

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

Part 22.1: Standard to Enhance Harmful Interference Protection for Low-Power Licensed Devices Operating in TV Broadcast Bands

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

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

IEEE Std 802.22.1™-2010

1 November 2010

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

**Part 22.1: Standard to Enhance
Harmful Interference Protection for
Low-Power Licensed Devices
Operating in TV Broadcast Bands**

Sponsor

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

Approved 30 September 2010

IEEE-SA Standards Board

Abstract: This standard defines the protocol and data formats for communication devices forming a beaconing network that are used to protect low-power, licensed devices operating in television broadcast bands from harmful interference generated by license-exempt devices, such as Wireless Regional Area Networks (WRAN), intended to operate in the same bands. The devices being protected are devices licensed as secondary under Title 47, Part 74, Subpart H in the USA and equivalent devices in other regulatory domains.

Keywords: ad hoc network, beacons, TV white space, Wireless Regional Area Networks, WRAN

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

Copyright © 2010 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 1 November 2010. Printed in the United States of America.

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

PDF: ISBN 978-0-7381-6497-7 STD97037
Print: ISBN 978-0-7381-6498-4 STDPD97037

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

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

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

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

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

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

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

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration. 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 interpretation of the 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, explanation, or interpretation of the IEEE. Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Recommendations to change the status of a stabilized standard should include a rationale as to why a revision or withdrawal is required.

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

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

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

Introduction

This introduction is not part of IEEE Std 802.22.1-2010, IEEE Standard for Information Technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 22.1: Standard to Enhance Protection for Low-Power, Licensed Devices Operating in Television Broadcast Bands.

This standard defines the protocol and data formats for communication devices offering enhanced protection for low-power, licensed devices, such as those used in the production and transmission of broadcast programs (i.e., devices licensed as secondary under Title 47, Part 74, Subpart H in the USA and equivalent devices in other regulatory domains), operating in television broadcast bands. Protection is provided through the use of a beacon, which contains information relevant to the licensed device, including its physical location and estimated duration of TV channel occupancy. The standard uses the ALOHA medium access mechanism, and all transmissions are broadcast.

The physical layer uses direct sequence spread spectrum (DSSS) with differential quadrature phase-shift keying (DQPSK). A synchronization word and countdown mechanism (i.e., time until the next beacon transmission) is transmitted continuously on the I channel, while beacons and inter-device communications are transmitted on the Q channel. Frequency, modulation rate, and transmit power vary from region to region and shall adhere to local regulations.

Notice to users

Laws and regulations

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

Copyrights

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

Updating of IEEE documents

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

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

Errata

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

Downloads

Portions of this standard can be downloaded from the Internet. Materials include PICS tables, data tables, and code. URLs are listed in the text in the appropriate sections.

Interpretations

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

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE 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 standard was submitted to the IEEE-SA for approval, the IEEE P802.22 Working Group had the following voting members:

Apurva N. Mody, *Chair*
Gerald Chouinard, *Vice Chair*

Greg Buchwald, *P802.22.1 Chair*
Monique Bourgeois Brown, *P802.22.1 Editor-in-Chief*

Chee Wei Ang	Baowei Ji	Edgar Reihl
Kwok Shum Au	Jerome J. Kalke	Jon Rosdahl
John Benko	Ramon Khalona	William Rose
Winston Caldwell	Chang-Joo Kim	Shigenobu Sasaki
Ed Callaway	Kihong Kim	Cheng Shan
Dave Cavalcanti	Gwangzeen Ko	Steve Shellhammer
Soo-Young Chang	Stephen Kuffner	Kirk Skeba
Chris Clanton	Jeong Suk Lee	Eli Sofer
Johnny Dixon	Zhongding Lei	Myung Sun Song
Charles Einolf	Kyutae Lim	Srikathyayani Srikanteswara
Wen Gao	Jiezhen Lin	Carl R. Stevenson
Ingo Gaspard	David Mazzaresse	Victor Tawil
Monisha Ghosh	Peter Murray	James Tomcik
Joanna Guenin	Paul Nikolich	Jungsun Um
Tom Gurley	Moh Nouroozian	George Vlantis
Anh Tuan Hoang	Juha Pihlaja	Kelly Williams
Wendong Hu	Jeff Poston	Yu-chun Wu
Sung Hyun Hwang	Ranga Reddy	Changlong Xu
Tae-In Hyon	Ivan Reede	Yonghong Zeng
		Jianwei Zhang

Major contributions were received from the following individuals:

Greg Buchwald	Mingwei Jie	K. Sivanesan
Ed Callaway	Jerome J. Kalke	Yang Tang
Soo-Young Chang	Mark Kenkel	Victor Tawil
Gerald Chouinard	Stephen Kuffner	Ke Wang
Chris Clanton	Linjun Lv	Kexue Wang
Charles Einolf	David Mazzaresse	Yu-chun Wu
Paul Gorday	Tom Messerges	Zhou Wu
Garret Heath	Zhixue Shi	Jianwei Zhang
Baowei Ji	Dave Silk	Xuesheng Zhu

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

Thomas Alexander	Randall Groves	Apurva N. Mody
Butch Anton	Joanna Guenin	Ronald Murias
Reza Arefi	Thomas Gurley	Peter Murray
Kwok Shum Au	Hiroshi Harada	Michael S. Newman
Alexander Awuviri	William Hayes	Charles Ngethe
Taehan Bae	Robert F. Heile	Paul Nikolich
Youngkyo Baek	Oliver Hoffmann	John Notor
Stephen Berger	Seock Deock Hong	Satoshi Obara
Gregory Best	Victor Hou	Jisung Oh
Harry Bims	David Hunter	Okundu Omeni
Gennaro Boggia	Chan-Soo Hwang	Chris Osterloh
Achim Brandt	Tae-In Hyon	Satoshi Oyama
Nancy Bravin	Tetsushi Ikegami	Jeongho Park
Monique Bourgeois Brown	Atsushi Ito	Jungshin Park
Timothy Brown	Raj Jain	William Pechey
William Byrd	Jaehyuk Jang	Venkatesha Prasad
Sean Cai	Hongkyu Jeong	Michael Probasco
Peter J. Calderon	Suryong Jeong	Henry Ptasinski
Edgar Callaway	Baowei Ji	Ivan Reede
James Carlo	Bobby Jose	Benjamin Rolfe
Juan Carreon	Woochul Jung	William Rose
Youngbin Chang	Tal Kaitz	Jaeyoung Ryu
Clint Chaplin	Hyunjeong Kang	Randall Safier
Yung-Mu Chen	Noh-Gyoung Kang	John Santhoff
Kichun Cho	Efthymios Karabetsos	Shigenobu Sasaki
Hokyuu Choi	Piotr Karocki	Bartien Sayogo
Seung-Hoon Choi	Shuzo Kato	Cheng Shan
Gerald Chouinard	Stuart J. Kerry	Stephen Shellhammer
Keith Chow	Eunkyung Kim	Jaejeong Shim
jinyong chung	Gil Kim	Changyong Shin
Jules Cohen	Sangbum Kim	Wonjae Shin
Charles Cook	Yongbum Kim	Gil Shultz
Todor Cooklev	Youngdoo Kim	Jaeseung Son
Jose Costa	Youngsoo Kim	Jung Je Son
Manoj Das	Patrick Kinney	Kapil Sood
Russell Dietz	Dongkeon Kong	Amjad Soomro
Thomas Dineen	Bruce Kraemer	Kenneth Stanwood
Mi-Sun Do	Joseph Kwak	Thomas Starai
Carlo Donati	Edwin Kwon	Adrian Stephens
Paul Eastman	Jeremy Landt	Carl Stevenson
Peter Ecclesine	Insun Lee	Rene Struik
Richard Eckard	Jae Min Lee	Walter Struppler
Charles Einolf	Jeongho Lee	Jun Ichi Takada
Marc Emmelmann	Mi Hyun Lee	Masahiro Takagi
Bernard Eydt	Myung Lee	Michael Johas Teener
Shulan Feng	Seong-hee Lee	David Tepen
Michael Fischer	Sungjin Lee	Adam Toner
C. Fitzgerald	Tae Hoon Lee	Wen Tong
David Fort	Joseph Levy	Masahiro Umehira
Andre Fournier	Jan-Ray Liao	Dmitri Varsanofiev
Robert Frazier	Arthur Light	Prabodh Varshney
Avraham Freedman	Chiwoo Lim	George Vlantis
Devon Gayle	Hyoung-Kyu Lim	Stanley Wang
Theodore Georgantas	Daniel Lubar	Robert Weller
James Gilb	William Lumpkins	Harry Worstell
Jaganathan Gnanavelu	Mark Maloney	Xuyong Wu
Mariana Goldhamer	David Mazzaresse	Oren Yuen
	Steven Methley	Juan Zuniga

When the IEEE-SA Standards Board approved this standard on 30 September 2010, it had the following membership:

Robert M. Grow, *Chair*
Richard H. Hulett, *Vice Chair*
Steve M. Mills, *Past Chair*
Judith Gorman, *Secretary*

Karen Bartleson
Victor Berman
Ted Burse
Clint Chaplin
Andy Drozd
Alexander Gelman
Jim Hughes

Young Kyun Kim
Joseph L. Koepfinger*
John Kulick
David J. Law
Hung Ling
Oleg Logvinov
Ted Olsen

Ronald C. Petersen
Thomas Prevost
Jon Walter Rosdahl
Sam Sciacca
Mike Seavey
Curtis Siller
Don Wright

*Member Emeritus

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

Satish Aggarwal, NRC Representative
Richard DeBlasio, DOE Representative
Michael Janezic, NIST Representative

Catherine Berger
IEEE Project Editor

Michael Kipness
IEEE Standards Program Manager, Technical Program Development

Contents

1.Overview.....	1
1.1 General.....	1
1.2 Scope.....	2
1.3 Purpose.....	2
2.Normative references.....	3
3.Definitions.....	5
4.Acronyms and abbreviations.....	7
5.Functional overview.....	9
5.1 Introduction.....	9
5.2 Architecture.....	9
5.3 Superframe structure.....	9
5.3.1 Synchronization burst structure.....	11
5.3.2 Beacon frame structure.....	12
5.4 Data transfer model.....	12
5.5 Security.....	14
5.5.1 Public-key approach.....	15
5.5.2 Security considerations.....	15
5.5.3 Beaconing device life cycle.....	16
6.PHY specification.....	17
6.1 General requirements and definitions.....	17
6.1.1 TV channel/licensed auxiliary service (LAS) channel information.....	17
6.1.2 Modulation rates and beacon offset location for ATSC DTV regions.....	17
6.1.3 RF power measurement information.....	18
6.1.4 Receiver sensitivity definitions.....	18
6.2 PHY service specifications.....	18
6.2.1 PHY data service.....	19
6.2.2 PHY management service.....	20
6.2.3 PHY enumerations description.....	27
6.3 Synchronization burst.....	27
6.4 PPDU format.....	28
6.5 Inter-device communication interval (ICI).....	28
6.5.1 Receive (Rx) period.....	29
6.5.2 Acknowledgement/no acknowledgement period (ANP).....	31
6.6 PHY constants and PIB attributes.....	32
6.6.1 PHY constants.....	32
6.6.2 PHY PIB attributes.....	32
6.7 PHY specifications.....	33
6.7.1 Modulation and spreading.....	33
6.7.2 Forward error correction (FEC).....	35
6.8 Radio specifications.....	37
6.8.1 Transmit center frequency tolerance.....	37
6.8.2 Transmit power.....	37
6.8.3 Transmit PSD mask.....	37
6.8.4 Modulation accuracy.....	38
6.8.5 Spurious harmonically-related emission suppression.....	39

6.8.6	Receiver sensitivity.....	39
6.8.7	Adjacent LAS channel rejection.....	39
6.8.8	Receiver maximum input level of compliant beacon signal (blocking).....	40
6.8.9	Link quality indicator (LQI).....	40
7.	MAC sublayer specification.....	43
7.1	MAC sublayer service specification.....	43
7.1.1	MAC management service.....	43
7.1.2	Enumeration description.....	51
7.2	MAC beacon frame.....	52
7.2.1	MSF 1.....	52
7.2.2	MSF 2.....	58
7.2.3	MSF 3.....	61
7.3	MAC constants and MIB attributes.....	62
7.3.1	MAC constants.....	62
7.3.2	MIB attributes.....	62
7.4	MAC functional description.....	65
7.4.1	Transmission protocol.....	65
7.4.2	Retry procedure for an SPD.....	65
7.4.3	Frame reception and rejection.....	66
7.4.4	Device initialization procedure.....	67
7.4.5	Inter-device communications.....	68
7.4.6	Beaconing.....	71
7.4.7	Ceasing transmissions.....	75
7.4.8	Setting the internal operating states of the transceiver.....	77
7.5	Security suite specifications.....	79
7.5.1	Notation and representation.....	79
7.5.2	Representation of time.....	79
7.5.3	Elliptic-curve building blocks.....	80
7.5.4	Signature scheme.....	81
7.5.5	Certificate scheme.....	83
7.6	Message sequence charts (MSC) illustrating MAC-PHY interaction.....	85
Annex A	(informative) Example block decoding method.....	93
A.1	Introduction.....	93
A.2	Method details.....	93
A.3	Example.....	94
Annex B	(informative) Recommended deployment in the United States.....	97
B.1	Introduction.....	97
B.2	Physical positioning of the beacon relative to other Part 74 wireless systems.....	97
B.3	Additional beacon antenna deployment considerations.....	101
B.4	Recommended wireless system frequencies of operation.....	102
Annex C	(informative) Receiver sensitivity and adjacent channel protection justification.....	103
C.1	Introduction.....	103
C.2	Analysis.....	104
C.2.1	Link 1 Rx signal.....	104
C.2.2	Microphone interfering signal.....	105

Annex D (informative) Next higher layer (NHL) operation.....	107
D.1 Introduction.....	107
D.2 Next higher layer (NHL).....	107
D.3 Beaconing device types	108
D.4 Generic PD behavior.....	108
D.4.1 Searching for existing PDs	108
D.4.2 Deciding to become a PPD or an SPD after an initial search	109
D.4.3 Beacon frame construction and transmission	110
D.4.4 Beacon transmission failure-related enumeration values	111
D.5 PPD behaviors.....	112
D.5.1 Insertion of Rx period and ANP	112
D.5.2 Data aggregation	112
D.5.3 NPD selection	114
D.5.4 SPD lost indication	115
D.5.5 NPD lost indication.....	115
D.5.6 Planned termination of PPD transmission	115
D.5.7 Behavior following termination of SPD transmissions without warning	115
D.5.8 Behavior following planned termination of SPD transmissions.....	116
D.5.9 Behavior following termination of NPD transmissions without warning	116
D.5.10 Behavior following planned termination of NPD transmissions	116
D.5.11 PPD-specific MIB attributes	116
D.6 SPD behaviors.....	116
D.6.1 PPD beacon reception	117
D.6.2 Inter-device communication	117
D.6.3 Inter-device communication—retry failure handling	117
D.6.4 Inter-device communication—selection of the same RTS codeword	118
D.6.5 PPD protection of an SPD	118
D.6.6 Transmission of more than one beacon frame	119
D.6.7 Behavior on SPD receiving a PPD lost indication.....	119
D.6.8 Behavior on SPD receiving an NPD lost indication	119
D.6.9 Behavior if both PPD and NPD are lost	120
D.6.10 Behavior when selected as the NPD	121
D.6.11 Planned termination of SPD transmission	121
D.6.12 Behavior following termination of PPD transmissions without warning	121
D.6.13 Behavior following planned termination of PPD transmissions.....	122
D.6.14 Behavior following termination of NPD transmissions without warning	122
D.6.15 Behavior following planned termination of NPD transmissions	122
D.6.16 SPD-specific MIB attributes	122
D.7 NPD behaviors.....	122
D.7.1 Behavior on reception of the PPD NPD Indication subfield	123
D.7.2 Behavior on NPD receiving a PPD lost indication	123
D.7.3 Planned termination of NPD transmissions	123
D.7.4 Behavior following termination of PPD transmissions without warning	123
D.7.5 Behavior following planned termination of PPD transmissions.....	124
D.7.6 NPD-specific MIB attributes	124

Annex E (informative) Acquisition of time and location information	125
E.1 Introduction.....	125
E.2 Time	125
E.3 Location information	126
Annex F (informative)Concept of service primitives	129
Annex G (informative) Bibliography	131

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

Part 22.1: Standard to Enhance Harmful Interference Protection for Low-Power Licensed Devices Operating in the TV Broadcast Bands

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

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

1. Overview

1.1 General

National regulators are advancing regulations that allow license-exempt devices to operate on a non-interfering basis within the portions of the TV spectrum that are not used for broadcasts or required to remain unused in order to protect broadcast stations from interference. The Federal Communications Commission (FCC) in the United States of America has proposed to allow license-exempt devices to operate on a non-interfering basis within the portions of the TV spectrum that are not used for broadcasts or required to remain unused in order to protect broadcast stations from interference. It is expected that other regulatory bodies will take similar actions. Although the TV channels in these portions are not used for TV broadcasts, low-power, licensed devices, such as wireless microphones operated by broadcasters, do use these channels, and are entitled to protection by regulation to avoid disrupting incumbent services.

Furthermore, national regulators in many regions of the world are advancing regulations that allow license-exempt devices to operate on a non-interfering basis within the portions of the TV spectrum that are not used for broadcasts or required to remain unused in order to protect broadcast stations from interference.

This standard describes the air interface of devices that serve as warning beacons to protect the operation of incumbent licensed low-power devices. The beacons should be installed by the operator of the licensed device at a location appropriate to afford protection to the protected service. The beacons transmit identifiable sync bursts as well as, optionally, information about locations and operational parameters of the protected devices. The unlicensed system should include an appropriate receiver to receive and decode the information and should have an operation policy that would avoid inflicting any interference to the protected device.

1.2 Scope

This standard specifies methods to provide enhanced protection to protected devices such as those used in the production and transmission of broadcast programs [e.g., devices licensed as secondary under Title 47 of the Code of Federal Regulations (CFR) in the USA and equivalent devices in other regulatory domains] from harmful interference caused by license-exempt devices (e.g., IEEE P802.22TM^{1, 2}) that also are intended to operate in the TV Broadcast Bands.

1.3 Purpose

This standard provides a standard and efficient method for license-exempt devices to provide enhanced protection to low-powered licensed devices that are entitled to protection from harmful interference, and that share the same spectrum. This standard may be applicable in global regulatory environments.

¹Information on references can be found in Clause 2.

²This IEEE standards project was not approved by the IEEE-SA Standards Board at the time this publication went to press. For information about obtaining a draft, contact the IEEE.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used, so each referenced document is cited in text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

ANSI X9.63-2001, Public Key Cryptography for the Financial Services Industry—Key Agreement and Key Transport Using Elliptic Curve Cryptography, American Bankers Association, November 20, 2001.³

IEEE P802.22/Draft 5.0, Draft Standard for Local and metropolitan area networks, Part 22: Cognitive Wireless RAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Policies and procedures for operation in the TV Bands, October 2010.⁴

IEEE Std 1363TM-2000, IEEE Standard Specifications for Public-Key Cryptography.^{5, 6}

IEEE Std 1363aTM-2004, IEEE Standard Specifications For Public Key Cryptography—Amendment 1: Additional Techniques.

NIST Special Publication 800-57, “Recommendation for Key Management—Part 1: General.”⁷

Standards for Efficient Cryptography, SEC 4: Elliptic Curve Cryptography, Version 1.1r1, Certicom Research, June 9, 2006.⁸

U.S. Code of Federal Regulations, Title 47, Part 74, Subpart H.⁹

³ANSI publications are available from the Sales Department, American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (<http://www.ansi.org/>).

⁴Numbers preceded by P are IEEE authorized standards projects that were not approved by the IEEE-SA Standards Board at the time this publication went to press. For information about obtaining drafts, contact the IEEE.

