

IEEE Standard for
Local and metropolitan area networks—

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

Amendment 1: MAC sublayer

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

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

IEEE Std 802.15.4e™-2012
(Amendment to
IEEE Std 802.15.4™-2011)

16 April 2012

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

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

Amendment 1: MAC sublayer

Sponsor

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

Approved 6 February 2012

IEEE-SA Standards Board

Abstract: IEEE Std 802.15.4-2011 is amended by this standard. The intention of this amendment is to enhance and add functionality to the IEEE 802.15.4 MAC to (a) better support the industrial markets and (b) permit compatibility with modifications being proposed within the Chinese WPAN.

Keywords: asynchronous multi-channel adaptation (AMCA), deterministic and synchronous multi-channel extension (DSME), enhanced beacon request (EBR), enhancements for security, fast association (FastA), IEEE 802.15.4, IEEE 802.15.4e, low energy, low latency deterministic networks (LLDN), MAC performance metrics (metrics), RFID Blink, timeslotted channel hopping (TSCH)

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

Copyright © 2012 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 16 April 2012. Printed in the United States of America.

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

The Bluetooth trademark is owned by Bluetooth SIG, Inc. USA.

WirelessHART is a registered trademark of the HART Communication Foundation.

ZigBee is a registered trademark of the ZigBee Alliance.

Print: ISBN 978-0-7381-7204-0 STD97211
PDF: ISBN 978-0-7381-7225-5 STDPD97211

IEEE prohibits discrimination, harassment and bullying. For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.
No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.

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

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

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

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

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

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

Official Statements: A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered the official position of IEEE or any of its committees and shall not be considered to be, nor be relied upon as, a formal position of IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position of IEEE.

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

Comments on standards should be submitted to the following address:

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

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

Notice to users

Laws and regulations

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

Copyrights

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

Updating of IEEE documents

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

Errata

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

Patents

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

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

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

Seong-Soon Joo, *Task Group 4e Chair*
Patrick W. Kinney, *Task Group 4e past Chair and Secretary*
Ludwig Winkel, *Task Group 4e Chief-Editor*
Gahng-Seop Ahn, *Task Group 4e Co-Editor*
Jonathan Simon, *Task Group 4e Co-Editor*
Benjamin Rolfe, *Task Group 4e Co-Editor*
Wun-Cheol Jeong, *Task Group 4e Co-Editor*
Liang Li, *Task Group 4e Co-Editor*

Emad Afifi
Roberto Aiello
Arthur Astrin
Taehan Bae
Michael Bahr
John Barr
Anuj Batra
Tuncer Baykas
Philip Beecher
Ashutosh Bhatia
Ghulam Bhatti
Gary Birk
Mathew Boytim
Peter David Bradley
Nancy Bravin
David Britz
Monique Brown
Sverre Brubk
Brian Buchanan
John Buffington
Kiran Bynam
Brent Cain
Edgar Callaway
Chris Calvert
Ruben Cardozo
Douglas Castor
Jaesang Cha
Russell Chandler
Soo-Young Chang
Kuor-Hsin Chang
Hind Chebbo
Chang-Soon Choi
Sangsung Choi
Ciaran Connell
David Cypher
Matthew Dahl
David Davenport

Mark Dawkins
Hendricus De Ruijter
Upkar Dhaliwal
Gang Ding
Paul Dixon
Guido Dolmans
Igor Dotlic
Michael Dow
Dietmar Eggert
David Evans
Charles Farlow
John Farserotu
Jeffrey Fischbeck
Mike Fischer
George Flammer
Ryosuke Fujiwara
Noriyasu Fukatsu
Kiyoshi Fukui
John Geiger
Gregory Gillooly
Tim Godfrey
Paul Gorday
Elad Gottlib
Robert Hall
Shinsuke Hara
Hiroshi Harada
Timothy Harrington
Rodney Hemminger
Marco Hernandez
Garth Hillman
Jin-Meng Ho
Wei Hong
Srinath Hosur
David Howard
Taeho Hwang
Jung-Hwan Hwang
Ichirou Ida

Tetsushi Ikegami
Akio Iso
Yeong Min Jang
Adrian Jennings
Steven Jillings
Tae Kang
Noh-Gyoung Kang
Tae-Gyu Kang
Tatsuya Kato
Shuzo Kato
Jeritt Kent
Prithpal Khakuria
Dae Ho Kim
Dong-Sun Kim
Dukhyun Kim
Jaehwan Kim
Jeffrey King
Ryuji Kohno
Fumihide Kojima
Bruce Kraemer
Raymond Krasinski
Masahiro Kuroda
John Lampe
Zhou Lan
Khanh Le
Cheolhyo Lee
Hyungsoo Lee
Myung Lee
Daniel Lewis
Huan-Bang Li
Sang-Kyu Lim
Jeremy Link
Michael Lynch
Robert Mason
Tomokuni Matsumura
Jeff McCullough
Michael McGillan
Michael McInnis

Michael McLaughlin
 Charles Millet
 Siamak Mirnezami
 Rishi Mohindra
 Emmanuel Monnerie
 Rajendra Moorti
 Robert Moskowitz
 Hamilton Moy
 Peter Murray
 Theodore Myers
 Chiu Ngo
 Paul Nikolich
 Hirohito Nishiyama
 David Olson
 Okundu Omeni
 Ryoji Ono
 Laurent Ouvry
 James Pace
 Jahng Park
 Seung-Hoon Park
 Taejoon Park
 Hyung-Il Park
 Ranjeet Patro
 Al Petrick
 Dalibor Pokrajac
 Daniel Popa
 Stephen Pope
 Richard Powell
 Clinton Powell
 Chang-Woo Pyo
 Mohammad Azizur Rahman

Sridhar Rajagopal
 Jayaram Ramasastry
 Marc Reed
 Ivan Reede
 Richard Roberts
 Craig Rodine
 June Chul Roh
 Seung-Moon Ryu
 Didier Sagan
 Kentaro Sakamoto
 Filippo Will San
 H. Sanderford
 Kamran Sayrafian
 Timothy Schmidl
 Michael Schmidt
 Jean Schwoerer
 Cristina Seibert
 Neal Seidl
 Kunal Shah
 Steve Shearer
 Stephen Shellhammer
 Shusaku Shimada
 CheolHo Shin
 Chang Sub Shin
 Michael Sim
 Jaeseung Son
 Paul Stadnik
 René Struik
 Chin-Sean Sum
 Hui-Hsia Sung
 Gu Sungi
 Kenichi Takizawa

Hirokazu Tanaka
 Larry Taylor
 Mark Thompson
 James Tomcik
 Ichihiko Toyoda
 David Tracey
 Khanh Tran
 Jerry Upton
 Jana van Greunen
 Hartman van Wyk
 Michel Veillette
 Billy Verso
 Bhupender Virk
 Joachim Walewski
 Junyi Wang
 Quan Wang
 Xiang Wang
 Andy Ward
 Scott Weikel
 Nicholas West
 Mark Wilbur
 Eun Tae Won
 Alan Chi Wai Wong
 Tao Xing
 Yang Yang
 Wen-Bin Yang
 Kazuyuki Yasukawa
 Kamya Yazdandoost
 Kaoru Yokoo
 Mu Zhao
 Bin Zhen

Major contributions were received from the following individuals:

Micheal Bahr
 Ghulam Bhatti
 Kuor-Hsin Chang
 Tim Godfrey
 Hiroshi Harada
 Wei Hong

Chol Su Kang
 Jeffrey King
 Fumihide Kojima
 Myung Lee
 HaiTao Liu
 Tae Rim Park
 Kentaro Sakamoto

