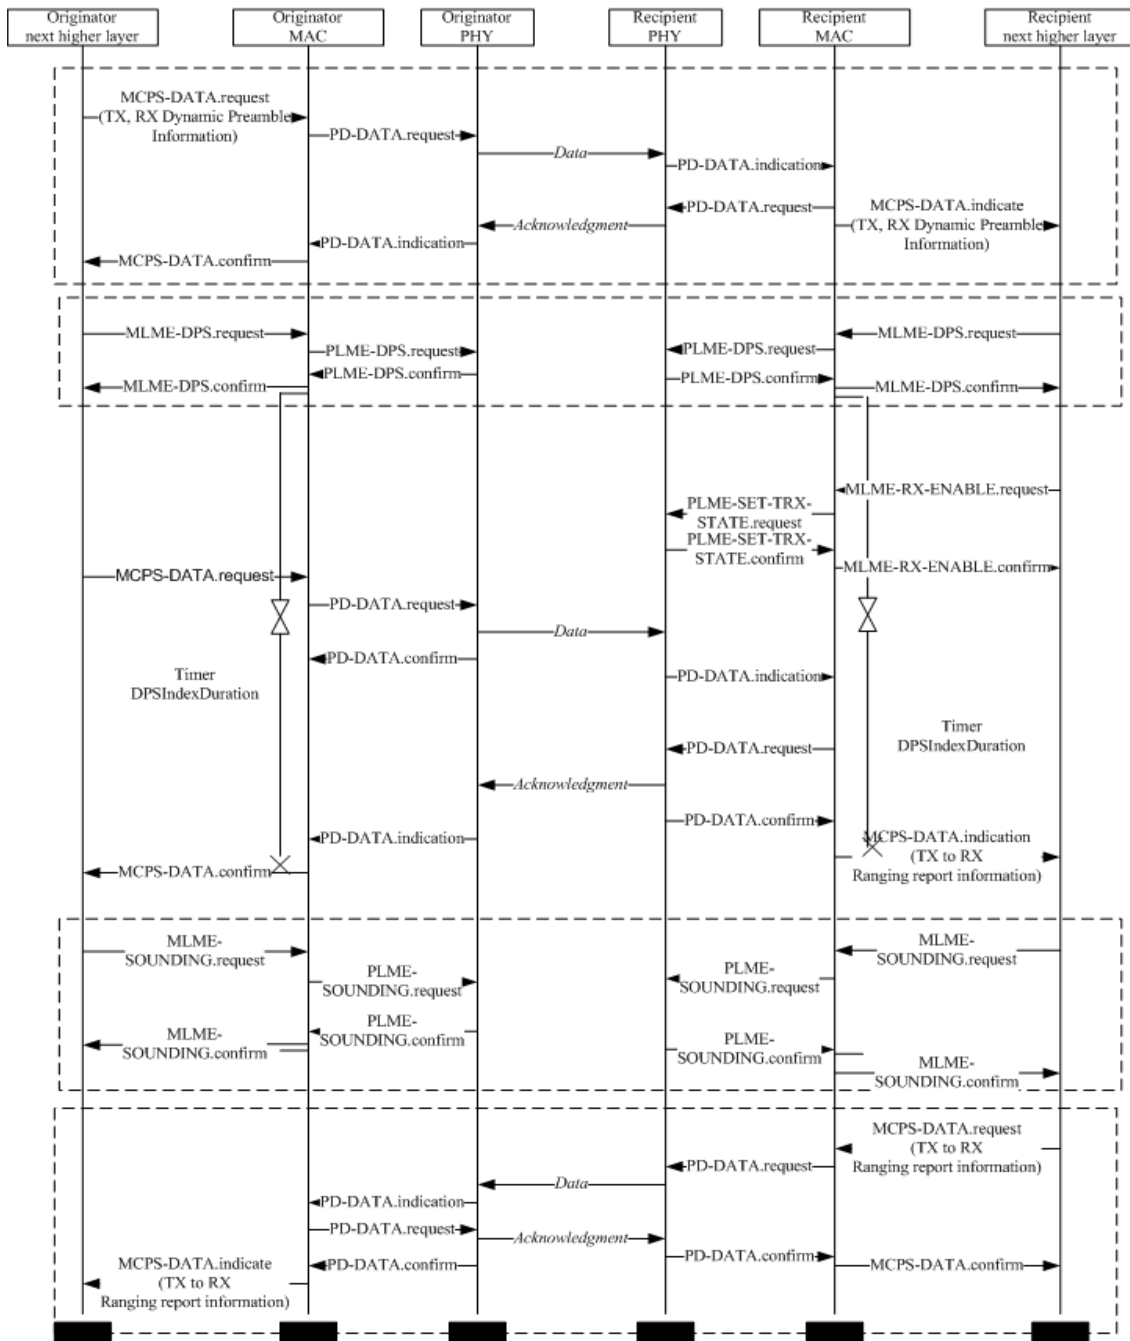


Figure 73a—A message sequence for ranging

7.5.7a.4 The ranging exchange

The essential core of the ranging exchange is shown in Figure 73a starting just below the middle dotted box. The application is responsible for initiating the MLME-RX-ENABLE.request primitive (described in 7.1.10.1) with RangingRxControl equal to RANGING_ON. That primitive in turn causes the MAC sublayer to initiate PLME-SET-TRX-STATE.request primitive (described in 6.2.2.7.1) with state equal to RX_WITH_RANGING_ON. Once the RDEV has received the MLME-RX-ENABLE.request primitive

with RangingRxControl equal to RANGING_ON, all future RFRAMES received by the RDEV shall generate timestamp reports until ranging is disabled.

At the initiator, the application is responsible for initiating a MCPS-DATA.request primitive with Ranging equal to ALL_RANGING. Upon receipt of a MCPS-DATA.request primitive with Ranging equal to ALL_RANGING, RDEV shall generate timestamp reports for all RFRAMES after the transmit frame is transmitted. The timestamp reports will continue until ranging is disabled. The Tx-to-Rx turnaround enabling the originator to receive the acknowledgment frame is necessary and is not shown in Figure 73a. This turnaround is the normal turnaround that is done for any exchange expecting an acknowledgment. The turnaround happens without any action required by the originator next higher layer. Timestamp reports are generated to the next higher layer independent of the state of the acknowledge request bit in the MAC header of received RFRAMES.

As shown in Figure 73a, the first timestamp report to the originator next higher layer shall come back as elements of the MCPS-DATA.confirm. The first timestamp report to the recipient next higher layer shall come back as elements of the MCPS-DATA.indication primitive. All subsequent timestamp reports on either side of the link shall come back as elements of MCPS-DATA.indication primitives. The potential additional MCPS-DATA.indication primitives that would be due to unexpected stray RFRAMES are not shown in Figure 73a for simplicity. The timestamp reports due to any strays shall continue until ranging is disabled. The reporting of timestamps for a stream of “strays” is the behavior that enables the RDEV to be used as an infrastructure RDEVs in one-way ranging applications.

In all of these timestamp reports, the timestamp itself shall consist of the 12 octets defined for the timestamp report in 6.8a.15.1, 6.8a.15.2, and 6.8a.15.3. Use of nonzero timestamp reports is limited to RDEVs. Only UWB PHYS may have phyRangingCapabilities bit 1 set. (That is actually what makes those devices RDEVs.) Only devices that have phyRangingCapabilities bit 1 set shall return a nonzero timestamp report to a next higher layer.

Annex D

(normative)

Protocol implementation conformance statement (PICS) proforma

D.1 Introduction

D.1.2 Scope

Change the text of D.1.1 as shown:

This annex provides the PICS proforma for IEEE Std 802.15.4-2006 and IEEE Std 802.15.4a-2007 in compliance with the relevant requirements, and in accordance with the relevant guidance, given in ISO/IEC 9646-7:1995.

D.1.2 Purpose

Change the first paragraph of D.1.2 as shown:

The supplier of a protocol implementation claiming to conform to IEEE Std 802.15.4-2006 and IEEE Std 802.15.4a-2007 shall complete the following PICS proforma and accompany it with the information necessary to identify fully both the supplier and the implementation.

D.5 Identification of the protocol

Change the text of D. 5 as shown:

This PICS proforma applies to IEEE Std 802.15.4-2006 and IEEE Std 802.15.4a-2007.

D.7 PICS proforma tables

Change the third sentence in D.7 as shown:

The first subclause contains the major roles for a device compliant with IEEE 802.15.4-2006 and IEEE 802.15.4a-2007.

D.7.2 Major capabilities for the PHY

D.7.2.1 PHY functions

Insert after the PLF8.3 row the following new rows in Table D.2 (the entire table is not shown) and change the last row as indicated:

Table D.2—PHY functions

Item number	Item description	Reference	Status	Support		
				N/A	Yes	No
<u>PLF8.4</u>	<u>Mode 4</u>	<u>6.9.9</u>	<u>RF4:O.6</u>			
<u>PLF8.5</u>	<u>Mode 5</u>	<u>6.9.9</u>	<u>RF4:O.6</u>			
<u>PLF8.6</u>	<u>Mode 6</u>	<u>6.8a.14, 6.9.9</u>	<u>RF4:O.6</u>			
<u>PLF9</u>	<u>Ranging</u>	<u>5.5.7, 6.8a.15</u>	<u>RF4:O</u>			
<u>PLF9.1</u>	<u>Crystal characterization</u>	<u>5.5.7.5, 6.8a.15.2</u>	<u>O</u>			
<u>PLF9.2</u>	<u>Dynamic preamble selection (DPS)</u>	<u>5.5.7.8.2, 6.2.2.11, 6.2.2.12</u>	<u>O</u>			
<u>PLF10</u>	<u>Ultra-wide band (UWB) pulse shape</u>	<u>5.5.8, 6.8a.13</u>	<u>RF4:M</u>			
<u>PLF10.1</u>	<u>Default</u>	<u>6.8a.12</u>	<u>M</u>			
<u>PLF10.2</u>	<u>Chirp on UWB (CoU)</u>	<u>6.8a.13.1</u>	<u>O</u>			
<u>PLF10.3</u>	<u>Continuous spectrum (CS)</u>	<u>6.8a.13.2</u>	<u>O</u>			
<u>PLF10.4</u>	<u>Linear combination of pulses (LCP)</u>	<u>6.8a.13.3</u>	<u>O</u>			
<u>PLF10.5</u>	<u>Chaotic</u>	<u>Annex H</u>	<u>O</u>			
<u>PLF11</u>	<u>Able to bin for “detect and avoid” (DAA) procedures</u>	<u>6.2.2.4</u>	<u>RF4:O</u>			
<u>PLF12</u>	<u>Supports 2450 MHz CSS 1 Mb/s</u>	<u>6.5a</u>	<u>RF3:M</u>			
<u>PLF12.1</u>	<u>Supports 2450 MHz CSS 250 kb/s</u>	<u>6.5a</u>	<u>O</u>			
<u>PLF13</u>	<u>Supports UWB 850 kb/s (rate 01)</u>	<u>6.8a.4, 6.8a.7.1</u>	<u>RF4:M</u>			
<u>PLF13.1</u>	<u>Supports rate 00</u>	<u>6.8a.4, 6.8a.7.1</u>	<u>O</u>			
<u>PLF13.2</u>	<u>Supports rate 02</u>	<u>6.8a.4, 6.8a.7.1</u>	<u>O</u>			
<u>PLF13.3</u>	<u>Supports rate 03</u>	<u>6.8a.4, 6.8a.7.1</u>	<u>O</u>			
<u>PLF13.4</u>	<u>Supports rate 04</u>	<u>6.8a.4, 6.8a.7.1</u>	<u>O</u>			
<p>O.2 At least one of these features shall be supported. O.6 At least one of these features shall be supported.</p>						

D.7.2.3 Radio frequency (RF)

Change the RF2 row in Table D.4 (the entire table is not shown), insert after the RF2 row the following new rows, and change the last row as indicated:

Table D.4—Radio frequency (RF)

Item number	Item description	Reference	Status	Support		
				N/A	Yes	No
RF2	2450 MHz <u>DSSS</u> PHY	5.4.1, Clause 6, Table 1, 6.6.1, 6.5	O.3			
<u>RF3</u>	<u>2450 MHz CSS</u> <u>PHY</u>	<u>Table 1,</u> <u>Table 1a, 6.5a</u>	<u>O.3</u>			
<u>RF4</u>	<u>UWB PHY</u>	<u>Table 1,</u> <u>Table 1b, 6.8a</u>	<u>O.3</u>			
<u>RF4.1</u>	<u>250–750 MHz</u> <u>UWB PHY</u>	<u>Table 1, 6.8a.11</u>	<u>O.5</u>			
<u>RF4.2</u>	<u>3244–4742 MHz</u> <u>UWB PHY</u>	<u>Table 1, 6.8a.11</u>	<u>O.5</u>			
<u>RF4.3</u>	<u>5944–10 234 MHz</u> <u>UWB PHY</u>	<u>Table 1, 6.8a.11</u>	<u>O.5</u>			
O.3 At least one of these features shall be supported. O.5 At least one of these features shall be supported.						

Insert after D.7.2.3 the following new subclause (D.7.2.4):

D.7.2.4 Channel capabilities for UWB PHY

Table D.4a—UWB channels

Item number	Item description	Reference	Status	Support		
				N/A	Yes	No
PCH1	Channel number 0	Table 1b	RF4.1:M			
PCH2	Channel number 1	Table 1b	RF4.2:O			
PCH3	Channel number 2	Table 1b	RF4.2:O			
PCH4	Channel number 3	Table 1b	RF4.2:M			
PCH5	Channel number 4	Table 1b	RF4.2:O			
PCH6	Channel number 5	Table 1b	RF4.3:O			
PCH7	Channel number 6	Table 1b	RF4.3:O			
PCH8	Channel number 7	Table 1b	RF4.3:O			

Table D.4a—UWB channels (continued)

Item number	Item description	Reference	Status	Support		
				N/A	Yes	No
PCH9	Channel number 8	Table 1b	RF4.3:O			
PCH10	Channel number 9	Table 1b	RF4.3:M			
PCH11	Channel number 10	Table 1b	RF4.3:O			
PCH12	Channel number 11	Table 1b	RF4.3:O			
PCH13	Channel number 12	Table 1b	RF4.3:O			
PCH14	Channel number 13	Table 1b	RF4.3:O			
PCH15	Channel number 14	Table 1b	RF4.3:O			
PCH16	Channel number 15	Table 1b	RF4.3:O			

D.7.3 Major capabilities for the MAC sublayer

D.7.3.1 MAC sublayer functions

Insert after the MLF13 row the following new rows in Table D.5:

Table D.5—MAC sublayer functions

Item number	Item description	Reference	Status	Support		
				N/A	Yes	No
MLF14	Ranging	7.5.7a	RF4:O			
MLF14.1	DPS	7.5.7a.3, 7.1.16a	O			

Insert after Annex D the following new annex (Annex D1):

Annex D1

(informative)

Location topics

D1.1 Overview

This informative annex provides supplemental material to the normative clauses of the standard in the application of the ranging capabilities of the devices for location services. The material refers to techniques and algorithms to reconstruct the location of devices in a network from range (or direction) data between units. It also considers implementation issues such as channel sounding, leading-edge detection, and the error induced by drift in timing crystal offsets.

Subclause D1.2 presents several methods and associated implementations to extract the time of arrival of a message; the latter is used to estimate the range between two devices. Subclause D1.3 introduces synchronous and asynchronous ranging and their application to different network architectures. It also discusses the effect of finite crystal tolerances on ranging precision. Subclause D1.4 outlines the aforementioned different network structures and furnishes the corresponding equations to transform gathered ranges data into location. Subclause D1.5 describes a class of network location algorithms suited in particular to large sensor networks where the location estimation takes place in a distributed fashion across the whole network to render optimal results.

D1.2 Time-of-arrival estimation from channel sounding

The range between a pair of transmitter and receiver devices can be estimated from the measured multipath profile characterizing the wireless channel between them. The peaks of the profile correspond to the arrivals, the first denoted as the time of arrival. Given τ and knowing that the signal travels at the speed of light c , the range between the two devices can be estimated as $c \cdot \tau$. The multipath profile appears in the form of a cross-correlation function of the received signal and the transmitted pseudo-random template sequence. In the proposed circuit in Figure D1.1 to estimate the delay in an AWGN channel, the receiver first samples the received signal through an ADC and then digitally correlates it with the template to generate a cross-correlation function.