⁵IEEE Publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, Piscataway, NJ 08854, USA (<http://standards.ieee.org>).

⁶The IEEE standards or products referred to in this clause are trademarks of the Institute of Electrical and Electronics Engineers, Inc.

⁷National Institute of Standards and Technology publications are available from <http://csrc.nist.gov/>.

⁸Standards for Efficient Cryptography publications are available from <http://www.secg.org/>.

⁹U.S. Regulatory Guides are available from the Superintendent of Documents, U.S. Government Printing Office, P.O. Box 37082, Washington, DC 20013-7082, USA (<http://www.access.gpo.gov/>).

3. Definitions

For the purposes of this document, the following terms and definitions apply. *The IEEE Standards Dictionary: Glossary of Terms & Definitions* should be referenced for terms not defined in this clause.¹⁰

3.1 beacon certificate: A certificate used by a beaconing device to convey its public key in a verifiable way to other entities.

3.2 beacon channel: The logical channel on which the PHY protocol data unit (PPDU) is sent. The PPDU comprises the 120-octet length beacon frame followed by 4 octets of all zeros.

3.3 beaconing device: A device that complies with IEEE Std 802.22.1.

3.4 beaconing network: A collection of devices consisting of a single primary protecting device (PPD) and the secondary protecting devices (SPDs) the PPD protects.

3.5 beacon private key: A key used by the beaconing device to create a digital signature of the beacon frame.

3.6 beacon public key: A key used by a receiving device to verify the authenticity of a digital signature that was created using the corresponding private key.

3.7 certificate authority: An entity responsible for provisioning asymmetric authentication data (public keys, digital certificates, etc.) to beaconing devices.

3.8 FCC Part 74 device: Low-power auxiliary stations as defined in the U.S. Code of Federal Regulations, Title 47, Part 74, Subpart H.

3.9 inter-device communication interval (ICD): A length of time, equal to one slot, consisting of a receive period (Rx period) and an acknowledgement/no acknowledgement period (ANP).

3.10 licensed auxiliary service (LAS) channel: A 200 kHz-wide portion of a TV channel. The lowest frequency LAS channel is centered 100 kHz above the lower edge of the TV channel, and the highest frequency LAS channel is centered 100 kHz below the upper edge of the TV channel. The number of LAS channels within a given TV channel depends on the width of the TV channel.

3.11 license-exempt device: A device that is tested and authorized by a regulatory body and is allowed to operate according to certain technical specifications without acquiring a license.

3.12 next-in-line protecting device (NPD): An secondary protecting device (SPD) that will become a primary protecting device (PPD) in the event that the already-existing PPD stops transmitting periodic beacon frames.

3.13 primary protecting device (PPD): A device that uses periodic beacons to protect its corresponding licensed device. Its protection may be extended to other licensed devices in the area [i.e., secondary protecting devices (SPDs)].

3.14 protected device: A low-power, licensed device that is being protected by a beaconing device.

3.15 protecting device (PD): A beaconing device that is protecting a low-power, licensed device. The PD may be either the primary protecting device (PPD) or a secondary protecting device (SPD).

¹⁰*The IEEE Standards Dictionary: Glossary of Terms & Definitions* is available at <http://shop.ieee.org/>.

3.16 secondary protecting device (SPD): A device that shares the responsibility of protecting its corresponding licensed device with the primary protecting device (PPD). An SPD occasionally sends beacons for the sole purpose of communicating with the PPD.

3.17 slot: An interval of time equal to 32 symbol times, or the duration of one sync burst.

3.18 synchronization channel: The logical channel consisting of a succession of 32-bit long synchronization bursts.

3.19 television (TV) channel: A contiguous segment of spectrum within the television broadcast bands, which may be 6, 7, or 8 MHz wide, depending on relevant regulations.

3.20 wireless microphone: A wireless microphone licensed for operation in the TV broadcast bands (e.g., an FCC Title 47, Part 74, Subpart H device in the United States).

3.21 wireless regional area network (WRAN) device: A device that is compliant with IEEE P802.22 or, once published, the latest published version of the standard.

4. Acronyms and abbreviations

ACK	acknowledgement
AGL	above ground level
ANP	acknowledgement/no acknowledgement period
ARIB	Association of Radio Industries and Businesses
ATSC	Advanced Television Systems Committee
CRC	cyclic redundancy check
DSSS	direct sequence spread spectrum
DQPSK	differential quadrature phase-shift keying
DTV	digital television
ECC	elliptic-curve cryptography
ECSSR-PV	elliptic-curve signature scheme with recovery, Pintsov-Vanstone version
ECSP-NR2/PV	elliptic curve pre-signature primitive, Nyberg-Rueppel and Pintsov-Vanstone version
ECSP-PV	elliptic curve signature primitive, Pintsov-Vanstone version
ECSSR-PV	elliptic-curve signature scheme with recovery, Pintsov-Vanstone version
ECVP-PV	elliptic curve verification primitive, Pintsov-Vanstone version
ED	energy detection
EIRP	effective isotropic radiated power
EMSR2	encoding method for signatures giving message recovery
ETSI	European Telecommunications Standards Institute
EVM	error vector magnitude
FCC	Federal Communications Commission
HDTV	high definition television
I2OSP	integer to octet string conversion primitive
ICI	inter-device communication interval
KDF2	key derivation function 2
LAS channel	licensed auxiliary service channel
LQI	link quality indicator
LSB	least significant bit
MAC	medium access control sublayer
MFR	MAC footer
MHR	MAC header
MIB	MAC information base
MIC	message integrity code
MLME	MAC sublayer management entity
MLME-SAP	MAC sublayer management entity service access point
MPDU	MAC protocol data unit
MSB	most significant bit
MSF	MAC subframe
MSI	manufacturer-specific information
NACK	no acknowledgement
NED	National Elevation Dataset
NHL	next higher layer
NPD	next-in-line protecting device
NST	next SPD superframe to transmit

NTSC	National Television System Committee
OSI	open systems interconnection
PD	protecting device
PD-SAP	PHY data service access point
PER	packet error rate
PHR	PHY header
PHY	physical layer
PIB	PHY information base
PLME	PHY layer management entity
PLME-SAP	PHY layer management entity service access point
PN	pseudo-random noise
PPD	primary protecting device
PPDU	PHY protocol data unit
PSDU	PHY service data unit
RATSC	symbol rate defined by the Advanced Television Systems Committee
RTS	request to send
Rx period	receive period
SAP	service access point
SHR	synchronization header
SPD	secondary protecting device
SRTM	Shuttle Radar Topography Mission
USGS	United States Geological Survey
UTC	universal coordinated time

5. Functional overview

5.1 Introduction

This standard defines the protocol and data formats for communication devices that form a beaconing network that offers an enhanced protection method to help prevent interference with low-power, licensed devices operating in television broadcast bands.

A brief introduction to the beaconing device architecture is given in 5.2, which is followed by a brief overview of the general functions of a beaconing network in 5.3 through 5.5. The general functions include information on the superframe structure, the data transfer model, and security.

Upon initialization, a protecting device (PD) is neither a primary protecting device (PPD) nor a secondary protecting device (SPD). Each PD monitors its TV channel for a random number of superframes to determine the presence or absence of a PPD. If no PPD is heard, the PD should act as a PPD and begin beaconing; if one or more PPDs are heard, under control of an upper layer, the PD may determine either to act as a PPD and initiate its own beacon, or as an SPD and attempt to contact a PPD. The device may take into consideration factors such as location and keep-out zone radius when making this decision. See 7.4.4 for more details on the device initialization procedure.

When a PPD is in operation, it should choose an SPD as the next-in-line protecting device (NPD). In the event that the PPD stops beaconing, the NPD shall promote itself from SPD to PPD. The remaining SPDs will eventually contact the new PPD, and the new PPD should choose an SPD to become the new NPD.

5.2 Architecture

The beaconing device architecture, shown in Figure 1, is a multilayer structure based on the open systems interconnection (OSI) seven-layer model. Each layer is responsible for one part of the protocol and offers services to the layers above it.

An IEEE 802.22.1 device comprises the physical layer (PHY), which contains the radio frequency (RF) transceiver along with its low-level control mechanism, and the medium access control sublayer (MAC), which provides access to the physical channel for the purpose of transferring information. The PHY and MAC are fully described in Clause 6 and Clause 7, respectively.

The upper layers provide services such as selecting an operating channel, deciding on an operating mode (i.e., PPD or SPD), starting and stopping beacon frame transmissions, processing incoming beacon frame information, aggregating data, and handling security errors. The definition of the upper layers is outside the scope of this standard. However, the intended behavior of the layer just above the MAC sublayer, which is known here as the next higher layer (NHL), is specified in Annex D.

The interfaces between the layers are called service access points (SAPs). Each SAP provides a means of passing information between two adjacent layers. More information on each particular SAP can be found in the text describing its associated layer (see 6.2 for information on PD-SAP, PLME-SAP, and RF-SAP; see 7.1 for information on MLME-SAP).

5.3 Superframe structure

The superframe structure repeats without interruption on a given TV channel occupied by an incumbent licensed wireless microphone or other licensed low-power devices. The superframe format is shown in Figure 2.

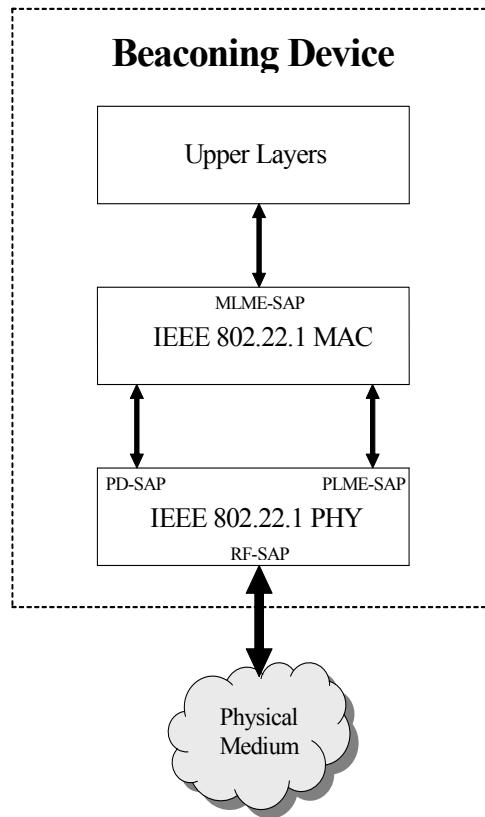


Figure 1—Beaconing device architecture


A large portion of the superframe structure is divided over two logical channels, which are transmitted in parallel. The synchronization logical channel consists of a succession of synchronization bursts (6.3), and the beacon logical channel consists of the PHY protocol data unit (PPDU), which contains the MAC beacon frame (7.2).

The superframe structure consists of a succession of 31 slots. Each slot is comprised of 32 DQPSK symbols, where one symbol has a duration of 1/9609.1 s. Each slot contains one synchronization burst, as well as a fixed number of PPDU bits. An inter-device communication interval (ICI) (6.5) may also be included. This format repeats on the TV channel without interruption as long as at least one PD is in operation. Note that the synchronization bursts and PPDU may be sent by either a PPD or an SPD.

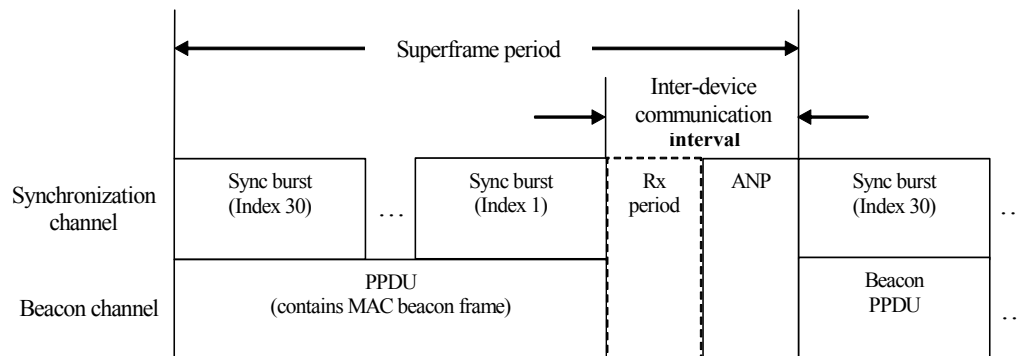
The synchronization bursts, which consists of a synchronization word followed by a decrementing index value, enable a receiver asynchronously sampling the radio channel to quickly determine when the next beacon will be sent. Parity bits follow the index value and provide error detection and correction on the index value.

The PPDU consists largely of the MAC beacon frame. The MAC beacon frame contains information relevant to the device or devices protected by the PD, including the physical location of the beaconing device and the estimated duration of TV channel occupancy.

Following thirty synchronization bursts and the PPDU, there may be an ICI, which includes a receive (Rx) period, during which the PPD pauses to monitor the TV channel for a request to send (RTS) burst transmitted by an SPD or an NPD codeword transmitted by the NPD, and an acknowledgement/no acknowledgement period (ANP), during which the PPD indicates whether it will transmit its own beacon frame during the next superframe or allow an SPD to transmit instead. During the initial transmission period, which lasts for *aInitializationPeriod* superframes (6.6.1), no ICI shall be included in the superframe structure.



(a) Initial transmission period



(b) Following the initial transmission period

Figure 2—Superframe logical format

If the PPD is in its initial transmission period and, therefore, the ICI is not included, the final index value in the series of synchronization bursts shall be zero and the next superframe shall start immediately after the end of this final synchronization burst [as illustrated in Figure 2 (a)]. While this synchronization burst with index zero is being transmitted on the synchronization channel, the beacon channel shall transmit all zeros. Because an SPD that has just joined the beaconing network shall not attempt to communicate on the channel for 100 superframes, no SPD will interfere with the PPD during its initial transmission period. Once the PPD has concluded its initial transmission period and, therefore, the ICI is included, the final index value in the series of synchronization bursts shall be one and the next superframe shall start after the ICI [as illustrated in Figure 2 (b)].

5.3.1 Synchronization burst structure

Figure 3 shows the structure of the synchronization burst sequence, which originates from within the PHY layer of a PD. Each synchronization burst contains a 15-bit synchronization word, an 8-bit parity field for detecting and correcting errors on the subsequent index value, a 7-bit index value that decrements with each

burst transmission, and a 2-bit reserved field. The synchronization burst sequence enables fast detection of the PPD or SPD that transmitted the sequence, while the decrementing index value identifies the start time of the next superframe transmission in multiples of slots. Each synchronization burst occupies one 32-bit-long synchronization channel slot of duration $32 \text{ bits}/9609.1 \text{ Hz} = 3.3301 \text{ ms}$.




Figure 3—Schematic view of synchronization burst sequence

5.3.2 Beacon frame structure

Figure 4 shows the structure of the beacon frame, which originates from within the MAC sublayer of either the PPD or an SPD. The beacon frame contains three MAC subframes (MSF). Each MSF is composed of a MAC header (MHR) and a MAC footer (MFR). The MHR in MSF 1 contains the three MAC parameter fields, the Source Address field, and the Location field. The MHR in MSF 2 contains the Map and Signature fields, while the MHR in MSF 3 contains the Certificate field; the Signature and Certificate fields are part of the public-key cryptography security solution (5.5.1). The MFR 1, MFR 2, and MFR 3 each contain a 2-octet cyclic redundancy check (CRC). The three MSFs together form the MAC protocol data unit (MPDU), which has the same contents as the unencoded PHY service data unit (PSDU).

The MAC beacon frame is passed to the PHY as the unencoded PSDU (or unencoded PHY payload). The MSF 1 in the PHY payload is protected by a half-rate convolutional code (6.7.2.2), and the addition of this code increases the payload length by 17 octets.

Because the length of a synchronization burst is 4 octets, zero padding is added following MSF 3 to ensure that the PPDU length is a multiple of four. Therefore, the overall length of the PPDU is 120 octets. The PHY payload and zero padding together form the PPDU (i.e., the PHY packet).

5.4 Data transfer model

All transmissions are broadcast. Transmitted data may be received and processed by any device in the area, including WRAN devices and other devices compliant with IEEE Std 802.22.1.

Two types of data transfer transactions exist for a device compliant with IEEE Std 802.22.1. The first type is the data transfer from the PPD to an SPD, in which the PPD transmits the data. The second type is the data transfer from an SPD to the PPD, in which the SPD transmits the data.

When the PPD wants to send data to an SPD, it does so by placing the information in its beacon PSDU. The SPD, which may be one of many in range of the PPD, monitors the beacon PSDU of the PPD, and decodes its address and recovers the message. No acknowledgement of data reception is provided at the MAC level. See Figure 5.

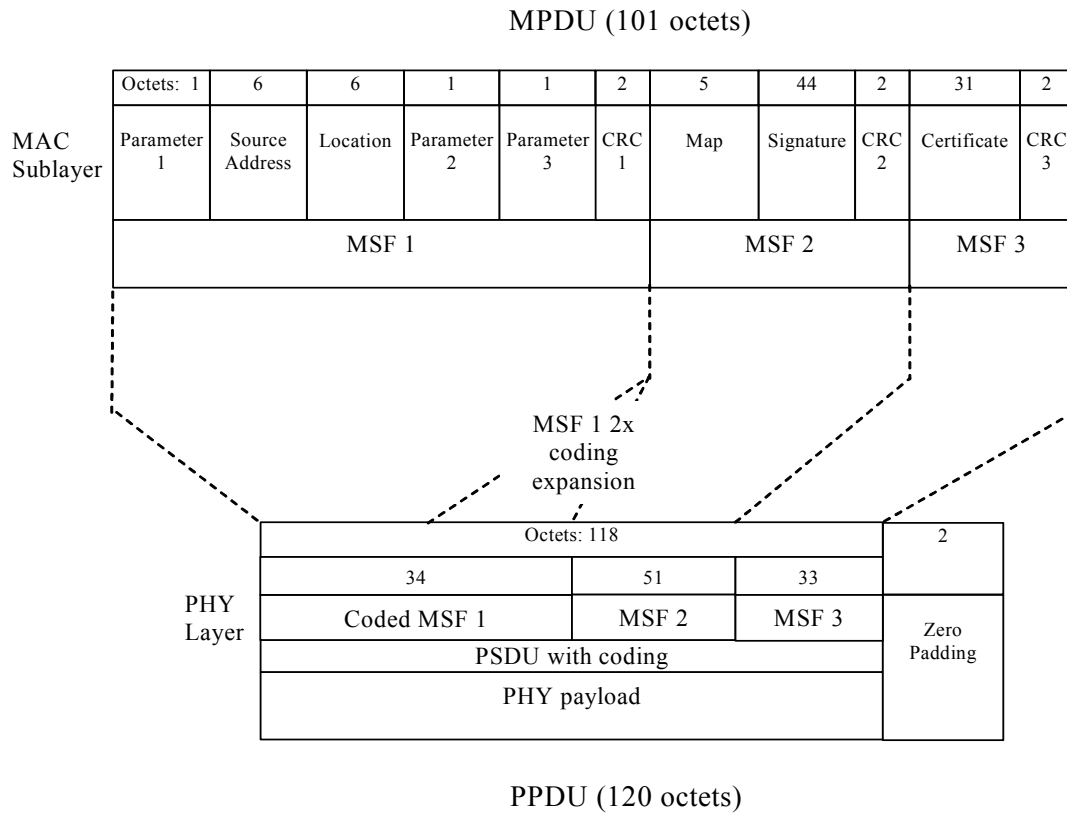


Figure 4—Schematic view of the beacon frame and the PHY packet (PPDU)




Figure 5—Communication from primary to secondary protecting devices

When an SPD wants to send data to the PPD, it first sends an RTS burst to the PPD during the Rx period of the PPD’s superframe. When the PPD receives the RTS burst, it may send the corresponding acknowledgement (ACK) during the ANP (6.5.2). Assuming the PPD does send the corresponding ACK, the PPD then yields the following superframe to the SPD, which then transmits its own superframe, containing its data and the synchronization channel, during this time. Following the transmission by the SPD, the PPD monitors the radio channel during the Rx period in order to check for the appearance of an NP codeword. The PPD then sends a no acknowledgement (NACK) during the subsequent ANP, since it is required to transmit at least every other beacon frame. Following that, the PPD resumes its normal transmission of the beacon frames. See Figure 6.