Chang Sub Shin
 René Struik
 Tao Xing
 Yang Yang
 Kazuyuki Yasukawa
 Mu Zhao

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

Jon Adams	Ichirou Ida	Ivan Reede
Emad Afifi	Tetsushi Ikegami	Maximilian Riegel
Rick Alfvin	Noriyuki Ikeuchi	Robert Robinson
Butch Anton	Akio Iso	Craig Rodine
Lee Armstrong	Atsushi Ito	Benjamin Rolfe
Takamasa Asano	Raj Jain	Carl Royer
Arthur Astrin	Oyvind Janbu	Randall Safier
Michael Bahr	Junghoon Jee	Kentaro Sakamoto
Jonathan Bailey	Jorjeta Jetcheva	Kazuyuki Sakoda
Edward Baranski	Tal Kaitz	Shigenobu Sasaki
Tuncer Baykas	Shinkyō Kaku	Katsuyoshi Sato
Philip Beecher	Ruediger Kays	Bartien Sayogo
Harry Bims	Raymond H. Kelley	Timothy Schmidl
Mathew Boytim	Jeritt Kent	Cristina Seibert
Nancy Bravin	Stuart J. Kerry	Kunal Shah
Vern Brethour	Yongbum Kim	Jie Shen
Monique Brown	Jeffrey King	Shusaku Shimada
Sverre Brubk	Patrick W. Kinney	Gil Shultz
John Buffington	Fumihide Kojima	Jonathan Simon
William Byrd	Bruce Kraemer	Kapil Sood
Edgar Callaway	Yasushi Kudoh	Amjad Soomro
Radhakrishna Canchi	Jeremy Landt	Robert Soranno
Dave Cavalcanti	Khanh Tuan Le	Thomas Starai
Kuor-Hsin Chang	David G. K. Li	René Struik
Clint Chaplin	Jan-Ray Liao	Walter Struppler
Keith Chow	Arthur Light	Mark Sturza
Charles Cook	Lu Liru	Chin-Sean Sum
Joseph Decuir	HaiTao Liu	Chen Sun
Patrick Diamond	William Lumpkins	Larry Taylor
Thomas Dineen	Greg Luri	Ichihiko Toyoda
Sourav Dutta	Michael Lynch	Ha Nguyen Tran
Peter Ecclesine	Elvis Maculuba	Anna Urra
Richard Edgar	Stephen Makgill	Jana van Greunen
John Egan	Robert Mason	Dmitri Varsanofiev
Stanislav Filin	Michael McInnis	Prabodh Varshney
Will San Filippo	Paul Mobus	John Vergis
Kimberly Francisco	Apurva Mody	Billy Verso
James P. K. Gilb	Emmanuel Monnerie	Dalton Victor
Gregory Gillooly	Kenichi Mori	Bhupender Virk
Tim Godfrey	Robert Moskowitz	Jerome Viviano
Randall C. Groves	Peter Murray	George Vlantis
Michael Gundlach	Michael S. Newman	Khurram Waheed
Henrik Gustafsson	Paul Nikolich	Junyi Wang
Jose Gutierrez	John Notor	Stanley Wang
C. Guy	Okundu Omeni	Xiang Wang
Rainer Hach	Masayuki Oodo	Andy Ward
Hiroshi Harada	Chris Osterloh	Hung-Yu Wei
Timothy Harrington	Satoshi Oyama	Scott Weikel
Rodney Hemminger	Kevyn Pair	Mark Wilbur
Marco Hernandez	Dalibor Pokrajac	Ludwig Winkel
Ryoichi Higashi	Daniel Popa	Andreas Wolf
Wei Hong	Clinton Powell	Tao Xing
Srinath Hosur	Venkatesha Prasad	Yang Yang
David Howard	Mohammad Azizur Rahman	Kazuyuki Yasukawa
Heqing Huang	Jayaram Ramasastry	Jianming Yuan
David Hunter	Mark Ranta	Oren Yuen

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

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

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

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

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

*Member Emeritus

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

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

Catherine Berger
IEEE Standards Program Manager, Document Development

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

Introduction

This introduction is not part of IEEE Std 802.15.4e-2012, IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)—Amendment 1: MAC sublayer.

This amendment to IEEE Std 802.15.4-2011 specifies additional media access control (MAC) behaviors and frame formats that allow IEEE 802.15.4 devices to support a wide range of industrial and commercial applications that were not adequately supported prior to the release of this amendment. Specifically, these changes to the MAC sublayer facilitate industrial applications such as Process Control and Factory Automation in addition to the MAC behaviors that support the Chinese Wireless Personal Area Network (CWPAN) standard.

Contents

3.	Definitions, acronyms, and abbreviations.....	2
3.1	Definitions	2
3.2	Acronyms and abbreviations	3
4.	General description	5
4.1	General.....	5
4.3	Network topologies.....	5
4.3.1	Star network formation	5
4.3.2	Peer-to-peer network formation.....	5
4.5	Functional overview	5
4.5.1	Superframe structure.....	5
4.5.2	Data transfer model.....	7
4.5.3	Frame structure	8
4.5.4	Improving probability of successful delivery	9
4.5.5	Power consumption considerations	9
4.5.6	Security	10
5.	MAC protocol	11
5.1	MAC functional description	11
5.1.1	Channel access	11
5.1.1a	Channel hopping.....	11
5.1.2	Starting and maintaining PANs	22
5.1.3	Association and disassociation	26
5.1.4	Synchronization	26
5.1.6	Transmission, reception, and acknowledgment.....	30
5.1.9	LLDN transmission states.....	32
5.1.10	Deterministic and synchronous multi-channel extension (DSME)	37
5.1.11	LE transmission, reception and acknowledgment	51
5.1.12	Asynchronous multi-channel adaptation (AMCA).....	57
5.2	MAC frame formats.....	59
5.2.1	General MAC frame format.....	59
5.2.2	Format of individual frame types.....	63
5.2.3	Frame compatibility	79
5.2.4	Information Element	80
5.3	MAC Command frames.....	95
5.3.1	Association request command	95
5.3.2	Association response command.....	97
5.3.7	Beacon request command	97
5.3.10	LL commands	98
5.3.11	DSME commands	104
5.3.12	LE commands	117
5.3.13	AMCA commands	118
6.	MAC services	123
6.2	MAC management service.....	123
6.2.0a	General.....	123
6.2.0b	TSCH MAC management service	123

6.2.0c	LL-MAC management service	123
6.2.0d	DSME MAC management service	123
6.2.1	Common requirements for MLME primitives	124
6.2.2	Association primitives	124
6.2.4	Communications notification primitives	131
6.2.10	Primitives for channel scanning	133
6.2.12	Primitives for updating the superframe configuration	137
6.2.18	Primitives for Beacon Generation	140
6.2.19	Primitives for TSCH	143
6.2.20	Primitives for LLDN	149
6.2.21	Primitives for DSME	153
6.3	MAC data service	167
6.3.1	MCPS-DATA.request	167
6.3.2	MCPS-DATA.confirm	169
6.3.3	MCPS-DATA.indication	170
6.4	MAC constants and PIB attributes	171
6.4.2	MAC PIB attributes	171
6.4.3	Calculating PHY dependent MAC PIB values	172
7.	Security	187
7.3	Security operations	187
7.3.2	CCM* Nonce	187
7.4	Auxiliary security header	187
7.4.1	Security Control field	187
7.4.2	Frame Counter field	188
7.5	Security-related MAC-PIB attributes	188
10.	O-QPSK PHY	189
10.1	PPDU format	189
10.1.3	Frame Length field	189
Annex A (informative)	Bibliography	191
Annex D (informative)	Protocol implementation conformance statement (PICS) proforma	192
D.7	PICS proforma tables	192
D.7.3	Major capabilities for the MAC sublayer	192
D.7.3.1	MAC sublayer functions	192
Annex I (informative)	MAC behaviors for industrial and commercial application domains	195
I.1	General	195
I.2	Timeslotlotted Channel Hopping (TSCH)	196
I.2.1	Typical application domains addressed by the TSCH mode	196
I.2.2	TSCH mode overview	196
I.2.3	MAC behaviors unique to TSCH	196
I.3	Low Latency Deterministic Networks	197
I.3.1	Typical application domains and requirements addressed by LLDNs	197
I.3.2	LLDN application domain overview	197
I.3.3	LLDN assumptions	198
I.3.4	LLDN mode characteristics	198
I.3.4.1	General	198

	I.3.4.2	Time Division Multiple Access (TDMA)	199
	I.3.4.3	Addressing.....	199
	I.3.4.4	Network topology.....	199
	I.3.5	MAC behaviors unique to the LLDN mode.....	200
I.4		Deterministic and synchronous multi channel extension (DSME).....	200
	I.4.1	Typical application domains and requirements.....	200
	I.4.1.1	Features	200
	I.4.2	Recommended DSME parameter settings for various application types	201
	I.4.3	MAC behaviors unique to the DSME mode	201
	I.4.4	Channel diversity.....	202
	I.4.4.1	Overview	202
	I.4.4.2	Channel diversity methods.....	202
	I.4.4.3	MAC behaviors unique to Channel Diversity	204
I.5		Blink.....	204
	I.5.1	Blink mode overview	204
	I.5.2	MAC behaviors unique to the Blink mode.....	204
I.6		Asymmetric multi-channel adaptation (AMCA)	204
	I.6.1	AMCA mode overview	204
	I.6.2	MAC behaviors unique to AMCA	205
I.7		Low energy	205
	I.7.1	LE overview	205
	I.7.2	MAC behaviors unique to LE	205
I.8		Information elements	206
	I.8.1	Overview	206
	I.8.2	IE types.....	206
	I.8.3	IE Identification.....	206
	I.8.4	IE Construction	206
	I.8.5	MAC behaviors unique to IE	206
I.9		Enhanced beacons (EB) and enhanced beacon requests (EBR)	206
	I.9.1	EB and EBR overview	206
	I.9.2	MAC behaviors unique to EB and EBR.....	207
I.10		Multipurpose frame.....	207
	I.10.1	Multipurpose frame overview	207
	I.10.2	MAC behaviors unique to the multipurpose frame	207
I.11		MAC performance metrics	208
	I.11.1	Requirements of routing (such as in mesh networks)	208
	I.11.2	Requirements of long transmit durations due to large frames or low data rates	208
	I.11.3	Metrics.....	208
	I.11.4	MAC behaviors unique to MAC performance metrics	208
I.12		Fast association	208
	I.12.1	Overview	208
	I.12.2	Operation	209
	I.12.3	MAC behaviors unique to FastA.....	209

IEEE Standard for Local and metropolitan area networks—

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

Amendment 1: MAC sublayer

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

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

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.¹

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. Editing instructions, change markings, and this NOTE will not be carried over into future editions because the changes will be incorporated into the base standard.