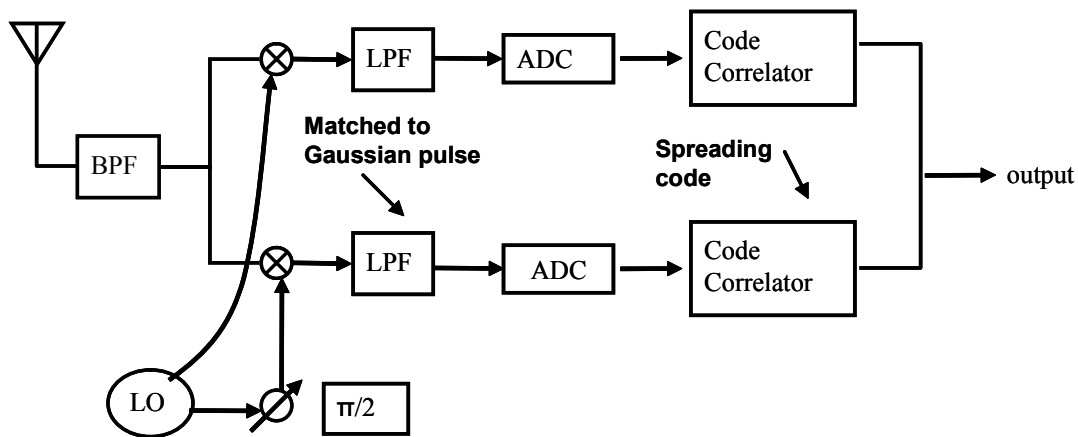


Figure D1.1—Circuit to compute multipath profile at receiver

Narrowband communication systems can afford a sampling rate equal to and up to five times the Nyquist rate² in the conventional approach to furnish good estimation accuracy. However, such a high sampling rate is difficult to implement with UWB devices demanding low cost and low power consumption. Motivated by low sampling rate, Qi and Kohno³ propose an approach tailored to such UWB devices that resorts to linear interpolation of the cross-correlation function through a second-order approximation of the maximum likelihood estimate. This estimate exploits both the given autocorrelation function of the template sequence and the given statistical characteristics of the noise, as shown in Figure D1.2.

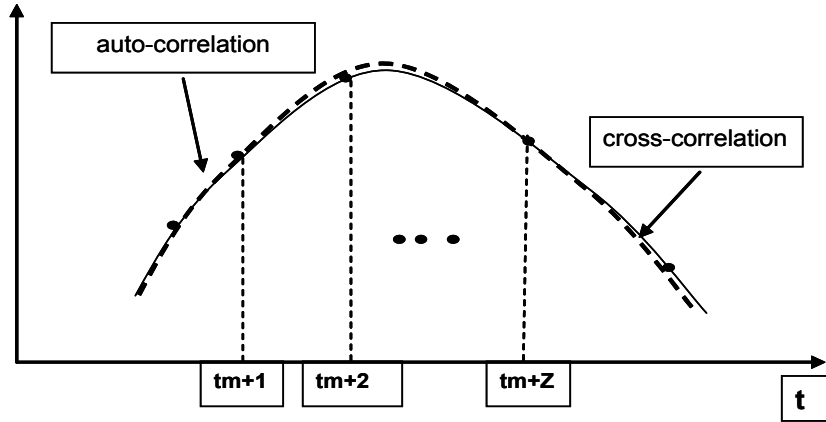


Figure D1.2—Autocorrelation and cross-correlation functions for simplified maximum likelihood estimator

Let the three largest adjacent correlation samples be denoted as $h(t_3) = [h(t_1) h(t_2) h(t_3)]^T$, the corresponding time instants as $t_3 = [t_1 t_2 t_3]^T$, and the inverse of the correlation matrix as

$$W_3 = \begin{bmatrix} g(0) & g(T_s) & g(2T_s) \\ g(T_s) & g(0) & g(2T_s) \\ g(2T_s) & g(T_s) & g(0) \end{bmatrix}^{-1} \quad \text{where } T_s \text{ is the sampling period; let } g(a) = \int s(t) \times s(t-a) dt \text{ with } s(t)$$

being a ternary pseudo-random sequence of length 31. The delay estimate can be expressed as a simple

algebraic solution, $\hat{\tau} = \frac{t_3^T \times W_3 \times h(t_3)}{1_3^T \times W_3 \times h(t_3)}$, where $1_3 = [1 \ 1 \ 1]^T$. This solution can be shown to be optimal

in the sense that the estimate is approaching the theoretical lower limit as the sampling rate grows sufficiently large.

Figure D1.3a shows simulation results for a Gaussian UWB pulse of width 500 MHz, a PRF of 30.875 MHz, and an ADC sampling rate of 494 MHz (or $16 \times$ PRF). The conventional approach denotes choosing the sample time with largest peak, the interpolation approach denotes simple interpolation without the autocorrelation function, and the simplified maximum likelihood denotes the proposed approach. Figure D1.3a shows a decreasing RMS estimation error with increasing SNR, and Figure D1.3b shows the same decreasing error with increasing ADC sampling rate. In both plots, the simplified maximum likelihood outperforms the other two approaches.

²Twice the bandwidth of the transmitted signal.

³Y. Qi and R. Kohno, "Mitigation of sampling-induced errors in delay estimation," *Proceeding of IEEE International Conference on UWB 2005 (ICU2005)*, Zurich, Switzerland, Sept. 2005).




Figure D1.3—Performance comparison between the simplified maximum likelihood estimator and other approaches

D1.2.1 Time-of-arrival estimation in non-line-of-sight (NLOS) conditions

In line-of-sight (LOS) conditions, the dominant peak in the cross-correlation function generated at the receiver corresponds to the first arrival. Since its strength is generally much greater than the subsequent peaks in the profile, it proves easy to isolate. However, in NLOS conditions, the first peak corresponding to the direct path is seldom the strongest, attenuated by transmission through walls and other objects; the strongest peak often corresponds to a reflected path whose travel time is greater than the direct path.

In the latter case, a sequential linear cancellation scheme⁴ is devised for leading edge detection based on the aforementioned simplified maximum likelihood scheme. It can cope with the accuracy degradation when the first arriving signal component is weak compared to a dominant multipath component. This scheme reduces to an iterative algorithm. In each step, the amplitude \hat{A} of the present strongest component in the cross-correlation function is estimated based on a sliding delay τ :

$$\hat{A} = \frac{\hat{g}(\hat{\tau}) \times W_3 \times h(t_3)}{\hat{g}(\hat{\tau}) \times W_3 \times \hat{g}(\hat{\tau})}$$

The autocorrelation samples, scaled to amplitude \hat{A} and time delay of the strongest component, are subtracted from the cross-correlation samples, effectively eliminating this component as in Figure D1.4a. Since only the delay of the first arrival is of interest, components with delays greater than τ are subsequently removed in the following iterations, as shown in Figure D1.4b, until no such components above a certain threshold exist. Guvenc and Sahinoglu^{5,6} and Lee and Scholtz⁷ cover threshold estimation techniques for UWB systems based on correlation.

⁴Y. Qi, H. Kobayashi, and H. Suda, "On time-of-arrival positioning in a multipath environment," *IEEE Trans. on Vehicular Technology*, 2006.

⁵I. Guvenc and Z. Sahinoglu, "Threshold-Based TOA Estimation for Impulse Radio UWB Systems," *IEEE International Conference on Ultra Wideband Systems and Technologies*, pp. 420–425, Sept. 2005.

⁶I. Guvenc and Z. Sahinoglu, "Threshold Selection for UWB TOA Estimation Based on Kurtosis Analysis," *IEEE Communication Letters*, vol. 9, no. 12, pp. 1025–1027, Dec. 2005.

⁷J.-Y. Lee and R.A. Scholtz, "Ranging in Sense Multipath Environment Using an UWB Radio Link," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 9, Dec. 2002.




Figure D1.4—Leading-edge detection in NLOS conditions

D1.3 Asynchronous ranging

As described in the previous subclause, the time of arrival is extracted from the cross-correlation function generated at the receiver. It is used for the computation of the time of flight t_p , defined as the time for propagation of the signal between the transmitter and receiver. The latter is found through an exchange of messages between the two devices in order to estimate range. The number of messages depends on the backbone structure of the network, described in detail in D1.4. With clock synchronization between the devices in the network, a single message suffices in one-way ranging to estimate t_p ; in the absence of such synchronization, more messages are required. The finite crystal tolerance of the clocks is susceptible to drift and, therefore, has an effect on the number of messages required. This subclause considers two schemes for asynchronous ranging.

D1.3.1 Two-way ranging (TWR)

In the absence of clock synchronization between two ranging devices, request device A uses its own clock as a time reference, as depicted in Figure D1.5.




Figure D1.5—Exchange of message in two-way ranging

Device A begins the session by sending a range request message to device B. While device B can measure the absolute time of arrival of the message, lacking synchronization with device A, it does not know the time of departure of the message and, therefore, cannot extract t_p . Rather device B waits a time t_{replyB} , known to both devices, to send a request back to device A. Now device A can measure the round-trip time $t_{roundA} = 2t_p + t_{replyB}$ and extract t_p with respect to its own reference time.

Concern should be given to the finite tolerance of the device crystal reference frequency since frequency error introduces error in the measurement of t_p . In reference to Figure D1.5, the true value of t_p is computed in terms of the transmitted and received times (denoted by subscripts T and R, respectively) at devices A and B.

$$2t_p = \left(\frac{\tau_{BR} - \tau_{AT}}{t_p} \right) + \left(\frac{\tau_{AR} - \tau_{BT}}{t_p} \right) = (\tau_{AR} - \tau_{AT}) + (\tau_{BR} - \tau_{BT})$$

Therefore, the estimated value \hat{t}_p follows as

$$2\hat{t}_p = \left(\frac{\tau_{AR} - \tau_{AT}}{t_p} \right) \times (1 + e_A) + \left(\frac{\tau_{BR} - \tau_{BT}}{t_p} \right) \times (1 + e_B)$$

where e_A and e_B represent the crystal tolerances of the respective devices expressed in parts per million. Substituting for $\tau_{AR} - \tau_{AT} = 2t_p + t_{replyB}$ and $\tau_{BR} - \tau_{BT} = -t_{replyB}$ in the above equation and simplifying gives

$$\hat{t}_p - t_p = \frac{1}{2}(t_{replyB} \times e_A - t_{replyB} \times e_B + 2t_p \times e_A)$$

Note the t_{replyB} is not the turnaround time between the received message from device A and the sent message from device B, but rather includes both the packet duration and this turnaround time. Since the packet duration is on the order of several milliseconds according to the normative subclauses, this duration implies $t_{replyB} \gg t_p$ and, therefore,

$$\hat{t}_p - t_p \approx \frac{1}{2} \times t_{replyB} \times (e_A - e_B)$$

Table D1.1 presents some typical values for $\hat{t}_p - t_p$ according to the other system parameters.

Table D1.1—Typical errors in time-of-flight estimation using TWR

$t_{replyB}/(e_A - e_B)$	2 ppm	20 ppm	40 ppm	80 ppm
100 μs	0.1 ns	1 ns	2 ns	4 ns
5 ms	5 ns	50 ns	100 ns	200 ns

The project authorization request specifies a ranging precision of 1 m; therefore, the estimated \hat{t}_p must lie within 3 ns of the true time of flight given the speed of light. Obviously for the normative packet duration, even with high-quality crystals with tolerance of 2 ppm, the measurement error is greater than the required resolution of the ranging system.

D1.3.2 Symmetric double-sided two-way ranging (SDS-TWR)

In order to compensate for the shortcomings of simple two-way ranging, Hach⁸ proposes an additional message exchange in the ranging session to reduce the effect of the finite crystal tolerances of the devices. Figure D1.6 shows the message exchange.




Figure D1.6—Exchange of message in SDS-TWR

The diagram shows that the round-trip times t_{roundA} and t_{roundB} can be expressed in terms of t_p and the respective t_{replyA} and t_{replyB} as follows:

$$t_{roundA} = 2t_p + t_{replyB}$$

$$t_{roundB} = 2t_p + t_{replyA}$$

Combining the two equations above allows for isolating the true value of t_p as

$$4t_p = t_{roundA} - t_{replyA} + t_{roundB} - t_{replyB}$$

and the estimated \hat{t}_p follows by introducing the finite crystal tolerance e_A and e_B :

$$4\hat{t}_p = (t_{roundA} - t_{replyA}) \times (1 + e_A) + (t_{roundB} - t_{replyB}) \times (1 + e_B)$$

Without loss of generality, replacing t_{replyA} and t_{replyB} in the above equation with

$$t_{replyA} = t_{reply}$$

$$t_{replyB} = t_{reply} + \Delta_{reply}$$

⁸R. Hach, "Symmetric Double Sided Two-Way Ranging," IEEE P802.15 Working Group for Wireless Personal Area Networks (WPAN), Doc. IEEE P.802.15-05-0334-00-004a, June 2005.

reduces it to

$$\hat{t}_p - t_p = \frac{1}{2} \times t_p \times (e_A + e_B) + \frac{1}{4} \times \Delta_{reply} \times (e_A - e_B)$$

Assuming that $t_p \ll \Delta_{reply}$, the above simplifies further to

$$\hat{t}_p - t_p \approx \frac{1}{4} \times \Delta_{reply} \times (e_A - e_B)$$

Table D1.2 shows the typical errors in the SDS-TWR time-of-flight estimation versus frequency tolerance.

Table D1.2—Typical errors in time-of-flight estimation using SDS-TWR

$\Delta_{reply}/(e_A - e_B)$ (μs)	2 ppm (ns)	20 ppm (ns)	40 ppm (ns)	80 ppm (ns)
1	0.0005	0.005	0.01	0.02
10	0.005	0.05	0.1	0.2
100	0.05	0.5	1	2

The extra message in the SDS-TWR accommodates a much smaller error margin even with low-quality crystals of 80 ppm.

D1.4 Location estimation from range data

The ranging capabilities of the devices can be used to estimate their locations through network collaboration. A device wishing to determine its location gathers at least three or four ranges to neighboring devices with known location in a two- or three-dimensional network. The technique used to triangulate the estimated ranges to an estimated location of the device depends largely on the network topology and communication protocols. The subclause considers the two main approaches to simple triangulation.

D1.4.1 Time of arrival

Consider device A initiating a ranging session by sending a request message to device B. Device B can estimate the time of arrival τ_A from the message through one-way ranging. If the two clocks are synchronized to the same time reference, device B can also extract the time of departure τ_D included in the message by device A and hence compute the time of flight $t_p = \tau_D - \tau_A$. In practice, it proves difficult or inefficient to establish and/or maintain synchronization between two mobile devices; therefore, a network lacking a wired backbone must resort to asynchronous ranging described in D1.3.

The time-of-arrival technique for triangulation of ranges applies to the general network lacking synchronization between devices and/or a priori organization. The technique assumes that three (or more) ranges $c \times t_1$, $c \times t_2$, and $c \times t_3$ are gathered from anchor devices $i = 1, 2, \text{ and } 3$, respectively, with known locations (x_i, y_i) , as in Figure D1.7.




Figure D1.7—Time-of-arrival triangulation of ranges to determine location

$$c \times t_1 = \sqrt{(x_1 - x_M)^2 + (y_1 - y_M)^2}$$

$$c \times t_2 = \sqrt{(x_2 - x_M)^2 + (y_2 - y_M)^2}$$

$$c \times t_3 = \sqrt{(x_3 - x_M)^2 + (y_3 - y_M)^2}$$

Solving the set of equations above translates to finding the intersection of the three circles (or spheres in three dimensions), yielding the unknown coordinates of the mobile device (x_M, y_M) .

Time of arrival is appealing due to its application to the general network architecture; however, the associate asynchronous ranging requires two or more messages per ranging session. This requirement may potentially increase the network traffic considerably.

D1.4.2 Time difference of arrival

A preinstalled network can be configured so that the mobile devices within a deployment area can maintain connectivity with at least three anchor devices positioned at known locations and connected through wire to maintain clock synchronization. Clock synchronization enables one-way ranging in conjunction with the time-difference-of-arrival technique, as opposed to time of arrival. In order to carry out a range request, the devices in the network classified as stationary anchor devices and mobile devices operate in one of the two following modes: Mode 1 or Mode 2.