Although it is possible for more than one SPD to select and transmit the same RTS codeword during a given receive period and, therefore, possibly cause a collision, the probability is reduced, since there are multiple RTS codewords from which each SPD may choose. However, if this does happen and each SPD receives the ACK corresponding to this RTS codeword, each SPD shall execute the retry procedure (7.4.2). The retry procedure shall also be executed when an SPD sends an RTS codeword but either receives a NACK or an ACK that does not match the transmitted RTS codeword. If the SPD still does not receive the corresponding ACK following the conclusion of the retry procedure, the SPD shall not send a beacon frame.

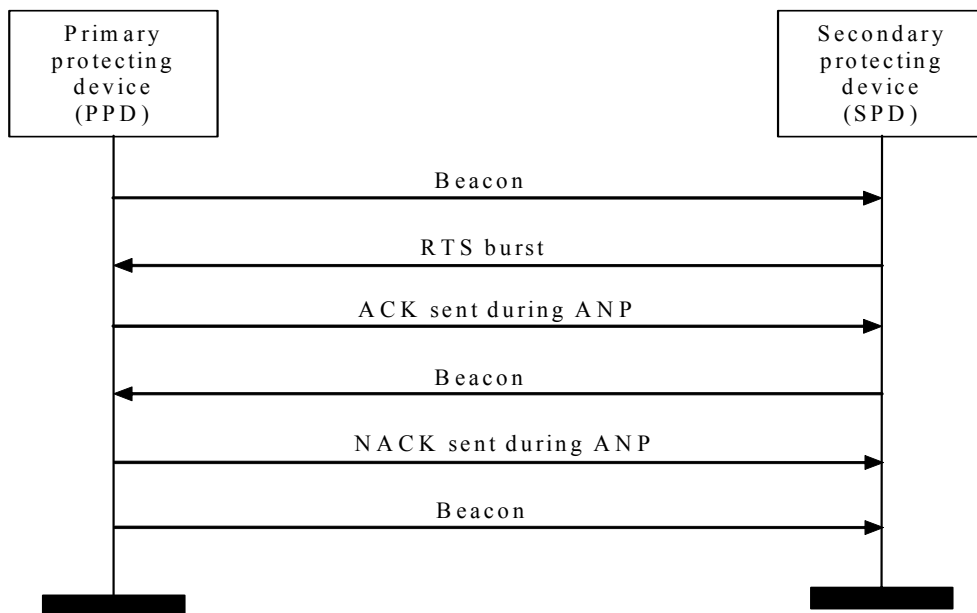


Figure 6—Communication from secondary to primary protecting devices

5.5 Security

It is important that the authenticity of a beacon frame can be determined by a device receiving the beacon frame. If authenticity cannot be determined, security issues could arise. For example, a rogue device may send out an illegitimate beacon to cause another device to leave a particular TV channel (a denial-of-service attack).

A device that operates in the radio frequency (RF) spectrum managed via beacons should, on a regular basis, listen for the potential presence of beacon frames. The security scheme chosen for this standard is based on asymmetric cryptography (i.e., public-key cryptography). Although this choice leads to a larger beacon frame over a symmetric-key approach, the effect of this larger frame size on system performance can be

minimized. The receiving device can trade off network downtime to intercept a beacon, but each MSF must be received sequentially and contiguously. For example MSF1 can be copied alone, or MSF1 plus MSF2, or MSF1 plus MSF2 plus MSF3, but each preceding MSF element must be copied in order to copy the following MSF element. Also, the burden of a larger beacon frame can be offset by a significant simplification in the area of key management. For example, because the beacon frame is comprised of three successive subframes, the receiving device can trade off network downtime and the level of trust to intercept a beacon, recognizing that by copying only MSF1 or MSF1 plus MSF2, the operator is accepting the risk that the beacon that may turn out to be illegitimate. Rather than relying on complex on-line verification (e.g., where symmetric keys needed for verification are kept within the infrastructure) or relying on complex key distribution schemes via tamper-resistant hardware, the asymmetric key approach requires minimal key management for receiving devices. Furthermore, since asymmetric cryptography leads to beacon frames that can be authenticated off-line, it is well-suited for securing inter-beacon communications, where the beaconing devices do not have a backend connection to the infrastructure.

5.5.1 Public-key approach

In this standard, each beaconing device, identified via its MAC address, shall be associated with a public key and corresponding private key. The beaconing device shall use its private key to generate a signature. Any receiving device can use the public key corresponding to a beaconing device's private key to verify a signature. However, a receiving device cannot create a valid signature simply through knowledge of the public key.

In one approach, a trusted authority would assign each beaconing device a unique private/public key pair and give each beaconing device a certificate for the public key. A receiving device would need the appropriate public key to verify a particular beaconing device's signatures. The appropriate public key can be obtained either via a database of MAC address/public key pairs or via a beacon's public-key certificate that a beaconing device shall send over-the-air as part of the beacon frame. When sent over-the-air, the public key is encapsulated into a certificate that is signed by an authority trusted by the receiving device. The authenticity of the received public key can be verified through the help of this certificate.

Figure 4 shows that MSF 2 shall contain the beaconing device's digital signature. This digital signature enables a receiving device to verify the authenticity and timeliness of the contents of the MHR in MSF 1 and the Map field in MSF 2. The signature shall be generated using the beaconing device's private key and can be verified using the receiving device's public key, as specified in 7.5.4. Figure 4 also shows that MSF 3 shall contain a public-key certificate. This certificate enables a receiving device to verifiably acquire the purported public key of the beaconing device. The certificate scheme used in this approach is specified in 7.5.5.

5.5.2 Security considerations

There are a variety of threats that need to be mitigated in a beaconing system. For example, a legitimate beacon could be replayed in an attempt to trick a receiving device into leaving a TV channel. To mitigate a replay attack, the signature is generated not only over the MHR in MSF 1, but also over time and date information. Relay attacks, where a device relays a valid beacon from one geographic location to another may also be possible. A relay attack is mitigated because the MHR contains location information, which is checked during the signature verification. A cryptanalysis attack may also be possible if weak encryption is used. Table 4 in NIST Special Publication 800-57¹¹ has recommended key lengths of cryptosystems for various time frames. It is desirable for this standard to be secure through the year 2030 time frame, so according to NIST, the security level should be at a minimum of 112 bits of strength. NIST also recommends that the cryptoperiod for a private signature key should be limited to 1–3 years (see Table 1 in

¹¹For information on references, see Clause 2.

NIST Special Publication 800-57). As cryptanalysis technology progresses, updates to these recommendations may be needed.

5.5.3 Beaconsing device life cycle

At time of manufacture, the beaconsing device shall be assigned a unique MAC address that will be embedded into its hardware. This MAC address shall resist tampering, i.e., manipulation and change, by any means (e.g., physical, electrical and software). The beaconsing device alone will not be able to generate beacon frames with legitimate signatures, so the end user will need to get a license to operate it. The end user will need to go to a licensing authority (e.g., an industry consortium or an appropriate regulatory body) and obtain a license for that particular beaconsing device (identified via its MAC address) that may be valid for a number of years. This licensing authority would be responsible for provisioning a private key and a public-key certificate to the beaconsing device. It is important that the private key remain a secret, so a good approach would be to provision this private key embedded in a tamper-resistant device such as a secure smart card. The end user would get this smart card, insert it into the beaconsing device, and the smart card would be used to create the beacon frame signatures. If a smart card for a beaconsing device was ever lost, stolen, or decommissioned, the end user could report this to the licensing authority so that a blacklist of devices could be maintained and periodically queried by receiving devices. The certificate would contain an expiration date so that end users would periodically need to renew their licenses to get new certificates (e.g., to enforce NIST's recommended cryptoperiod recommendation). Also, the certificate should be cryptographically bound to a specific beaconsing device's MAC address so that, if the security of a smart card is compromised, this smart card cannot be replicated and easily used with other beaconsing devices.

6. PHY specification

The PHY is responsible for the following tasks:

- Activation and deactivation of the radio transceiver
- Link quality indicator (LQI) for received packets
- TV channel frequency selection
- Data transmission and reception
- Synchronization of PDs
- Generating codewords during the inter-device communication interval (ICI)

Constants and attributes that are specified and maintained by the PHY are written in the text of this clause in italics. Constants have a general prefix of “a”, e.g., *aRegion*, and are listed in Table 18. Attributes have a general prefix of “phy”, e.g., *phyCurrentChannel*, and are listed in Table 19.

6.1 General requirements and definitions

This subclause describes the TV channel/licensed auxiliary service (LAS) channel information, modulation rates and beacon offset location, RF power measurement information, and receiver sensitivity definitions.

6.1.1 TV channel/licensed auxiliary service (LAS) channel information

The region of operation determines the width of the protected TV channels. Protected TV channels may be 6 MHz wide, 7 MHz wide, or 8 MHz wide. Each TV channel is divided into 200 kHz-wide LAS channels, with the lowest frequency LAS channel being centered 100 kHz above the lower edge of the TV channel and the highest frequency LAS channel being centered 100 kHz below the upper edge of the TV channel. A 6 MHz-wide TV channel has 30 LAS channels, while a 7 MHz-wide and an 8 MHz-wide TV channel have 35 and 40 LAS channels, respectively.

Equation (1) gives the location of the center frequency of LAS channel *M*.

$$\text{LAS channel } M \text{ center frequency} = f_{TV_N} - BW_{TV}/2 + (2M - 1) \times 100 \text{ kHz} \quad (1)$$

where

f_{TV_N} is the center frequency of TV channel *N*

BW_{TV} is the bandwidth of the TV channel

6.1.2 Modulation rates and beacon offset location for ATSC DTV regions

A beaconing device shall operate using the parameters in Table 1.

Table 1—Modulation rates and beacon offset location for ATSC DTV regions

Offset from lower TV channel edge (kHz)	Chip rate (kchips/s)	Symbol rate (kBaud)
309.4	76.873	9.6091

The beacon offset location was chosen due to its proximity to the DTV pilot frequency; the exact frequency of the DTV pilot frequency is 309.4406 kHz. An explanation of the origin of the chip and symbol rate are given in 6.7.1.6.

Figure 7 illustrates the position of the beacon offset frequency with respect to the LAS channel boundaries.




Figure 7—The position of the beacon offset frequency with respect to LAS channel boundaries

6.1.3 RF power measurement information

Unless otherwise stated, all RF power measurements, either transmit or receive, shall be made at the appropriate transceiver-to-antenna connector. The measurements shall be made with the transceiver terminated with the characteristic impedance of the antenna. For devices without an antenna connector, the measurements shall be interpreted as effective isotropic radiated power (EIRP) (i.e., employing a 0 dBi gain antenna), and any radiated measurements shall be corrected to compensate for the antenna gain in the implementation (ANSI C63.17-1998 [B1]).

6.1.4 Receiver sensitivity definitions

The definitions related to the receiver sensitivity and referred to in later subclauses are given in Table 2.

6.2 PHY service specifications

The PHY provides an interface between the MAC sublayer and the physical radio channel. The PHY conceptually includes a management entity called the PHY Layer Management Entity (PLME). The PLME provides a means of passing information between the MAC sublayer and the PHY layer (but not across the air interface). The PLME is also responsible for maintaining a database of variables pertaining to the PHY, which is called the PHY Information Base (PIB).

Figure 8 depicts the components and interfaces of the PHY.

The PHY provides two services, accessed through two service access points (SAPs): the PHY data service, accessed through the PHY data SAP (PD-SAP), and the PHY management service, accessed through the PLME-SAP. The RF-SAP provides a means to pass information across the air interface between the PHY of the present device and that of another device.

Table 2—Receiver sensitivity definitions

Term	Definition of term	Conditions
Packet	One synchronization word, one index value, one MAC subframe of the beacon channel, or one RTS/ANP burst.	<ul style="list-style-type: none"> – Synchronization word packet = 15 bits – Index value packet = 15 bits – MSF1 beacon packet = 17 uncoded octets (coded: 34 octets using 1:2 encoding) – MSF2 beacon packet = 51 octets – MSF3 beacon packet = 33 octets – RTS/ANP burst packet = 2 octets
Packet error rate (PER)	Average fraction of transmitted packets that are not correctly received.	Averaged over 10 000 superframes in a Gaussian channel with structured MAC synchronization channel data and random MAC beacon channel data.
Sensitivity	Threshold input signal power that yields a specified PER.	<ul style="list-style-type: none"> – PER = 1%. – Power measured at antenna terminals. – Sensitivities of different packet types will differ due to packet sizes and coding strength (6.8.6).

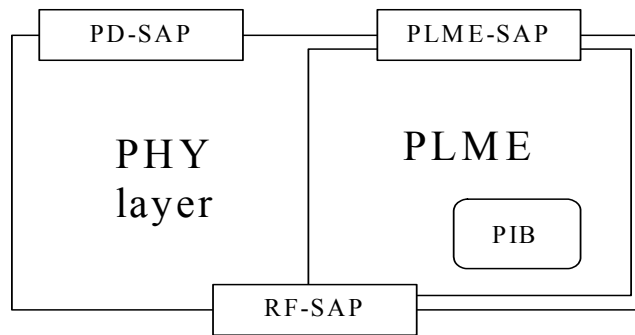


Figure 8—PHY reference model

6.2.1 PHY data service

The PD-SAP is an interface that provides a means of passing MPDUs between the MAC sublayer and the PHY layer. In the case of this standard, the MPDU always contains a beacon frame. Table 3 lists the primitives¹² supported by the PD-SAP. These primitives are discussed in 6.2.1.1 through 6.2.1.3.

Table 3—PD-SAP primitives

PD-SAP primitive	Request	Confirm	Indication
PD-DATA	6.2.1.1	6.2.1.2	6.2.1.3

¹²See Annex F for an overview of the concept of service primitives.

6.2.1.1 PD-DATA.request

The PD-DATA.request primitive is generated by a MAC sublayer entity and issued to its PHY entity to request the transfer of an MPDU (i.e., PSDU). Table 4 specifies the parameters for the PD-DATA.request primitive.

Table 4—PD-DATA.request parameters

Name	Type	Valid range	Description
psdu	Set of octets	—	The set of octets forming the PSDU to be transmitted by the PHY entity.

On receipt of the PD-DATA.request primitive by the PHY entity, the PHY will attempt to transmit the supplied PSDU. Provided the transmitter is enabled, the PHY will construct and transmit a PPDU containing the supplied PSDU. When the PHY entity has completed the transmission, it will issue the PD-DATA.confirm primitive with a status of SUCCESS.

6.2.1.2 PD-DATA.confirm

The PD-DATA.confirm primitive is generated by the PHY entity and issued to its MAC sublayer entity in response to a PD-DATA.request primitive, in order to confirm the end of the transmission attempt of an MPDU (i.e., PSDU) from a local PHY layer entity to a peer PHY layer entity. This primitive has no parameters.

6.2.1.3 PD-DATA.indication

The PD-DATA.indication primitive is generated by the PHY entity and issued to its MAC sublayer entity to transfer a received MPDU (i.e., PSDU). Table 5 specifies the parameters for the PD-DATA.indication primitive.

Table 5—PD-DATA.indication parameters

Name	Type	Valid range	Description
psdu	Set of octets	—	The set of octets forming the PSDU received by the PHY entity.
ppduLinkQuality	Integer	0x00–0xff	The LQI value measured during reception of the PPDU (6.8.9).

6.2.2 PHY management service

The PLME-SAP is an interface that provides a means of passing information between the MAC sublayer and the PHY layer via the MAC sublayer management entity (MLME) and the PLME. Table 6 lists the primitives supported by the PLME-SAP. These primitives are discussed in the clauses referenced in the table.

Table 6—PLME-SAP primitives

PLME-SAP primitive	Request	Indication	Confirm
PLME-ANP-DECISION	6.2.2.1	—	6.2.2.2
PLME-ANP-RESPONSE	—	6.2.2.3	—
PLME-GET	6.2.2.4	—	6.2.2.5
PLME-INITIALIZE	6.2.2.6	—	6.2.2.7
PLME-INITIATE-RTS-BURST	6.2.2.8	—	6.2.2.9
PLME-NPD-ACTIVE	—	6.2.2.10	—
PLME-NPD-HEARTBEAT	6.2.2.11	—	6.2.2.12
PLME-SET	6.2.2.13	—	6.2.2.14
PLME-SET-TRX-STATE	6.2.2.15	—	6.2.2.16

6.2.2.1 PLME-ANP-DECISION.request

The PLME-ANP-DECISION.request primitive is generated by the MLME of the PPD and issued to its PLME to request that the PHY entity send a particular response during the ANP. Table 7 specifies the parameter for the PLME-ANP-DECISION.request primitive.

Table 7—PLME-ANP-DECISION.request parameter

Name	Type	Valid range	Description
ANPResponse	Enumeration	ACK, NACK, GO_ON	The response that the MAC sublayer requests to be sent during the ANP.

On receipt of the PLME-ANP-DECISION.request primitive, the PLME will consider both the ANPResponse parameter received in PLME-ANP-DECISION.request primitive and what was received over the air during the Rx period before making a final decision on which response to send (6.5.2).

6.2.2.2 PLME-ANP-DECISION.confirm

The PLME-ANP-DECISION.confirm primitive is generated by the PLME of the PPD and issued to its MLME in response to a PLME-ANP-DECISION.request primitive. This primitive contains the response transmitted by the PHY layer during the ANP. Table 8 specifies the parameter for the PLME-ANP-DECISION.confirm primitive.

On receipt of the PLME-ANP-DECISION.confirm primitive, the MLME of the PPD is notified of the response transmitted by the PHY layer during the ANP. An ANPResponse parameter equal to either ACK or GO_ON indicates that the PPD will allow an SPD to transmit its beacon frame during the upcoming superframe. An ANPResponse parameter equal to NACK indicates that the PPD will transmit its own beacon frame during the upcoming superframe.

Table 8—PLME-ANP-DECISION.confirm parameter

Name	Type	Valid range	Description
ANPResponse	Enumeration	ACK, NACK, GO_ON	The response transmitted by the PHY layer during the ANP.

6.2.2.3 PLME-ANP-RESPONSE.indication

The PLME-ANP-RESPONSE.indication primitive is generated by the PLME of an SPD and issued to its MLME as a notification of the response received from the PPD during the ANP. This primitive is only issued if the SPD did not send an RTS burst in the Rx period immediately preceding the ANP; see 7.4.5.2 for details on the usage of this primitive. Note that when an SPD does send an RTS burst, it will receive the response via the PLME-INITIATE-RTS-BURST.response primitive (6.2.2.9).

Table 9 specifies the parameter for the PLME-ANP-RESPONSE.indication primitive.

Table 9—PLME-ANP-RESPONSE.indication parameter

Name	Type	Valid range	Description
ANPResponse	Enumeration	ACK, NACK, GO_ON	The response received during the ANP.

6.2.2.4 PLME-GET.request

The PLME-GET.request primitive is generated by the MLME and issued to its PLME to obtain information from the PIB about one or more PIB attributes. Table 10 specifies the parameters for the PLME-GET.request primitive.

Table 10—PLME-GET.request parameters

Name	Type	Valid range	Description
PIBAttribute1	Integer	See Table 19	The identifier of the PIB attribute being requested.
PIBAttribute2	Integer	See Table 19	The identifier of the PIB attribute being requested.
...			
PIBAttributeN	Integer	See Table 19	The identifier of the PIB attribute being requested.

On receipt of the PLME-GET.request primitive, the PLME attempts to retrieve the requested PIB attribute or attributes from its database. Each successfully retrieved PIB attribute will have a corresponding status of SUCCESS in the PLME-GET.confirm primitive. Any PIB attribute value(s) that could not be retrieved will have a corresponding status indicating the appropriate error status (6.2.2.5).