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

3. Definitions, acronyms, and abbreviations

3.1 Definitions

Insert in alphabetical order the following definitions:

blink: A multipurpose frame with a one-octet Frame Control field used by unassociated devices.

channel hopping: Periodically switching the channel using a sequence known to both sending and receiving devices where the entire frame is sent on a single channel.

channel offset: A number used in the channel calculation of a slotted channel hopping system to allow for different channels to be used in the same slot.

coordinated sampled listening (CSL): A low-energy mode of media access control specifying how receiving devices periodically monitor the channel(s) for incoming transmissions at low duty cycles.

coordinated sampled listening channel sample: The operation of performing energy detection on a channel and attempting to receive a wake-up frame if energy is detected.

coordinated sampled listening payload frame: A beacon, data, or command frame carrying the application payload following the wake-up sequence in a coordinated sampled listening period.

coordinated sampled listening period: The time interval in which receiving devices monitor the channel(s) for incoming transmissions.

coordinated sampled listening phase: The length of time between a given time and the next channel sample, as a fraction of the coordinated sampled listening period.

coordinated sampled listening rendezvous time: A timestamp in the coordinated sampled listening wake-up frame payload indicating the expected length of time in milliseconds between the end of the wake-up frame transmission and the beginning of the payload frame transmission.

coordinated sampled listening wake-up frame: A special short frame repeatedly transmitted before the payload frame to ensure the reception of the payload frame by a coordinated sampled listening receiving device.

coordinated sampled listening wake-up frame sequence: A sequence of back-to-back wake-up frames up to the duration of the coordinated sampled listening period.

downlink: Data communication from the personal area network (PAN) coordinator to the PAN device.

encapsulated service data unit: A unit of data that has been passed down from a higher layer that has been encapsulated.

enhanced active scan: An active scan using one or more enhanced beacon requests.

enhanced beacon (EB): A beacon that contains information elements (IEs) (rather than a superframe specification) to describe various properties of the network.

enhanced beacon request (EBR): A beacon request with optional filter elements to limit the number of coordinator responses.

information element (IE): A formatted data object consisting of an ID, a length, and a data payload used to pass data between layers or devices.

low latency deterministic network (LLDN): A personal area network (PAN) organized as a star network with a superframe structure and using LLDN frames.

low latency deterministic network (LLDN) device: A device in an LLDN that is associated to an LLDN coordinator.

receiver initiated transmission (RIT): An alternative low-energy mode to coordinated sampled listening in which receiving devices periodically broadcast data request frames, and transmitting devices only transmit to a receiving device upon receiving a data request frame.

slotframe: A collection of timeslots repeating in time, analogous to a superframe in that it defines periods of communication opportunities.

slot owner: A low latency deterministic network (LLDN) device that is assigned exclusive access rights at the beginning of a timeslot in an LLDN.

timeslot: A defined period of time during which a frame and an acknowledgement may be exchanged between devices.

uplink: data communication from the personal area network (PAN) device to the PAN coordinator.

3.2 Acronyms and abbreviations

Insert in 3.2 the following abbreviations and acronyms in alphabetical order:

ACK	positive acknowledgment
AMCA	asymmetric multi-channel adaptation
ASN	absolute slot number
CSL	coordinated sampled listening
CTS	clear to send
DSME	deterministic and synchronous multi-channel extension
EB	enhanced beacon
EBR	enhanced beacon request
ECFP	extended contention-free period
ESDU	encapsulated service data unit
ESOR	enhanced security and overhead reduction
EUI	extended unique identifier
FastA	fast association

GACK	group acknowledgment
IE	information element
LE	low energy
LL	low latency
LLDN	low latency deterministic network
NACK	negative acknowledgment
NHL	next higher layer
PA	process automation
RIT	receiver initiated transmission
RTS	request to send
SAB	slot allocation bitmap
SD	superframe duration
TSCH	timeslotted channel hopping
WLAN	wireless local area network

4. General description

4.1 General

Insert before 4.2 the following text:

Supplemental information for different industrial domains is contained in Annex I.

4.3 Network topologies

4.3.1 Star network formation

Insert before 4.3.2 the following paragraphs:

A low latency deterministic network (LLDN) operates in a star topology. More information on the star topology of LLDNs is given in Annex I.

Timeslotted channel hopping (TSCH) PANs are topology independent and can be used in star topologies as well as partial or full mesh topologies.

4.3.2 Peer-to-peer network formation

Insert before 4.4 the following paragraphs:

Deterministic and synchronous multi-channel extension (DSME) supports both star and peer-to-peer network topologies.

TSCH devices support peer-to-peer topologies.

4.5 Functional overview

4.5.1 Superframe structure

Insert the following before the existing text:

4.5.1.1 General

This standard allows the optional use of different superframe structures, as follows:

- Superframe structure described in 4.5.1.1 (as defined in IEEE Std 802.15.4™-2011).
- Superframe structure described in 4.5.1.2 based on beacons defined in 5.2.2.1 (Enhanced Beacon) with an information element (IE) defined in 5.2.4.9 (DSME PAN Descriptor).
- Superframe structure described in 4.5.1.3 based on LL-Beacons defined in 5.2.2.5.2.

A Slotframe structure, as described in 4.5.1.4, takes the place of superframes in a TSCH PAN.

Delete the following existing text:

~~This standard allows the optional use of a superframe structure.~~

Insert before 4.5.2 the following subclauses and figures:

4.5.1.2 DSME multi-superframe structure

The DSME-enabled (i.e., *macDSMEenabled* is TRUE) PANs use the DSME multi-superframe structure. The format of the DSME multi-superframe structure is defined by coordinators that periodically transmit an Enhanced Beacon (EB) with the DSME PAN descriptor IE. A multi-superframe is a cycle of repeated superframes, each of which consists of a beacon frame, a Contention Access Period (CAP) and a Contention-Free Period (CFP). An example of a multi-superframe structure is shown in Figure 4a.




Figure 4a—General DSME Multi-superframe Structure

The single common channel, the logical channel number used in the successful association, is used to transmit the enhanced beacon frames and the frames transmitted during the CAP. The channel diversity method, either channel adaptation or channel hopping as selected by *macChannelDiversityMode*, is applied to transmit data frames and acknowledge frames during the CFP. Beacons are scheduled as defined in 5.1.10.6. Frames sent during the CFP are transmitted using the allocated channel for DSME-GTS. A DSME-GTS can be allocated on any of the available channels in the current ChannelPage.

Details on the DSME Multi-superframes structure are described in 5.1.10.1.

4.5.1.3 Superframe structure based on LL Beacons

LLDN PANs (i.e., *macLLEnabled* is TRUE) use the LLDN superframe structure as described in 5.1.1.6. The superframe is divided into a beacon slot, 0 or 2 management timeslots (i.e., 2 if *macLLDNmgmtTS* is TRUE), and *macLLDNnum TimeSlots* number of timeslots of equal length as shown in Figure 4b.




Figure 4b—LLDN Superframe with dedicated timeslots

The first timeslot of each superframe contains an LL-Beacon frame. The LL-Beacon frame is used for synchronization with the superframe structure. It is also used for re-synchronization of devices that, for instance, went into power save or sleep mode.

The beacon timeslot may be followed by two management timeslots, one for downlink and one for uplink.

The remaining timeslots are assigned to the LLDN devices in the network; there is no explicit addressing necessary inside the frames provided that there is exactly one device assigned to a timeslot as per 5.1.1.6.6. The determination of the sender is achieved through the indexing of timeslots. If there is more than one device assigned to a timeslot, the timeslot is referred to as shared group timeslot, and a simple addressing scheme with 8-bit addresses, *macSimpleAddress*, is used as described in 6.3.

4.5.1.4 Slotframes

In a TSCH PAN, the concept of the superframe is replaced with a slotframe. The slotframe also contains defined periods of communications between peers that may be either CSMA-CA or guaranteed, but the slotframe automatically repeats based on the participating devices' shared notion of time, and does not require beacons to initiate communications. Unlike the superframe, slotframes and a device's assigned timeslot(s) within the slotframe may initially be communicated by beacon, but are typically configured by a higher layer as the device joins the network. Because all devices share common time and channel information, devices may hop over the entire channel space to minimize the negative effects of multipath fading and interference, and do so in a slotted way to avoid collisions, minimizing the need for retransmissions. Both of these features are desirable for operation in a harsh industrial environment.

4.5.2 Data transfer model

4.5.2.1 Data transfer to a coordinator

Insert before 4.5.2.2 the following paragraph and figure:

When a device wishes to transfer data to a PAN coordinator in an LLDN, it first listens for the network beacon. When the beacon is found, the device synchronizes to the superframe structure. At its assigned timeslot, the device transmits its data frame to the LLDN PAN coordinator. If the device transmits its data frame in a dedicated timeslot or as slot owner of a shared group timeslot, the data frame is transmitted without using CSMA-CA. If the device transmits its data frame in a shared group timeslot and is not the slot owner, the data frame is transmitted using slotted CSMA-CA as described in 5.1.1.4.4, or ALOHA described in 4.5.4.3, depending on the used PHY. The LLDN PAN coordinator may acknowledge the successful reception of the data by transmitting an optional acknowledgment frame. Successful data transmissions in dedicated timeslots or by the slot owner are acknowledged by the LLDN PAN coordinator with a Group Acknowledgment either in the next beacon or as a separate group acknowledgment (GACK) frame. This sequence is summarized in Figure 4c.




Figure 4c—Communication to a PAN coordinator in an LLDN

4.5.2.2 Data transfer from a coordinator

Insert before 4.5.2.3 the following paragraphs and figure:

A data transfer from an LLDN PAN coordinator is only possible in the *macLLDNumBidirectionalTS* timeslots described in 5.1.1.6.5 and if the Transmission Direction field in the Flags field of the beacon indicates downlink direction.

When the LLDN PAN coordinator wishes to transfer data to an LLDN device assigned to a bidirectional timeslot in an LLDN, it indicates in the network beacon that the transmission direction is downlink. At the appropriate time, the LLDN PAN coordinator transmits its data frame to the device without using CSMA-CA. The device may acknowledge the successful reception of the data by transmitting an acknowledgment frame to the LLDN PAN coordinator in the same timeslot of the next superframe. In order to do so, the transmission direction has to be uplink in that superframe. This sequence is summarized in Figure 4d.