D1.4.2.1 Mode 1

The synchronized anchor nodes jointly send range requests to the mobile node at the same time instant; therefore, the number of messages equals the number of anchors. Although the mobile node lacks synchronization with the three anchors indexed through $i = 1, 2, 3$, it can still measure the time difference of arrivals $t_{32} = \tau_3 - \tau_2$ and $t_{31} = \tau_3 - \tau_1$ within its own time reference, where τ_i denotes the arrival time of

the message from anchor i at the mobile M . Figure D1.8 depicts the locations of anchors i and respective locations (x_i, y_i) . Solving the following equations translates to finding the intersection of two hyperbolae and yields the unknown location (x_M, y_M) of the mobile device. The burden of the calculation lies on the mobile device, which may have limited resources with respect to the anchor devices due to its smaller dimensions to preserve battery life.




Figure D1.8—Time-difference-of-arrival triangulation of ranges to determine location

D1.4.2.2 Mode 2

This mode operates in a similar fashion to Mode 1; however, the mobile device initiates the range request by sending a single message received by all the neighboring anchor nodes. The request message arrives at the anchor nodes at different times of arrival measured individually. These values are either distributed throughout the wired backbone or sent to a central controller to compute the time difference of arrivals t_{32} and t_{31} and in turn the location of the mobile device. This mode offers the advantage of reducing the wireless network traffic to just one message per request and also does not place the burden of computation on the mobile node with potentially limited resources.

D1.5 Network location algorithms

In the basic approach to location estimation described in D1.4, a mobile device gathers range measurements from a number of anchor devices with known locations and triangulates these ranges to a single point (or area) through simple geometrical relationships. This approach assumes that the mobile device receives at least three or four estimates in a two- or three-dimensional coordinate system.

Interest in dense sensor networks due to falling price and reduced size has motivated research in network location algorithms in recent years, where the number of mobile devices vastly outnumbers the number of fixed stations. Consider a rapidly deployable network whose nodes are scattered about an area of interest and self-organize in an ad hoc fashion to determine their locations through simple messaging and ranging. Most of the sensors in such networks lack connectivity to fixed stations. Rather, the high sensor-to-sensor connectivity allows them to infer their locations from the locations of other sensors that do have connectivity to the fixed stations. This class of algorithms to process large amounts of range data is commonly known as

network location algorithms. These algorithms can render good location accuracy despite significant errors in range estimates between sensors.

This subclause provides a survey of the benchmark network location algorithms developed in recent years and divides them into three main classes: ad hoc algorithms, centralized algorithms, and convex optimization algorithms.

D1.5.1 Ad hoc algorithms

Network location algorithms arose from the need to locate nodes in ad hoc networks characterized by high mobility and dynamic architecture, with nodes joining and leaving the network at random times. Here greater importance is placed on continuously updating location in a distributed manner rather than furnishing precision. Niculescu and Nath⁹ contributed to this pioneering effort. Their work describes several algorithms to this end based on the range measurements available to the nodes. The strength of these algorithms lies in their simplicity and in turn their applications to ad hoc networks. In general they foster reduced complexity for larger networks.




Figure D1.9—Performance comparison between the simplified maximum likelihood estimator and other approaches

The first algorithm known as *DV-Hop* generates location based on mere connectivity quantified as the minimum number of hops between two nodes, or the minimum-hop distance. Hence the nodes in the network lack any ranging capabilities. The network is categorized into two types of nodes: *anchor nodes* whose locations are known and *sensor nodes* whose locations are unknown. Rather than estimate range between each other, the nodes simply determine the number of hops between each other. Given the ground-truth ranges between the anchors through simple messaging of their known locations through the ad hoc network, each anchor node i computes a factor f_i , which is simply the sum of the ground-truth ranges between all the anchors in the network divided by the sum of the respective minimum hops.

Consider the network in Figure D1.9a. Anchor L1 computes factor f_1 by summing the ground-truth ranges from L2 (40 m) and L3 (100 m), divided by the sum of the minimum hops 6 and 2, or $f_1 = \frac{40 + 100}{2 + 6}$ m.

Anchor L2 and L3 compute their factors f_2 and f_3 in the same manner as $f_2 = \frac{40 + 75}{2 + 5}$ m and $f_3 = \frac{75 + 100}{5 + 6}$ m, respectively, and then distribute them to all the sensors in the network. Ultimately the sensor node A, pictured in Figure D1.9b, finds its location from its three closest neighboring anchor nodes; however, rather

⁹D. Niculescu and B. Nath, “Ad Hoc Positioning System (APS),” *IEEE Conf. on Global Communications*, pp. 2926–2931, Nov. 2001.

than through the triangulation of three measured ranges R_1 , R_2 , and R_3 to them, respectively, it finds its location through DV-ranges $R_1 = 3 \times f_1$, $R_2 = 2 \times f_2$, and $R_3 = 3 \times f_3$ given through the minimum hops and the factors.

In the same paper, the authors present an alternative algorithm known as *DV-Distance* as an adaptation of the ad hoc positioning system to accommodate nodes with ranging technology (such as those in IEEE Std 802.15.4a-2007) for enhanced precision without compromising simplicity. Rather than scaling the ground-truth distance by the minimum-hop distance between two anchors in computing the factors of each anchor node, it is scaled by the sum of the measured distances on the multihop path between two anchors. It directly follows in the triangulation step that the distances between the unknown sensor and the three closest anchor nodes are scaled by the measured distances on the multihop path between the two.

Saverese et al.¹⁰ present another algorithm designed for ad hoc networks, which also incorporates range estimates. Knowing its coordinates (x_1, y_1) , anchor L1 orients a local coordinate system pictured in Figure D1.10a in the direction of an arbitrary neighbor, e.g., L2. Given the measured range r_{12} , the coordinates $(x_2, y_2) = (x_1 + r_{12}, y_1)$ of sensor 2 are easily computed in the same coordinate system. Given further the range estimates r_{13} between anchor L1 and sensor L3 and r_{23} between sensors L2 and L3, the coordinates of $(x_3, y_3) = (x_2 + dx, y_2 + dy)$ are computed through

$$dx = \frac{r_{12}^2 + r_{13}^2 + r_{23}^2}{2 \times r_{12}}$$

$$dy = \sqrt{r_{12}^2 - dx^2}$$

In this manner, anchor L1 propagates the coordinates of its successive neighbors throughout the network as an initialization step. The other anchor nodes independently do likewise.




Figure D1.10—Ad hoc network location algorithm proposed by Savarese et al.

In propagation through the individual neighbors of the n_A anchor nodes in the network, node 2 receives n_A estimates of its coordinates. It can then refine its estimate (\hat{x}_2, \hat{y}_2) as the equally weighted average of all the accumulated estimates. In sequence, sensor 3 refines its estimate through the propagating of the coordinates from anchor L1 and sensor 2 as in the initialization step; however, it replaces (x_2, y_2) with (\hat{x}_2, \hat{y}_2) . Propagation then continues and eventually converges to an acceptable solution after a number of iterations, although the authors provide no proof. In the network simulations, the algorithm furnishes a location error of about 5% even with range errors up to 50%.

¹⁰C. Saverese, J. M. Rabaey, and J. Beutel, “Location in Distributed Ad-Hoc Networks,” *IEEE Conf. on Acoustics, Speech, and Signal Processing*, pp. 2037–2040, May 2001.

D1.5.2 Centralized algorithms

Centralized algorithms gather all the data available from the network to process it collectively. As expected, these algorithms in general render better results than the ad hoc algorithms, which consider only local data and process it independently of the data available to other parts of the network. The drawback of the former involves designating a central controller, a condition which may be unsuitable for some applications. Even so, the dynamic links of nodes in motion may require rapid updating; alternatively, relaying information across a large network sanctions the centralized processing of obsolete data at the controller, and this condition limits scalability. It is worth noting that most of these algorithms have distributed versions, which, however, compromise the quality of the results. This subclause considers the architecture of two centralized algorithms commonly referred to in literature.

Many centralized algorithms estimate the locations of the unknown nodes by minimizing an objective function such as in Savvides et al.¹¹ A popular function, as used here, is the least-squared sum of the residuals; the residual is defined as the difference between a measured range and the range estimated through the algorithm. The minimization is performed through Kalman filtering, which ultimately finds only a local minimum. In fact, an initialization step is required to estimate the positions of the nodes from the measured ranges and the anchor nodes. Consider sensor node C in Figure D1.11 as an example. The simplistic initialization step finds a bounding rectangle for it as the intersection of the individual bounding squares of each anchor node (here A and B) to the sensor: the length of half the side of the square is the sum of the measured distances on the multihop between the anchor and the sensor.




Figure D1.11—Autocorrelation and cross-correlation functions for the maximum likelihood estimator

¹¹A. Savvides, H. Park, and M. B. Srivastava, “The Bits and Flops of the N-hop Multilateration Primitive for Node Localization Problems,” *ACM Conf. on Wireless Sensor Networks and Applications*, pp. 112–121, Sept. 2002.

Shang et al.¹² have designed another centralized algorithm to minimize the least-squared sum of the residuals, but use a more powerful technique called *multidimensional scaling* to find the global maximum. As in Savvides, the initial step consists of computing the minimum-hop distances between all nodes in the network, but here they are stored in a range matrix R having the following structure:

$$R = \begin{bmatrix} r_{11} & r_{12} & r_{13} & & \\ r_{21} & r_{22} & r_{23} & \cdots & \\ r_{31} & r_{32} & r_{33} & & \end{bmatrix}$$

The minimum-hop distance between two nodes is just the sum of the estimated ranges between them.¹³ Performing the Singular Value Decomposition of R^2 yields a matrix of eigenvectors in V and respective squared eigenvalues on the diagonal matrix A . The relative coordinates of the nodes are subsequently reconstructed as $X = V\sqrt{A}$. These coordinates are then scaled, translated, and rotated to fit the anchor nodes.

D1.5.3 Convex optimization algorithms

The most powerful class of network location algorithms places the geometrical constraints of the physical world on the nodes while minimizing an objective function similar to the one in D1.5.2. The advantage of the technique over the ones in D1.5.2 lies in its ability to achieve the global maximum independent of any initial estimates. Doherty et al.¹⁴ first proposed the technique in defining the following optimization problem with quadratic constraints to minimize a linear objective function:

$$\begin{aligned} \min \quad & \sum |\alpha_{ij}| \\ \text{st} \quad & \|\bar{x}_i - \bar{x}_j\|_2 = d_{ij}, \quad \forall (i, j) \in N \\ & \|\bar{x}_i - \bar{x}_j\|_2 \geq R \quad \forall (i, j) \notin \bar{N} \end{aligned}$$

where

- $d_{ij} = \hat{d}_{ij} + \alpha_{ij}$ is the estimated range resultant of the algorithm
- \hat{d}_{ij} is the measured range
- α_{ij} is the residual between the two
- N is the set containing the index of neighboring nodes
- \bar{N} is the set of non-neighboring nodes
- R is the radio range, i.e., the range at which a node loses connectivity with another node in the network

Since many of the constraints are nonconvex, the paper relaxes the problem simply by removing these constraints. The corresponding solution not only offers less precision, but also forces the sensors to lie within the convex hull of the anchors. An alternative approach proposed by Biswas et al.¹⁵ relaxes the

¹²Y. Shang, W. Rumi, Y. Zhang, and M. P. J. Fromherz, "Localization from Mere Connectivity," *ACM Conf. on Mobile Ad Hoc Networking and Computing*, pp. 201–212, June 2003.

¹³The estimated ranges are actually normalized to fit the framework of the multidimensional scaling.

¹⁴L. Doherty, K. S. J. Pister, and L. El Ghaoui, "Convex Position Estimation in Wireless Sensor Networks," *IEEE Conf. on Information Theory and Communications*, pp. 1655–1663, April 2001.

¹⁵P. Biswas and Y. Ye, "Semidefinite Programming for Ad Hoc Wireless Sensor Network Localization," *IEEE Conf. on Information Processing in Sensor Networks*, pp. 46–54, April 2004.

problem above to a semi-definite program that yields an average and standard variation for the positions of the unknown nodes.

Rather than relaxing the problem from the original, Gentile¹⁶ directly applies linear constraints given through the triangle inequality:

$$\begin{aligned} & \min \sum |\alpha_{ij}| \\ & \text{st } \left. \begin{aligned} d_{ij} + d_{jk} &\geq d_{ik} \\ d_{ij} + d_{ik} &\geq d_{jk} \\ d_{jk} + d_{ik} &\geq d_{ij} \end{aligned} \right\} \begin{cases} \forall (i, j) \in N \\ \forall (j, k) \in N \\ \forall (i, k) \in N \end{cases} \end{aligned}$$

The original convex constraints necessitate no relaxation and hence render a much tighter solution than the other two approaches mentioned. A distributed version of the algorithm yields the same results as the centralized version with no compromise in optimality.¹⁷

In order to substantiate the effectiveness of the network location algorithms, Table D1.3 presents the results from Gentile.¹⁶ The simulation platform consists of a network with 50 sensor nodes uniformly distributed in a unit area. The number of anchor nodes varies as a parameter {3, 5, 7} in the table and R as {0.20, 0.25, 0.30}. The noise parameter controls the percentage of Gaussian-distributed noise perturbing the measured radio range from the ground-truth range and varies as {0.0, 0.1, 0.2, 0.3}. The table shows that despite the measured range errors up to 30%, the location errors yielded by the algorithm can lie on the order of only 5%.

Table D1.3—Location results according to varying network parameters

noise	R=0.20			R=0.25			R=0.30		
	3	5	7	3	5	7	3	5	7
0.0	0.043	0.042	0.041	0.007	0.006	0.006	< 1e-6	< 1e-6	< 1e-6
0.1	0.075	0.064	0.064	0.053	0.044	0.027	0.045	0.036	0.025
0.2	0.085	0.080	0.068	0.077	0.065	0.049	0.057	0.057	0.046
0.3	0.107	0.088	0.087	0.095	0.083	0.077	0.077	0.074	0.046

D1.5.4 Location estimation using multipath delays

For location estimation in a multipath environment, the conventional approach is based on leading-edge detection, where the time-of-arrival estimate of the first arriving signal is taken as the distance between the transmitter and the receiver of interest up to a constant and the delays of other multipath signals are completely ignored. This approach works well when the first arrival signal is sufficiently strong and via a LOS propagation path. However, in a typical UWB channel, the first arrival signal is usually weak, e.g., 6 dB lower than a dominant multipath component, and can be subject to NLOS propagation. Hence the conventional approach can cause severe degradation of the positioning accuracy. To address this problem,

¹⁶C. Gentile, "Sensor Location through Linear Programming with Triangle Inequality Constraints," *IEEE Conf. on Communications*, pp. 3192–3196, May 2005.

¹⁷C. Gentile, "Distributed Sensor Location through Linear Programming with Triangle Inequality Constraints," *IEEE Conf. on Communications*, June 2006.

one method is to utilize time-of-arrival estimates of multipath components in addition to the first arriving signals for location estimation. Although subject to NLOS propagation, the second and later arriving signals should also carry information regarding the position of interest. Hence the method incorporating the multipath delays can improve the positioning accuracy under certain conditions.