6.2.2.5 PLME-GET.confirm

The PLME-GET.confirm primitive is generated by the PLME and issued to its MLME in response to a PLME-GET.request primitive, and it reports the results of an information request from the PIB. Table 11 specifies the parameters for the PLME-GET.confirm primitive.

Table 11—PLME-GET.confirm parameters

Name	Type	Valid range	Description
Status1	Enumeration	SUCCESS or ATTRIBUTE_NOT_FOUND	The result of the request for PIB attribute information.
PIBAttribute1	Integer	See Table 19	The identifier of the PIB attribute that was requested.
PIBAttributeValue1	Various	Attribute specific	The value of the indicated PIB attribute that was requested. This parameter has zero length when the Status parameter is set to ATTRIBUTE_NOT_FOUND.
Status2	Enumeration	SUCCESS or ATTRIBUTE_NOT_FOUND	The result of the request for PIB attribute information.
PIBAttribute2	Integer	See Table 19	The identifier of the PIB attribute that was requested.
PIBAttributeValue2	Various	Attribute specific	The value of the indicated PIB attribute that was requested. This parameter has zero length when the Status parameter is set to ATTRIBUTE_NOT_FOUND.
...			
StatusN	Enumeration	SUCCESS or ATTRIBUTE_NOT_FOUND	The result of the request for PIB attribute information.
PIBAttributeN	Integer	See Table 19	The identifier of the PIB attribute that was requested.
PIBAttributeValueN	Various	Attribute specific	The value of the indicated PIB attribute that was requested. This parameter has zero length when the Status parameter is set to ATTRIBUTE_NOT_FOUND.

On receipt of the PLME-GET.confirm primitive, the MLME is notified of the results of its request to read one or more PIB attributes. If the request to read a PIB attribute was successful, the corresponding Status parameter is set to SUCCESS. If the identifier of a PIB attribute was not found in the database, the corresponding Status parameter is set to ATTRIBUTE_NOT_FOUND.

6.2.2.6 PLME-INITIALIZE.request

The PLME-INITIALIZE.request primitive is generated by the MLME of the PPD and issued to its PLME upon receipt of the MLME-START-BEACON.request primitive with the Initialize parameter set to TRUE to request that the PD enter the initial transmission period. If the PD is a PPD, the PHY entity will start to transmit a continuous series of superframes, which shall not include inter-device communication intervals, for a period of *aInitializationPeriod* superframes. If the PD is an SPD, the PHY entity will wait for a period of *aInitializationPeriod* superframes before attempting to send an RTS burst. This primitive has no parameters.

6.2.2.7 PLME-INITIALIZE.confirm

The PLME-INITIALIZE.confirm primitive is generated by the PLME of the PPD and issued to its MLME in response to a PLME-INITIALIZE.request primitive once the initial transmission period is over. This primitive has no parameters.

6.2.2.8 PLME-INITIATE-RTS-BURST.request

The PLME-INITIATE-RTS-BURST.request primitive is generated by the MLME of an SPD and issued to its PLME to request that the PHY entity generate an RTS burst packet. The PLME-INITIATE-RTS-BURST.request primitive has no parameters.

Once the burst packet is transmitted and the ANP concludes, the PLME will issue the PLME-INITIATE-RTS-BURST.confirm primitive with the appropriate status (6.2.2.9).

6.2.2.9 PLME-INITIATE-RTS-BURST.confirm

The PLME-INITIATE-RTS-BURST.confirm primitive is generated by the PLME of an SPD and issued to its MLME in response to a PLME-INITIATE-RTS-BURST.request primitive. This primitive confirms that an RTS burst packet was sent by the PHY entity and indicates whether the SPD has permission to send a beacon frame. Table 12 specifies the parameters for the PLME-INITIATE-RTS-BURST.confirm primitive.

Table 12—PLME-INITIATE-RTS-BURST.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	ACK, INCORRECT_ACK, or NACK	The result of the request to send a beacon frame.

On receipt of the PLME-INITIATE-RTS-BURST.confirm primitive, the MLME is notified whether its request to send a beacon frame was accepted or rejected. The Status parameter shall be set equal to ACK if the request was accepted. If the request was rejected due to the receipt of an acknowledgement codeword that does not match the transmitted RTS codeword (e.g., a competing SPD transmitted a different RTS codeword during the same Rx period), the Status parameter shall be set equal to INCORRECT_ACK. If the request was rejected due to the receipt of a NACK, the Status parameter shall be set equal to NACK.

6.2.2.10 PLME-NPD-ACTIVE.indication

The PLME-NPD-ACTIVE.indication primitive is generated by the PLME of a PD and issued to its MLME as a notification that an NPD codeword has been detected. This primitive has no parameters.

6.2.2.11 PLME-NPD-HEARTBEAT.request

The PLME-NPD-HEARTBEAT.request primitive is generated by the MLME of the SPD that was selected as the NPD and issued to its PLME to request that the PHY entity send the NPD codeword during the upcoming receive period. The PLME-NPD-HEARTBEAT.request primitive has no parameters.

Once the NPD codeword is transmitted, the PLME will issue the PLME-NPD-HEARTBEAT.confirm primitive.

6.2.2.12 PLME-NPD-HEARTBEAT.confirm

The PLME-NPD-HEARTBEAT.confirm primitive is generated by the PLME of the SPD that was selected as the NPD and issued to its MLME in response to a PLME-NPD-HEARTBEAT.request primitive. This primitive confirms that the NPD codeword was sent by the PHY entity. This primitive has no parameters.

On receipt of the PLME-NPD-HEARTBEAT.confirm primitive, the MLME is notified that its request to send the NPD codeword was granted.

6.2.2.13 PLME-SET.request

The PLME-SET.request primitive is generated by the MLME and issued to its PLME to attempt to set the indicated PIB attribute(s) to the given value(s). This primitive may have several pairs of parameters, with each pair specifying a PIB attribute and its corresponding value. Table 13 specifies the parameters for the PLME-SET.request primitive.

Table 13—PLME-SET.request parameters

Name	Type	Valid range	Description
PIBAttribute1	Enumeration	See Table 19	The identifier of the PIB attribute to set.
PIBAttributeValue1	Various	Attribute specific	The value of the indicated PIB attribute to set.
PIBAttribute2	Enumeration	See Table 19	The identifier of the PIB attribute to set.
PIBAttributeValue2	Various	Attribute specific	The value of the indicated PIB attribute to set.
...			
PIBAttributeN	Enumeration	See Table 19	The identifier of the PIB attribute to set.
PIBAttributeValueN	Various	Attribute specific	The value of the indicated PIB attribute to set.

On receipt of the PLME-SET.request primitive, the PLME attempts to write the given value or values to the indicated PIB attributes in its database. Each successfully written PIB attribute will have a corresponding status of SUCCESS in the PLME-SET.confirm primitive. Any PIB attribute value(s) that could not be written will have a corresponding status indicating the appropriate error status (6.2.2.14).

6.2.2.14 PLME-SET.confirm

The PLME-SET.confirm primitive is generated by the PLME and issued to its MLME in response to a PLME-SET.request primitive, and it reports the results of the attempt to write one or more values to their corresponding PIB attributes. Table 14 specifies the parameters for the PLME-SET.confirm primitive.

Table 14—PLME-SET.confirm parameters

Name	Type	Valid range	Description
Status1	Enumeration	SUCCESS, ATTRIBUTE_NOT_FOUND, or INVALID_PARAMETER	The status of the attempt to set the requested PIB attribute.
PIBAttribute1	Enumeration	See Table 19	The identifier of the PIB attribute being confirmed.
Status2	Enumeration	SUCCESS, ATTRIBUTE_NOT_FOUND, or INVALID_PARAMETER	The status of the attempt to set the requested PIB attribute.
PIBAttribute2	Enumeration	See Table 19	The identifier of the PIB attribute being confirmed.
...			
StatusN	Enumeration	SUCCESS, ATTRIBUTE_NOT_FOUND, or INVALID_PARAMETER	The status of the attempt to set the requested PIB attribute.
PIBAttributeN	Enumeration	See Table 19	The identifier of the PIB attribute being confirmed.

On receipt of the PLME-SET.confirm primitive, the MLME is notified of the result of its request to set the value of one or more PIB attributes. If a requested value was successfully written to the indicated PIB attribute, the Status parameter is set to SUCCESS.

If a PIBAttribute parameter specifies an attribute that was not found in the database, the corresponding Status parameter is set to ATTRIBUTE_NOT_FOUND. If a PIBAttributeValue parameter specifies a value that is out of the valid range for the given attribute, the corresponding Status parameter is set to INVALID_PARAMETER.

6.2.2.15 PLME-SET-TRX-STATE.request

The PLME-SET-TRX-STATE.request primitive is generated by the MLME and issued to its PLME to request that the PHY entity change the internal operating state of the transceiver during a superframe period or periods. Table 15 specifies the parameters for the PLME-SET-TRX-STATE.request primitive.

The transceiver has the following three main states:

- Transceiver disabled (TRX_OFF)
- Transmitter enabled (TX_ON)
- Receiver enabled (RX_ON)

On receipt of the PLME-SET-TRX-STATE.request primitive, the PLME will cause the PHY to change the operating state of the transceiver, according to the parameters of this primitive, during the appropriate periods of the following superframe(s). Once the state change is accepted, the PLME will issue the PLME-SET-TRX-STATE.confirm primitive.

For the special case of device initialization (7.4.4), the concept of a superframe period does not yet exist. In order to start the search portion of the device initialization procedure (i.e., following the receipt of the MLME-SEARCH.request primitive by the MLME), a PLME-SET-TRX-STATE.request primitive will be issued with the parameters set in the following way: TRX_State1 = RX_ON, TRX_State2 = RX_ON,

Table 15—PLME-SET-TRX-STATE.request parameters

Name	Type	Valid range	Description
TRX_State1	Enumeration	RX_ON, TX_ON, or TRX_OFF	The new state in which to configure the transceiver during the beaconing portion of the superframe.
TRX_State2	Enumeration	RX_ON, TX_ON, or TRX_OFF	The new state in which to configure the transceiver during the Rx period.
TRX_State3	Enumeration	RX_ON, TX_ON, or TRX_OFF	The new state in which to configure the transceiver during the ANP.
Periodic	Boolean	TRUE or FALSE	If this parameter is TRUE, the operating states of the transceiver will be valid for each superframe that follows, unless a new PLME-SET-TRX-STATE.request primitive is received from the MLME. Otherwise, the operating states will only apply to the upcoming superframe.

TRX_State3 = RX_ON, and Periodic = TRUE. The receiver will remain on until the search portion is over and a new PLME-SET-TRX-STATE.request primitive is issued.

If, upon conclusion of the search portion, the PD promotes itself to PPD, a new PLME-SET-TRX-STATE.request primitive will be issued in order to start the initial transmission period with the parameters set in the following way: TRX_State1 = TX_ON, TRX_State2 = TX_ON, TRX_State3 = TX_ON, and Periodic = TRUE. The transmitter will remain on until the device initialization procedure is over and a new PLME-SET-TRX-STATE.request primitive is issued with the parameters set in the following way: TRX_State1 = TX_ON, TRX_State2 = RX_ON, TRX_State3 = TX_ON, and Periodic = TRUE.

6.2.2.16 PLME-SET-TRX-STATE.confirm

The PLME-SET-TRX-STATE.confirm primitive is generated by the PLME and issued to its MLME upon receipt of a PLME-SET_TRX_STATE.request primitive. This primitive has no parameters.

6.2.3 PHY enumerations description

The enumeration values used by the PHY data and management primitives are given in Table 37.

6.3 Synchronization burst

A synchronization burst is composed of four fields: the Sync field, the Parity field, the Index field, and the Reserved field. The format of the synchronization burst is shown in Figure 9.




Figure 9—Format of the synchronization burst

The Sync field is used by a receiver to detect the presence of the synchronization burst and to synchronize to the slot timing. The Sync field shall have the value shown in Table 16.

Table 16—Format of the Sync field

Bits	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Name	s ₀	s ₁	s ₂	s ₃	s ₄	s ₅	s ₆	s ₇	s ₈	s ₉	s ₁₀	s ₁₁	s ₁₂	s ₁₃	s ₁₄
Value	1	1	1	1	0	1	0	1	1	0	0	1	0	0	0

The Index field is used to obtain frame synchronization with an incoming beacon. It contains a numerical value equal to the number of slots remaining (not including the present slot) before the start of the next superframe. The Index field shall be decremented by one each time the data contained within a slot is transmitted until the index reaches either zero or one, depending on whether the PPDU will be followed by a Rx period (6.5.1) and a corresponding acknowledgement/no acknowledgement period (ANP) burst (6.5.2).

If the PPDU will not be followed by a Rx period and a corresponding ANP, the final index shall be zero, and the next superframe shall start immediately thereafter. If the PPDU is to be followed by a Rx period and an ANP, the final index shall be one, and the next superframe shall start immediately after the Rx period and corresponding ANP.

To improve the receive performance of the Index field, a (15, 7) linear block code shall be applied. The Parity field shall be used for both error-correction and error-detection of the index (6.7.2).

The two bits in the Reserved field shall be set to zero.

6.4 PPDU format

Each PPDU packet consists of a constant length payload of length 120 octets, which carries the MAC sublayer beacon frame. The beacon frame is divided into three MAC subframes (MSF), as shown in Figure 4 of Clause 5, and is passed to the PHY layer for inclusion in the PHY payload.

A half-rate convolutional code with puncturing shall be applied by the PHY layer to the first MAC subframe (MSF 1), which includes the three MAC parameter fields, the Source Address field, the Location field, and a 2-octet CRC. The addition of this convolutional code with puncturing increases the PHY payload length by 17 octets, bringing the total length to 118 octets. Then, because the length of a synchronization burst is 4 octets, zero padding shall be added to the end of the PPDU (i.e., following MSF 3) to ensure that its length is a multiple of four. Therefore, the overall length of the PPDU is 120 octets.

For the PPDU packet structure, the multiple octet field shall be transmitted or received least significant octet first and each octet shall be transmitted or received least significant bit (LSB) first.

6.5 Inter-device communication interval (ICI)

The ICI is only excluded from the superframe if the PPD is executing the device initialization procedure (7.4.4). This period consists of the Rx period and the ANP. It is 32 symbols in duration, corresponding to the slot in the superframe immediately following the 30 synchronization bursts and the PPDU. The order of symbols within the ICI is as follows:

- a) 3 symbols of turnaround time
- b) An 11 symbol Rx period composed of
 - 2 symbols of AGC training (since the RTS will be coming from an unfamiliar transmitter)
 - A phase reference symbol for the RTS burst
 - The 8 symbol RTS burst
- c) 4 symbols of turnaround time

- d) A 9 symbol ANP burst composed of
 - A phase reference symbol for the ANP burst
 - The 8 ANP symbols
- e) 4 symbols of turnaround time
- f) A phase reference symbol for the start of the next superframe

The 3 symbol turn-around time is composed of 1 symbol (8 chips) of PPD superframe ramp-down, 1 symbol of propagation delay, and 1 symbol of RTS ramp-up. The first 4 symbol turn-around time has 1 symbol of RTS ramp-down, 2 symbols for propagation delay and processing, and 1 symbol of ANP ramp-up. The second 4 symbol turn-around time has 1 symbol of ANP ramp-down, 2 symbols of propagation delay and processing, and 1 symbol of ramp-up for the next superframe. Figure 10 illustrates an ICI during which an RTS burst is sent by an SPD and an ANP burst is sent by the PPD.

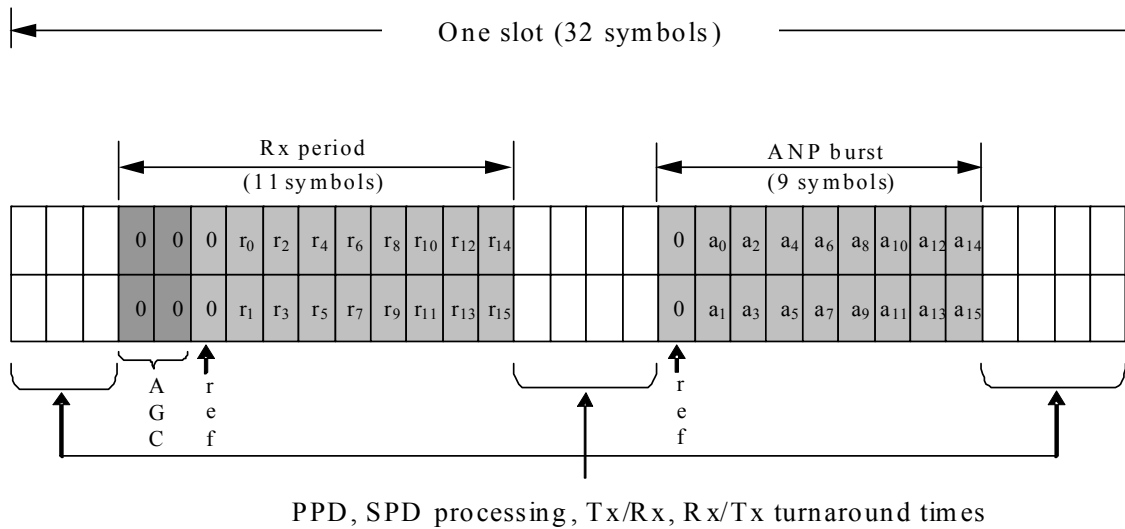


Figure 10—ICI structure

6.5.1 Receive (Rx) period

The Rx period is a length of time equal to 11 symbols that has two purposes. The first is for the PPD to listen for an RTS burst transmission from an SPD. The second is for the PPD to listen for an NPD codeword transmission from the NPD in order to learn whether the NPD is still active.

For convenience, the Rx period structure is presented so that the leftmost field as written in this standard shall be transmitted or received first. All multiple octet fields shall be transmitted or received least significant octet first and each octet shall be transmitted or received LSB first.

The RTS burst is used by an SPD to reserve the next superframe to transmit its beacon frame to the PPD. Each RTS burst consists of an RTS codeword field, and the list of available RTS codewords is given in Table 17, with the exception of the codewords with indices 1 and 32. Note that these two codewords shall not be available as RTS codewords but shall instead be reserved for special purposes that will be explained in the following paragraphs. An SPD shall randomly select an RTS codeword from the list to send in the Rx period; this process shall be repeated prior to each transmission (including retransmissions).