Figure 4d—Communication from a PAN coordinator to a device in an LLDN

4.5.3 Frame structure

Change in 4.5.3 the first paragraph and dashed list as follows:

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

- A Beacon and Enhanced beacon frame, used by a coordinator to transmit beacons
- A ~~d~~Data frame, used for ~~all~~ transfers of data
- An unsecured and enhanced acknowledgment frame, used for confirming successful frame reception
- A MAC command frame, used for handling ~~all~~ MAC peer entity control transfers
- A multipurpose frame, used for transfers of data where additional MHR options are needed
- A LLDN frame, used to minimize overhead in latency critical networks

Insert at the end of 4.5.3 the following paragraph:

This standard makes use of IEs to transfer formatted data between layers and between devices. IEs consist of an identification (ID) (or tag), a length, and the IE content. Devices can accept or discard a particular element if the ID is known, and skip over unknown ID elements. IDs are categorized into managed spaces, whose content is defined in this standard, and unmanaged spaces, whose content is left to a higher layer to define. As defined in this standard, IEs provide extensibility of the MAC protocol: future versions of the protocol can add functionality by adding new IE types in a backwards compatible way, without redefining the structure of the MAC frame or new frame types. This removes the limitations of fixed length MAC Header fields such as the Frame Version and Addressing Mode fields.

4.5.4 Improving probability of successful delivery

4.5.4.1 CSMA-CA mechanism

Insert before 4.5.4.2 the following:

LLDNs use a slotted CSMA-CA channel access mechanism for management timeslots and shared group timeslots, where the backoff slots are aligned as follows:

- With the start of the beacon transmission in management timeslots
- With tSlotTxOwner in shared group timeslots

Each time a device wishes to transmit data frames with CSMA-CA at the appropriate places, it locates the boundary of the next backoff slot and then waits for a random number of backoff slots. If the channel is busy, following this random backoff, the device waits for another random number of backoff slots before trying to access the channel again. If the channel is idle, the device begins transmitting on the next available backoff slot boundary. Acknowledgment and beacon frames are sent without using a CSMA-CA mechanism.

Devices in a TSCH PAN also use a slotted CSMA-CA scheme for shared slots, but use a slotted backoff scheme as described in 5.1.1.4.3.

4.5.4.3 Frame acknowledgment

Insert in 4.5.4.3 after the first paragraph the following text:

The receiving device may insert additional content in an enhanced acknowledgment encapsulated as IEs. If the originator does not understand the IE content of the acknowledgment, it is ignored, but the transmission is considered successful.

Insert before 4.5.5 (after 4.5.4.4) the following subclause:

4.5.4.5 Asynchronous multi-channel adaptation

The single common channel approach may not be able to connect all devices in the PAN. The variance of channel condition can be large and channel asymmetry between two neighboring device can happen. Asynchronous multi-channel adaptation (AMCA) is a solution to handle such cases. The AMCA is performed in non-beacon mode, refer to 5.1.12 for details.

4.5.5 Power consumption considerations

Insert the following header before the existing text:

4.5.5.1 General

Insert before 4.5.6 the following new subclause:

4.5.5.2 Low-energy mechanisms

Two low-energy mechanisms are provided to further reduce energy consumption by allowing devices to communicate while maintaining low duty cycles. They are coordinated sampled listening (CSL) and receiver initiated transmission (RIT).

CSL allows receiving devices to periodically sample the channel(s) for incoming transmissions at low duty cycles. The receiving device and the transmitting device are coordinated to reduce transmitting overhead.

RIT allows receiving devices to periodically broadcast data request frames, and transmitting devices only transmit to a receiving device upon receiving a data request frame. RIT is suitable for the following application scenarios:

- Low data traffic rate and loose latency requirement, where a few seconds of latency is allowable by application.
- Local regulations restricting the duration of continuous radio transmissions (e.g., 950 MHz band in Japan).

4.5.6 Security

Insert before the last paragraph the following paragraph:

Implementations should include mechanisms to prevent unauthorized access to locally stored keys.

5. MAC protocol

5.1 MAC functional description

5.1.1 Channel access

Insert before 5.1.1.1 the following subclause (5.1.1a) and figure:

5.1.1a Channel hopping

Networks may support channel hopping where permitted or when required using any multichannel PHY. Hopping sequences are defined in the Hopping Sequence PIB (as defined in Table 52f) and may be of arbitrary length up to 511 elements, and cover all or a subset of channels available to the PHY. Devices may hop in a slotted mode (e.g., TSCH or DSME) or in an unslotted mode.

The Hopping Sequence PIB includes an ID (*macHoppingSequenceID*), fields describing operation of the PHY, the length of the sequence (*macHoppingSequenceLength*), an ordered list of channels (*macHoppingSequenceList*), and the dwell time for unslotted hopping, if applicable.

All hopping sequences are referred to by an ID, *macHoppingSequenceID*, with ID = 0 denoting the default sequence for a particular PHY (or PHY configuration if the PHY supports more than one channel list).

This default sequence is a pseudo-randomly shuffled list of all of the channels available to the PHY. The mechanism to generate the default sequence is defined as follows:

- a) SHUFFLE is a *macHoppingSequenceLength*-sized array. The contents of this array are equivalent to the first *macHoppingSequenceLength* outputs of a 9-bit linear feedback shift register (LFSR) with polynomial $x^9 + x^5 + 1$ and a starting seed of 255. Each LFSR output is modulo *macHoppingSequenceLength*, so that each entry of SHUFFLE is between 0 and (*macHoppingSequenceLength* – 1), inclusive.
- b) CHANNELS is a *macHoppingSequenceLength*-sized array that is initially populated with the monotonically increasing list of channels available to the PHY.
- c) CHANNELS is shuffled as per Figure 7a. Elements may wind up being swapped multiple times in this process.
- d) The default sequence (i.e., *macHoppingSequenceList* for *macHoppingSequenceID* = 0) is equivalent to the shuffled CHANNELS array.

The use of other sequences (*macHoppingSequenceID* > 0) may be defined by a particular channel hopping system. The *macHoppingSequenceList* for a *macHoppingSequenceID* > 0 may be longer or shorter than the default sequence and may be specified algorithmically or set as a predefined channel list. Two hopping devices cannot communicate unless their PHYs support the same number of channels, and they either use the default hopping sequence or agree upon the hopping sequence being used, either through carrying this information in an enhanced beacon frame, or through pre-configuration.




Figure 7a—CHANNELS shuffle algorithm

For cases where *macHoppingSequenceLength* is greater than the number of channels available to the PHY, this implies that some channels will appear multiple times in the array. For cases where *macHoppingSequenceLength* is less than or equal to the number of channels available to the PHY, some channels available to the PHY may be excluded from the array. The selection of channels (the subset of available PHY channels, and which, if any, channels are used multiple times in the hopping sequence) is implementation-specific.

In general, any optional MAC mode as described in Annex I can calculate the channel as

$$CH = \text{macHoppingSequenceList} [\text{COUNTER} \% \text{macHoppingSequenceLength}]$$

where the COUNTER is some appropriate shared counter for a pair of devices communicating using that mode, and % indicates modular division.

5.1.1.2 Superframe structure

Insert in 5.1.1.1 after the first paragraph the following text:

For LLDN applications an additional superframe structure with LLDN beacons is required, as described in 5.1.1.6.

For LE-applications an additional and optional superframe structure is adopted, as described in 5.1.1.7.

5.1.1.4 CSMA-CA algorithm

Insert after the heading of 5.1.1.4 the following subheading followed by the existing text:

5.1.1.4.1 General

Replace existing Figure 11 with the following new Figure 11:




Figure 11—CSMA-CA Algorithm

Insert before 5.1.2 the following subclauses (5.1.1.4.2–5.1.1.7.4):

5.1.1.4.2 TSCH CCA algorithm

When a device is operating in TSCH mode as described in 5.1.2.6, CCA may be used to promote coexistence with other users of the radio channel. For other devices in the same network, the start time of transmissions, *macTsTxOffset*, is closely aligned making intra-network collision avoidance using CCA ineffective. TSCH devices also use channel hopping so there is no backoff period used when CCA prevents a transmission.

When a device has a packet to transmit, it shall wait for a link to the destination device. If CCA has been enabled, the MAC requests the PHY to perform a CCA at the designated time in the timeslot, *macTsCCAOOffset*, without any backoff delays. Figure 11a extends Figure 11 for the TSCH mode.




Figure 11a—TSCH CSMA-CA algorithm

5.1.1.4.3 TSCH CSMA-CA algorithm

Shared links (links with the linkOptions Bitmap set to shared transmission) are intentionally assigned to more than one device for transmission. This can lead to collisions and result in a transmission failure detected by not receiving an acknowledgment. To reduce the probability of repeated collisions when the packets are retransmitted, a retransmission backoff algorithm shall be implemented for shared links.

When a packet is transmitted on a shared link for which an acknowledgment is expected and none is received, the transmitting device shall invoke the TSCH CA retransmission algorithm. Subsequent retransmissions may be in either shared links or dedicated links as retransmission occurs in the next link to the destination. This backoff algorithm has the following properties:

The retransmission backoff wait applies only to the transmission on shared links. There is no waiting for transmission on dedicated links. The retransmission backoff is calculated in the number of shared transmission links. The backoff window increases for each consecutive failed transmission in a shared link. A successful transmission in a shared link resets the backoff window to the minimum value. The backoff window does not change when a transmission is a failure in a dedicated link.