For simplicity, consider that a mobile node is synchronized with B anchor nodes, whose locations $\{(x_b, y_b), b = 1, 2, \dots, B\}$ are known. Each anchor node receives radio signals transmitted from the mobile node via multipath propagation. A received signal at the b -th anchor node is expressed as shown in Equation (D1.1).

$$r_b(t) = \sum_{i=1}^{N_b} A_{bi} \times s(t - \tau_{bi}) + n_b(t) \quad (\text{D1.1})$$

where τ_{bi} is the delay of the i -th multipath component, given by $\tau_{bi} = \frac{1}{c} \sqrt{(x - x_b)^2 + (y - y_b)^2} + l_{bi}$

which consists of the LOS delay corresponding to the distance between the mobile node and the anchor node, and the NLOS induced path length error l_{bi} . The quantity $l_b = (l_{b1}, l_{b2}, \dots, l_{bN_b})^T$ is usually modeled as a multivariate random variable, which can be determined by field experiments¹⁸ or theoretical models. Noise $n_b(t)$'s are independent white Gaussian processes, and A_{bi} is the signal amplitude. Estimation of the multipath delays yields

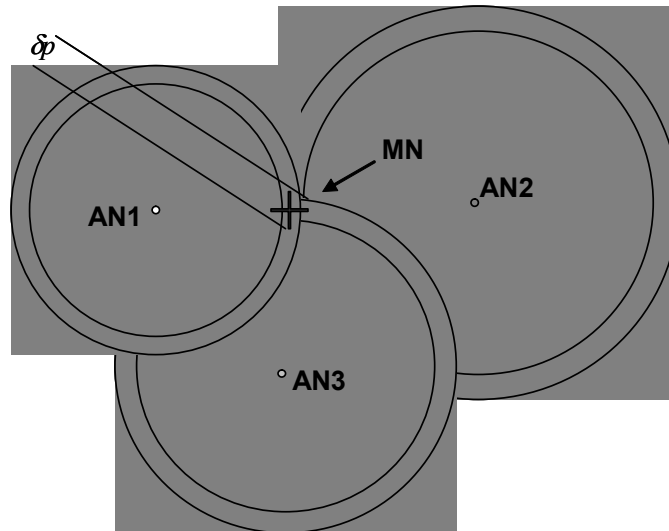
$$\hat{\tau}_{bi} = \tau_{bi} + \xi_{bi}$$

where $\xi_b = (\xi_{b1}, \xi_{b2}, \dots, \xi_{bN_b})^T$ is a multivariate Gaussian random variable with zero mean and an explicit covariance matrix F_{ξ} . Based on Equation (D1.1) and the probability density functions of l_b and ξ_b , location estimation of (x, y) can be formulated as the maximum a priori estimation. It is shown that the positioning accuracy enhancement depends on two principal factors, the strength of multipath components and the variance of the NLOS induced errors. In certain situations, significant accuracy improvement, e.g., above 50%, can be obtained. The limitation of this approach is that its computation complexity is higher than the conventional approach. The exact formula of accuracy improvement and detailed discussion on this approach can be found in Qi et al.¹⁸

D1.5.5 Reduced dimension approach to bad geometric dilution of precision (GDOP) problem

The bad GDOP problem concerns severe degradation of the positioning accuracy when anchor nodes and the mobile node to be located are lined up. The GDOP index corresponding to such an anchor-node–mobile-node geometric layout is infinitely large. To see this clearly, first consider an ordinary anchor-node–mobile-node layout as illustrated in Figure D1.12, where the positions of the anchor nodes and the mobile node are distributed evenly on a plane. An anchor node is located at the center of a circle. The radius and the width of the circle represent a time-of-arrival estimate and the corresponding estimation error, respectively. It is seen that the positioning error, denoted by symbol δp is comparable to the time-of-arrival estimate error. In contrast, in the bad GDOP case as illustrated in Figure D1.13, where the mobile node and the anchor nodes are almost lined up, although the estimation errors of the three time-of-arrival estimates are same as in the previous case, the positioning error is considerably increased, much larger than the scale of the time-of-arrival errors. The problem is that the location estimation is set to be a two-dimensional problem, yet the unfavorable mobile-node–anchor-node layout is essentially a one-dimensional configuration and thus cannot provide accurate two-dimensional location estimation.

¹⁸Y. Qi, H. Kobayashi, and H. Suda, "On time-of-arrival positioning in a multipath environment," *IEEE Trans. on Vehicular Technology*, 2006.



MN = mobile node; AN = anchor node

Figure D1.12—An example of the good GDOP case




Figure D1.13—An example of the bad GDOP case

A so-called reduced dimension method is proposed to solve this problem. Consider the specific Euclidean coordinate system where the y -axis is set to be along the line where anchor nodes and the mobile node are (almost) lined up, as illustrated in Figure D1.14. By decomposing the positioning accuracy (in terms of MSE) into two orthogonal components along the x and y axes, it is seen that the y -axis defines a “good” dimension in the sense that location estimation in this dimension is a regular one-dimensional estimation problem, and the corresponding position error denoted by δp_y is “normal,” i.e., comparable to the time-of-arrival estimation errors; and the x -axis renders a “bad” dimension in which the positioning error δp_x is large, which is the main cause of the large aggregated positioning error. In fact, however, δp_x is not difficult to minimize once the orientation of y -axis is known. Based on this observation, the basic idea of the reduced dimension approach is to separate the “good” dimension from the “bad” dimension and then to perform location estimation separately in each dimension. In practice, a strict mobile-node–anchor-node line-up rarely happens. Hence, the “good” dimension needs to be approximated. Simulation results are presented in Qi et al.¹⁹




Figure D1.14—Decomposition of the positioning error into the “good” and “bad” dimensions

¹⁹See Footnote 17.

Annex E

(informative)

Coexistence with other IEEE standards and proposed standards

E.1 Introduction

Insert the following new paragraph at the end of E.1:

This annex also considers IEEE 802.15.4a devices. For UWB devices specifically, additional consideration is given to certain non-IEEE standards.

E.2 Standards and proposed standards characterized for coexistence

Change the first paragraph of E.2 as shown:

This clause enumerates IEEE-compliant devices that are characterized and the devices that are not characterized for operation in proximity to IEEE 802.15.4 devices or IEEE 802.15.4a devices.

Insert the following new paragraphs at the end of E.2:

IEEE 802.15.4a CSS PHYs for the 2400 MHz ISM band are specified for operation in 14 channels. Channel 0 through channel 13 reside in frequencies from 2412–2484 MHz bands and, therefore, may interact with other IEEE-compliant devices operating in those frequencies.

Standards and proposed standards characterized in this annex for coexistence are as follows:

- IEEE Std 802.11-2007 (ERP)
- IEEE Std 802.11-2007 (2400 MHz DSSS)
- IEEE Std 802.15.1TM-2005 [2400 MHz frequency hopping spread spectrum (FHSS)]
- IEEE Std 802.15.3TM-2003 (2400 MHz DSSS)
- IEEE Std 802.15.4-2006 (2400 MHz DSSS)
- IEEE Std 802.15.4a-2007 (2400 MHz CSS)

Standards not characterized in this annex for coexistence are as follows:

- IEEE Std 802.11-2007, frequency hopping (FH) (2400 MHz FHSS)
- IEEE Std 802.11-2007, infrared (IR) [333 GHz amplitude modulation (AM)]
- IEEE Std 802.16TM-2004 (2400 MHz OFDM)
- IEEE Std 802.11-2007 (5.2 GHz DSSS)

IEEE 802.15.4a UWB PHYs for the 250–750 MHz band reside in frequencies that may interact with other IEEE standards in development. UWB PHYs for the 3244–4742 MHz and 5944–10 234 MHz bands may interact with both IEEE-compliant devices and non-IEEE-compliant devices.

Standards and proposed standards characterized in this annex for coexistence are as follows:

- IEEE Std 802.16-2004
- IEEE P802.22
- ECMA 368²⁰

E.3 General coexistence issues

Insert the following new paragraphs at the end of E.3:

In addition, IEEE Std 802.15.4a-2007 provides several mechanisms that enhance coexistence of UWB PHYs with other wireless devices operating in the same spectrum. This subclause describes the mechanisms that are defined in this standard:

- UWB modulation with extremely low PSD
- Low duty cycle
- Low transmit power
- Dynamic channel selection
- Coordinated piconet capabilities

These mechanisms are each described briefly in the following subclauses.

E.3.2 Modulation

Insert after E.3.2.2 the following new subclause (E.3.2.3):

E.3.2.3 Direct sequence UWB modulation

The UWB PHY specified for IEEE Std 802.15.4a-2007 uses a UWB direct sequence modulation. This power-efficient modulation method achieves low requirements for signal-to-noise ratio (SNR) and signal-to-interference ratio (SIR) through the use of a signal bandwidth that is significantly larger than the symbol rate. A defining feature of systems that use UWB modulation is that they are less likely to cause interference in other devices due to their reduced PSD. In fact, even the least restrictive regulations for UWB devices today require the emission PSD levels to be at or below the levels allowed for unintentional emissions by other electrical or electronic devices. In some cases, the UWB PSD limits are as much as 35 dB below these same unintentional emissions limits. For the same reason, UWB devices have some degree of immunity from interfering emitters, making them a good choice for environments where coexistence may be an issue.

E.3.4 Low duty cycle

Insert the following new paragraph at the end of E.3.4:

An important contribution of IEEE Std 802.15.4a-2007 is the definition of a new UWB PHY with higher optional bits rates. In the UWB bands, the data rates have been increased to a nominal mandatory rate of 850 kb/s. Although not designed to provide continuous higher throughputs, the UWB PHY also provides for optional data rates as high as 27 Mb/s. These rates are not designed to support high-rate applications such as video transport, but instead are provided to allow devices in close proximity to shorten their transmission duty cycle by as much as a factor of 32 relative to the mandatory rate to further reduce the likelihood that these devices will interfere with or be subject to interference by other devices when conditions allow.

²⁰ECMA 368, High Rate Ultra Wideband PHY and MAC Standard (December 2005) (www.ecma-international.org).


Insert after E.3.4 the following new subclause (E.3.4.1):

E.3.4.1 Low-duty-cycle considerations for UWB PHYs

Low-duty-cycle piconet scenarios are used to model the following situations:

- IEEE 802.15.4a devices are deployed in high density in a limited area, e.g., hot-spot deployment scenarios.
- Some UWB victim systems cover a much larger area than the coverage of a typical IEEE 802.15.4a piconet, are located well above the local cluster (e.g., IEEE 802.16, radio astronomy service, and satellite service), or are closely located with a piconet coordinator (e.g., devices placed at the same desk or even within the same computer).

In such cases, transmissions from every device in the piconet can affect the victim receiver. For reasons of less complexity, lower power consumption, as well as physical limitations, it is difficult for simple IEEE 802.15.4a devices to detect victim systems reliably. The aggregate interference from the piconet increases with piconet members. Given 1% average device duty cycle and pure ALOHA protocol, the aggregate interference is 17.6% from a piconet with 18 members. See Figure E.1a. Besides, the channel idle periods are randomly segmented into small pieces. Therefore, it is hard to use the channel effectively. Analyzing the interference in the channel is similar to the collision analysis of a pure ALOHA system.



LDC = low duty cycle

Figure E.1a—Aggregate normalized interference

The maximal interference level to such kinds of victim systems can be limited by controlling the duty cycle of the piconet through general active/inactive periods (see Figure E.1b). The traffic can occur only in the active period. Victim systems are free of interference in the inactive period. The distribution of active/inactive periods is controlled by the piconet coordinator. This can be implemented by a clock in the application layer. The piconet coordinator defines global time of the piconet and duration of the active period. When a device joins a piconet, it synchronizes its clock with that of coordinator.




Figure E.1b—Generalized active/inactive periods

The interference level is restricted by the ratio of active period to the total period. The possible packet collision in the active period can be mitigated as follows:

- Adopt CSMA-CA mechanism.
- Adopt channel-dependent ALOHA: The channel-dependent ALOHA is used to set transmission probability related with the channel quality, which can be obtained through listening to a beacon from the coordinator by means of LQI and receiver ED. The function to map channel quality to transmission probability is defined at application layer. A simple way is to set a threshold and only enable transmission when the channel quality is above the threshold.
- Limit the number of piconet members through association.
- Use traffic shaping, e.g., a combination of short packet to large packet.

Considering the applications for which the UWB PHY is designed, in application scenarios where a greater number of nodes can be expected, duty cycle (aggregate and individual) can be expected to be orders of magnitude less than the 1% used above. Consider, for example, a sensor application where low-cost sensor nodes are deployed in large number (typically indoors). An individual node may be “awake” only milliseconds per hour. In such scenarios, the aggregate duty cycle would be under the control of the higher layer protocols and very low compared to the 1% used in the above analysis. This observation has two important implications:

- ALOHA is well suited to this application where probability of collision is small and controllable; therefore, the complexity advantage is a good trade-off.
- There is low impact on coexistence due to a large number of IEEE 802.15.4a nodes as the aggregate duty cycle remains very low.

E.3.5 Low transmit power

Insert after E.3.5.3 the following new subclause (E.3.5.4):

E.3.5.4 UWB PHYs

The UWB PHY defined by IEEE Std 802.15.4a-2007 operates under strict regulations for unlicensed UWB devices worldwide. The least restrictive regulations for UWB are available under the Federal Communications Commission (FCC) rules, US 47 CFR Part 15, subpart F. Under these rules, the highest allowable limits for UWB emissions are based on an equivalent emission PSD of -41.3 dBm/MHz. Other future UWB regulations in other regions will likely be at this same level or even lower. Under these limits, the allowable transmit power for a 500 MHz bandwidth UWB device would be less than -14 dBm, or about 37 μ W transmit power. This transmit power level is at or below the limits for unintentional emissions from

other electrical or electronic devices, as well as less than the out-of-band emission limits for other unlicensed devices operating in designated bands such as the 2.4 GHz ISM or 5 GHz UNII bands. Additionally, since this transmit power is spread over at least 500 MHz of bandwidth, the highest power in the operating bandwidth of a typical narrowband 20 MHz victim system is less than -28 dBm, or about 1.5 μ W of transmit power per 20 MHz. These very low power levels emitted into the operating band of any potential victim system will reduce the likelihood that these devices might interfere with other systems.

E.3.6 Channel alignment

Insert the following new paragraphs at the end of E.3.6:

The alignment between IEEE 802.11 HR/DSSS (nonoverlapping sets) and IEEE 802.15.4a CSS channels (overlapping sets) is shown in Figure E.1c. There are 14 IEEE 802.15.4a CSS channels ($n = 0, 2, \dots, 13$). Operating an IEEE 802.15.4a network on one of these channels will minimize interference between systems.

When performing dynamic channel selection, either at network initialization or in response to an outage, an IEEE 802.15.4a CSS device will scan a set of channels specified by the ChannelList parameter. For IEEE 802.15.4a networks that are installed in areas known to have high IEEE 802.11 HR/DSSS activity, the ChannelList parameter can be defined from the above set in order to enhance the coexistence of the networks.




Figure E.1c—IEEE 802.15.4a CSS channel selection

E.3.7 Dynamic channel selection

Insert the following new paragraph at the end of E.3.7:

When performing dynamic channel selection, either at network initialization or in response to an outage, a IEEE 802.15.4a UWB device will scan a set of channels specified by the ChannelList parameter. For IEEE 802.15.4a UWB networks that are installed in areas known to have spectrum restrictions, the ChannelList parameter can be defined as the above sets in order to enhance the coexistence of the networks.

Change the title of E.4 as shown:

E.4 2400 MHz band coexistence performance (except for CSS PHYs)

E.5 800/900 MHz bands coexistence performance

Insert after E.5.5.6.3 the following new subclauses (E.6 through E.7.11) (instructions for renumbering the existing E.6 are given after this new text):

E.6 2400 MHz band coexistence performance for CSS PHYs

Subclauses E.3.2 and E.3.4 also describe the assumptions made for individual standards and quantify their predicted performance when coexisting with IEEE 802.15.4a CSS devices.