When more than one SPD simultaneously transmits an RTS codeword during the Rx period, the PPD may be able to detect one or more SPDs by their RTS codewords, since each SPD may select a different RTS codeword (6.5.2). In the unlikely event that more than one SPD simultaneously transmits the same RTS

Table 17—List of available codewords

Index	Codeword															
	r ₀ / a ₀	r ₁ / a ₁	r ₂ / a ₂	r ₃ / a ₃	r ₄ / a ₄	r ₅ / a ₅	r ₆ / a ₆	r ₇ / a ₇	r ₈ / a ₈	r ₉ / a ₉	r ₁₀ / a ₁₀	r ₁₁ / a ₁₁	r ₁₂ / a ₁₂	r ₁₃ / a ₁₃	r ₁₄ / a ₁₄	r ₁₅ / a ₁₅
1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
3	0	0	0	0	1	1	1	1	0	0	0	0	1	1	1	1
4	0	0	0	0	1	1	1	1	1	1	1	1	0	0	0	0
5	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1
6	0	0	1	1	0	0	1	1	1	1	0	0	1	1	0	0
7	0	0	1	1	1	1	0	0	0	0	1	1	1	1	0	0
8	0	0	1	1	1	1	0	0	1	1	0	0	0	0	1	1
9	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1
10	0	1	0	1	0	1	0	1	1	0	1	0	1	0	1	0
11	0	1	0	1	1	0	1	0	0	1	0	1	1	0	1	0
12	0	1	0	1	1	0	1	0	1	0	1	0	0	1	0	1
13	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1	0
14	0	1	1	0	0	1	1	0	1	0	0	1	1	0	0	1
15	0	1	1	0	1	0	0	1	0	1	1	0	1	0	0	1
16	0	1	1	0	1	0	0	1	1	0	0	1	0	1	1	0
17	1	0	0	1	0	1	1	0	0	1	1	0	1	0	0	1
18	1	0	0	1	0	1	1	0	1	0	0	1	0	1	1	0
19	1	0	0	1	1	0	0	1	0	1	1	0	0	1	1	0
20	1	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1
21	1	0	1	0	0	1	0	1	0	1	0	1	1	0	1	0
22	1	0	1	0	0	1	0	1	1	0	1	0	0	1	0	1
23	1	0	1	0	1	0	1	0	0	1	0	1	0	1	0	1
24	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
25	1	1	0	0	0	0	1	1	0	0	1	1	1	1	0	0
26	1	1	0	0	0	0	1	1	1	1	0	0	0	0	1	1
27	1	1	0	0	1	1	0	0	0	0	1	1	0	0	1	1
28	1	1	0	0	1	1	0	0	1	1	0	0	1	1	0	0
29	1	1	1	1	0	0	0	0	0	0	0	0	1	1	1	1
30	1	1	1	1	0	0	0	0	1	1	1	1	0	0	0	0
31	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0
32	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

codeword, the PPD will probably correctly detect the RTS codeword from one of the SPDs, which will lead to the simultaneous transmission of the SPD's beacon frames during the following superframe. However, this will not cause any lasting problems for either the SPD or PPD (7.4.5.2).

The NPD codeword is sent by the NPD during the Rx period and is used to inform the PD that the NPD is still active. The NPD codeword sequence is the codeword with index 1 in Table 17.

6.5.2 Acknowledgement/no acknowledgement period (ANP)

The ANP is a length of time equal to 9 symbols that is used by the PPD to respond to information received during previous Rx periods and, more generally, to communicate with other PDs. For convenience, the ANP structure is presented so that the leftmost (least significant) bit as written in this standard shall be transmitted or received first.

The ANP burst is used by the PPD to indicate whether it will transmit its own beacon frame during the next superframe, or whether an SPD will be allowed to transmit its beacon frame during this time. The PPD may allow transmission of an SPD beacon frame if an RTS burst was detected by the PPD during the Rx period immediately preceding the ANP burst. If it has approved a request, the PPD shall transmit an ACK. Otherwise, it shall transmit a NACK. The ACK sequences are given in Table 17, with the exception of the codewords with indices 1 and 32. The NACK sequence is the codeword with index 32 in Table 17.

The list of possible ACKs is identical to the list of possible RTS codewords, and each ACK is associated with an ACK index. If the PPD receives an RTS codeword from an SPD and the PPD decides to grant permission to that SPD to transmit a beacon frame, the PPD shall transmit the ACK that is identical to the received RTS codeword. For example if the PPD receives an RTS codeword with the RTS index equal to 23, it will transmit the ACK corresponding to the ACK index equal to 23. If the PPD is able to detect more than one RTS codeword from different SPDs within the same Rx period, the PPD can grant permission to at most one SPD to transmit a beacon frame. The PPD shall grant permission to an SPD by transmitting the ACK corresponding to the RTS codeword that was sent by the randomly chosen SPD.

A special acknowledgement message called the Go-On is sent by the PPD to allow an SPD to transmit an additional beacon frame without needing to issue another RTS burst. This message is useful in the case when an SPD wants to send both TV channel information and LAS channel information, because an SPD can only send one or the other in any given beacon frame (7.2.2.1). If the PPD receives an SPD's beacon frame that has the NST subfield set to one (7.2.1.4), the PPD can allow the SPD to transmit one additional beacon frame by sending a Go-On during the ANP. Note that there shall be at least one PPD beacon frame between the two consecutive SPD beacon frames in order to allow the PPD to continue its protection of the channel. The Go-On sequence is the codeword with index 1 in Table 17.

Before deciding what to transmit during the ANP, the PHY layer of the PPD shall consider both the ANPResponse parameter in PLME-ANP-DECISION.request primitive sent by the MLME and what was received over the air during the Rx period.

- If the ANPResponse parameter is equal to GO_ON, the Go-On sequence shall be transmitted by the PHY during the ANP. The PLME will then issue the PLME-ANP-DECISION.confirm primitive with the ANPResponse parameter set to GO_ON.
- If the ANPResponse parameter is equal to NACK, the NACK sequence shall be transmitted by the PHY during the ANP. The PLME will then issue the PLME-ANP-DECISION.confirm primitive with the ANPResponse parameter set to NACK. The PHY shall transmit its beacon frame during the next superframe.
- If the ANPResponse parameter is equal to ACK and the PHY detected at least one RTS codeword during the preceding Rx period, one of the detected RTS codewords shall be randomly selected by the PHY and the corresponding ACK codeword shall be transmitted during the ANP. The PLME

will then issue the PLME-ANP-DECISION.confirm primitive with the ANPResponse parameter set to ACK.

- If the ANPResponse parameter is equal to ACK and the PHY did not detect an RTS codeword, the NACK codeword shall be transmitted by the PHY during the ANP. The PLME will then issue the PLME-ANP-DECISION.confirm primitive with the ANPResponse parameter set to NACK. The PHY shall transmit its beacon frame during the next superframe.

All the SPDs shall listen during the ANP in order to learn whether the upcoming beacon frame will be sent by the PPD or an SPD. This knowledge is useful when an SPD wants to send a beacon frame, since it is only possible for the PPD to acknowledge an RTS burst transmitted during the Rx period following its own beacon frame. The PLME of the SPD shall send the responses received during each ANP to the MAC sublayer by repeatedly issuing the PLME-ANP-RESPONSE.indication primitive.

6.6 PHY constants and PIB attributes

This subclause specifies the constants and attributes required by the PHY layer.

6.6.1 PHY constants

The constants that define the characteristics of the PHY layer are presented in Table 18.

Table 18—PHY constants

Constant	Description	Value
<i>aInitializationPeriod</i>	The length of time, measured in superframes, that the initial transmission period, which is part of the device initialization procedure, will last.	100
<i>aRegion</i>	The geographical region of operation.	Dependent on physical location (no units)

6.6.2 PHY PIB attributes

The PIB attributes required to manage the PHY layer are presented in Table 19.

Table 19—PHY PIB attributes

Attribute	Identifier	Type	Range	Description
<i>phyCurrentChannel</i>	0x00	Integer	Band and region dependent.	The RF channel (i.e., TV channel) to use for all following transmissions and receptions. The channel specified shall be a valid channel for the geographic region specified by <i>aRegion</i> .

6.7 PHY specifications

6.7.1 Modulation and spreading

The IEEE 802.22.1 PHY shall employ direct sequence spread spectrum (DSSS) with differential quadrature phase-shift keying (DQPSK).

6.7.1.1 Reference modulator diagram

The functional block diagram in Figure 11 is provided as a reference for specifying the PHY modulation and spreading functions. The number in each block refers to a subclause that describes that function. Each bit in the synchronization burst and beacon PPDU shall be processed through the differential encoding, bit-to-chip mapping and modulation functions in octet-wise order. Within each octet, the LSB, b_0 , is processed first and the most significant bit (MSB), b_7 , is processed last.




Figure 11—Modulation and spreading functions

6.7.1.2 Mapping of logical channels to physical channels

Data bits either belong to the synchronization logical channel, the beacon logical channel, the RTS burst, the NPD codeword, or the ANP burst. These bits are parsed between the physical I channel and the physical Q channel, which are used as input for DQPSK encoding (6.7.1.3).

Parsing of the RTS burst bits, NPD codeword bits, and ANP burst bits to the I and Q channels is described in Table 17.

Bits of the beacon frame belong to the beacon logical channel and shall first be parsed into consecutive 4-octet words in the same order as they have been passed from the MAC. Bits of the synchronization burst belong to the synchronization logical channel and are naturally parsed into consecutive 4-octet synchronization bursts.

The bits of the beacon logical channel shall directly be mapped to the bits of the physical Q channel. The bits of the synchronization logical channel shall directly be mapped to the bits of the physical I channel. Both physical channels shall be referenced to a common time reference, meaning that they are simultaneous channels, and as such one bit dI from the I channel and one bit dQ from the Q channel shall be transmitted simultaneously in a single modulation symbol.

6.7.1.3 Differential QPSK encoding

Differential QPSK encoding is a phase change applied to the previous DQPSK symbol according to the two raw data bits from the I and Q channels being encoded. The DQPSK symbols belong to the constellation $(1 + j, -1 + j, 1 - j, -1 - j)$. DQPSK encoding is performed by the transmitter and can be described by Table 20.

Table 20—DQPSK encoding table

Input bits (<i>dI, dQ</i>)	Phase change (φ)
0 0	0
1 0	$\pi/2$
0 1	$3\pi/2$
1 1	π

Differential QPSK encoding can equivalently be performed with a complex multiplication, and it is described by Equation (2):

$$E_n = E_{n-1} \times e^{j\varphi_n} \quad (2)$$

where

φ_n is the phase change according to the two raw data bits being encoded

E_n is the corresponding differentially encoded symbol

E_{n-1} is the previous differentially encoded symbol

For each packet transmitted, φ_1 corresponds to the first two raw data bits to be encoded and E_0 is assumed to be $1 + j$.

Conversely, for each packet received, φ_1 corresponds to the first two raw data bits to be decoded, and E_0 is assumed to be $1 + j$.

The DQPSK modulator shall be initialized at the beginning of each of the following instances:

- New superframe
- RTS burst
- NPD codeword
- ACK burst
- NACK burst
- Go-On burst

6.7.1.4 Symbol-to-chip mapping

Each DQPSK symbol shall be mapped into an 8-chip, complex, pseudo-random noise (PN) sequence as specified in Table 21. During each symbol period, the least significant chip, c_0 , is transmitted first, and the most significant chip, c_7 , is transmitted last.

The DQPSK symbol-to-chip mapping operation can be equivalently viewed as the complex multiplication of the DQPSK symbol with the complex spreading sequence $(1-j, -1-j, 1+j, 1-j, 1-j, 1-j, -1+j, -1-j)$, followed by a division by two.

For ease of implementation, the DQPSK chips are rotated by $\pi/4$ before being transmitted, so that the transmitted chips belong to the constellation $(0.707+0.707j, -0.707+0.707j, 0.707-0.707j, -0.707-0.707j)$.

Table 21—DQPSK symbol-to-chip mapping

Input		Chip values							
DQPSK symbol	DQPSK symbol phase	c ₀	c ₁	c ₂	c ₃	c ₄	c ₅	c ₆	c ₇
1 + j	$\pi/4$	1	-j	j	1	1	1	-1	-j
-1 + j	$3\pi/4$	j	1	-1	j	j	j	-j	1
1 - j	$-\pi/4$	-j	-1	1	-j	-j	-j	j	-1
-1 - j	$-3\pi/4$	-1	j	-j	-1	-1	-1	1	j

The despreading of the complex signal is performed, after synchronization with the chip pulses and rotation of the chips by $-\pi/4$, by summing the element-wise product of the eight consecutive received chips with the complex conjugate of the spreading sequence, and then dividing by the processing gain of eight.

6.7.1.5 Pulse shaping

The chip sequence is modulated onto the carrier with square-root-raised-cosine pulse shaping (roll-off factor = 0.5) applied separately to the in-phase and quadrature components of the complex modulation chips.

Equation (3) (Kuffner [B8]) is a time domain equation for a square root raised cosine pulse with roll-off factor α and chip duration T . This equation is for an infinite duration pulse, which should be truncated and windowed/filtered for practical application. The windowing of the pulse shall be sufficient to satisfy the spectral mask (6.8.3) and EVM (6.8.4) requirements.

$$h(t) = \frac{4\alpha \cos[(1 + \alpha)((\pi t)/T)] + T \sin[(1 - \alpha)((\pi t)/T)] / (4\alpha T)}{\pi \sqrt{T} [1 - ((4\alpha t)/T)^2]} \quad (3)$$

6.7.1.6 Chip and symbol rates

Both chip rate and symbol rate are related to the Advanced Television Systems Committee (ATSC) digital television (DTV) symbol rate (RATSC = 10.7622378 MHz). However, a tighter tolerance of the chip rate and symbol rate is utilized for a device that is compliant with IEEE Std 802.22.1. The chip rate, R_c , shall be equal to the nominal RATSC value divided by 140 (approximately 76.8731 kchip/s), and the symbol rate, R_s , shall be equal to $R_c/8$ (approximately 9.6091 ksymbols/s). The chip rate stability shall be better than ± 2 ppm.

For other regions, different chip and symbol rates may apply.

6.7.2 Forward error correction (FEC)

FEC is applied to portions of both the synchronization burst and the beacon frame. A (15, 7) linear block code shall be applied to the Index field of the synchronization burst. A half-rate, binary convolutional code shall be applied to the first MAC subframe (MSF 1) of the beacon frame.

6.7.2.1 Encoding the Index field of the synchronization burst

The 7-bit Index field of the synchronization burst (6.3) shall be encoded using a (15, 7) linear block code. The eight parity bits shall be used for both error-correction and error-detection of the Index field. The parity bits are determined using the generator polynomial shown in Equation (4).

$$g(D) = D^8 + D^7 + D^6 + D^4 + 1 \tag{4}$$

The 15-bit code word shall be generated in systematic form as shown in Table 22. The bit order is reversed from conventional ordering to avoid an instance of false synchronization.

Table 22—(15, 7) Linear block code in systematic form

Parity Field							Index Field							
P7	P6	P5	P4	P3	P2	P1	P0	i6	i5	i4	i3	i2	i1	i0

The Index and Parity fields are represented in polynomial form as shown in Equation (5) and Equation (6), respectively.

$$i(D) = i_0D^6 + i_1D^5 + i_2D^4 + i_3D^3 + i_4D^2 + i_5D + i_6 \tag{5}$$

$$p(D) = p_0D^7 + p_1D^6 + p_2D^5 + p_3D^4 + p_4D^3 + p_5D^2 + p_6D + p_7 \tag{6}$$

The parity bits shall be computed as shown in Equation (7).

$$p(D) = D^8 \times i(D) \times \text{mod}[g(D)] \tag{7}$$

For example, Table 23 shows the Index and Parity fields for an index of 25.

Table 23—(15, 7) Linear block coding example (index = 25)

Parity Field							Index Field							
P7	P6	P5	P4	P3	P2	P1	P0	i6	i5	i4	i3	i2	i1	i0
0	0	0	1	0	0	1	1	1	0	0	1	1	0	0

The index bits in polynomial form are shown in Equation (8).

$$i(D) = D^6 + D^3 + D^2 \tag{8}$$

The parity bits in polynomial form are shown in Equation (9).

$$p(D) = (D^{14} + D^{11} + D^{10}) \times \text{mod}[g(D)] = D^4 + D + 1 \tag{9}$$

See Annex A for an example decoding method.

6.7.2.2 Encoding the MSF 1 of the beacon frame

The MSF 1 of the beacon frame is encoded using a half-rate, binary convolutional encoder. The constraint length of this coder is equal to 7, and its generator polynomials are 171₀ and 133₀. Figure 12 shows the

pictorial depiction of the generator polynomials. Output A and output B represent the first and second output bits, respectively, of this encoder.

The encoder memory of the convolutional coder shall be initialized with an all-zero state at the beginning of the MSF 1.

The data length shall be 17 octets or 136 bits before encoding and $(136 \times 2 + 12)$ bits after encoding. After encoding, the output data of the encoder is punctured from $(272 + 12)$ bits to 272 bits. The following 12 bits shall be punctured: 1 23 45 67 89 111 133 155 177 199 221 243. The output data starts at bit 0.

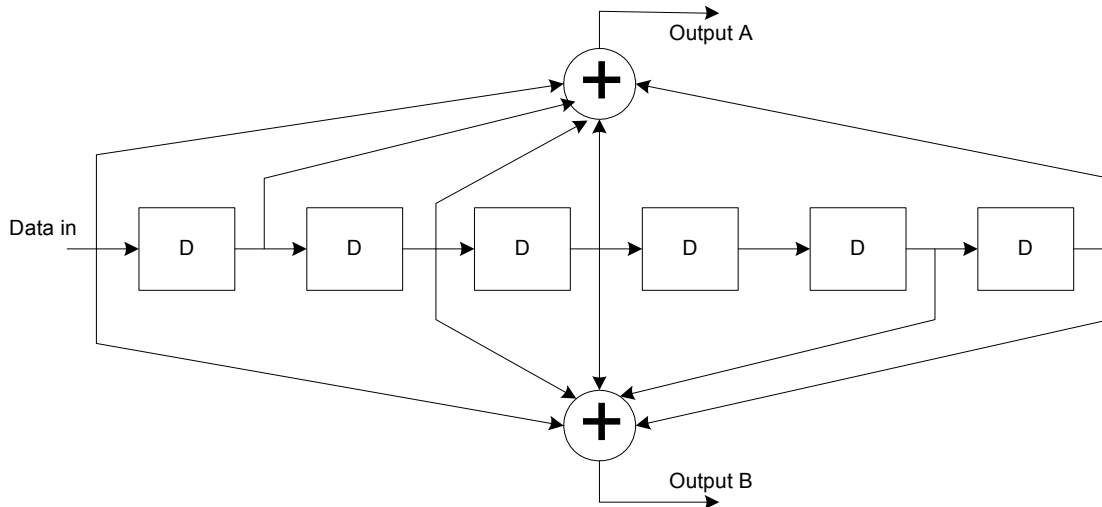


Figure 12—Half-rate convolutional coder with generator polynomials 171_0 , 133_0
(The delay element represents a delay of 1 bit)

6.8 Radio specifications

6.8.1 Transmit center frequency tolerance

The transmitted center frequency tolerance shall be ± 2 ppm maximum. The device shall maintain the center frequency tolerance over a temperature range of -30 °C to $+60$ °C.

6.8.2 Transmit power

The maximum transmitter output power of beaconing devices shall be limited to 250 mW for UHF bands and 50 mW for VHF bands, unless superseded by the local regulatory environment. The transmit power shall comply with the RF measurement information given in 6.1.3.

6.8.3 Transmit PSD mask

Transmitter emissions shall comply with the piece-wise linear mask with breakpoints indicated in Table 24 and shown in Figure 13. This mask shall be valid for non-harmonic spurious emissions. For offsets less than ± 1 MHz, the measurements shall be made with 1 kHz resolution and video bandwidth and a peak hold detector. The units “dBr” indicate “dB relative to the total average signal power.” For offset frequencies beyond ± 1 MHz, the transmitter emissions shall be below -90 dBr/kHz and shall be made using an average

detector. Adherence to local regulations, which may impose emissions stricter than indicated in Table 24, is required.




Figure 13—Piece-wise linear transmit mask described by Table 24
 [The mask is symmetric about the center frequency (0 kHz offset in this plot)]

Table 24—Beacon transmitter emissions mask

Offset frequency (kHz)	Attenuation (dBm/kHz)
± 50	0
± 70	20
± 100	60
± 200	80
≥ ± 1000	90

6.8.4 Modulation accuracy

An error vector magnitude (EVM) measurement is used to determine modulation accuracy. In order to calculate the EVM measurement, a time record of N received signal coordinate pairs (I_j, Q_j) is captured. For each received chip, a decision is made as to which chip was transmitted. The ideal position of the chosen chip is represented by the vector (I_j, Q_j) . The error vector $(\delta I_j, \delta Q_j)$ is defined as the distance from this ideal position to the actual position of the received chip.