The backoff window does not change when a transmission is successful in a dedicated link and the transmission queue is still not empty afterwards.

The backoff window is reset to the minimum value if the transmission in a dedicated link is successful and the transmit queue is then empty.

In the TSCH mode, backoff is calculated in shared links, so the CSMA-CA *aUnitBackoffPeriod* is not used.

For the *macMaxBE* and *macMinBE* values when the device is in TSCH mode refer to Table 52b.

The device shall use an exponential backoff mechanism analogous to that described in 5.1.1.4.1. A device upon encountering a transmission failure in a shared link shall initialize the BE to *macMinBE*. The MAC sublayer shall delay for a random number in the range 0 to $2^{BE} - 1$ shared links (on any slotframe) before attempting a retransmission on a shared link. Retransmission on a dedicated link may occur at any time. For each successive failure on a shared link, the device should increase the backoff exponent until the backoff exponent = *macMaxBE*. Successful transmission on a shared link resets the backoff exponent to *macMinBE*.

If an acknowledgment is still not received after *macMaxFrameRetries* retransmissions, the MAC sublayer shall assume the transmission has failed and notify the next higher layer of the failure. This is illustrated in Figure 11b.

5.1.1.4.4 LLDN simplified CSMA-CA

A simplified CSMA-CA algorithm is used during Management timeslots and Shared Group timeslots in LLDNs.

The simplified CSMA-CA is a slotted CSMA-CA mechanism and follows the same algorithm as described in 5.1.1.4.1.

The backoff slots of *aUnitBackoffPeriod* symbols are aligned with the start of the beacon transmission in management timeslots and with *tSlotTxOwner* in shared group timeslots.




Figure 11b—TSCH CSMA-CA backoff

5.1.1.5 TSCH Slotframe structure

5.1.1.5.1 General

A slotframe is a collection of timeslots repeating in time. Each timeslot allows enough time for a pair of devices to exchange a frame and an acknowledgment. It is possible, although usually undesirable, to define a timeslot that is not long enough for a pair of devices to exchange a maximum length frame and an acknowledgment. The number of timeslots in a given slotframe (slotframe size) determines how often each timeslot repeats, thus setting a communication schedule for nodes that use the timeslots. When a slotframe is created, it is associated with a slotframe handle for identification. Each slotframe repeats on a cycle dependent on its length. Each timeslot is an opportunity for a device to send or receive a single frame, and optionally receive or transmit an acknowledgment to that frame. Slotframes and timeslots are configured by a higher layer.

5.1.1.5.2 Absolute Slot Number (ASN)

The total number of timeslots that has elapsed since the start of the network or an arbitrary start time determined by the PAN coordinator is called the Absolute Slot Number (ASN). It increments globally in the network every *macTsTimeslotLength* μ s (refer to Table 52e). It may be beamed by devices already in a TSCH PAN, allowing new devices to synchronize. It is used globally by devices in a TSCH PAN as the frame counter (thus allowing for time-dependent security) and is used to compute the channel for any given pairwise communication as described in 5.1.1.5.3.

5.1.1.5.3 Links

Figure 11c shows how nodes may communicate in a sample three-timeslot slotframe. Nodes A and B communicate during timeslot 0, nodes B and C communicate during timeslot 1, and timeslot 2 is not being used. Every three timeslots, the schedule repeats, but note that ASN increments continuously. The pairwise assignment of a directed communication between devices in a given timeslot on a given channelOffset is a link. Physical channel, CH, in a link is made according to the following formula:

$$CH = \text{macHoppingSequenceList}[(\text{macASN} + \text{channelOffset}) \% \text{macHoppingSequenceLength}]$$

Use of a channelOffset allows for different channels to be used at a given *macASN* for a given *macHoppingSequenceList*. There are *macNumberOfChannels* channelOffsets that will result in a unique channel for that combination of *macASN* and *macHoppingSequenceList*.




Figure 11c—Example of a three time-slot slotframe

5.1.1.5.4 Multiple slotframes

A given network using timeslot-based access may contain several concurrent slotframes of different sizes. Multiple slotframes may be used to define a different communication schedule for various groups of nodes or to run the entire network at different duty cycles by giving some devices many active timeslots in a slotframe, and others few or none.

A network device may participate in one or more slotframes simultaneously, and not all devices need to participate in all slotframes. By configuring a network device to participate in multiple overlapping slotframes of different sizes, it is possible to establish different communication schedules and connectivity matrices that all work at the same time.

Slotframes can be added, removed, and modified while the network is running. Even though this is the case, all slotframes are aligned to timeslot boundaries, and timeslot 0 of the first repetition of every slotframe is projected back to *macASN* = 0, which is determined by the PAN coordinator (or other network device that starts the network). Because of this, timeslots in different slotframes are always aligned, even though the

beginning and end of a particular repetition of that slotframe may not be as illustrated in Figure 11d. When, for any given timeslot, a device has links in multiple slotframes, transmissions take precedence over receives, and lower slotframeHandle slotframes takes precedence over higher slotframeHandle slotframes.

	ASN=0	ASN=1	ASN=2	ASN=3	ASN=4	ASN=5	ASN=6	ASN=7	
Slotframe 1 5 slots	TS 0	TS 1	TS 2	TS 3	TS 4	TS 0	TS 1	TS 2	...
Slotframe 2 3 slots	TS 0	TS 1	TS 2	TS 0	TS 1	TS 2	TS 0	TS 1	...

Figure 11d—Multiple slotframes in the network

5.1.1.6 LLDN Superframe structure

5.1.1.6.1 General structure of superframe

The LLDN superframe is divided into a beacon slot, management timeslots if present, and *macLLDNumTimeSlots* base timeslots of equal length as illustrated in Figure 11e.




Figure 11e—Superframe with dedicated timeslots

The first timeslot of each superframe contains a beacon frame. The beacon frame is used for synchronization with the superframe structure. It is also used for re-synchronization of devices that went into power save or sleep mode.

The remaining timeslots are assigned to specific devices of the network. Each timeslot may have assigned a so-called slot owner. The slot owner has access privileges in the timeslot (dedicated timeslot). There is no explicit addressing necessary inside the frames if the slot owner transmits in its timeslot. The determination of the sender is achieved through the number of the timeslot. More than one device can be assigned to a timeslot (shared group timeslot). The devices use a contention-based access method (modified CSMA-CA as specified in 5.1.1.4.4) and a simple addressing scheme with 8-bit addresses in shared group timeslots.

Multiple adjacent base timeslots can be concatenated to a single, larger timeslot, as illustrated in Figure 11f.




Figure 11f—Usage and order of slots in a superframe

As shown in Figure 11f, there is a specific order in the meaning or usage of the timeslots, as follows:

- Beacon Timeslot: always present.
- Management Timeslots: one timeslot downlink, one timeslot uplink, presence is configurable in *macLLDNmgmtTS* during the Configuration state.
- Uplink timeslots for LLDN devices: *macLLDNnumUplinkTS* timeslots uplink (unidirectional communication), *macLLDNnumRetransmitTS* timeslots at the beginning can be reserved for retransmissions according to the Group Acknowledgement field contained in the LL-beacon as described in 5.2.2.5.2 and 5.1.9.4.
- Bidirectional timeslots for LLDN devices: *macLLDNnumBidirectionalTS* timeslots uplink/downlink (bidirectional communication).

It is also possible to use a separate Group Acknowledgement (GACK) frame as described in 5.2.2.5.4 in order to facilitate retransmissions of failed transmissions in the uplink timeslots within the same superframe. The use of a separate GACK is configurable during configuration mode. If the use of a separate GACK is configured, the structure of the superframe is as depicted in Figure 11g.




Figure 11g—Usage and order of slots in a superframe with configured use of separate GACK

Descriptions of the following configuration parameters and intervals for the superframe with a separate GACK are only different for the Uplink Timeslots:

- Beacon Timeslot
- Management Timeslots
- Uplink Timeslots: *macLLDNnumUplinkTS* denotes the total number of timeslots available for uplink (unidirectional) communication. Typically, one timeslot is allocated to each LLDN device. In this case, *M* denotes the number of LLDN devices, *macLLDNnumRetransmitTS* denotes the number of timeslots allocated for LLDN devices that failed their original transmissions prior to the GACK and

need to retransmit their message, and N denotes the number of LLDN devices that are allowed to retransmit. One timeslot is allocated for each retransmitting LLDN device.

- GACK: It contains an M bit bitmap to indicate successful and failed uplink transmissions in the same order as the uplink transmissions.
- Bidirectional Timeslots

The LL Beacon frame in the LLDN mode always carries the GACK bitmap even if a separate GACK frame is used. The GACK bitmap is used for acknowledging the successful retransmissions in timeslots $R1$, $R2$, ..., RN since some of the retransmitted frames (in $R1$, $R2$, ..., RN timeslots) may fail.

5.1.1.6.2 Beacon timeslot

The beacon timeslot is reserved for the LLDN PAN coordinator to indicate the start of a superframe with the transmission of a beacon. The beacon is used to synchronize the devices and to indicate the current transmission mode. The beacon also contains acknowledgments for the data transmitted in the last superframe.

The beacon timeslot is available in every superframe.

5.1.1.6.3 Management timeslots

The first portion of a superframe after the beacon timeslot is formed by the management timeslots, i.e., the downlink and uplink management timeslots.