E.6.1 Assumptions for coexistence performance

E.6.1.1 Receiver sensitivity

The receiver sensitivity assumed is the reference sensitivity specified in each standard as follows:

- -76 dBm for IEEE 802.11 HR/DSSS 11 Mb/s CCK
- -82 dBm for IEEE 802.11 ERP 6 Mb/s OFDM
- -74 dBm for IEEE 802.11 ERP 24 Mb/s OFDM
- -65 dBm for IEEE 802.11 ERP 54 Mb/s OFDM
- -70 dBm for IEEE 802.15.1 devices
- -75 dBm for IEEE P802.15.3 22 Mb/s DQPSK
- -85 dBm for IEEE 802.15.4 devices
- -85 dBm for IEEE 802.15.4a 1 Mb/s CSS

E.6.1.2 Transmit power

The transmit power for each coexisting standard has been specified as follows:

- 14 dBm for IEEE Std 802.11 HR/DSSS
- 0 dBm for IEEE Std 802.15.1-2005
- 8 dBm for IEEE Std 802.15.3-2003
- 0 dBm for IEEE Std 802.15.4-2006
- 0 dBm for IEEE 802.15.4a CSS

E.6.1.3 Bit error rate (BER) calculations

BER for IEEE 802.15.4a CSS

$$[(M-2) \times Q(\sqrt{SNR_0 \times \log_2(M)}) + Q(\sqrt{SN(R_0 \times 2 \log_2(M))})] / 2$$

where

SNR₀ = SNR × 14 × 1.6667, M = 8 for 1 Mb/s

SNR₀ = SNR × 14 × 1.6667 × 4, M = 64 for 250 kb/s

BER for IEEE 802.11 ERP

For the three IEEE 802.11 ERP data rates to be considered, the following are assumed:

- 1) OFDM, 6 Mb/s: M-PSK, M = 2, 5.7 dB coding gain

$$BER_{11g}(M = 2) = Q\left(\sqrt{2 \times \frac{E_b}{N_0} \times 10^{\frac{5.7}{10}}}\right)$$

- 2) OFDM, 24 Mb/s: M-ary QAM, M = 16, coding gain Cg = 5.7 dB
- 3) OFDM, 54 Mb/s: M-ary QAM, M = 64, coding gain Cg = 3.8 dB

For 2) and 3) equations for M-ary QAM can be applied.

$$BER_{11g}(M > 2, Cg) = 1 - \left[1 - 2 \left(\left(1 - \frac{1}{\sqrt{M}} \right) \times Q \left(\sqrt{\frac{3}{M-1} \times \frac{\log_2(M) \cdot E_b}{N_0} \times 10^{\frac{Cg}{10}}} \right) \right) \right]^2 \cdot \frac{1}{\log_2(M)}$$

The relationship between E_b/N_0 and SNR is assumed to be computable from the subcarrier spacing $F_s = 0.3125$ MHz and the OFDM symbol rate $R_s = 0.25$ Msymbol/s as follows:

$$SNR = \frac{E_b}{N_0} \times \frac{F_s}{R_s}$$

E.6.1.4 Packet error rate (PER)

- IEEE 802.11 HR/DSSS average frame length: 1500 octets
- IEEE 802.11 HR/DSSS average duty cycle: 50%
- IEEE 802.11 ERP average frame length: 1500 octets
- IEEE 802.11 ERP average duty cycle: 50%
- IEEE 802.15.1 average frame length: 2871 bits
- IEEE 802.15.1 average duty cycle: 50%
- IEEE 802.15.3 average frame length: 1024 octets
- IEEE 802.15.3 average duty cycle: 50%
- IEEE 802.15.4 average frame length: 22 octets
- IEEE 802.15.4 normal duty cycle: 1%
- IEEE 802.15.4 rare (aggregated) duty cycle: 10%
- IEEE 802.15.4a CSS average frame length: 32 octets
- IEEE 802.15.4a CSS normal duty cycle: 0.25%, 1%
- IEEE 802.15.4a CSS rare (aggregated) duty cycle: 2.5%, 10%

E.6.1.5 BER model for IEEE Std 802.15.4a-2007

Figure E.16 illustrates also the relationship between BER and SNR for IEEE 802.11 HR/DSSS, IEEE 802.15.3 base rate, IEEE 802.15.1, IEEE 802.15.4, and IEEE 802.15.4a CSS PHYs.




Figure E.16—BER results of IEEE 802.11 HR/DSSS, IEEE 802.15.1, IEEE 802.15.3, IEEE 802.15.4 (2400 MHz) and IEEE 802.15.4a CSS PHYs

E.6.2 Coexistence simulation results

The shapes of the assumed transmit spectra and receive filter shapes are defined in Table E.5.

Table E.5—Transmit spectra and receiving filter shapes

IEEE 802	Transmit		Receive	
	Frequency offset (MHz)	Attenuation (dB)	Frequency offset (MHz)	Attenuation (dB)
15.1	0	0	0	0
	0.25	0	0.25	0
	0.75	38	0.75	38
	1	40	1	40
	1.5	55	1.5	55

Table E.5—Transmit spectra and receiving filter shapes (continued)

IEEE 802	Transmit		Receive	
	Frequency offset (MHz)	Attenuation (dB)	Frequency offset (MHz)	Attenuation (dB)
11 HR/DSSS	0	0	0	0
	4	0	4	0
	6	10	6	10
	9	30	9	30
	15	50	15	50
	20	55	20	55
11 ERP	0	0	0	0
	5	0	5	0
	8	4	8	4
	9	10	9	10
	10	25	10	25
	15	40	15	40
	40	43	40	43
15.3	0	0	0	0
	8	0	8	0
	8	30	8	30
	15	30	15	30
	15	40	15	40
	22	50	22	50
15.4	0	0	0	0
	0.5	0	0.5	0
	1	10	1	10
	1.5	20	1.5	20
	2	25	2	25
	2.5	30	2.5	30
	3	31	3	31
	3.5	33	3.5	33
	4	34	4	34
	5	40	5	40
	6	55	6	55

Table E.5—Transmit spectra and receiving filter shapes (continued)

IEEE 802	Transmit		Receive	
	Frequency offset (MHz)	Attenuation (dB)	Frequency offset (MHz)	Attenuation (dB)
15.4a CSS	0	0	0	0
	6	0	6	0
	12	32	12	32
	15	55	15	55

E.6.3 Low-duty-cycle assumption

In general, IEEE 802.15.4 and IEEE 802.15.4a devices address low-duty-cycle applications. The assumption of 1% duty cycle for IEEE 802.15.4 devices was introduced in E.2.4. Under the assumption that IEEE 802.15.4a devices are battery-powered and have a lifetime of at least one year, the 1% assumption can be hardened by taking into account state-of-the-art numbers: A typical AA battery has a capacity of 1.8 Ah. A typical IEEE 802.15.4 device operating at 2.4 GHz has a transmit current of 30 mA. If the device only transmits during its entire lifetime, the result would be $30/1800 = 60$ h of operation. Over a lifetime of one year (365×24 h = 8760 h), the duty cycle would be 0.0068, which is clearly below 1%. In reality, traffic generated by several nodes might accumulate. On the other hand, a significant part of the battery power will be spent in receive mode (which requires more current than the transmit mode for many implementations). Thus the 1% duty cycle also is valid for networks of IEEE 802.15.4 devices. In some rare cases, traffic might aggregate in proximity to coordinator nodes. Thus an aggregated duty cycle of up to 10% can be assumed in rare cases.

E.6.4 Impact of increased data rate

It should be noted that IEEE 802.15.4 and IEEE 802.15.4a devices will serve applications with similar low required data traffic. Since IEEE 802.15.4a devices offer a significantly increased data rate (1 Mb/s versus 250 kb/s), the duty cycle of IEEE 802.15.4a devices can be expected to be significantly below the duty cycle of IEEE 802.15.4 devices. Since the 2.4 GHz ISM band has become an extremely busy medium, a low duty cycle achieved by high data rates is crucial for reasonable coexistence performance.

E.6.5 Co-channel scenario

Operating any two systems at the same location and at the same center frequency is obviously not a desirable situation. As long as no active interference cancellation is provided, the coexistence performance will be determined by the duty cycle behavior of both systems. Applying the duty cycle assumptions on IEEE 802.15.4a devices as stated above will result in reasonable performance. However, whenever possible, it is recommended that this situation be avoided by using a nonoverlapping channel. When a nonoverlapping channel is not available to the CSS PHY, because other networks (for example, IEEE 802.11 networks) are themselves already using the nonoverlapping channels, the recommendation is to select for the CSS PHY a channel between the channels already in use. It is further recommended that in the case of IEEE 802.11 networks, the CSS center frequency be selected so that the spatially closer IEEE 802.11 network has a frequency offset of at least 15 MHz.

Figure E.17 through Figure E.36 show the computed PER versus separation distances (in meters) for co-channel pairings of systems when those systems use the spectra and filter properties given in Table E.1.




Figure E.17—IEEE 802.15.4a CSS receiver, IEEE 802.11 HR/DSSS interferer




Figure E.18—IEEE 802.11 HR/DSSS receiver, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.19—IEEE 802.11 HR/DSSS receiver, IEEE 802.15.4a CSS interferer with rare duty cycle




Figure E.20—IEEE 802.15.4a CSS receiver, IEEE 802.11 ERP interferer




Figure E.21—IEEE 802.11 ERP receiver, 6 Mb/s, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.22—IEEE 802.11 ERP receiver, 6 Mb/s, IEEE 802.15.4a CSS interferer with rare duty cycle




Figure E.23—IEEE 802.11 ERP receiver, 24 Mb/s, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.24—IEEE 802.11 ERP receiver, 24 Mb/s, IEEE 802.15.4a CSS interferer with rare duty cycle




Figure E.25—IEEE 802.11 ERP receiver, 54 Mb/s, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.26—IEEE 802.11 ERP receiver, 54 Mb/s, IEEE 802.15.4a CSS interferer with rare duty cycle




Figure E.27—IEEE 802.15.4a CSS receiver, IEEE 802.15.1 interferer




Figure E.28—IEEE 802.15.1 receiver, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.29—IEEE 802.15.1 receiver, IEEE 802.15.4a CSS interferer with rare duty cycle




Figure E.30—IEEE 802.15.4a CSS receiver, IEEE 802.15.3 interferer

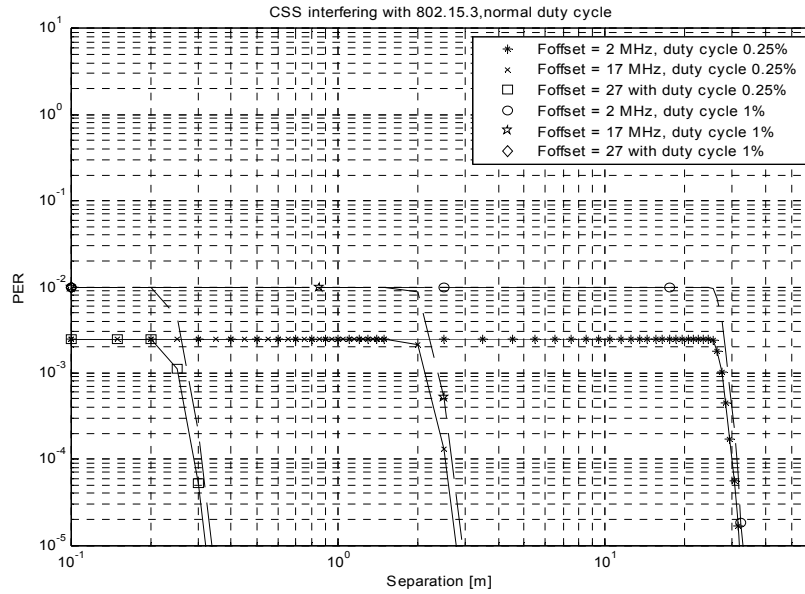


Figure E.31—IEEE 802.15.3 receiver, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.32—IEEE 802.15.3 receiver, IEEE 802.15.4a CSS interferer with rare duty cycle




Figure E.33—IEEE 802.15.4a CSS receiver, IEEE 802.15.4 interferer with normal duty cycle




Figure E.34—IEEE 802.15.4a CSS receiver, IEEE 802.15.4 interferer with rare duty cycle




Figure E.35—IEEE 802.15.4 receiver, IEEE 802.15.4a CSS interferer with normal duty cycle




Figure E.36—IEEE 802.15.4 receiver, IEEE 802.15.4a CSS interferer with rare duty cycle

E.7 UWB coexistence performance

E.7.1 Specific regulatory requirements for UWB coexistence

Surprisingly, despite the wide bandwidth of the UWB PHY, there is only one other IEEE standard waveform that may occupy the same frequency bands, namely, IEEE 802.16 systems below 10 GHz. Cognizant of the potential for coexistence issues, regulators in the parts of the world where IEEE 802.16 systems (such as WiMAX) may be deployed in bands overlaid by UWB spectrum are creating specific regulatory requirements to further reduce the likelihood of any coexistence problems. In both Asia and the European Union, regulators are creating rules for unlicensed UWB operation that will require specific active mitigation mechanisms to ensure peaceful coexistence with IEEE 802.16 systems or other similar systems used for fixed or mobile wireless access.

Additionally, a proposed IEEE standard, P802.22, proposes to occupy parts of the bandwidth in the UWB PHY 150–650 MHz band. In the regulatory domains where this is presently allowed (FCC), the maximum transmit power is specified an additional (approximately) 35 dB lower compared the limits for the 3.1–10 GHz bands. Some regulatory domains (including FCC) have suggested that certain applications, specifically those involving personnel location in emergency response situations, would be allowed at higher PSD levels under specific conditions, where other factors such as operating limitations would provide required protection of incumbent services. Clearly it is beyond the scope of this standard to anticipate specific future regulatory actions. However, in considering the application scenarios presented in the call for applications and responding to specific guidance from regulators in the United States, it can be observed that coexistence with the IEEE P802.22 systems and other known incumbent systems is assured through operating conditions. As a primary mitigation factor, it is unlikely such systems will be operating in near physical proximity at the same time as emergency response teams. Such conditions are the scope of regulatory agencies to define, and it is the responsibility of implementers of this standard to conform with applicable regulations and conditions.

In considering other personnel location scenarios, the mitigations factors described for other UWB applications apply equally to all UWB bands.

E.7.2 Mitigation of interference from UWB devices using low PAN duty cycles

One proposal made to the Task Group 4a is to use a lower duty cycle within a UWB piconet to reduce potential interference effects. Low-duty-cycle piconet scenarios could be used in the following situations:

- IEEE 802.15.4a devices are deployed in high density in a limited area, e.g., hot-spot deployment scenarios.
- UWB victim systems cover much larger area than the coverage of a typical IEEE 802.15.4a PANs.

In these cases, transmissions from every device in the PAN can affect the victim receiver. For reasons of less complexity, lower power consumption, as well as physical limitations, it is difficult for simple IEEE 802.15.4a devices to detect victim system reliably. The aggregate interference from the PAN increases with increment in number of PAN members. The interference to victim systems could be limited by controlling duty cycle of the PAN through general active/inactive periods. The UWB traffic can occur only in the active period. Victim systems would then be free of interference in the inactive period. The interference level could be controlled by the ratio of active period to the total period.

E.7.3 Coexistence assurance: methodology and assumptions

In order to quantify the coexistence performance of the IEEE 802.15.4a UWB PHY, the techniques described by Shellhammer²¹ have been adapted.

The coexistence assurance methodology predicts the PER of an affected wireless network (AWN, or victim) in the presence of an interfering wireless network (IWN, or assailant). In its simplest form, the methodology assumes an Awn and an Iwn, each composed of a single transmitter and a receiver. The methodology takes as input a path loss model, a quantitative model for the BER of the Awn, and predicted temporal models for packets generated by the Awn and for “pulses,” i.e., packets generated by the Iwn. Based on these inputs, the methodology predicts the PER of the Awn as a function of the physical spacing between the Iwn transmitter and the Awn receiver.

The appeal of the coexistence assurance methodology is that multiple networking standards can be characterized and compared with just a few parameters, notably,

- Bandwidth of Awn and Iwn devices
- Path loss model for the networks
- BER as a function of SIR of Awn devices
- Temporal model for Awn packets and Iwn “pulses” (interfering packets)

Subclauses E.7.4 through E.7.7 describe the general assumptions made across all of the PHYs covered under this standard.

E.7.4 UWB PHY coexistence

E.7.4.1 Victims and assailants

At present, IEEE Std 802.15.4a-2007 for UWB systems is the only wireless networking standard in the UWB bands covered under IEEE Std 802[®]. The only other IEEE wireless standard waveforms that overlap this same spectrum are IEEE 802.16 systems occupying 3400–3800 MHz licensed frequency bands in some regions (parts of Europe and Asia). In addition, the proposed standard IEEE P802.22 would occupy parts of the band between 150 MHz to 650 MHz.

In addition to IEEE standardized wireless systems, another UWB standard produced by ECMA is specified in ECMA 368. A limited analysis of the coexistence between this system and IEEE 802.15.4a waveform is given here.

In this analysis, the assumption is made that the PHYs will serve as both victims (i.e., participants in Awns) and as assailants (i.e., participants in Iwns).

E.7.4.2 Bandwidth for UWB systems

The IEEE 802.15.4a UWB PHYs that operate in any of the three UWB bands have one or more channels, approximately 500 MHz wide or, optionally, 1300 MHz wide. The ECMA 368 PHY has a nominal bandwidth of 1500 MHz. In contrast to these UWB systems, the narrowband IEEE 802.16 PHYs that operate in the 2–10 GHz band have multiple defined channels, each 20 MHz wide or less. IEEE P802.22 would have multiple defined channels, each 6 MHz to 8 MHz wide. The coexistence methodology assumes that any UWB device in an Awn or Iwn will have a much greater bandwidth than a narrowband device in a corresponding Awn or Iwn (so BUWB >> BNB).

²¹S. J. Shellhammer, “Estimating Packet Error Rate Caused by Interference—A Coexistence Assurance Methodology,” IEEE 802.19-05/0029r0, Sept. 14, 2005.