Thus, the received vector is the sum of the ideal vector and the error vector as shown in Equation (10).

$$(\tilde{I}_j, \tilde{Q}_j) = (I_j, Q_j) + (\delta I_j, \delta Q_j) \tag{10}$$

The EVM for this standard is defined in Equation (11) through Equation (13).

$$\bar{I} = \frac{1}{N} \sum_{j=1}^N |\tilde{I}_j| \quad (11)$$

$$\bar{Q} = \frac{1}{N} \sum_{j=1}^N |\tilde{Q}_j| \quad (12)$$

$$EVM = \sqrt{\frac{\sum_{i=1}^N [(|\tilde{I}_i| - \bar{I})^2 + (|\tilde{Q}_i| - \bar{Q})^2]}{N(\bar{I}^2 + \bar{Q}^2)}} \quad (13)$$

The transmit rms error vector magnitudes shall be less than 14% averaged over one superframe ($N = 124$ octets $\times 8 = 992$ symbols).

The EVM measurement shall be made on baseband I and Q data after recovery through an ideal reference receiver system. The ideal reference receiver shall perform carrier lock, chip timing recovery, and amplitude adjustment while making the measurements. The ideal reference receiver shall have a data filter impulse response that approximates that of an ideal root raised cosine filter with 50% excess bandwidth.

6.8.5 Spurious harmonically-related emission suppression

Harmonically-related spurious emissions shall be suppressed a minimum of 60 dB below the equivalent unmodulated carrier power level. The reference carrier frequency shall be the unmodulated carrier frequency or the geometric center frequency of the modulated beacon RF emission. Measurement of harmonically-related spurious emissions shall be measured at the antenna terminal of the beaconing device or, if the antenna is permanently fixed to the device, by means of a calibrated sense antenna and appropriate measurement equipment. Local regulations may supersede the requirements of this subclause.

6.8.6 Receiver sensitivity

Under the conditions specified in 6.1.4, 1% PER Gaussian channel sensitivity for the synchronization word, the index value, and MSF 1 shall be no worse than -107 dBm, and shall be no worse than -106 dBm for the RTS/ANP burst. The 1% PER Gaussian channel sensitivity for both MSF 2 and MSF 3 shall be no worse than -100 dBm. This difference in sensitivities is due to the omission of error protection redundancies in MSF 2 and 3 to minimize the overall sensing dwell time.

6.8.7 Adjacent LAS channel rejection

The scenario is for two IEEE 802.22.1-compliant devices attempting to intercommunicate while one of the devices is in close proximity (both spatially and spectrally) to one of its low-powered protected devices (e.g., a wireless microphone). See C.2 for a complete link analysis. The adjacent LAS channel is defined here to be the first usable LAS channel (per Annex B) at an offset of Δf within the selected TV channel, where the low-power, licensed link can operate with minimal sensitivity degradation while in close proximity to the beacon transmitter. See Figure 14.

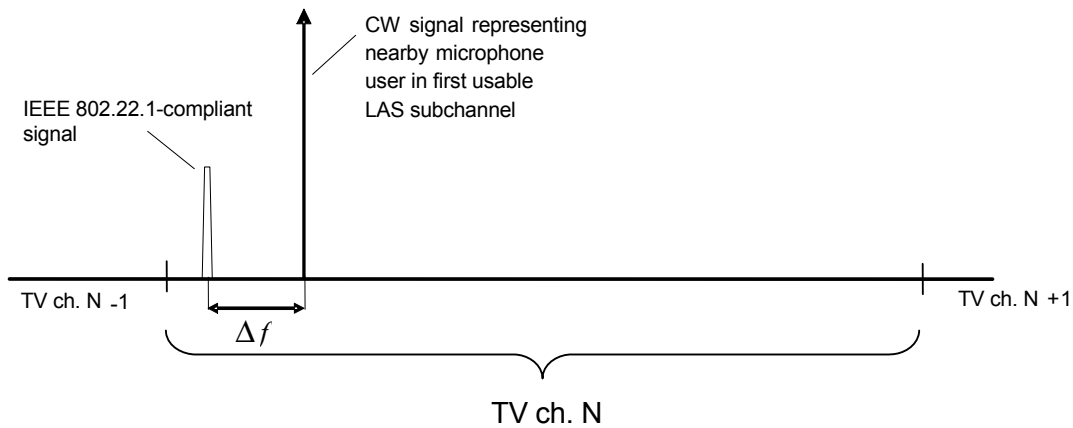


Figure 14—Pictorial definition of adjacent LAS channel within a TV channel

Annex B defines Δf to be $1.075 \text{ MHz} - 309.4 \text{ kHz} = 765.6 \text{ kHz}$. The beacon receiver sensitivity shall be at or below -104 dBm in the presence of a -36 dBm continuous wave interferer located at an offset Δf of $+765.56 \text{ kHz}$ (the first available Title 47, Part 74, Subpart H LAS channel) for MSF 1, and at or below -97 dBm for MSF 2 and MSF 3. The difference in sensitivity is due to the rate-1/2 error correcting code on MSF 1.

6.8.8 Receiver maximum input level of compliant beacon signal (blocking)

An IEEE 802.22.1-compliant device shall be capable of receiving a compliant beacon signal greater than or equal to 0 dBm .

6.8.9 Link quality indicator (LQI)

The LQI measurement is a characterization of the quality of a received beacon frame. The use of the LQI result by the higher layers is not specified in this standard.

The LQI measurement shall use the average of the absolute value of the radian phase error from each received differential phase symbol spanning a single frame. The range of the measurement will be between 0 and 0.4, with 0 corresponding to a very good channel and 0.4 corresponding to a very bad channel. The resolution of the measurement shall be $1/256$ of the range.

The LQI measurement shall be performed for each received beacon frame, and the result shall be reported to the MAC sublayer using the PD-DATA.indication primitive (6.2.1.3) as an integer ranging from $0x00$ to $0xff$, according to Table 25.

An LQI measurement can be converted to a decimal LQI measurement as shown in Equation (14):

$$dM = \text{round}(640M) \tag{14}$$

where

- M is the captured LQI measurement
- dM is the decimal LQI measurement

Table 25—LQI mapping

LQI measurement	LQI mapping	Corresponding SNR _C ^a (dB)
0.0	0x00	∞
0.0015625	0x01	45.2
0.0031250	0x02	39.2
0.0046875	0x03	35.6
...
0.10	0x40	9.1
...
0.20	0x80	3.0
...
0.3984375	0xff	< -6

^aNOTE—The chip SNR shown here is only for an AWGN channel with no channel impairments, like delay spread or frequency offset. For channels with impairments, the LQI will reflect these impairments, but the actual SNR is not as easily related to the LQI measure.

The equation for mapping a decimal LQI measurement to a hexadecimal value is as follows in Equation (15).

$$xM = 0x\left(\text{int}\left[\frac{dM}{16}\right]_X, \text{rem}\left[\frac{dM}{16}\right]_X\right) \tag{15}$$

where

xM is the hexadecimal mapping

dM is the decimal LQI measurement

int is the integer portion, rem is the remainder

X subscript indicates representation of the digit in hex format (e.g., if $dM = 197$, $\text{int}[197/16] = 12$, so $\text{int}(197/16)_X = C$. If $dM > 255$, then $dM = 255$; otherwise, $dM = dM$)

Equation (16) shows the relationship of the LQI measurement value and the chip SNR value.

$$E\{|\Delta\phi_{err}|\} = 0.284 \times 10^{-SNR_C/20}, \text{ for } SNR_C > -1\text{dB} \tag{16}$$

where

$|\Delta\phi_{err}|$ is the time series of the absolute value of the phase deviation from the ideal DQPSK constellation point per symbol

$E\{|\Delta\phi_{err}|\}$ is the expected value of the time series and gives the LQI measure M

SNR_C is the chip SNR in dB

7. MAC sublayer specification

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

- Generating beacon frames
- Supporting message authentication and integrity
- Employing the radio channel access mechanism
- Providing a reliable link between two peer MAC entities

Constants and attributes that are specified and maintained by the MAC sublayer are written in the text of this clause in italics. Constants have a general prefix of “a”, e.g., *aAddress*, and are listed in Table 46. Attributes have a general prefix of “mac”, e.g., *macNumSyncBursts*, and are listed in Table 47.

7.1 MAC sublayer service specification

The MAC sublayer provides an interface between the next higher layer (NHL) and the PHY. The MAC sublayer conceptually includes a management entity called the MAC Sublayer Management Entity (MLME). The MLME provides a means of passing information between the NHL and the MAC sublayer. The MLME is also responsible for maintaining a database of variables pertaining to the MAC sublayer, which is called the MAC Information Base (MIB).

Figure 15 depicts the components and interfaces of the MAC sublayer.

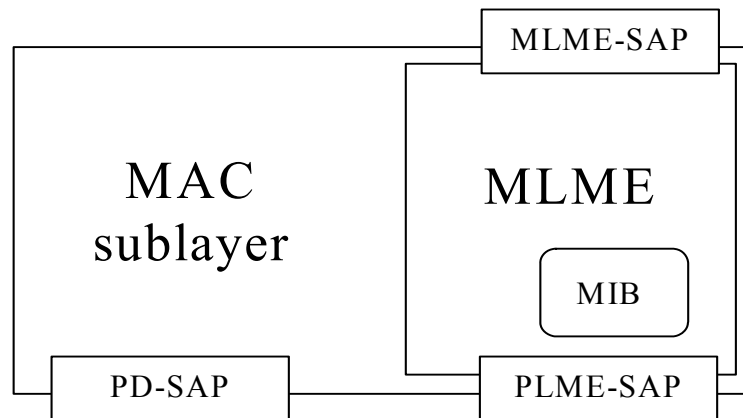


Figure 15—MAC sublayer reference model

The MAC sublayer provides a MAC management service, accessed through the MLME-SAP.

7.1.1 MAC management service

The MLME-SAP is an interface that provides a means of passing information between the NHL and MAC sublayer (using the MLME). Table 26 summarizes the primitives¹³ supported by the MLME through the MLME-SAP interface. The primitives are discussed in the subclauses referenced in the table.

¹³See Annex F for an overview of the concept of service primitives.

Table 26—Primitives supported by the MLME-SAP

Name	Request	Indication	Confirm
MLME-PPD-LOST	—	7.1.1.1	—
MLME-GET	7.1.1.2	—	7.1.1.3
MLME-INCOMING-BEACON	—	7.1.1.4	—
MLME-NPD	7.1.1.5	—	7.1.1.6
MLME-NPD-LOST	—	7.1.1.7	—
MLME-SEARCH	7.1.1.8	—	7.1.1.9
MLME-SET	7.1.1.10	—	7.1.1.11
MLME-SPD-LOST	—	7.1.1.12	—
MLME-START-BEACON	7.1.1.13	—	7.1.1.14

7.1.1.1 MLME-PPD-LOST.indication

The MLME-PPD-LOST.indication primitive is generated by the MLME of an SPD or NPD and issued to its NHL as a notification that either the PPD was not heard during the last *macMaxMissedPPDBeacons* superframes or the PPD was heard during this time but its beacon frames were considered “invalid” during the frame reception and rejection procedure (7.4.3). This primitive has no parameters.

7.1.1.2 MLME-GET.request

The MLME-GET.request primitive is generated by the NHL and issued to its MLME to request information about one or more MIB attributes. Table 27 specifies the parameters for the MLME-GET.request primitive.

Table 27—MLME-GET.request parameters

Name	Type	Valid range	Description
MIBAttribute1	Integer	See Table 47	The identifier of the MIB attribute to read.
MIBAttribute2	Integer	See Table 47	The identifier of the MIB attribute to read.
...			
MIBAttributeN	Integer	See Table 47	The identifier of the MIB attribute to read.

On receipt of the MLME-GET.request primitive, the MLME attempts to retrieve the requested MIB attribute or attributes from its database. Each successfully retrieved MIB attribute will have a corresponding status of SUCCESS in the MLME-GET.confirm primitive. Any MIB attribute value(s) that could not be retrieved will have a corresponding status indicating the appropriate error status (see 7.1.1.3).

7.1.1.3 MLME-GET.confirm

The MLME-GET.confirm primitive is generated by the MLME and issued to its NHL in response to an MLME-GET.request primitive, and it reports the results of an information request from the MIB. Table 28 specifies the parameters for the MLME-GET.confirm primitive.

Table 28—MLME-GET.confirm parameters

Name	Type	Valid range	Description
Status1	Enumeration	SUCCESS or ATTRIBUTE_NOT_FOUND	The result of the request for MIB attribute information.
MIBAttribute1	Integer	See Table 47	The identifier of the MIB attribute that was read.
MIBAttributeValue1	Various	Attribute specific; see Table 47	The value of the indicated MIB attribute that was read. This parameter has zero length when the Status parameter is set to ATTRIBUTE_NOT_FOUND.
Status2	Enumeration	SUCCESS or ATTRIBUTE_NOT_FOUND	The result of the request for MIB attribute information.
MIBAttribute2	Integer	See Table 47	The identifier of the MIB attribute that was read.
MIBAttributeValue2	Various	Attribute specific; see Table 47	The value of the indicated MIB attribute that was read. This parameter has zero length when the Status parameter is set to ATTRIBUTE_NOT_FOUND.
...			
StatusN	Enumeration	SUCCESS or ATTRIBUTE_NOT_FOUND	The result of the request for MIB attribute information.
MIBAttributeN	Integer	See Table 47	The identifier of the MIB attribute that was read.
MIBAttributeValueN	Various	Attribute specific; see Table 47	The value of the indicated MIB attribute that was read. This parameter has zero length when the Status parameter is set to ATTRIBUTE_NOT_FOUND.

On receipt of the MLME-GET.confirm primitive, the NHL is notified of the results of its request to read one or more MIB attributes. If the request to read an MIB attribute was successful, the corresponding Status parameter is set to SUCCESS. If the identifier of an MIB attribute was not found in the database, the corresponding Status parameter is set to ATTRIBUTE_NOT_FOUND.

7.1.1.4 MLME-INCOMING-BEACON.indication

The MLME-INCOMING-BEACON.indication primitive is generated by the MLME and issued to its NHL in order to transfer the parameters contained within a received beacon frame. The primitive also transfers a measure of the LQI, the time the beacon frame was received, and the channel number on which it was received. Table 29 specifies the parameters for the MLME-INCOMING-BEACON.indication primitive.

On receipt of the MLME-INCOMING-BEACON.indication primitive, the NHL is notified of the arrival of a beacon frame at the MAC sublayer.

Table 29—MLME-INCOMING-BEACON.indication parameters

Name	Type	Valid range	Description
Parameter1	Integer	See 7.2.1.1	The parameter contains the information from the MHR field of the same name (i.e., frame version number, priority level, antenna height, and device rank).
Parameter2	Integer	See 7.2.1.4	The parameter contains the information from the MHR field of the same name (i.e., TV channel width, cease tx, time parity, next-in-line protecting device, next SPD superframe to transmit, sub-group aggregation data, keep out zone data).
Parameter3	Integer	See 7.2.1.5	The parameter contains the information from the MHR field of the same name (i.e., antenna location, required need data timer).
Address	IEEE address	A valid 48-bit IEEE address	The address of the originator of the beacon frame.
Location	Integer	See 7.2.1.3	The location of the originator of the beacon frame.
Map	Bitmap	40-bit field	Either the occupied TV channels or the LAS channels within an occupied TV channel that are protected by the device that transmitted the beacon frame or manufacturer-specific information.
Channel	Integer	Region dependent.	The TV channel on which the beacon frame was received.
RxTime	Integer	See 7.5.2	The time the beacon frame was received relative to the clock of the receiving device.
LinkQuality	Integer	0x00–0xff	The LQI at which the beacon frame was received. Lower values represent lower LQI (6.8.9).
SecurityStatus	Enumeration	SIGNATURE_NOT_CHECKED, SIGNATURE_VALID, SIGNATURE_INVALID, or CERTIFICATE_INVALID	The security-related status of the received beacon frame.

The signature of the incoming beacon frame is only checked if *macSignatureCheckEnabled* is set to TRUE (see Table 47). If the signature was checked but was found to be invalid (7.4.3), the *SecurityStatus* parameter is set to *SIGNATURE_INVALID*. Otherwise, if the signature was found to be valid, the *SecurityStatus* parameter is set to *SIGNATURE_VALID*. If the signature of the incoming beacon frame was not checked because *macSignatureCheckEnabled* is set to FALSE, the *SecurityStatus* parameter is set to *SIGNATURE_NOT_CHECKED*.

If the certificate of the received beacon frame was found to be invalid (7.4.3), the *SecurityStatus* parameter is set to *CERTIFICATE_INVALID*.

7.1.1.5 MLME-NPD.request

The MLME-NPD.request primitive is generated by the NHL of a PPD and issued to its MLME with the address of the SPD that was chosen to be the next-in-line protecting device (NPD). Table 30 specifies the parameter for the MLME-NPD.request primitive. This primitive shall only be issued by a PPD.

Table 30—MLME-NPD.request parameter

Name	Type	Valid range	Description
NPDAddress	IEEE address	A valid 48-bit IEEE address	The address of the SPD that was chosen as the NPD.

On the receipt of the MLME-NPD.request primitive, the MLME is notified of the address of the chosen NPD.

7.1.1.6 MLME-NPD.confirm

The MLME-NPD.confirm primitive is generated by the MLME of a PPD and issued to its NHL in response to an MLME-NPD.request primitive following the transmission of a beacon frame with the NPD Indication subfield set to 01. This primitive has no parameters.

7.1.1.7 MLME-NPD-LOST.indication

The MLME-NPD-LOST.indication primitive is generated by the MLME of a PD other than the NPD and issued to its NHL as a notification that it has not received either an NPD beacon frame or an NPD codeword within the last ($aMaxMissedNPDCodes \times aNPDPeriod$) superframes. This primitive has no parameters.

7.1.1.8 MLME-SEARCH.request

The MLME-SEARCH.request primitive is generated by the NHL and issued to its MLME to initiate a listening period on a given TV channel or channels. Table 31 specifies the parameters for the MLME-SEARCH.request primitive. Note that although this primitive allows more than one TV channel to be searched, the ChannelList parameter is limited to one TV channel for this version of the standard.

Table 31—MLME-SEARCH.request parameters

Name	Type	Valid range	Description
ChannelList	Bitmap	69-bit field	The list of TV channels to be searched. The 69 bits (b_0, b_1, b_{69}) indicate which channels are to be searched (1 = search, 0 = do not search).
Duration	Integer	0–90	A value used to calculate the length of time to spend searching each TV channel. The length of time spent searching each channel is $5 + m$ superframe periods, where m is the value of the Duration parameter.

On receipt of the MLME-SEARCH.request primitive, the MLME enables the receiver and listens to each TV channel specified by the ChannelList parameter for a finite amount of time, as specified by the Duration parameter.

7.1.1.9 MLME-SEARCH.confirm

The MLME-SEARCH.confirm primitive is generated by the MLME and issued to its NHL when the TV channel search operation initiated with the MLME-SEARCH.request primitive has completed. This primitive has no parameters.

7.1.1.10 MLME-SET.request

The MLME-SET.request primitive is generated by the NHL and issued to its MLME to attempt to write the given value or values to the indicated MIB attributes. This primitive may have several pairs of parameters, with each pair specifying the MIB attribute and its corresponding value. Table 32 specifies the parameters for the MLME-SET.request primitive.

Table 32—MLME-SET.request parameters

Name	Type	Valid range	Description
MIBAttribute1	Integer	See Table 47	The identifier of the MIB attribute to write.
MIBAttributeValue1	Various	Attribute specific; see Table 47	The value to write to the indicated MIB attribute.
MIBAttribute2	Integer	See Table 47	The identifier of the MIB attribute to write.
MIBAttributeValue2	Various	Attribute specific; see Table 47	The value to write to the indicated MIB attribute.
...			
MIBAttributeN	Integer	See Table 47	The identifier of the MIB attribute to write.
MIBAttributeValueN	Various	Attribute specific; see Table 47	The value to write to the indicated MIB attribute.

On receipt of the MLME-SET.request primitive, the MLME attempts to write the given value or values to the indicated MIB attributes in its database. Each successfully written MIB attribute will have a corresponding status of SUCCESS in the MLME-SET.confirm primitive. Any MIB attribute value(s) that could not be written will have a corresponding status indicating the appropriate error status (7.1.1.11).