The downlink direction is defined as sending data to the LLDN device. The uplink direction is defined as sending data from the LLDN device to the LLDN Coordinator.

Management timeslots provide a mechanism for bidirectional transmission of management data in downlink and uplink direction. Downlink and uplink timeslots are provided in equal number in a superframe. There are two management timeslots per superframe at maximum. Management timeslots are implemented as shared group access timeslots.

Management downlink and uplink timeslots are used in the Discovery state and the Configuration state and are optional in the Online state. These states are described in 5.1.9.

5.1.1.6.4 Uplink timeslots

After the management timeslots, timeslots for the transmission of data are contained in a superframe. Uplink timeslots allow for unidirectional communication (uplink) only.

The first *macLLDNnumRetransmitTS* of the *macLLDNnumUplinkTS* uplink timeslots are dedicated timeslots for retransmissions of failed uplink transmission attempts in dedicated timeslots of the previous superframe. The dynamic assignment of nodes to retransmission timeslots is described in 5.1.9.4.

5.1.1.6.5 Bidirectional timeslots

Bidirectional timeslots allow for bidirectional communication between the LLDN PAN coordinator and the LLDN device. The direction of the communication is signaled in the beacon as described in 5.2.2.5.2. Bidirectional timeslots are used for the transmission of device data to the LLDN PAN coordinator (uplink) as well as of data from the LLDN PAN coordinator to the LLDN device (downlink).

5.1.1.6.6 Channel access within timeslots

Each timeslot is described by four time attributes as illustrated in Figure 11h and described in Table 0a.




Figure 11h—Time attributes of timeslots

Table 0a—Time attributes of timeslots

Attribute	Description
tSlotStart	Starting time of timeslot
tSlotTxOwner	End time of privileged access by device that owns the timeslot
tSlotTxGW	If timeslot is unused, LLDN PAN coordinator can use the timeslot
tSlotEnd	End time of timeslot

From tSlotStart till tSlotTxOwner, the device that owns the slot, the slot owner, has exclusive access to the timeslot.

From tSlotTxOwner till tSlotTxGW, any device other than the LLDN PAN Coordinator may use the timeslot for data transmission with a modified CSMA-CA access scheme as described in 5.1.1.4.4 if the timeslot is not used by the slot owner. If the timeslot is not used by the slot owner, the LLDN PAN Coordinator shall indicate this by broadcasting a Clear To Send (CTS) Shared Group frame (5.3.10.4). To reduce the chances of collisions from other LLDN devices trying to use this timeslot, an LLDN device should send a Request To Send (RTS) frame (5.3.10.5) and wait for the receipt of the corresponding CTS frame (5.3.10.6) that identifies this LLDN device, from the LLDN PAN Coordinator before it transmits its data with a modified CSMA-CA access scheme as described in 5.1.1.4.4.

From tSlotTxGW till tSlotEnd, the LLDN PAN coordinator may use the timeslot, if the timeslot is still unused.

Dedicated timeslots are reserved for a single device (slot owner). This is achieved by setting tSlotTxOwner and tSlotTxGW to tSlotEnd. A dedicated timeslot allows the transmission of exactly one packet. Dedicated timeslots are only used during online mode as described in 5.1.9.4.

Shared group timeslots with contention-based access for every allowed device can be achieved by setting tSlotTxOwner to tSlotStart.

5.1.1.7 LE Functional description

This subclause specifies functionalities of devices supporting the following PIB attributes:

- *macCSLPeriod*
- *macRITPeriod*
- *macCSLMaxPeriod*
- *macLowEnergySuperframeSupported*
- *macLowEnergySuperframeSyncInterval*

5.1.1.7.1 LE Contention access period (LE CAP)

When *macCSLPeriod* is nonzero, CSL is deployed in CAP. CSL behavior is defined in 5.1.11.1. The *macRITPeriod* shall be set to zero in a beacon-enabled PAN.

When *macLowEnergySuperframeSupported* is TRUE, the transaction shall be ensured to be completed one IFS period before the end of the inactive period. If a device senses a frame in the CAP that does not end within the CAP when *macLowEnergySuperframeSupported* is TRUE, the device may continue receiving the frame until it ends before the end of the inactive period. When *macLowEnergySuperframeSupported* is TRUE, the coordinator shall not allocate GTSSs in order to avoid the interference from the frames exceeding the CAP and going into the CFP. When *macLowEnergySuperframeSupported* is TRUE, the coordinator shall notify the devices that already associated or intend to associate the condition of *macLowEnergySuperframeSupported* in the beacon frames.

5.1.1.7.2 LE Superframe structure

If *macLowEnergySuperframeSupported* is TRUE and *macLowEnergySuperframeSyncInterval* is not zero, the coordinator shall transmit beacon frames not in every beacon interval, but once in every beacon interval time *macLowEnergySuperframeSyncInterval*, except when requested to do so. If *macLowEnergySuperframeSupported* is TRUE and *macLowEnergySuperframeSyncInterval* is zero, the coordinator shall transmit beacon frames only when it is requested to do so.

5.1.1.7.3 LE-incoming and outgoing superframe timing

If a device supports *macLowEnergySuperframeSupported*, the beacon order and superframe order may be equal for all superframes on a PAN.

5.1.1.7.4 LE scan

When *macCSLPeriod* is nonzero, CSL is deployed in channel scans. When *macCSLMaxPeriod* is nonzero, each coordinator broadcasts beacon frames with wake-up frame sequence. This allows devices to perform channel scans with low duty cycles.

5.1.2 Starting and maintaining PANs

5.1.2.1 Scanning through the channels

Change the first paragraph as follows:

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

5.1.2.1.2 Active and passive channel scan

Change the fourth paragraph as follows:

~~An active or passive scan or an active scan or an enhanced active scan~~ over a specified set of channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate ~~an active or passive scan, an active scan or enhanced active scan respectively~~. For each channel, the device shall first switch to the channel, by setting *phyCurrentChannel* and *phyCurrentPage* accordingly. For an active scan, the device shall send a beacon request command, as described in 5.3.7. For UWB and CSS PHYs, the scan process shall be repeated for each mandatory preamble code, setting the *phyCurrentCode* appropriately. For an enhanced active scan the device shall send an enhanced beacon request as described in 5.3.7.2. Upon successful transmission of the beacon request command for an active scan or after switching to the channel for a passive scan, the device shall enable its receiver for at most $[aBaseSuperframeDuration \times (2^n + 1)]$, where n is the value of the ScanDuration parameter. During this time, the device shall reject all nonbeacon frames and record the information contained in all unique beacons in a PAN descriptor structure, as described in Table 17, including the channel information and, if required, the preamble code.

Insert the following new text, table, and figure in 5.1.2.1.2 after the eighth paragraph:

If a coordinator capable of responding to the enhanced beacon request (EBR) command and capable of filtering as described in 5.3.7.2.2 receives an EBR command, it shall perform the filtering as indicated in the EBR. If the filter conditions are satisfied, it shall transmit the appropriate enhanced beacon as per Table 0b.

Table 0b—Channel access for EBR response

Mode of operation	Response required	Access method	When to respond	Notes
Beacon PAN	Y	Slotted CSMA-CA	Next available CAP	
Nonbeacon PAN	Y	Unslotted CSMA-CA	ASAP	E.G., unslotted hopping
DSME beacon	Y	Slotted CSMA-CA	Next available EB	
DSME non-beacon	Y	Unslotted CSMA-CA	ASAP	
TSCH	Y	Slotted CSMA-CA	Next available time slot on same channel	

The messages exchanged when using an Enhanced Active Scan with an Enhanced Beacon Request are shown in Figure 13a.




Figure 13a—Enhanced Active scan with Enhanced Beacon request

Insert before 5.1.3 (after 5.1.2.5) the following subclause and figure:

5.1.2.6 TSCH PAN formation

A TSCH PAN is formed when a device, usually the PAN coordinator, advertises the presence of the network by sending Enhanced Beacons upon receipt of a MLME BEACON.request from a higher layer. The Enhanced Beacons contain the following:

- Time information so new devices can synchronize to the network (as described in 5.2.4.13).
- Channel hopping information (as described in 5.1.1.6).
- Timeslot information describing when to expect a frame to be transmitted and when to send an acknowledgment (as described in 5.2.4.15)
- Initial link and slotframe information so new devices know when to listen for transmissions from the advertising device and when they can transmit to the advertising device (as described in 5.2.4.14).

The device wishing to join the network begins passively (preferred) or actively scanning for the network as the result of receiving an MLME SCAN.request from a higher layer. Once the listening device has heard a valid Enhanced Beacon, it generates an MLME-BEACON-NOTIFY.indication to a higher layer. The higher layer may initialize the slotframe and links contained in the Enhanced Beacon and switch the device into TSCH mode with a TSCH MODE.request or wait for additional MLME-BEACON-NOTIFY.indication primitives before doing so.

At this point the device is synchronized to the network and may optionally send in an Association Request. If the device uses Association, it may request a 2-octet address. The sequence of messages exchanged to synchronize a device to the networks is shown in Figure 16a, and the process of synchronization is described in 5.1.4.2a.

Typically at this point the device will go through a procedure to allocate additional communication resources (slotframes and links) to the joining device. This procedure may include a security handshake to mutually authenticate the joining device, configure encryption keys, and configure routing information. The mechanism and rules for setting up these additional communication links and configure other policies would normally be defined in a higher layer standard—the content of these messages is beyond the scope of this document.