E.7.5 Path loss model

The coexistence methodology uses a variant of the path loss model described by Shellhammer,²² which stipulates a two-segment function with a path loss exponent of 2.0 for the first 8 meters and then a path loss model of 3.3 thereafter. The formula given by Shellhammer is as follows:

$$pl(d) = \begin{cases} 40.2 + 20\text{Log}_{10}(d) & d \leq 8\text{ m} \\ 58.5 + 33\text{Log}_{10}\left(\frac{d}{8}\right) & d > 8\text{ m} \end{cases}$$

The constants in this formula are based on a 2.4 GHz center frequency. To adapt the model to a typical center frequency in the 3100–4800 MHz frequency band, it can be generalized as follows:

$$pl(d) = \begin{cases} pl(1) + 10\gamma_1\text{Log}_{10}(d) & d \leq 8\text{ m} \\ pl(8) + 10\gamma_8\text{Log}_{10}\left(\frac{d}{8}\right) & d > 8\text{ m} \end{cases}$$

where $pl(1)$ is the path loss at 1 m (in decibels), 1 is the path loss exponent at 1 m (2.0), and 8 is the path loss exponent at 8 m (3.3). The initial condition of $pl(1)$ is computed as follows:

$$pl(1) = 10\gamma_1\text{Log}_{10}\left(\frac{4\pi f}{C}\right)$$

With $1 = 2.0$, $f = 3400$ MHz, and $C = \text{speed of light} = 299792458 \text{ ms}^{-1}$, then $pl(1) = 43.08$ and $pl(8) = 61.14$. The path loss function modified for 3400 MHz is, therefore,

$$pl(d) = \begin{cases} 43.03 + 20\text{Log}_{10}(d) & d \leq 8\text{ m} \\ 61.09 + 33\text{Log}_{10}\left(\frac{d}{8}\right) & d > 8\text{ m} \end{cases}$$

With $f = 400$ MHz for the sub-gigahertz UWB band, then $pl(1) = 24.49$ and $pl(8) = 78.75$. The path loss function for 400 MHz center frequency is the same as for 3400 MHz with the substitution of the following constants:

$$pl(d) = \begin{cases} 24.49 + 20\text{Log}_{10}(d) & d \leq 8\text{ m} \\ 78.75 + 33\text{Log}_{10}\left(\frac{d}{8}\right) & d > 8\text{ m} \end{cases}$$

A plot of the path loss as a function of device separation distance is shown in Figure E.37.

²²S. J. Shellhammer, "Estimation of Packet Error Rate Caused by Interference using Analytic Techniques—A Coexistence Assurance Methodology," IEEE 802.19-05/0028r0, Sept. 14, 2005.




Figure E.37—Path loss function

E.7.6 BER as a function of SIR

For the PHY specifications analyzed in this standard, there are no analytic expressions for the BER or symbol error rate (SER) of the signal due to the use of FEC methods to improve reliability.

In this analysis, a method is used that is equivalent to using interpolation of table values. In order to simplify the calculations and still provide meaningful results, the relationship is approximated between the changes in BER (on a logarithmic scale) and varying SNR as a linear with a slope of 0.6 dB per order of magnitude (10x) change in BER over the range of BER that is relevant to this analysis (about $1e-8$ to $1e-5$ BER). This approximation is reasonable for the FEC methods used for IEEE Std 802.16-2004 (Reed-Solomon block code), ECMA 368, IEEE P802.22, and IEEE Std 802.15.4a-2007 (convolutional coding).

For each of the systems, the effect of the IWN on the AWN is characterized by computing the rise in the effective operating noise floor of the AWN by the interference of the IWN (modeled as uncorrelated wideband noise). The analysis will assume a baseline operating effective noise floor (including effects of thermal noise floor, noise figure, and operating margin to account for other real-world effects such as multipath propagation effects and co-channel or adjacent channel interference). This approach allows the characterization of the effect of the IWN on the AWN as the IWN is moved from a large separation distance (when the AWN has a baseline nominal PER) to a very close distance where the interference effect of the IWN dominates the PER during periods of operation (subject to duty cycle assumptions).

Although this analysis approach is perhaps not as elegant as the use of an analytic expression (not possible in these cases), it will provide a good characterization of the coexistence of these systems under real-world conditions and can be used to estimate a range of effects for an equivalent range of assumptions about operating margin.

E.7.7 Temporal model

In IEEE Std 802.15.4a-2007, packet overhead is kept to minimum. The maximum PSDU size is 128 bytes, and a typical packet may be only 32 bytes, including PSDU and synchronization bytes. For this coexistence methodology, all packets, whether belonging to the AWN or IWN, are assumed to be 32 bytes.

Although there is no duty-cycle limitation in the authorized UWB bands at this point, many IEEE 802.15.4a-based networks are expected to operate at well under 5% duty cycle, particularly devices that are battery-powered. This 5% duty cycle level has also been used by regulators as a high value for expected UWB communications device operating levels on various coexistence studies. In addition, IEEE Std 802.15.4a is based on the use of an ALOHA contention-based access mechanism that is intended to support only lower duty cycle applications. Based on these factors, it is reasonable to expect that IEEE 802.15.4a piconets used for many applications will operate at duty cycles as high as 10%. For purposes of modeling coexistence, the assumption is made that all UWB devices operating in piconets will have a shared duty cycle of 10% and that such piconets will operate within a range of a few tens of meters. Based on this and a typical active device population of five devices per piconet, an average operating duty cycle of 2% is assumed for any particular device within a piconet.

For the other wireless systems considered in this analysis (IEEE 802.16, IEEE P802.22, and ECMA 368), anticipated applications are focused on higher bandwidth connectivity over wide areas for IEEE 802.16 and IEEE P802.22 systems and over short WPAN ranges for ECMA 368 systems. Because these systems are not deployed in great numbers, it is not possible to qualify typical operating duty cycle. For this analysis, therefore, the initial assumption is a very conservative continuous operation as a baseline worst-case scenario.

E.7.8 Coexistence analysis

This subclause details the assumptions for the coexistence analysis and presents the results for each of the cases analyzed.

E.7.8.1 Impact of IEEE 802.15.4a devices on IEEE 802.16 networks

Assumptions

- The IEEE 802.16 receiver is the victim (AWN) and is an indoor fixed or nomadic client node of the network. The base station node will not be susceptible to IEEE 802.15.4a UWB interference due to site positioning. The AWN operates in 3.4–3.8 GHz licensed bands (available in most of the world except the United States).
- The IEEE 802.16 receiver is operating in a real-world environment in the presence of multipath fading and interference, and a 3–10 dB margin above sensitivity functions well. The baseline PER is $1e-6$ at 3 dB above sensitivity in the absence of any UWB device effects, and the receiver noise floor is 6 dB.
- UWB interference is wideband uncorrelated noise since the bandwidth is much wider than victim receiver. The difference in antenna gains is 10 dB since the indoor or outdoor IEEE 802.16 antenna will have gain in the direction of the desired base station downlink signal. The UWB device will not directly block the LOS.

E.7.8.1.1 Coexistence methodology results

Table E.6 shows the calculation of the allowable path loss that would result in an IEEE 802.15.4a UWB emission level at the AWN equal to the effective operating noise floor. Base on this path loss, the effect on AWN PER is computed as a function of separation distance, shown in Figure E.38.

Table E.6—Computation of acceptable levels of IEEE 802.15.4a device emissions for an operating IEEE 802.16 client node

Quantity	Value	Units	Notes
UWB transmit PSD limit (PLIM)	-41.3	dBm/MHz	Set by regulatory authority.
Average margin to limit (MBO)	1.7	dB	Transmit power back-off due to spectral ripple (0.5+ dB) and ~1 dB margin for manufacturing tolerance, etc.
Average UWB antenna gain (GUWB)	-2	dBi	Average gain from small, low-cost UWB antenna to arbitrary victim receiver over 360°.
Average emissions PSD (PLIM – MBO + GUWB) seen by IEEE 802.16 device receiver	-45	dBm/MHz	Average PSD seen in direction of arbitrary victim receiver.
IEEE 802.16 thermal noise floor (kTB)	-114	dBm/MHz	Thermal noise floor (room temperature).
IEEE 802.16 NF	6	dB	Noise figure for indoor IEEE 802.16 terminal.
Average IEEE 802.16 antenna gain in direction of interfering UWB	-4	dBi	Gain of IEEE 802.16 antenna in main beam (to desired IEEE 802.16 base station) is 6–7 dBi and to nearby UWB interferer (not blocking antenna main beam) -4 dBi.
IEEE 802.16 operating margin (M16)	3–10	dB	Operating margin for acceptable performance in presence of multipath fading and adjacent cell/channel interference.
IEEE 802.16 effective operating noise floor for UWB interference susceptibility: (kTB + NF16 – G16 + MOP)	-101 to -94	dBm/MHz	The effective operating noise floor level for the IEEE 802.16 operating receiver.
Level of wideband IEEE 802.15.4a UWB interference that result in a 3 dB rise in IEEE 802.16 effective operating noise floor	-101 to -94	dBm/MHz	For 3 dB rise, wideband UWB emissions in-band can be at the same level as effective operating noise floor for indoor IEEE 802.16 node receiver.

Table E.6—Computation of acceptable levels of IEEE 802.15.4a device emissions for an operating IEEE 802.16 client node (continued)

Quantity	Value	Units	Notes
Path loss (range) from UWB to IEEE 802.16 receiver (average case) for 3 dB rise in effective operating noise floor	49 to 56 (2 to 4.5)	dB (m)	For 3 dB rise, wideband UWB emissions in-band can be at the same level as effective operating noise floor for indoor IEEE 802.16 node receiver.
Path loss (range) from UWB to IEEE 802.16 receiver (average case) for 1 dB rise in effective operating noise floor	55 to 61 (4 to 8)	dB (m)	For 1 dB rise, wideband UWB emissions in-band must be 6 dB below effective operating noise floor for indoor IEEE 802.16 node receiver.




Figure E.38—Effect on IEEE 802.16 AWN as a function of separation distance from IEEE 802.15.4a UWB device

E.7.8.2 Impact of an IEEE 802.16 device on IEEE 802.15.4a UWB networks

Assumptions

- The IEEE 802.15.4a UWB device is the affected device (AWN). The IEEE 802.16 device is the interferer (IWN) and is an indoor fixed or nomadic client node of the network. The base station node will have less interference effects on IEEE 802.15.4a UWB devices due to UWB device deployment much closer to subscriber or mobile IEEE 802.16 devices. The IWN operates in 3.4–3.8 GHz licensed bands (available in most of the world except the United States). For this analysis, the IWB operates at a conservative 50% duty cycle (IEEE 802.16 subscriber uplink).
- The IEEE 802.15.4a UWB receiver is operating in a real-world environment in the presence of multipath fading and interference, and the margin above sensitivity is 3 dB during operation. The baseline PER is 1e-7 at 3 dB above sensitivity in the absence of any UWB device effects, and the receiver noise floor is 10 dB.

- UWB interference is wideband uncorrelated noise since the bandwidth is much wider than victim receiver. The difference in antenna gains is 10 dB since the indoor or outdoor IEEE 802.16 antenna will have gain in the direction of the desired base station downlink signal. The UWB device will not directly block the LOS.

E.7.8.3 Coexistence methodology results

Table E.7 shows the calculation of the allowable path loss that would result in a IEEE 802.15.4a UWB emission level at the AWN equal to the effective operating noise floor. Base on this path loss, the effect on AWN PER is computed as a function of separation distance, shown in Figure E.39.

Table E.7—Computation of acceptable levels of IEEE 802.15.4a device emissions for an operating IEEE 802.16 client node

Quantity	Value	Units	Notes
IEEE 802.16 client device transmit power (P16)	17	dBm	Assumes subscriber station in small cell.
IEEE 802.16 client device bandwidth	5	MHz	
IEEE 802.15.4a UWB device bandwidth	500	MHz	
Average IEEE 802.16 antenna gain (G16)	-2	dBi	Average gain from antenna to arbitrary victim receiver over 360° (IWN typically not in main beam).
Average emissions PSD (P16 + G16 - 10Log(BUWB) seen by IEEE 802.15.4a UWB device receiver	-12	dBm/MHz	Average PSD seen in direction of arbitrary victim receiver (assumes that UWB receiver can spread interference power into receiver bandwidth).
IEEE 802.15.4a UWB thermal noise floor (kTB)	-114	dBm/MHz	Thermal noise floor (room temperature).
IEEE 802.15.4a UWB NF	10	dB	Noise figure for low-cost IEEE 802.15.4a device.
IEEE 802.15.4a UWB operating margin (MUWB)	3	dB	Operating margin for acceptable performance in presence of multipath fading (assumes no interference other than IWN).
IEEE 802.15.4a UWB effective operating noise floor for UWB interference susceptibility: (kTB + NFUWB + MUWB)	-101	dBm/MHz	The effective operating noise floor level for the IEEE 802.15.4a operating receiver.
Level of interference power density to achieve a 3 dB rise in IEEE 802.15.4a UWB effective operating noise floor	-101	dBm/MHz	For 3 dB rise, IEEE 802.16 power emissions in-band can be at the same level as effective operating noise floor for UWB receiver.

Table E.7—Computation of acceptable levels of IEEE 802.15.4a device emissions for an operating IEEE 802.16 client node (continued)

Quantity	Value	Units	Notes
Path loss (range) from IEEE 802.16 to UWB receiver (average case) for 3 dB rise in effective operating noise floor	89 (48)	dB (m)	For 3 dB rise, IEEE 802.16 power emissions in-band can be at the same level as effective operating noise floor for UWB receiver.
Path loss (range) from IEEE 802.16 to UWB receiver (average case) for 1 dB rise in effective operating noise floor	95 (75)	dB (m)	For 1 dB rise, wideband UWB emissions in-band must be 6 dB below effective operating noise floor for indoor IEEE 802.16 node receiver.




Figure E.39—Effect on IEEE 802.15.4a UWB AWN as a function of separation distance from IEEE 802.16 IWN device

E.7.8.4 Low-duty-cycle UWB assaulting a WiMAX link

These results are an extract from a French contribution to Electronic Communications Committee (ECC) Task Group 3 meeting #15.

The impact of UWB on a fixed broadband wireless access (FBWA) system is measured on video streaming (see Table E.8), which is considered as a relevant service in term of vulnerability, bandwidth use, and timing constraint.

The methodology used is the following:

- Set the WiMAX received signal strength at equipment at minimum sensitivity level (−98 dBm).
- Get a reference measure without UWB (depending on each test case).
- Measure the degradation with low-duty-cycle UWB emission [for any considered activity factor (AF) and distances]. Degradation is, in decibels, the increase of power needed by the WiMAX receiver to reestablish the reference link quality.

Table E.8—Impact of UWB on FBWA system measured on video streaming

Degradation (dB)		Distance (m)			
		0.5	1	2	4
AF (Ton/Toff ms)		0.5	1	2	4
2%	075/38	1	N/A	1	N/A
	5/245	0	N/A	0	N/A
	10/490	1	N/A	1	N/A
5%	2/38	2	1	0	N/A
	5/95	1	N/A	1	N/A
	10/190	0	N/A	1	N/A
10%	2/18	3	N/A	0	0
	5/45	3	2	0	0
	10/90	2	N/A	1	0.5

Table E.9 shows the evolution of the lowest needed receive signal strength indicator (RSSI) to achieve a reliable 1 Mb/s throughput with respect to UWB activity. The reference level is -98 dBm (i.e., without UWB activity).

Table E.9—Lowest RSSI to achieve reliable 1 Mb/s throughput

RSSI needed to achieve 1 Mb/s data rate (dBm)		Distance (m)		
AF (Ton/Toff ms)		0.5	2	4
2%	075/38	-98 (-98)	-98 (-98)	N/A
	5/245	-98 (-98)	-98 (-98)	N/A
	10/490	-97 (-98)	-97 (-98)	N/A
5%	2/38	-98 (-98)	-98 (-98)	N/A
	5/95	-98 (-98)	-98 (-98)	N/A
	10/190	-97 (-98)	-98 (-98)	N/A
10%	2/18	-97 (-98)	-98 (-98)	N/A
	5/45	-98 (-98)	-98 (-98)	N/A
	10/90	-97 (-98)	-98 (-98)	N/A

E.7.9 Impact of IEEE 802.15.4a devices on ECMA 368 networks

Assumptions

- The ECMA 368 receiver is the victim (AWN). The AWN operates using frequency hopping in bands across the 3.1–4.8 GHz unlicensed UWB bands (available only in the United States at this time), but the IEEE 802.15.4a device operates only in band 3 (mandatory).
- The ECMA 368 receiver is operating in a real-world environment in the presence of multipath fading and interference, and a 5 dB margin above sensitivity functions well. The baseline PER is $8e-2$ at sensitivity ($8e-7$ at 3 dB above sensitivity) in the absence of any UWB device effects, and the receiver noise floor is 6 dB.