7.1.1.11 MLME-SET.confirm

The MLME-SET.confirm primitive is generated by the MLME and issued to its NHL in response to an MLME-SET.request primitive, and it reports the results of an attempt to write one or more values to their corresponding MIB attributes. Table 33 specifies the parameters for the MLME-SET.confirm primitive.

On receipt of the MLME-SET.confirm primitive, the NHL is notified of the result of its request to set the value of one or more MIB attributes. If a requested value was written to the indicated MIB attribute, the corresponding Status parameter is set to SUCCESS.

Table 33—MLME-SET.confirm parameters

Name	Type	Valid range	Description
Status1	Enumeration	SUCCESS, ATTRIBUTE_NOT_FOUND, or INVALID_PARAMETER	The result of the request to write the MIB attribute.
MIBAttribute1	Integer	See Table 47	The identifier of the MIB attribute that was written.
Status2	Enumeration	SUCCESS, ATTRIBUTE_NOT_FOUND, or INVALID_PARAMETER	The result of the request to write the MIB attribute.
MIBAttribute2	Integer	See Table 47	The identifier of the MIB attribute that was written.
...			
StatusN	Enumeration	SUCCESS, ATTRIBUTE_NOT_FOUND, or INVALID_PARAMETER	The result of the request to write the MIB attribute.
MIBAttributeN	Integer	See Table 47	The identifier of the MIB attribute that was written.

If a MIBAttribute parameter specifies an attribute that is not found in the database, the corresponding Status parameter is set to ATTRIBUTE_NOT_FOUND. If a MIBAttributeValue parameter specifies a value that is out of the valid range for the given attribute, the corresponding Status parameter is set to INVALID_PARAMETER.

7.1.1.12 MLME-SPD-LOST.indication

The MLME-SPD-LOST.indication primitive is generated by the MLME of the PPD and issued to its NHL as a notification that the beacon frames from a particular SPD were not heard in the last *macMissedSPDBeacons* superframes. Table 34 specifies the parameter for the MLME-SPD-LOST.indication primitive.

Table 34—MLME-SPD-LOST.indication parameters

Name	Type	Valid range	Description
SPDAddress	IEEE Address	A valid 48-bit IEEE address	The address of the lost SPD.

7.1.1.13 MLME-START-BEACON.request

The MLME-START-BEACON.request primitive is generated by the NHL and issued to its MLME to either start or stop beacon frame transmissions or to change the content of the transmitted beacon frames. Table 35 specifies the parameters for the MLME-START-BEACON.request primitive.

Table 35—MLME-START-BEACON.request parameters

Name	Type	Valid range	Description
Parameter1	Integer	See 7.2.1.1	The parameter contains the information for the MHR field of the same name (i.e., frame version number, priority level, antenna height, and device rank).
Parameter2	Integer	See 7.2.1.4	The parameter contains most of the information for the MHR field of the same name (i.e., TV channel width, cease tx, next-in-line protecting device, next SPD superframe to transmit, sub-group aggregation data, keep out zone data). The only information included in the MHR field that is not included here is the time parity.
Parameter3	Integer	See 7.2.1.5	The parameter contains the information for the MHR field of the same name (i.e., antenna location, required need data timer).
SourceAddress	IEEE address	A valid 48-bit IEEE address	The address of the originator of the beacon frame.
Location	Integer	See 7.2.1.3	The location of the originator of the beacon frame.
Map	Bitmap	40-bit field	This parameter contains one of the following: the occupied TV channels that are to be protected, the LAS channels within an occupied TV channel that are to be protected, or manufacturer-specific information.
Channel	Integer	Region dependent.	The TV channel on which to transmit the beacon frame(s).
Start	Boolean	TRUE or FALSE	If this parameter is TRUE, the device is to begin beacon frame transmission. Otherwise, the device is to stop beacon frame transmissions.
Initialize	Boolean	TRUE or FALSE	If this parameter is TRUE, the PPD will enter the initial transmission period. If this parameter is FALSE, the PPD will not enter the initial transmission period. This parameter only applies to the PPD.

On receipt of the MLME-START-BEACON.request primitive, the MLME attempts to carry out the requested action from the NHL. If the requested action was successfully executed, the MLME will issue the MLME-START-BEACON.confirm primitive with a status of SUCCESS. Otherwise, the MLME will issue the MLME-START-BEACON.confirm primitive with the appropriate error status (see 7.1.1.14).

If the MLME-START-BEACON.request primitive is received by the MLME of the PPD, the PPD sends a continuous series of beacon frames, except when interrupted by an SPD via inter-device communications. If the MLME-START-BEACON.request primitive is received by the MLME of an SPD, the SPD sends one beacon frame.

7.1.1.14 MLME-START-BEACON.confirm

The MLME-START-BEACON.confirm primitive is generated by the MLME and issued to its NHL in response to an MLME-START-BEACON.request primitive, and it reports the results of the attempt to start or stop beacon frame transmissions or to change the content of the transmitted beacon frame(s). Table 36 specifies the parameters for the MLME-START-BEACON.confirm primitive.

Table 36—MLME-START-BEACON.confirm parameters

Name	Type	Valid range	Description
Status	Enumeration	SUCCESS or TX_FAILURE	The result of the attempt to either start or stop the transmission of a beacon frame(s).

On receipt of the MLME-START-BEACON.confirm primitive, the NHL is notified of the results of its request. If the request primitive was issued by the NHL of an SPD and the beacon frame was not transmitted following repeated attempts (i.e., the retry procedure described in 7.4.2 failed), the Status parameter is set to TX_FAILURE. If the request was successful, the Status parameter is set to SUCCESS. In the case of a PPD, the status parameter is always set to SUCCESS.

7.1.2 Enumeration description

The enumeration values used by the MAC management primitives, as well as by the PHY data and management primitives, are given in Table 37.

Table 37—Enumeration values

Enumeration	Value	Description
SUCCESS	0x00	The requested operation was completed successfully.
ATTRIBUTE_NOT_FOUND	0x01	A SET/GET request was issued with the identifier of a MIB or PIB attribute that is not supported.
CERTIFICATE_INVALID	0x02	The certificate contained in the received beacon frame is invalid.
COMPLETE	0x03	The requested operation is complete.
GO_ON	0x04	The PPD will allow the last SPD that transmitted a beacon frame to transmit a second beacon frame during the upcoming superframe.
INCORRECT_ACK	0x05	A request to transmit a beacon frame by an SPD was rejected due to the receipt of an acknowledgement codeword not matching the transmitted RTS codeword.
INVALID_PARAMETER	0x06	A parameter in the primitive is either not supported or is out of the valid range.

Table 37—Enumeration values (continued)

Enumeration	Value	Description
NACK	0x07	A request to transmit a beacon frame by an SPD was rejected due to the reception of a NACK from the PPD.
RX_ON	0x08	The transceiver is either already in or requested to change to the received enabled state.
SIGNATURE_INVALID	0x09	The signature contained in the received beacon frame is invalid.
SIGNATURE_NOT_CHECKED	0x0a	The signature contained in the received beacon frame was not checked.
SIGNATURE_VALID	0x0b	The signature contained in the received beacon frame is valid.
TRX_OFF	0x0c	The transceiver is either already in or requested to change to the transceiver disabled state.
TX_FAILURE	0x0d	The beacon frame was not transmitted by the SPD, because the SPD was unable to gain permission from the PPD to access the channel following repeated attempts.
TX_ON	0x0e	The transceiver is in the transmitter enabled state.

7.2 MAC beacon frame

This subclause specifies the format of the MAC beacon frame (i.e., MPDU).

The beacon frame is described as a sequence of fields in a specific order. The beacon frame format is depicted in the order in which it is transmitted by the PHY, from left to right, where the leftmost bit is transmitted first in time. Bits within each field are numbered from 0 (leftmost and least significant) to $k - 1$ (rightmost and most significant), where the length of the field is k bits. Fields that are longer than a single octet are sent to the PHY in the order from the octet containing the lowest numbered bits to the octet containing the highest numbered bits.

All reserved bits shall be set to zero upon transmission and shall be ignored upon receipt.

The beacon frame format is composed of three MSFs, each having a MHR and a MFR. Together, these three MSFs form the MAC protocol data unit (MPDU). The beacon frame shall be formatted as illustrated in Figure 16.

Octets: 17	51	33
MSF 1	MSF 2	MSF 3

Figure 16— MAC beacon frame format (MPDU)

7.2.1 MSF 1

MSF 1 shall be formatted as illustrated in Figure 17. The fields of the MHR appear in a fixed order. The MHR contains three MAC parameter fields, the Source Address field, and the Location field. The MFR contains a 2-octet CRC.

Octets: 1	6	5 + 1 bit	1 + 7 bits	1	2
Parameter 1	Source Address	Location	Parameter 2	Parameter 3	CRC 1
MHR 1					MFR 1

Figure 17—MSF 1 format

7.2.1.1 Parameter 1 field

The Parameter 1 field includes the Frame Version Number, Priority Level, Antenna Height, and Rank subfields. The Parameter 1 field shall be formatted as illustrated in Figure 18.

Bits: 0–2	3–5	6	7
Frame Version Number	Priority Level	Antenna Height	Rank

Figure 18—Format of the Parameter 1 field

The Frame Version Number subfield specifies the version number of the transmitted frame. This subfield shall be set to 000 to indicate a frame compliant with this standard. All other subfield values shall be reserved for future use.

The Priority Level subfield specifies the priority of the service protected by the beacon frame transmission. This information may be used by IEEE 802.22.1 devices to establish a hierarchy of users and does not affect the actions of unlicensed users, since unlicensed users should vacate the channel for any licensed device. The value of the Priority Level subfield shall be numeric, in which the value 111 shall be defined as highest priority and the value 000 as lowest priority.

The Antenna Height subfield specifies the height above ground level (AGL) that the device (or devices) protected by the beaconing device has been deployed. Note that “ground level” shall be defined to be the highest point within the radius of protection;¹⁴ the radius of protection is defined by the Keep Out Zone subfield (7.2.1.4). The Antenna Height subfield shall be set to one if this height is greater than or equal to 10 m. It shall be set to zero if this height is less than 10 m. If the beaconing device is protecting more than one device, the collective antenna height shall be the greatest height among all the protected devices and the subfield shall be set accordingly.

The Rank subfield shall be set to one if the beaconing device is the PPD, or zero if it is an SPD.

7.2.1.2 Source Address field

The Source Address field is 6 octets in length and specifies the 48-bit IEEE address of the originator of the beacon frame.

7.2.1.3 Location field

The Location field is 41 bits in length and specifies the location of the originator of the beacon frame.

¹⁴See E.3 for a recommendation for obtaining altitude information.

The location information shall be represented in WGS84 datum coordinates and shall be acquired from geolocation latitude and longitude information.^{15,16} The latitude and longitude information are read from the geolocation message and formatted into a degrees/minutes/seconds format as shown in Table 38 and Table 39, respectively. The Location field is shown in Table 40.

Table 38—Latitude information taken from geolocation message

Bits: 7	6	6	1
Latitude Degrees	Latitude Minutes	Latitude Seconds	Latitude Direction (north = 0; south = 1)

Table 39—Longitude information taken from geolocation message

Bits: 8	6	6	1
Longitude Degrees	Longitude Minutes	Longitude Seconds	Longitude Direction (east = 0; west = 1)

Table 40—Location field (encoded data)

Bits: 20	21
Latitude	Longitude

7.2.1.4 Parameter 2 field

The Parameter 2 field includes the Channel Width, Cease Tx, Time Parity, Protected Radius, and Sub-group Channels subfields. If the beacon frame is sent by the PPD, the Parameter 2 field includes the NPD Indication subfield. If the beacon frame is sent by an SPD, the Parameter 2 field includes the NPD and NST subfields. The Parameter 2 field for the PPD shall be formatted as illustrated in Figure 19. The Parameter 2 field for an SPD shall be formatted as illustrated in Figure 20.

Bits: 0–2	3	4	5	6–12	13–14
Channel Width	Cease Tx	Time Parity	Keep Out Zone	Sub-group Channels	NPD Indication

Figure 19—Format of the Parameter 2 field in the PPD's beacon frame

The Channel Width subfield specifies whether cross-channel aggregation is employed and the width of the occupied TV channel or channels being protected by the transmitting device. The format of the Channel Width subfield is given in Table 41.

¹⁵See Annex E for possible methods for acquiring time and location information.

¹⁶If the geolocation-based location source is a global positioning system (GPS) or similar device that supports NMEA output strings, the data is most easily extracted from GPRMA or GPGGA strings (NMEA 0183 [B10]).

Bits: 0–2	3	4	5	6–12	13	14
Channel Width	Cease Tx	Time Parity	Protected Radius	Sub-group Channels	NPD	NST

Figure 20—Format of the Parameter 2 field in an SPD’s beacon frame

Table 41—Channel Width subfield

Bits: 0	1	2	Description
0	0	0	Cross-channel aggregation is not employed, and the occupied channel is 6 MHz wide.
0	0	1	Cross-channel aggregation is employed, and the occupied channels are 6 MHz wide.
0	1	0	Cross-channel aggregation is not employed, and the occupied channel is 7 MHz wide.
0	1	1	Cross-channel aggregation is employed, and the occupied channels are 7 MHz wide.
1	0	0	Cross-channel aggregation is not employed, and the occupied channel is 8 MHz wide.
1	0	1	Cross-channel aggregation is employed, and the occupied channels are 8 MHz wide.
1	1	0	Reserved.
1	1	1	Reserved.

The Cease Tx subfield specifies whether the transmitting device is planning to cease transmission. The subfield shall be set to one to indicate that the device plans to stop transmitting and shall be set to zero otherwise.

The Time Parity subfield indicates the parity of the time subfield *t* used when generating the signature. The value of the subfield shall be equal to $(t \bmod 2)$. For more information on the time subfield *t*, see 7.5.2.

The Keep Out Zone subfield specifies the radius of the protected area. The subfield shall be set to zero if the protected radius is 1.5 km. The subfield shall be set to one if the protected radius is 4.5 km.

Five UHF sub-groups are defined in the United States. These sub-groups are defined in Table 42.

The Sub-group Channels subfield shall be considered when the TV channel of operation falls within one of the five UHF sub-groups and when cross-channel aggregation is employed (i.e., the Channel Width subfield indicates cross-channel aggregation). If the TV channel of operation (i.e., the channel on which the beacon frames are sent) is not within one of the five UHF sub-groups or cross-channel aggregation is not employed, this subfield shall be ignored. The Sub-group Channels subfield is formatted as shown in Table 43.

Table 42—UHF sub-groups in the United States

UHF sub-group	TV channels included
1	14–20
2	21–28
3	29–36
4	38–44
5	45–51

Channel 1 represents the TV channel with the lowest channel number within the sub-group, while Channel 7 represents the TV channel with the highest channel number within the sub-group. A bit value of zero indicates that the TV channel shall not be protected, and a bit value of one indicates that the TV channel shall be protected. If the number of TV channels to be specified is less than seven, the unused channel bits shall be set to zero. The TV channel of operation shall not be included in this subfield, since its inclusion is implied.

Table 43—Sub-group Channels subfield

Bit: 6	7	8	9	10	11	12
Channel 1	Channel 2	Channel 3	Channel 4	Channel 5	Channel 6	Channel 7

The NPD Indication subfield is transmitted in every PPD beacon frame, and its format is shown in Table 44.

Table 44—NPD Indication subfield in the PPD’s beacon frame

Bit		NPD exists	Request for NPD	Description
6	7			
0	0	no	yes	No NPD currently exists in the beaconing network, and the PPD is actively seeking one.
1	1	no	no	No NPD currently exists in the beaconing network, and the PPD is not actively seeking one.
0	1	yes	n/a	An NPD currently exists in the beaconing network.
1	0	—	—	Reserved.

A value of 00 in the NPD Indication subfield indicates that there is no NPD in the beaconing network and the PPD is requesting one. A value of 11 in the NPD Indication subfield indicates that there is no NPD in the beaconing network and that the PPD is not requesting one. A value of 01 in the NPD Indication subfield indicates that there is an NPD present in the beaconing network. A value of 10 in the NPD Indication subfield indicates a reserved value.

The NPD subfield is transmitted in every SPD beacon frame. A value of one indicates that the beacon frame was sent by the NPD. A value of zero indicates that the beacon frame was sent by an SPD that is not the NPD.

The Next SPD Superframe to Transmit (NST) subfield is transmitted in every SPD beacon frame. The subfield indicates whether an SPD wants to send an additional beacon frame without needing to issue another RTS burst. If the SPD does not want to send an additional beacon frame beyond the current frame, the subfield shall be set to zero. If the SPD wants to send an additional frame, the subfield shall be set to one.

All other bits are reserved.

7.2.1.5 Parameter 3 field

The Parameter 3 field includes the Indoor/Outdoor and Required Need Timer subfields. The Parameter 3 field shall be formatted as illustrated in Figure 21.

Bits: 0	1–7
Indoor/Outdoor	Required Need Timer

Figure 21—Format of the Parameter 3 field

The Indoor/Outdoor subfield indicates the location of the receiver antenna of the protected device. It shall be set to one if the antenna is indoors and shall be set to zero if the antenna is outdoors. Zero (outdoors) shall be the default.

The Required Need Timer subfield shall be a numeric value indicating the estimated time remaining, in hours, that the TV channel will be occupied. A value of all zeros indicates that the channel will be occupied for an indeterminate amount of time.

7.2.1.6 CRC 1 field

The CRC 1 field is 2 octets in length and contains a 16-bit ITU-T CRC. The CRC is calculated over the MHR of MSF 1, which contains the three MAC parameter fields, the Source Address field, and the Location field.

The CRC shall be calculated using the standard generator polynomial of degree 16 shown in Equation (17).


$$G_{16}(x) = x^{16} + x^{12} + x^5 + 1 \quad (17)$$

The CRC shall be calculated for transmission using the following algorithm:

- Let $M(x) = b_0x^{k-1} + b_1x^{k-2} + \dots + b_{k-2}x + b_{k-1}$ be the polynomial representing the sequence of bits for which the checksum is to be computed.
- Multiply $M(x)$ by x^{16} , giving the polynomial $x^{16} \times M(x)$.
- Divide $x^{16} \times M(x)$ modulo 2 by the generator polynomial, $G_{16}(x)$, to obtain the remainder polynomial, $R(x) = r_0x^{15} + r_1x^{14} + \dots + r_{14}x + r_{15}$.
- The CRC field is given by the coefficients of the remainder polynomial, $R(x)$.

Here, binary polynomials are represented as bit strings, in highest polynomial degree first order.

A typical implementation is depicted in Figure 22, which was borrowed from IEEE Std 802.15.4-2006 [B5].



1. Initialize the remainder register (r_0 through r_{15}) to zero.
2. Shift the MHR into the divider in the order of transmission (LSB first).
3. After the last bit of the data field is shifted into the divider, the remainder register contains the CRC.
4. The CRC is appended to the data field so that r_0 is transmitted first.

Figure 22—Typical CRC implementation

7.2.2 MSF 2

MSF 2 shall be formatted as illustrated in Figure 23. The fields of the MHR appear in a fixed order. The MHR contains the Map and Signature fields. The MFR contains a 2-octet CRC.

Octets: 5	44	2
Map	Signature	CRC 2
MHR 2		MFR 2

Figure 23—MSF 2 format

7.2.2.1 Map field

The Map field is five octets in length and is a bitmap with three possible uses. It can be used to identify either a list of occupied TV channels being protected by the transmitting device or a list of the 200 kHz-wide LAS channels within an occupied TV channel being protected. The LAS channel information may be used by operators of protected devices to conserve spectrum by identifying vacant LAS channels. The Map field can also contain manufacturer-specific information (MSI). Such information is out of the scope of this standard and is not specified here.