Once synchronized and configured by a higher layer to do so, all FFDs that are already part of the network may send enhanced beacons announcing the presence of the network. The advertising rate and content is configured by a higher layer as appropriate to the density of devices, the desired rate of network formation, and the energy devoted to network formation.

After joining, the device may receive additional slotframes and links from a higher layer management entity or peer as required by the application, or the device may be instructed to remove slotframes and links obtained from the enhanced beacon.




Figure 16a—Message sequence chart for TSCH procedure to find an advertising device

5.1.3 Association and disassociation

5.1.3.1 Association

Insert at end of 5.1.3.1 the following paragraph:

Association is optional for devices operating in TSCH mode.

5.1.3.2 Disassociation

Insert at end of 5.1.3.2 the following paragraphs:

For devices using the optional TSCH mode, additional disassociation behavior is required. A device shall only disassociate from the PAN if it receives a disassociation notification command from either the PAN coordinator, or all of its timekeeping neighbors, i.e., neighbors to whom the device has receive links marked with the Timekeeping bit set in the *macLinkOptions*.

Upon determining that it should disassociate from the PAN, the device shall transmit disassociate notification command frames to all its neighbors on any available link for *macDisconnectTime* timeslots, after which it should clear all synchronization information and leave the PAN.

Insert the following subclause 5.1.3.3 before 5.1.4:

5.1.3.3 Fast association (FastA)

Fast association (FastA) is optional and is not defined in IEEE 802.15.4 devices.

A device is instructed to Fast associate with a PAN through the MLME ASSOCIATE.request primitive with the Association Type field of the Capability Information field set to one.

The MAC sublayer of an unassociated device shall initiate the FastA procedure by sending an association request command with the Association Type field of the Capability Information field set to one to the coordinator of an existing PAN. Because the association request command contains an acknowledgment request, the coordinator shall confirm its receipt by sending an acknowledgment frame.

On receipt of the acknowledgment to the association request command, the device shall wait for at most *macResponseWaitTime* symbols for the coordinator to make its association decision.

If the coordinator has sufficient resources, the next higher layer should allocate a 16-bit short address to the device, and the MAC sublayer shall generate an association response command containing the new address and a status indicating a successful FastA. If the Association Type field of the Capability Information field of the association request command is set to one, the MAC sublayer of the coordinator shall send the association response command to the device directly as described in Figure 60a. If the Association Type field of the Capability Information field of the association request command is set to zero, the association response command shall be sent as 5.1.3.1.

Because the association response command contains an acknowledgment request, the device requesting FastA shall confirm its receipt by sending an acknowledgment frame. If the Association Status field of the command indicates “FastA successful” (0x80), the device shall store the address contained in the 16-bit Short Address field of the command in *macShortAddress* for its communication use in the PAN.

5.1.4 Synchronization

Insert before 5.1.4.1 the following paragraph:

For devices using the optional TSCH mode, initial synchronization is performed by the use of enhanced beacons, and synchronization is maintained by slotted communication with other devices in the PAN.

5.1.4.2 Synchronization without beacons

Insert the following subclauses (5.1.4.2.a–5.1.4.2.a.2) before 5.1.4.3:

5.1.4.2a Synchronization in TSCH PAN

In a TSCH PAN, all communication happens in timeslots as described in 5.1.4.2a.1. To remain synchronized, the devices should have the same notion of when each timeslot begins and ends, within $\pm macTsRxWait/2$ μ s. In a typical TSCH PAN, time propagates outwards from the PAN coordinator. A device shall periodically synchronize its network time to at least one other network device. It may also provide its network time to one or more network devices. A network device uses a neighbor as a time source when it has received links to that neighbor with the Timekeeping bit set in the *macLinkOptions*. A higher layer may add or change time source neighbors at any time.

Note that a device sending Enhanced Beacons to advertise a TSCH PAN should set the Timekeeping bit in the Link Option field (as described in 5.2.4.14) for the joining devices' receive link so that joining devices can maintain time synchronization until additional time source neighbors are configured by a higher layer.

A network device may have more than one neighbor as its time source. In such cases, the device shall synchronize its time to all of the neighbors that are acting as its time source, synchronizing to the relative drift of all its time source neighbors.

Figure 22a shows an example of time propagation in a TSCH PAN. The arrows indicate the direction of time distribution. In this example, the PAN coordinator acts as the time source for the entire network. Device 1 synchronizes to the PAN coordinator only, and is the time source for device 3. Device 2 synchronizes its time to both 1 and the PAN coordinator, and device 4 synchronizes to the PAN coordinator, device 2, and device 3.

Synchronization is possible whenever a device exchanges a frame with a time source neighbor. This can either come from receiving an acknowledgment with time correction information (as described in 5.2.4.13) or from the arrival time of a frame from the time source neighbor. Both methods are described in 5.1.4.2a.2.

In order to ensure that it remains synchronized with the TSCH PAN, a network device shall ensure that it communicates with each of its timekeeping neighbors at least once per Keep Alive period. If a network device has not sent a packet to its time source neighbor within this interval, it shall send any empty acknowledged MAC frame and use the acknowledgement frame to perform acknowledgement based synchronization.

5.1.4.2a.1 Timeslot communication

During a timeslot in a slotframe, one node typically sends a frame, and another sends back an acknowledgment frame containing the ACK/NACK Time Correction IE as described in 5.2.4.11 if it successfully receives that frame and the frame passes authentication. The acknowledgment can be positive (ACK) or negative (NACK). A positive acknowledge indicates that the receiver has successfully received the frame and has taken ownership of it for further routing. A negative acknowledgment indicates that the receiver cannot accept the frame at this time, but has heard it with no errors. Both ACKs and NACKs carry timing information used by nodes to maintain network synchronization. Frames sent to a unicast node address require that a link-layer acknowledgment be sent in response during the same timeslot as shown in Figure 22b. If an acknowledgment is requested and not received within the time-out period, retransmission of the frame waits until the next assigned transmit timeslot to that address occurs.




Figure 22a—Example of possible time propagation in a TSCH PAN

As shown in Figure 22b, the timeslot starts at time $T = 0$ from the transmitting device’s perspective. The transmitter waits $macTsCCAOffset$ μs , and then performs CCA (if active). At $macTsTxOffset$ μs , the device begins transmitting the frame. If an acknowledgment is expected, the device waits $macTsRxAckDelay$ μs , and then enables the receiver to await the acknowledgment. If the acknowledgment does not arrive within $macTsAckWait$ μs the device may idle the radio and consider the transmission a failure. If no acknowledgment is expected, the transmitter may idle the radio after sending the frame.

On the receiver’s side, at its estimate of $T = 0$ it waits $macTsRxOffset$ μs and then goes into receive mode for $macTsRxWait$ μs . If the frame has not started by that time, it may idle the receiver. Otherwise, once the frame has been received, the receiver waits $macTsTxAckDelay$ μs and then sends an acknowledgment.

The transmitter or receiver may resynchronize clocks as described in 5.1.4.2a.2.




Figure 22b—Timeslot diagram of acknowledged transmission

5.1.4.2a.2 Node synchronization

Device-to-device synchronization is necessary to maintain connection with neighbors in a slotframe-based network. There are two methods for a device to synchronize to the network, Acknowledgment-based and Frame-based. Originator in this context is the device sending a frame, and receiver is the device receiving that frame and sending back an acknowledgment as is appropriate. Since timestamps are required to maintain synchronization in a TSCH PAN, all devices shall have *macTimestampSupported* = TRUE.

Acknowledged communication provides a basic method of time synchronization through the exchange of data and acknowledgment frames. The algorithm involves the receiver calculating the delta between the expected time of frame arrival and its actual arrival, and providing that information to the transmitting node in the subsequent acknowledgment.

The acknowledgment-based synchronization algorithm can be described as follows:

- Originator sends a frame, timing the start symbol to be sent at *macTsTxOffset* according to its clock, which would correspond to $macTsRxOffset + macTsRxWait/2$ in the receiver's clock if both clocks were perfectly synchronized.
- Receiver records the timestamp (e.g., from the MCPS DATA.indication) when it receives the start symbol of the frame.
- Receiver calculates a time correction = $macTsRxOffset + macTsRxWait/2 - \text{arrival timestamp}$.
- Receiver sends back the time correction in the IE field in the corresponding enhanced acknowledgment frame (assuming the incoming frame passes validation).
- Originator receives the acknowledgment. If the receiver node is a time source neighbor, the originator adjusts its own clock by incorporating the difference into an average of the drift to all its time source neighbors. The averaging method is implementation dependent. If the receiver is not a clock source, the time correction is ignored.

In Frame-based synchronization a node may synchronize its own network clock if it receives a frame from a time source neighbor. The receiver calculates the delta between expected time of frame arrival and its actual arrival time to use that information to adjust its own clock.

The Frame-based synchronization algorithm can be described as follows:

- Receiver records the timestamp (e.g., from the MCPS DATA.indication) when it receives the start symbol of the frame.
- The receiver calculates a time correction = $macTsRxOffset + macTsRxWait/2 - \text{arrival timestamp}$.
- If the originator was a time source neighbor, the receiver adjusts its own clock by incorporating the difference into an average of the drift to all its time source neighbors. The averaging method is implementation dependent. If the receiver is not a clock source, the time correction is ignored.

Figure 22c illustrates both time synchronization mechanisms. In both cases, the receiver calculates its time adjustment to either send back to the transmitting device or to use locally.