E.7.9.1 Coexistence methodology results

Table E.10 shows the calculation of the allowable path loss that would result in an IEEE 802.15.4a UWB emission level at the AWN equal to the effective operating noise floor. Base on this path loss, the effect on AWN PER is computed as a function of separation distance, shown in Figure E.40.

Table E.10—Computation of acceptable levels of IEEE 802.15.4a device emissions for an operating ECMA 368 device

Quantity	Value	Units	Notes
UWB Transmit PSD Limit (PLIM)	-41.3	dBm/MHz	Set by regulatory authority
Average margin to limit (MBO)	1.7	dB	Due to spectral ripple (0.5+ dB) and ~1 dB margin for manufacturing tolerance, etc.
Average UWB antenna gain (GUWB)	-2	dBi	Average gain from small, low-cost UWB antenna to arbitrary victim receiver over 360°
Average emissions PSD (PLIM - MBO + GUWB)	-45	dBm/MHz	Average PSD seen in direction of arbitrary victim receiver
UWB victim thermal noise floor (kTB)	-114	dBm/MHz	Thermal noise floor (room temperature)
UWB victim NF	6	dB	Noise figure for the ECMA 368 receiver
UWB victim frequency diversity	3	dB	ECMA UWB system uses 2x band frequency diversity for then encoding of each bit as part of its frequency hopping scheme
UWB victim operating margin (MECMA)	5	dB	Operating margin for acceptable performance in presence of multipath fading and RF interference
IEEE 802.16 effective operating noise floor for UWB interference susceptibility: (kTB + NFECMA368 + DFD + MOP)	-100	dBm/MHz	The effective allowable interference power level for the ECMA 368 operating receiver

Table E.10—Computation of acceptable levels of IEEE 802.15.4a device emissions for an operating ECMA 368 device (continued)

Quantity	Value	Units	Notes
Level of wideband UWB emissions that result in 3 dB rise in ECMA 368 effective operating noise floor	-100	dBm/MHz	For 3 dB rise, IEEE 802.15.4a UWB emissions in-band can be at the same level as effective operating noise floor for AWN device receiver
Path loss (range) from UWB to ECMA 368 receiver (average case) for 3 dB rise in effective operating noise floor	55 (3)	dB (m)	For 3 dB rise, wideband UWB emissions in-band can be at the same level as effective operating noise floor for AWN device receiver
Path loss (range) from UWB to ECMA 368 receiver (average case) for 1 dB rise in effective operating noise floor	61 (6)	dB (m)	For 1 dB rise, wideband UWB emissions in-band must be 6 dB below effective operating noise floor for indoor IEEE 802.16 node receiver




Figure E.40—Effect on ECMA 368 AWN as a function of separation distance from IEEE 802.15.4a UWB device

E.7.10 Impact of IEEE 802.15.4a devices on IEEE P802.22 networks

Based on the currently available draft of IEEE P802.22, the operating conditions are generally similar to IEEE Std 802.16-2004. The primary operating considerations include the following:

- The IEEE P802.22 network is a fixed-point-to-multipoint network, operating in a narrow band (6–8 MHz) widely spaced between 54 MHz and 862 MHz; the fixed node will not be susceptible to IEEE 802.15.4a interference due to positioning.
- The UWB PHY channel at 150 MHz to 650 MHz is operating, on average, at least –75 dBm (set by regulation, using current FCC limits), which is at approximately 34 dB lower power than the higher band UWB PHY (–41.3 dBm).
- UWB interference is wideband uncorrelated noise since the bandwidth is much wider than the victim receiver. A 10 dB difference in antenna gain is assumed in anticipation that the IEEE P802.22 antenna will require gain in the direction of the desired fixed node (base station) downlink signal, and it is also assumed that the UWB device will not directly block the LOS.

E.7.10.1 Coexistence methodology results

At the time of this analysis, the characteristics of the IEEE P802.22 AWN were not completely defined. Assuming similar characteristics to an IEEE 802.16 device with the operating frequencies specified above, note that the 150–650 MHz UWB PHY has a similar path loss curve to the 3100–4800 MHz UWB PHY with the noted 6–8 dB difference along the curve. Note further that the maximum radiated power is 34 dB lower and the effective interference seen by the AWN will be lower than shown for the IEEE 802.16 case.

E.7.11 Conclusions

These analyses characterize the expected coexistence behavior between IEEE 802.15.4a UWB devices and IEEE 802.16 devices. Also described are the expected effects of an IEEE 802.15.4a device on an ECMA 368 receiver and the proposed IEEE P802.22 devices. One conclusion that can be drawn is that the relative effects of the IEEE 802.15.4a device and IEEE 802.16 device to each other are quite different. The IEEE 802.15.4a device is impacted by the IEEE 802.16 device at much longer range than vice versa. The implication is that the IEEE 802.15.4a device would not be able to operate at all at ranges where its emissions would impact the IEEE 802.16 device because of the large asymmetry in the transmit power levels (+17 dB for the IEEE 802.16 device versus –15 dBm for the IEEE 802.15.4a device). In such a case, either the IEEE 802.15.4a device would accept the much higher PER, or else it could simply use a different channel or some other form of interference mitigation.

A similar conclusion can be reached regarding proposed IEEE P802.22 devices; there is an even greater asymmetry in power levels, as the sub-gigahertz band is operated at a substantially lower level than the higher UWB bands. One form of mitigation (in both directions) is to observe that when considering the application environment in which the sub-gigahertz UWB band has greatest advantage and is, therefore, most likely to be used, the operation of IEEE P802.22 devices in near proximity is unlikely. In application scenarios where it is expected that IEEE 802.15.4a sub-gigahertz devices may operate in proximity to IEEE P802.22 devices, the IEEE 802.15.4a devices may need to employ some other forms of interference mitigation. Additional mitigation is available to the IEEE P802.22 device. Note that a great number of potential channels are available above 650 MHz and provide the option to the IEEE P802.22 device to change to a channel outside the operating range of the sub-gigahertz UWB.

Change the following subclause number due to the insertion above of the new subclauses, E.6 and E.7:

E.8 ~~E.6~~Notes on the calculations

Annex F

(informative)

Change the title of Annex F as shown:

~~IEEE 802.15.4~~ **Regulatory requirements**

Insert the following new subclause (F.1) at the beginning of Annex F:

F.1 IEEE Std 802.15.4

Change the subclause numbering on the existing F.1 through F.8 as follows:

F.1.1 ~~F.1~~ Introduction

F.1.2 ~~F.2~~ Applicable U.S. (FCC) rules

F.1.2.1 ~~F.2.1~~ Section 15.35 of FCC CFR47

F.1.2.2 ~~F.2.2~~ Section 15.209 of FCC CFR47

F.1.2.3 ~~F.2.3~~ Section 15.205 of FCC CFR47

F.1.2.4 ~~F.2.4~~ Section 15.247 of FCC CFR47

F.1.2.5 ~~F.2.5~~ Section 15.249 of FCC CFR47

F.1.3 ~~F.3~~ Applicable European rules

F.1.3.1 ~~F.3.1~~ European 2400 MHz band rules

F.1.3.2 ~~F.3.2~~ European 868-870 MHz band rules

F.1.4 ~~F.4~~ Applicable Japanese rules

F.1.5 ~~F.5~~ Emissions specification analysis with respect to known worldwide regulations

F.1.5.1 ~~F.5.1~~ General analysis and impact of detector bandwidth and averaging rules

F.1.5.2 ~~F.5.2~~ Frequency spreading and averaging effects specific to IEEE Std 802.15.4

F.1.6 ~~F.6~~ Summary of out-of-band spurious emissions limits

F.1.7 ~~F.7~~ Phase noise requirements inferred from regulatory limits

F.1.8 ~~F.8~~ Summary of transmission power levels

Insert after F.1.8 the following new subclauses (F.2 through F.2.4.3):

F.2 IEEE 802.15.4a UWB

F.2.1 Introduction

Three regulatory environments are presented in F.2.2 through F.2.4: U.S. (FCC) rules, regulation in Europe, and regulation in Japan.

Worldwide UWB regulations are rapidly evolving, and the content in F.2 was current at the time this standard was published.

F.2.2 Applicable U.S. (FCC) rules

The FCC adopted “First report and order” in February 2002,²³ which allows UWB operation with the average emission limits given in Table F.13.

The peak EIRP limit being adopted in this report and order is 0 dBm when measured with a resolution bandwidth of 50 MHz and 20 log (RBW/50) dBm when measured with a resolution bandwidth ranging from 1 MHz to 50 MHz. RBW is the resolution bandwidth, in megahertz, actually employed. The minimum resolution bandwidth that may be employed is 1 MHz; the maximum resolution bandwidth that may be employed is 50 MHz.

Table F.13—UWB average emission limits, EIRP in dBm/MHz

Frequency band (MHz)	Imaging below 960 MHz	Imaging, mid-frequency	Imaging, high frequency	Indoor applications	Hand held, including outdoor	Vehicular radar
0.009–960	FCC §15.209	FCC §15.209	FCC §15.209	FCC §15.209	FCC §15.209	FCC §15.209
960–1610	–65.3	–53.3	–65.3	–75.3	–75.3	–75.3
1610–1990	–53.3	–51.3	–53.3	–53.3	–63.3	–61.3
1990–3100	–51.3	–41.3	–51.3	–51.3	–61.3	–61.3
3100–10 600	–51.3	–41.3	–41.3	–41.3	–41.3	–61.3
10 600–22 000	–51.3	–51.3	–51.3	–51.3	–61.3	–61.3
22 000–29 000	–51.3	–51.3	–51.3	–51.3	–61.3	–41.3
Above 29 000	–51.3	–51.3	–51.3	–51.3	–61.3	–51.3

²³“First Report and Order Regarding UWB Transmission,” issued by the FCC (Washington, DC 20554) {ET Docket 98-153} 14 February 2002.

F.2.3 Applicable European rules

F.2.3.1 Generic rules

In Europe, Draft ECC Decision ECC/DEC/(06)AA was under “public consultation” until 24 December 2005. A total of 67 comments were made, which were dealt with during the next meeting of the corresponding European Conference of Postal and Telecommunications Administration (CEPT) Task Group (TG3).

Decision ECC/DEC/(06)04, which was made 24 March, allows UWB operation in the upper band (6 GHz to 8.5 GHz) with the emission limits given in Table F.14. This decision is now to be implemented by national regulatory agencies of the 45 CEPT member countries.

The main points of this decision are :

1. that this ECC Decision defines general harmonised conditions for the use in Europe of devices using UWB technology in bands below 10.6 GHz;
2. that the devices permitted under this ECC Decision are exempt from individual licensing and operate on a non-interference, non-protected basis;
3. that this ECC decision is not applicable to:
 - a) flying models,
 - b) outdoor installations and infrastructure, including those with externally mounted antennas,
 - c) devices installed in road and rail vehicles, aircraft and other aviation;
4. that devices covered by the scope of this ECC Decision are not allowed to be used at a fixed outdoor location or connected to a fixed outdoor antenna;
5. that the technical requirements detailed in Annex 1 apply to devices permitted under this ECC Decision;
6. that this Decision enters into force on 24 March 2006;
7. that the preferred date for implementation of this Decision shall be 1 October 2006;
8. that CEPT administrations shall communicate the national measures implementing this Decision to the ECC Chairman and the Office when the Decision is nationally implemented.

A separate decision (ECC/DEC(06)12) has been made concerning low-duty-cycle limitation to enable access to the use of the lower band with the guarantee of an efficient protection of licensed services (see F.2.3.2).

Another separate decision concerning the lower band is still under consideration. Some issues concerning this band are still to be resolved. ECC Task Group 3 shall provide recommendations and specification of DAA procedures to enable the use of the lower band with the guarantee of an efficient protection of licensed services. The question of whether the ECC should adopt a “mitigation free period” until 2010 or 2012 is also under consideration. It should be noted that the ECC Decision intends to deliver a clear message that the 6–8.5 GHz band is identified in Europe for long-term UWB operation.

Table F.14—ECC Decision 06/04 on UWB emission limits, EIRP

Frequency range	Maximum mean EIRP density (dBm/MHz)	Maximum peak EIRP density (dBm/50MHz) (Note 2)
Below 1.6 GHz	−90	−50
1.6 to 3.8 GHz (Note 1)	−85	−45
3.8 to 4.8 GHz (Note 1)	−70	−30
4.8 to 6 GHz	−70	−30
6 to 8.5 GHz	−41.3	0
8.5 to 10.6 GHz	−65	−25
Above 10.6 GHz	−85	−45

Note 1—ECC is still considering whether to adopt a separate decision covering the 3.1–4.8 GHz frequency band.
 Note 2—The peak EIRP can be alternatively measured in a 3 MHz bandwidth. In this case, the maximum peak EIRP limits to be applied is scaled down by a factor of $20\log(50/3) = 24.4$ dB.

Two other technical requirements are also expressed:

- *Pulse repetition frequency (PRF)*. The PRF for UWB devices are not to be less than 1 MHz. This restriction does not apply to burst repetition frequency.
- *Transmission activity*. A communications system is allowed to transmit only when it is sending information to an associated receiver or attempting to acquire or maintain association. The device is required to cease transmission within 10 s unless it receives an acknowledgment from an associated receiver that its transmission is being received. An acknowledgment of transmission must continue to be received by the UWB device at least every 10 s, or it is required to cease transmitting. A device operating as a communication system is characterized by transmission between at least two devices. Noncommunication systems such as imaging systems are required to contain a manually operated switch that causes the transmitter to cease operation within 10 s of being released by the operator. In lieu of a switch located on the imaging system, it is permissible to operate an imaging system by remote control provided the imaging system ceases transmission within 10 s of the remote switch being released by the operator.

F.2.3.2 Mitigation by low-duty-cycle limitations

To permit uses of the low band (3.4 GHz to 4.8 GHz) to low-activity applications for which this band is essential, the European regulator has also defined a mitigation technique called *low duty cycle*. A device implementing low duty cycle is a UWB device as stated under the generic rules that also meets the following requirements:

$$T_{onmax} = 5ms$$

$$T_{offmean} \geq 38ms \text{ (averaged over 1 s)}$$

$$\Sigma T_{off} > 950ms \text{ per second}$$

$$\Sigma T_{on} < 5\% \text{ per second and } 0.5\% \text{ per hour}$$

T_{on} is defined as the duration of a burst irrespective of the number of pulses contained.

T_{off} is defined as the time interval between two consecutive bursts when the UWB emission is kept idle.

UWB devices implementing low duty cycle will be permitted to operate at a level of -41.3 dBm/MHz (instead of $-85/-70$ dBm/MHz) in the 3.4–4.8 GHz frequency band.

F.2.4 Applicable Japanese rules

F.2.4.1 Japanese spectrum mask

The spectrum mask regulated by the Telecommunication Council of Ministry of Internal Affairs and Communications (MIC) is shown in Figure F.1.




Figure F.1—Japanese spectrum mask (only indoor use)

The frequency bands of 3400 MHz through 4800 MHz and 7250 MHz through 10250 MHz are assigned for UWB operation. For 3400 MHz through 4800 MHz, interference mitigation techniques are required. However, for 4200 MHz through 4800 MHz, interference mitigation techniques are not required until the end of December 2008. UWB systems are not allowed to interrupt other radio systems operated in the same band. UWB systems cannot defer the operation of other radio systems.

F.2.4.2 General technical requirements

The general technical requirements of the Japanese rules are as follows:

- a) *UWB definition.* At the maximum radiation frequency, the 10 dB-down bandwidth (B-10) must be larger than 450 MHz, or the fractional bandwidth must be larger than 20%. Moreover, systems using frequency hopping or chirping are regarded as UWB systems as long as their instantaneous bandwidths meet the above UWB bandwidth definition.
- b) *UWB frequency band.* The frequency bands of 3400 MHz through 4800 MHz and 7250 MHz through 10250 MHz are assigned for UWB operation. For 3400 MHz through 4800 MHz, interference mitigation techniques are required. However, for 4200 MHz through 4800 MHz, interference mitigation techniques are not required until the end of December 2008. UWB systems are not allowed to interrupt other radio systems operated in the same band. UWB systems are not allowed to defer the operation of other radio systems.

- c) *Transmit power.* Average power and peak power are defined in Table F.15.

Table F.15—Transmit power

Frequency band (MHz)	Average power (dBm/MHz)	Peak power (dBm/50MHz)
3400–4800 ^a	< -41.3	< 0
7250–10250	< -41.3	< 0

^aAverage power and peak power must be -70 dBm/MHz and -64 dBm/MHz, respectively, if interference mitigation techniques are not installed. However, this is not applied for 4200 MHz through 4800 MHz until the end of December 2008.

- d) *Antenna gain.* Antenna gain is required to be smaller than 0 dBi. However, if the EIRP is below the power limit given in Table F.15, a large antenna gain can be used to reach the limit.
- e) *Transceiver and modulation.* Transceiver can be simplex, full-duplex or semi-duplex. There is no restriction on modulation types.
- f) *Spread bandwidth.* At the maximum radiation frequency, the 10 dB-down bandwidth (B-10) must be larger than 450 MHz, or the fractional bandwidth must be larger than 20%.
- g) *Data rate.* The data rate must be over 50 Mb/s. However, a lower data rate is permitted when the purpose of using the lower data rate is for interference avoidance from noise and noise-like sources to maintain quality of service.
- h) *Communication control.*
- 1) A UWB device is required to first detect the identification signals from other neighboring UWB devices before sending a new signal.
 - 2) A UWB device can send its own identification signal without detecting identification signals from other neighboring UWB devices.
- i) *Interference avoidance function.* Functions of automatically sending and receiving identification signals are required. UWB devices are required to be operated under the condition of no disruption or interference to other radio systems.
- j) *Communications between devices within the same terminal.*
- 1) Each UWB device must have an identification code with a length larger than 48 bits.
 - 2) Except in some special cases, a device is required to first do channel assessment. It will set up a link only when the channel is free.
- k) *Operation limitation.* Operation is limited to indoor. This should be guaranteed by the following methods:
- 1) A network coordinator device is required to be connected to an ac power supply. Other node devices are controlled by the network coordinator and are not necessarily connected to ac supplies.
 - 2) There must be a clear and easy-to-see note attached to the UWB device to remind the user of the indoor operation limitation.
- l) *Measures against illegal rebuild.* The devices are required to be built robust and be difficult to dismantle.
- m) *Electromagnetic compatibility (EMC) to medical devices.* The inter-EMC interference among UWB devices and electronic medical devices must be soundly taken into consideration.

F.2.4.3 Technical requirements on radio equipment

The technical requirements on radio equipment are as follows:

- a) *Transmitters.*
 - 1) *Occupied bandwidth.* For the purpose of conformity with the existing law, permitted occupied bandwidth instead of UWB transmission bandwidth is defined as follows for UWB transmitters:
 - i) It is required to be smaller than 1400 MHz for the 3400–4800 MHz frequency band.²⁴
 - ii) It is required to be smaller than 3000 MHz for the 7250–10250 MHz frequency band.
 - 2) *Unwanted emission level.* The unwanted emission level is given in Table F.16.

Table F.16—Permitted level of unwanted emission

Frequency (MHz)	Average power (dBm/MHz)	Peak power (dBm/MHz)
<1600	–90.0	–84.0
1600–2700	–85.0	–79.0
>2700	–70.0	–64.0
10 600–10 700 11 700–12 750	–85.0	–79.0

- 3) *Reference bandwidth.* The reference bandwidth for permitted unwanted emission is 1 MHz.
 - 4) *Deviation of transmit power.* The deviation within +20% is permitted for the transmit power.
 - 5) *Emission from chassis.* EIRP is required to be smaller than the permitted unwanted emission.
- b) *Receiver.* For 3400–4800 MHz and 7250–10 250 MHz frequency bands, the permitted unwanted emission is required to be smaller than 4 nW/MHz (–54 dBm/MHz). For other frequency bands, it is required to be below the level determined in Table F.16. This value is reduced to –70 dBm/MHz for the 3400–4800 MHz frequency band if interference mitigation techniques are not installed.
- However, the permitted unwanted emission may be smaller than 4 nW/MHz (–54 dBm/MHz) for the 4200–4800 MHz frequency band without interference mitigation techniques until the end of December 2008.

²⁴Interference mitigation techniques are required. However, this requirement is not applied for the 4200–4800 MHz frequency band until the end of December 2008. For this frequency band, the occupied bandwidth is required to be smaller than 600 MHz.

Insert after Annex G the following new annex (Annex H):

Annex H

(informative)

UWB PHY optional chaotic pulses

Another noncoherent optional pulse shape that may be used is a chaotic waveform. This optional pulse shape shall be used only when all other devices within the PAN are using a chaotic pulse. This mode can be used for low-power applications where long battery life is critically important. Since chaotic on-off keying (COOK) is noncoherent modulation, the receiver does not need to generate a corresponding chaotic signal for demodulation. For that reason, the technique chosen for generating a chaotic waveform can be freely determined by implementers.

When transmitting preamble fields, the duration of the chaotic pulses shall be the inverse of the mandatory peak PRF of the mandatory preambles ($T_c = 1/(2 \times \text{mean PRF})$), where the mean PRF is defined in Table 39a. In other words, a single chip of the ternary S code of the mandatory length 31 as shown in Table 39c shall be mapped into a noncoherent COOK pulse $c(t)$. Since COOK is noncoherent modulation, the receiver is not required to distinguish the sign of the chaotic pulse $c(t)$. This relationship is shown in Table H.1.

Table H.1—Ternary-code-to-COOK-pulse mapping

Ternary code chip	COOK chaotic pulse
+	$c(t)$
0	0
–	$c(t)$

After transmitting the preamble, the duration of the chaotic pulse shall be equal to the duration of pulse burst T_{burst} , which is shown in Figure 27c and Table 39a. Since the chaotic pulse is noncoherent modulation, as shown in Figure 27h, the receiver demodulates only the position of pulse. In other words, because pulse position modulation with a two-position symbol (2-PPM) is being used, position 0 or position 1 of the chaotic pulse can be distinguished by the receiver, but not the sign of the chaotic pulse $c(t)$. This relationship is shown in Table H.2.

Table H.2—UWB PHY bits-to-symbol mapping

Modulation symbols (g1g0)	Chaotic pulse COOK
–10	$c(t)$ PPM Position 0
–11	$c(t)$ PPM Position 1
10	$c(t)$ PPM Position 0
11	$c(t)$ PPM Position 1

Insert after Annex H the following new annex (Annex I):

Annex I

(informative)

Example UWB PHY transmit data frame encoding

I.1 Channel used in the example

This annex provides one example of how the PHY would encode a short sample PSDU received from the MAC sublayer. In the example, the PHY transmits at 850 kb/s on channel 3 using preamble code index 6. The annex shows how the data are changed by the PHY encoding steps that eventually lead to bursts of pulses for transmission.

I.2 Encoding progression

I.2.1 Transmit PSDU

In this example, the MAC sublayer presents a sample PSDU to the PHY via the PHY service access point (SAP). The sample that will be encoded is the following 17-octet PSDU:

UWB welcomes IEEE

converted from ASCII to decimal as follows:

85 87 66 32 119 101 108 99 111 109 101 115 32 73 69 69 69

which, in hexadecimal, is as follows:

55 57 42 20 77 65 6C 63 6F 6D 65 73 20 49 45 45 45

Note that the MAC sublayer would not usually present such a PSDU to the PHY because it does not contain a valid cyclic redundancy check (CRC). The purpose of this annex is to provide an example to help the reader to understand the PHY encoding process, not the MAC sublayer.

I.2.2 PSDU bits

For the rest of this annex, the ternary +, -, 0 notation will be used where + represents a one, 0 represents a zero, and - represents a -1.

The sample PSDU is converted to binary, LSB first and starting with the first bit in time first as follows:

+0+0+0+0+++0+0+00+0000+000000+00+++0+++0+0+00++000++0++0+++000++0++++0+++0+++0+++0+0+0+++0+0+0+++000000+00+00+00+0+0+000+0+0+000+0+0+000+0

1.2.3 Reed-Solomon encoded bits

Reed-Solomon encoding is then applied to these bits as described in 6.8a.9.1 to give the following bits:

+0+0+0+0+++0+0+00+0000+000000+00+++0+++0+0+00++000++0+0+++000++0++++0+++0+++0+++0+
0+00++0+++000000+00+00+00+0+0+000+0+0+000+0+0+000+000++0+++000+++00+0++++0+00+
++++++00+++0+++00+0+00

1.2.4 Convolutional encoder input bits

The next step is to prepend the PHR to this Reed-Solomon encoded data for input to the convolutional coder. Because of the data rate and length of PSDU, the PHR is as follows in Table I.1.

Table I.1—Example PHR

Rk1	Rk0	L6	L5	L4	L3	L2	L1	L0	Rng	Ext	P+1	P+0	C5	C4	C3	C2	C1	C0
0	1	0	0	1	0	0	0	1	0	0	0	1	1	1	0	0	1	1

Prepending this to the PSDU gives the following bit stream:

0+00+000+000+++00+++0+0+0+0+++0+0+00+0000+000000+00+++0+++0+0+00++000++0+++0+++000+
+0++++0+++0+0+0+0+00++0+++000000+00+00+00+0+0+000+0+0+000+0+0+000+000++0++
000+++00+0+++0+00++++++00+++0+++00+0+00

1.2.5 Convolutional encoder output bits

Two tail bits are appended to the input bits, and they are input to the convolutional encoder. The convolutional encoder produces two types of output bit: position encoding bits, denoted g_0 , and sign or parity bits, denoted g_1 .

For this example, the position bits g_0 will be simply a delayed version of the above with a zero tail bit or as follows:

00+00+000+000+++00+++0+0+0+0+++0+0+00+0000+000000+00+++0+++0+0+00++000++0+++0+++000+
++0++++0+++0+0+0+0+00++0+++000000+00+00+00+0+0+000+0+0+000+0+0+000+000++0+
+000+++00+0+++0+00++++++00+++0+++00+0+000

The sign bits g_1 will be as follows:

0+0++0+0+0+0+0++++0+000000+0+0000++0+00+0+0000+0+++0+0+0+0000++++0+++0+++0+++0+
++0+00+0+++000++0+++0000++++0++++0+++0000+0+++0+++00000+0+00000+0+00000+0+0+0+++0+
++0++0+++00+00+00+++00000++++0+0+0+++000+00

1.2.6 Scrambler output bits

The scrambler is now enlisted to help with the encoding process. For this complex channel, the scrambler initializer, starting at s_{-15} and going to s_{-1} , is as follows:

1 1 1 0 0 0 1 0 1 1 0 1 1 0 1

Table I.2—Example ternary output symbols 1

Symbol number	Burst Chip Position	Burst	Symbol number	Burst Chip Position	Burst
0	64	+++++-----+-----	51	16	+++++-----+-----
1	48	+++++-----+-----	52	288	+++++-----+-----
2	368	-----+-----+-----	53	368	-----+-----+-----
3	96	+++++-----+-----	54	320	-----+-----+-----
4	0	-----+-----+-----	55	48	-----+-----+-----
5	288	+++++-----+-----	56	336	+++++-----+-----
6	112	+++++-----+-----	57	336	+++++-----+-----
7	32	+++++-----+-----	58	368	+++++-----+-----
8	80	+++++-----+-----	59	0	+++++-----+-----
9	272	+++++-----+-----	60	320	+++++-----+-----
10	64	-----+-----+-----	61	16	+++++-----+-----
11	112	-----+-----+-----	62	256	+++++-----+-----
12	96	+++++-----+-----	63	64	-----+-----+-----
13	352	+++++-----+-----	64	80	+++++-----+-----
14	288	+++++-----+-----	65	304	+++++-----+-----
15	368	+++++-----+-----	66	304	+++++-----+-----
16	64	-----+-----+-----	67	80	+++++-----+-----
17	16	+++++-----+-----	68	48	+++++-----+-----
18	352	+++++-----+-----	69	112	+++++-----+-----
19	288	+++++-----+-----	70	352	+++++-----+-----
20	336	+++++-----+-----	71	256	+++++-----+-----
21	80	+++++-----+-----	72	32	+++++-----+-----
22	368	-----+-----+-----	73	368	+++++-----+-----
23	64	+++++-----+-----	74	320	-----+-----+-----
24	336	+++++-----+-----	75	80	+++++-----+-----
25	48	+++++-----+-----	76	272	+++++-----+-----
26	272	+++++-----+-----	77	320	+++++-----+-----
27	0	+++++-----+-----	78	16	+++++-----+-----
28	320	+++++-----+-----	79	96	+++++-----+-----
29	336	+++++-----+-----	80	0	+++++-----+-----
30	304	+++++-----+-----	81	352	+++++-----+-----
31	80	+++++-----+-----	82	352	+++++-----+-----
32	368	+++++-----+-----	83	64	+++++-----+-----
33	64	+++++-----+-----	84	272	+++++-----+-----
34	368	-----+-----+-----	85	352	+++++-----+-----
35	0	-----+-----+-----	86	288	+++++-----+-----
36	32	+++++-----+-----	87	336	+++++-----+-----
37	304	+++++-----+-----	88	48	+++++-----+-----
38	16	+++++-----+-----	89	272	+++++-----+-----
39	96	+++++-----+-----	90	320	+++++-----+-----
40	32	+++++-----+-----	91	80	+++++-----+-----
41	16	+++++-----+-----	92	336	+++++-----+-----
42	320	+++++-----+-----	93	112	+++++-----+-----
43	16	+++++-----+-----	94	352	+++++-----+-----
44	96	+++++-----+-----	95	352	+++++-----+-----
45	96	+++++-----+-----	96	0	+++++-----+-----
46	32	+++++-----+-----	97	288	+++++-----+-----
47	48	+++++-----+-----	98	336	+++++-----+-----
48	48	+++++-----+-----	99	112	+++++-----+-----
49	368	+++++-----+-----	100	320	+++++-----+-----
50	32	+++++-----+-----	101	112	+++++-----+-----

Table I.3—Example ternary output symbols 2

Symbol number	Burst Chip Position	Burst	Symbol number	Burst Chip Position	Burst
102	352	+-----+-----+	153	80	+-----+-----+
103	32	------+-----+	154	304	------+-----+
104	80	+-----+-----+	155	16	+-----+-----+
105	272	+-----+-----+	156	64	+-----+-----+
106	256	------+-----+	157	80	+-----+-----+
107	32	+-----+-----+	158	304	+-----+-----+
108	272	+-----+-----+	159	336	+-----+-----+
109	256	------+-----+	160	16	+-----+-----+
110	96	------+-----+	161	352	------+-----+
111	32	+-----+-----+	162	352	------+-----+
112	368	+-----+-----+	163	0	------+-----+
113	288	+-----+-----+	164	64	+-----+-----+
114	304	+-----+-----+	165	16	+-----+-----+
115	48	+-----+-----+	166	352	------+-----+
116	16	------+-----+	167	256	+-----+-----+
117	64	+-----+-----+	168	352	------+-----+
118	112	------+-----+	169	32	------+-----+
119	64	+-----+-----+	170	80	+-----+-----+
120	16	+-----+-----+	171	272	------+-----+
121	288	+-----+-----+	172	96	+-----+-----+
122	48	+-----+-----+	173	352	------+-----+
123	16	+-----+-----+	174	320	+-----+-----+
124	352	+-----+-----+	175	368	------+-----+
125	64	------+-----+	176	256	------+-----+
126	80	+-----+-----+	177	96	+-----+-----+
127	336	------+-----+	178	320	+-----+-----+
128	16	+-----+-----+	179	80	+-----+-----+
129	0	------+-----+	180	112	+-----+-----+
130	288	+-----+-----+	181	352	------+-----+
131	80	------+-----+	182	352	+-----+-----+
132	304	------+-----+	183	320	+-----+-----+
133	48	------+-----+	184	368	------+-----+
134	336	+-----+-----+	185	320	+-----+-----+
135	48	+-----+-----+	186	368	------+-----+
136	16	------+-----+	187	256	------+-----+
137	32	------+-----+	188	32	------+-----+
138	368	------+-----+	189	48	+-----+-----+
139	32	+-----+-----+	190	368	+-----+-----+
140	272	+-----+-----+	191	352	+-----+-----+
141	0	+-----+-----+	192	288	------+-----+
142	320	+-----+-----+	193	16	------+-----+
143	16	+-----+-----+	194	288	------+-----+
144	32	+-----+-----+	195	272	------+-----+
145	112	+-----+-----+	196	256	------+-----+
146	352	+-----+-----+	197	32	------+-----+
147	0	+-----+-----+	198	48	+-----+-----+
148	352	+-----+-----+	199	304	------+-----+
149	96	+-----+-----+	200	48	------+-----+
150	288	+-----+-----+	201	368	------+-----+
151	48	+-----+-----+	202	32	------+-----+
152	80	------+-----+	203	16	------+-----+
			204	64	+-----+-----+