Figure 22c—Time synchronization

5.1.6 Transmission, reception, and acknowledgment

5.1.6.2 Reception and rejection

Change the sixth paragraph in 5.1.6.2 as indicated:

If any of the third-level filtering requirements are not satisfied, the MAC sublayer shall discard the incoming frame without processing it further. If all of the third-level filtering requirements are satisfied, the frame shall be considered valid and processed further. For valid frames that are not broadcast, if the Frame Type field indicates a data or MAC command frame with the Frame Version field set to 0b00-0b01 and the AR field is set to request an acknowledgment, the MAC sublayer shall send an acknowledgment frame. Prior to the transmission of the acknowledgment frame, the sequence number included in the received data or MAC command frame shall be copied into the Sequence Number field of the acknowledgment frame. This step will allow the transaction originator to know that it has received the appropriate acknowledgment frame.

Insert the following paragraph after the sixth paragraph in 5.1.6.2:

For valid frames that are not broadcast, if the Frame Type field indicates a multipurpose frame, or a data or MAC command frame with the Frame Version field set to 0b10, and the AR field is set to request an acknowledgment, the MAC sublayer shall send an enhanced acknowledgement as described in 5.1.6.4.2.

5.1.6.4 Use of acknowledgments and retransmissions

5.1.6.4.2 Acknowledgment

Insert in 5.1.6.4.2 after the second paragraph the following text:

For enhanced acknowledgment frames (which are only sent after the incoming frame passes security filtering, if applicable), additional time may be required after *aTurnaroundTime* to complete incoming and outgoing security processing. If the enhanced acknowledgment is not sent before *macEnhAckWaitDuration* μ s, the sender will assume the frame was not successful.

The receiving device may include additional content in an enhanced acknowledgment encapsulated as IEs. If the originator does not understand the IE content of the acknowledgment, it is ignored, but the transmission is considered successful.

When returning Time Correction information in the acknowledgment (via the ACK/NACK Timing Correction IE as described in 5.2.4.11), the receiving device may acknowledge frames with an NACK to indicate that the frame could not be transferred to a higher layer. This is treated as a failure from the sender's perspective, and a notification of failure due to congestion is indicated to the higher layer through the MAC performance metrics.

LLDNs use several methods for the acknowledgment of data transmissions. The timings of these mechanisms are defined by the superframe structure of the LLDN. The transmission of an LL-Acknowledgment frame in response to an LL-data frame in an LLDN shall commence in the same bidirectional timeslot in the next superframe. The LL-Acknowledgment frame shall only be used with bidirectional timeslots.

Insert before 5.1.6.4.3 the following paragraph:

When in TSCH mode, incoming frames are acknowledged using the enhanced acknowledge frame as described in 5.2.2.3. Security of the acknowledgment frame should match that of the incoming frame. When operating in the TSCH mode (refer to 6.2.19.5), the enhanced acknowledgment frame is sent at the time specified by the *macTimeslotTemplate* being used (as described in 6.4.3.3.3 and 5.1.4.2a), i.e., *macEnhAckWaitDuration* should be set to a value corresponding to *macTsAckWait*.

5.1.6.4.3 Retransmissions

Change in 5.1.6.4.3 the second paragraph as follows:

A device that sends a data or MAC command frame with frame version 0b00-0b01 and with its AR field set to acknowledgment requested shall wait for at most *macAckWaitDuration* for the corresponding acknowledgment frame to be received. If an acknowledgment frame is received within *macAckWaitDuration* and contains the same DSN as the original transmission, the transmission is considered successful, and no further action regarding retransmission shall be taken by the device. If an acknowledgment is not received within *macAckWaitDuration* or an acknowledgment is received containing a DSN that was not the same as the original transmission, the device shall conclude that the single transmission attempt has failed.

Insert in 5.1.6.4.3 after the second paragraph the following paragraph:

A device that sends a data or MAC command frame with frame version 0b10 with its Acknowledgment Request field set to one shall wait for at most *macEnhAckWaitDuration* μ s for the corresponding acknowledgment frame to be received, as additional security processing is required. If an acknowledgment frame is received within *macEnhAckWaitDuration* μ s and contains the same DSN as the original transmission, the transmission is considered successful, and no further action regarding retransmission shall be taken by the device. If an acknowledgment is not received within *macEnhAckWaitDuration* μ s or an acknowledgment is received containing a DSN that was not the same as the original transmission, the device shall conclude that the single transmission attempt has failed.

Insert before 5.1.6.5 the following subclause (5.1.6.4.4):

5.1.6.4.4 TSCH retransmissions

A device that sends a data or MAC command frame with its Acknowledgment Request field set to one shall wait for *macTsRxAckDelay* μ s. If an acknowledgment frame begins within *macEnhAckWaitDuration* and

contains the same DSN as the original transmission, and otherwise has no errors, the transmission is considered successful, and no further action regarding retransmission shall be taken by the device. If an acknowledgment is not received within the appropriate time-out or an acknowledgment is received containing a DSN that was not the same as the original transmission, the device shall conclude that the single transmission attempt has failed.

All transmissions in TSCH are direct. If a single transmission attempt has failed, the device shall repeat the process of transmitting the data or MAC command frame and waiting for the acknowledgment, up to a maximum of *macMaxFrameRetries* times. In TSCH mode (6.2.19.5), retransmissions only occur on subsequent transmit links to the same recipient on any slotframe assigned to the device. If an acknowledgment is still not received after *macMaxFrameRetries* retransmissions, the MAC sublayer shall assume the transmission has failed and notify the next higher layer of the failure.

5.1.6.6 Transmission scenarios

Replace in 5.1.6.6 each occurrence of “macAckWaitDuration” with “macAckWaitDuration or macEnhAckWaitDuration as appropriate.”

Insert before 5.2 the following subclauses (5.1.9–5.1.12.5):

5.1.9 LLDN transmission states

5.1.9.1 General

The transitions between the different transmission states are illustrated in Figure 34a.




Figure 34a—Transitions between transmission states

The discovery state is the first step during network setup: the new devices are discovered and configured in the second step, the configuration state. After the successful completion of the configuration state, the network can go into online state. Data and readings from the devices can only be transmitted during online state. In order to reconfigure a network, the configuration state can be started again.

5.1.9.2 Discovery state

The Discovery state is the first step during network setup or for the addition of new devices to an existing network.

In the Discovery state, the superframe contains only the timeslot for the beacon described in 5.1.1.6.2 and two management timeslots, one downlink and one uplink (5.1.1.6.3).

A new device scans the different channels until it detects an LLDN PAN coordinator sending beacons that indicate Discovery state.

If a new device received a beacon indicating Discovery state, it attempts to access the medium in the uplink management timeslot in accordance with 5.1.1.4.4 in order to send a Discover Response frame to the LLDN PAN coordinator. The Discover Response frame is described in 5.3.10.1. The Discover Response frame contains the current configuration of the device. The new device shall repeat sending the Discover Response frame until it receives an acknowledgment frame for it or the Discovery state is stopped by the LLDN PAN coordinator. The acknowledgment frame is described in 5.2.2.5.4.

The LLDN coordinator changes from the Discovery state to the Configuration state if it did not receive any Discover Response frames within *macLLDNdiscoveryModeTimeout* seconds.

Figure 34b illustrates the Discovery state.




Figure 34b—Flow diagram of Discovery state

5.1.9.3 Configuration state

The Configuration state is the second step during network setup. It is also used for network reconfiguration.

In the Configuration state, the superframe contains only the timeslot for the beacon described in 5.1.1.6.2 and two management timeslots, one downlink and one uplink as described in 5.1.1.6.3.

If a device received a beacon indicating configuration state, it tries to get access to the transmission medium in the uplink management timeslot in order to send a Configuration Status frame to the LLDN PAN coordinator. The Configuration Status frame is described in 5.3.10.2. The Configuration Status frame contains the current configuration of the device. The new device shall repeat sending the Configuration Status frame until it receives a Configuration Request frame for it or the Configuration state is stopped by the LLDN PAN coordinator. The Configuration Request frame is described in 5.3.10.3. The Configuration Request frame contains the new configuration for the receiving device. After successfully receiving the

Configuration Request frame, the device sends an acknowledgment frame to the LLDN PAN coordinator. The acknowledgment frame is described in 5.2.2.5.4.

Figure 34c illustrates the Configuration state.




Figure 34c—Flow diagram of Configuration state

5.1.9.4 Online state

User data is only sent during Online state. The superframe starts with a beacon and is followed by several timeslots. The devices can send their data during the timeslots assigned to them during the Configuration state. The different types of timeslots are described in 5.1.1.6.

The existence and length of management timeslots in the Online state are contained in the Configuration Request frame.

The successful reception of data frames by the LLDN PAN coordinator is acknowledged in the Group Acknowledgment bitmap of the beacon frame of the next superframe described in 5.2.2.5.2 or in a separate Data Group Acknowledgment frame depicted in Figure 48h. This is the case for both uplink timeslots and bidirectional timeslots if the transmission direction is uplink. Figure 34d illustrates an example of the Online state for uplink transmissions. In this example, the network has three dedicated timeslots, and LLDN device 2 is assigned to timeslot 2.

If retransmission timeslots are configured (i.e., $macLLDNnumRetransmitTS > 0$), the retransmission slots are assigned to the owners of the first $macLLDNnumRetransmitTS$ with the corresponding bit in the group acknowledgment bitmap set to zero. Each LLDN device shall execute the algorithm as illustrated in Figure 34e in order to determine its retransmission timeslot. The LLDN PAN coordinator has to execute a similar algorithm in order to determine the senders of the frames in the retransmission slots.