The values of bits 0 and 1 determine how the remaining bits will be used. If bit 0 is one, the remaining bits shall identify protected LAS channels. If bit 0 is zero, the content of bit 1 shall be examined to determine the field's use. If bit 1 is zero, bits 2–39 shall identify protected TV channels. If bit 1 is one, bits 2–39 shall carry MSI.

7.2.2.1.1 Manufacturer-specific information (MSI)

Figure 24 illustrates the bit assignment for the case when bits 0 and 1 of the Map field are set to zero and one, respectively, indicating that the field is to be used to carry MSI.

Bit: 0	1	2–39
Channel/LAS channel (contents ignored)	MSI	Manufacturer-specific information

Figure 24—Field configured to carry MSI

7.2.2.1.2 Channel mapping

Figure 25 illustrates the bit assignment for the case when bits 0 and 1 of the Map field are both set to zero, indicating that the field is to be used as a channel map.

Bit: 0	1	2–6	7–12	13–18	19–24	25–30	31–36	37–39
Channel/ LAS channel	MSI	Region	Channel 1	Channel 2	Channel 3	Channel 4	Channel 5	Reserved

Figure 25—Field as a channel map (bits 0 and 1 set to zero)

Bits 2–6 specify the Region subfield, which is the geographical region of operation. Only fourteen region designators are necessary for TV channel mapping worldwide due to common sharing of band plans of the VHF and UHF segments among various countries, and in some cases, entire continents. For example, most of Western Europe utilizes a formalized band plan for UHF allocations, and this allocation plan is also utilized by most of Africa. Reserved values included in the list shall be reserved for future use, such as in the event of changes to band plans and allocations. See Table 45 for the list of region designators.

Table 45—Region designators for protected TV channels

Bit: 2 3 4 5 6	Region
0 0 0 0 0	United States of America with less than 64 TV channels (UHF and VHF)
0 0 0 0 1	The Americas (excluding USA) VHF
0 0 0 1 0	The Americas (excluding USA), Korea, Taiwan, Philippines, certain Pacific Islands UHF
0 0 0 1 1	Ireland VHF
0 0 1 0 0	Australia VHF
0 0 1 0 1	Australia UHF
0 0 1 1 0	Western Europe, Africa, Asia, Pacific Islands (excluding French Territories) UHF
0 0 1 1 1	France VHF
0 1 0 0 0	New Zealand UHF
0 1 0 0 1	Morocco
0 1 0 1 0	Eastern Europe/Russia VHF

Table 45—Region designators for protected TV channels (continued)

Bit: 2 3 4 5 6	Region
0 1 0 1 1	South Africa VHF
0 1 1 0 0	French Overseas Territories
0 1 1 0 1	China
0 1 1 1 0–1 1 1 1 0	Reserved
11111	Undefined

The five Channel subfields allow explicit entry of up to five protected TV channels. If the number of TV channels to be specified is less than five, the unused Channel subfields shall be filled with zeros. Note that it is unnecessary to include the TV channel of operation (i.e., the channel on which the beacon frames are sent), since its inclusion is implied.

Each Channel subfield is six bits long. If the number of a TV channel in a given region is less than 64, the channel number shall simply be converted to a 6-bit binary number before inserting it into the subfield. If the number of a TV channel in a given region equals or exceeds 64, a re-mapping procedure is needed in order to utilize the 6-bit subfield. The TV channel designations for these regions shall be re-mapped to a sequential numbering system, starting with “1”. This shall be accomplished by Equation (18).

$$C_m = C_a - C_1 \tag{18}$$

where

- C_m is the mapped TV channel number to be inserted into the Channel subfield
- C_a is the actual TV channel number to be protected
- C_1 is the TV channel number assigned to the first (lowest) channel designation for the region to be mapped

For example, if the actual TV channel to be protected is 68 and the first channel number assigned in the local region is 14, then the mapped channel number to be inserted into the Channel subfield is $68 - 14 = 54$.

Bits 37–39 are reserved for future use and shall be set to zero.

7.2.2.1.3 LAS channel mapping

A total of thirty 200 kHz LAS channels can be defined in a 6 MHz-wide TV channel, thirty-five 200 kHz LAS channels can be defined in a 7 MHz-wide TV channel, and forty 200 kHz LAS channels can be defined in an 8 MHz-wide TV channel. An 8 MHz-wide TV channel requires 40 bits to map all 40 available LAS channels, but only 39 bits are available. Therefore, in this case, the 200 kHz LAS channel utilized by the beacon shall not be mapped, thus allowing all other LAS channels to be mapped. The LAS channel map shall always apply to the TV channel of operation.

Figure 26 illustrates the bit assignment for the case when bit 0 of the Map field is set to one, indicating that the field is to be used as an LAS channel map. Bit 1 represents the 200 kHz LAS channel centered 100 kHz above the lower edge of the TV channel. The LAS channel centered 100 kHz below the upper edge of the TV channel corresponds to bit 30 for 6 MHz-wide TV channels, bit 35 for 7 MHz-wide TV channels, or bit 39 for 8 MHz-wide TV channels. See 6.1.1 for more information on TV channels and LAS channels.

Bit: 0	1	2	...	29	30
Channel/LAS channel	LAS channel 1	LAS channel 2	...	LAS channel 29	LAS channel 30

Figure 26—Field as an LAS channel map (bit 0 set to one) for a 6 MHz TV channel

When the protected device center frequency falls at the center of the LAS channel, the bit for that LAS channel shall be set to one. When the protected device center frequency falls such that a 200 kHz region around the center frequency overlaps two LAS channels, both of those LAS channel bits shall be set to one. If the occupied TV channel width is such that not all of the 40 bits are needed, the remaining MSBs (9 bits for a 6 MHz-wide TV channel and 4 bits for a 7 MHz-wide TV channel) shall be set to zero but shall not indicate unused LAS channels. The Channel Width subfield in the Parameter 2 field (7.2.1.4) indicates occupied TV channel width.

7.2.2.2 Signature field

The signature field is generated over the MHR of MSF 1, the Map field of MSF 2, and current time and date information. The algorithms to generate and verify this signature are defined in 7.5.4.

7.2.2.3 CRC 2 field

The CRC 2 field is 2 octets in length and contains a 16-bit ITU-T CRC. The CRC is calculated over the MHR of MSF 2, which contains the Map and Signature fields. See 7.2.1.6 for information on how to calculate the CRC.

7.2.3 MSF 3

MSF 3 shall be formatted as illustrated in Figure 27. The fields of the MHR appear in a fixed order. The MHR contains the Certificate field, and the MFR contains a 2-octet CRC.

Octets: 31	2
Certificate	CRC 3
MHR 3	MFR 3

Figure 27—MSF 3 format

7.2.3.1 Certificate field

The certificate field represents the implicit certificate used to transport the beaconing device’s public key. The certificate is stored as a MIB attribute, *macCertificate*, and the algorithm to generate and verify this certificate is defined in 7.5.5.

7.2.3.2 CRC 3 field

The CRC 3 field is 2 octets in length and contains a 16-bit ITU-T CRC. The CRC is calculated over the MHR of MSF 3, which contains the Certificate field. See 7.2.1.6 for information on how to calculate the CRC.

7.3 MAC constants and MIB attributes

This subclause specifies the constants and attributes required by the MAC sublayer.

7.3.1 MAC constants

The constants that define the characteristics of the MAC sublayer are presented in Table 46.

Table 46—MAC sublayer constants

Constant	Description	Value
<i>aAddress</i>	The 48-bit IEEE address assigned to the device.	Device-specific
<i>aMaxMissedNPDCodes</i>	The number of NPD codes that, when missed consecutively, will cause the PPD or an SPD to send a notification to its NHL.	5
<i>aNPDPeriod</i>	The absolute number of consecutive superframes between two consecutive NPD codeword transmissions.	10
<i>aNSTValidCount</i>	The maximum number of PPD beacon frames allowed between an initial SPD beacon frame and the subsequent SPD beacon frame permitted by a Go-On response before the SPD will terminate the NST process. If the initial SPD beacon frame is in superframe n , the Go-On response shall be received in either superframe $n+1$ or $n+2$, in order to continue the NST process.	2

7.3.2 MIB attributes

The MIB attributes required to manage the MAC sublayer are presented in Table 47. The list of *macPeerPublicKeyTable* table entries and the list of *macAuthorityPublicKeyTable* table entries are given in Table 48 and Table 49, respectively.

Table 47—MIB attributes

Attribute	ID	Type	Range	Description	Default
<i>macActivePeriodSPD</i>	0x00	Integer	1000–4999	The maximum interval, measured in superframes, between the transmission of beacon frames by an SPD.	2000
<i>macAuthorityPublicKeyTable</i>	0x01	Table of Authority -Public-Key entries (see Table 49)	—	Entries in this table correspond to authorities that are trusted to sign a public-key certificate.	(empty)
<i>macCertificate</i>	0x02	Integer	Any 31-octet value	A value representing the implicit certificate used to transport the beaconing device's public key.	Device-specific

Table 47—MIB attributes (continued)

Attribute	ID	Type	Range	Description	Default
<i>macMaxMissedPPDBeacons</i>	0x03	Integer	1–29	The number of consecutive missed PPD beacon frames that will cause an SPD or the NPD to send a notification to the NHL. If the device is operating as the PPD, this attribute shall not apply.	15
<i>macMissedSPDBeacons</i>	0x04	Integer	5000–6000	The number of SPD beacon frames that, when missed consecutively, will cause the PPD to send a notification to the NHL. If the device is operating as an SPD or NPD, this attribute shall not apply.	5000
<i>macNPAddress</i>	0x05	IEEE address	A valid 48-bit IEEE address	The 48-bit address of the NPD. This attribute only represents a valid NPD address if <i>macNPDPresent</i> is TRUE.	—
<i>macNPDPresent</i>	0x06	Boolean	TRUE or FALSE	An indication of whether an NPD is present in the beaconing network. TRUE indicates that an NPD is present and that its address is <i>macNPAddress</i> . FALSE indicates that no NPD is currently present.	FALSE
<i>macNumSyncBursts</i>	0x07	Integer	30–31	The number of synchronization bursts to be sent prior to transmitting the next beacon frame. A value of 31 shall be chosen when the device is operating in its initial transmission period (7.4.4). Otherwise, the value shall be 30.	30
<i>macPeerPublicKeyTable</i>	0x08	Table of Peer-Public-Key entries (see Table 48)	—	Entries in this table correspond to peer devices (e.g., other beaconing devices) that have been heard and whose public keys have been validated. Entries in this table are indexed by the peer device's address, the issuer, and the key identifiers. The expiration date of a certificate, from which a key in this table was derived, shall be verified before using this key to verify a newly received signature.	(empty)
<i>macPPDAddress</i>	0x09	IEEE address	A valid 48-bit IEEE address	The 48-bit address of the PPD.	—

Table 47—MIB attributes (continued)

Attribute	ID	Type	Range	Description	Default
<i>macPrivateKey</i>	0x0a	Integer	Any 14-octet value	A value representing the private key that a beaconing device uses to create the digital signature of an outgoing beacon frame.	Device-specific
<i>macSignatureCheckEnabled</i>	0x0b	Boolean	TRUE or FALSE	An indication of whether the device will perform a check of the signatures of incoming beacon frames. TRUE indicates that signature checking will be enabled, while FALSE indicates that signature checking will be disabled.	TRUE

Table 48—Elements of a PeerPublicKey entry

Name	Type	Range	Description
PeerDeviceAddress	IEEE address	A valid 48-bit IEEE address	The 48-bit IEEE address assigned to the peer device.
PublicKey	Octet string	Any valid 29-octet string representing a point on the elliptic curve	A point on the elliptic curve (7.5.3.2).
ExpirationDate	Integer	0x00–0xff	The expiration date of the certificate from which this key was derived (7.5.3.5).
KeyID	Integer	0x00–0xff	The identifier of the key (7.5.3.4).
KeyIssuerID	Integer	0x00–0xff	The trusted authority that assigned the public key to the peer device (7.5.3.5).

Table 49—Elements of AuthorityPublicKey entry

Name	Type	Range	Description
PublicKey	Octet string	Any valid 29-octet string representing a point on the elliptic curve	A point on the elliptic curve (7.5.3.2).
KeyIssuerID	Integer	0x00–0xff	The trusted authority that assigned the public key to the peer device (7.5.3.5).

7.4 MAC functional description

This subclause provides a detailed description of the MAC functionality.

7.4.1 Transmission protocol

Only the PPD shall transmit beacon frames unless an SPD has been granted permission to do so in place of the PPD. In order to gain permission to transmit a beacon frame, an SPD transmits an RTS burst in the Rx period immediately following the beacon frame transmitted by the PPD. In the event that two or more SPDs transmit an RTS burst simultaneously, at least one should be successfully received by the PPD. In this case if no RTS burst is received, the PPD will transmit a NACK during the ANP and then continue to transmit its own beacon frame. If one is received and the PPD does not elect to transmit its own beacon frame, the PPD will transmit the ACK corresponding to the RTS burst it heard. In the unlikely event that two or more SPDs chose the same RTS codeword and so received the corresponding ACK, both will proceed to transmit a beacon frame. If the purpose of the SPD's beacon frame is aggregation, the SPD will know whether the PPD received its beacon frame by examining the contents of the PPD's subsequent beacon frames. The PPD's beacon frame shall only contain the data from an SPD if it correctly received the SPD's beacon frame. If the purpose of the SPD's beacon frame is to notify the PPD that the SPD is still active, the SPD will be unaware of the collision but should have another chance to send the notification before the PPD considers the SPD to be inactive. For more details on inter-device communications, see 7.4.5.2.

7.4.2 Retry procedure for an SPD

The retry procedure is initiated by an SPD in the event that its request to transmit a beacon frame was denied (i.e., an SPD transmits an RTS burst and does not receive an ACK corresponding to the transmitted RTS codeword). The MLME of the SPD shall learn of the failed request through the receipt of the PLME-INITIATE-RTS-BURST.confirm primitive with the Status parameter set either to NACK, indicating that a NACK was received during the ANP, or INCORRECT_ACK, indicating that an ACK codeword that did not match the transmitted RTS codeword was received during the ANP. Note that the retry procedure applies to all SPDs, including the NPD.

Figure 28 illustrates the retry procedure. Following the failed request, the SPD shall choose a random number K in the range $[0, \dots, 15]$, where K is the number of SPD beacon frame transmission opportunities that the SPD shall wait before retransmitting the RTS burst. Note that SPD beacon frame transmission opportunities only occur following a superframe in which the PPD transmitted its beacon frame. If the MLME of the SPD does not receive a NACK via the PLME-ANP-RESPONSE.indication primitive during the last superframe (i.e., the last beacon frame was not sent by the PPD), it defers the contention. If the SPD receives a NACK during the last superframe and the value of K is not zero, the SPD shall decrement K by one and defer the contention. If the SPD receives a NACK during the last superframe and the value of K is zero, the SPD shall send an RTS burst during the Rx period in the current superframe. If the MLME of the SPD receives a matching ACK via the PLME-INITIATE-RTS-BURST.confirm primitive following its RTS burst transmission, the SPD shall send a beacon frame in the following superframe and the procedure shall be considered successful. If the SPD does not receive a matching ACK following its RTS burst transmission, the SPD shall select another random number and begin the countdown process again. Each time the SPD selects a random number and transmits an RTS burst, it counts as a retry attempt. The maximum number of retry attempts shall be three, meaning that the total number of transmission attempts shall not exceed four.

If the SPD reaches the maximum number of retry attempts and was still unable to transmit its beacon frame, the MLME of the SPD shall notify the NHL of the transmission failure via the MLME-START-BEACON.confirm primitive with the Status parameter set to TX_FAILURE. The action taken by the NHL on receipt of this primitive is out of the scope of this standard.

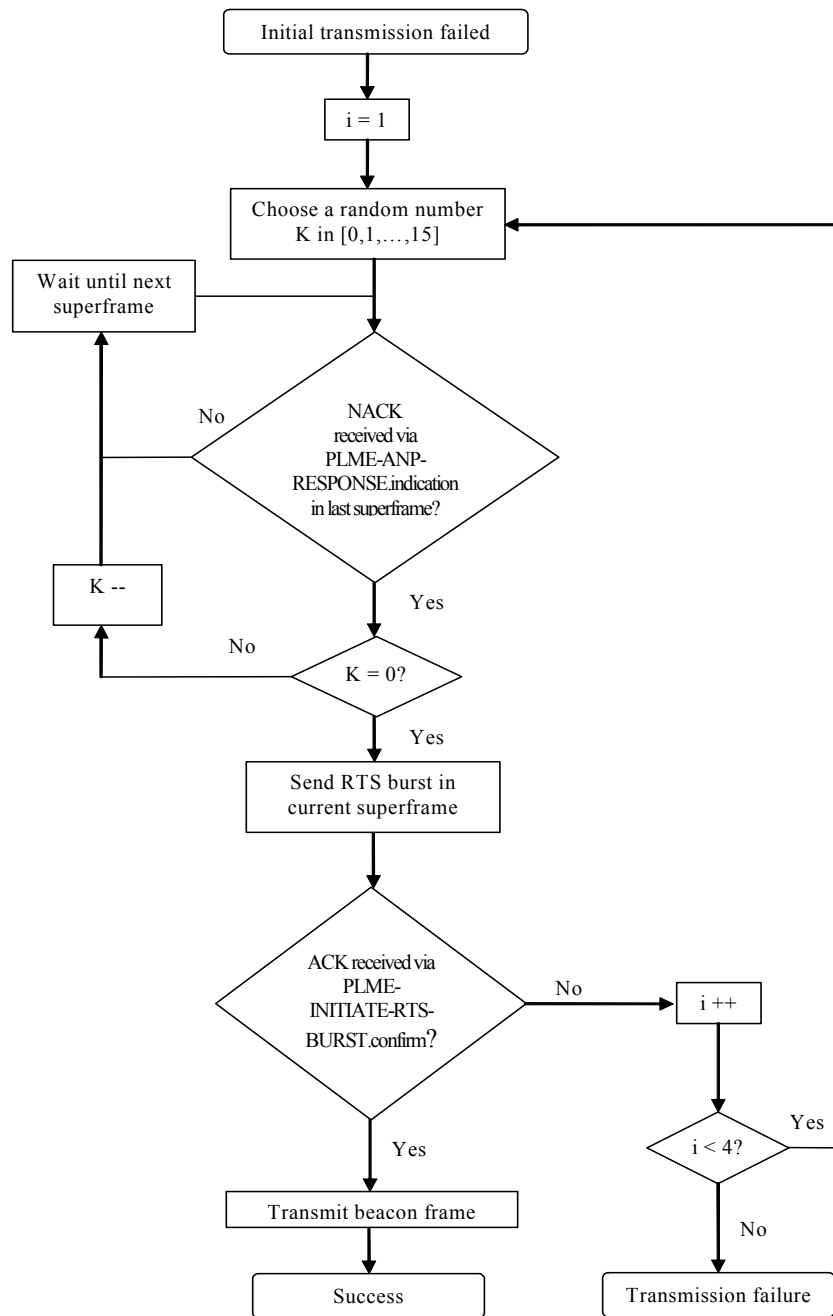


Figure 28—Retry procedure

7.4.3 Frame reception and rejection

The MAC sublayer shall filter incoming frames and retain only the frames that are of interest to the upper layers. For the first level of filtering, the MAC sublayer shall discard all received frames that do not contain a correct value in their CRC 1 and CRC 2 fields.

When signature checking is enabled (e.g., via the MIB attribute *macSignatureCheckEnabled*) and this checking requires use of the frame's certificate field (e.g., as opposed to using a previously received certificate), as a second level of filtering the MAC sublayer shall discard all received frames that do not contain a correct value in their CRC 3 field.

If the frame is not discarded by either the first or second levels of filtering, then an MLME-INCOMING-BEACON.indication primitive is issued with the SecurityStatus parameter set to one of the following four values: SIGNATURE_NOT_CHECKED, SIGNATURE_VALID, SIGNATURE_INVALID, or, in the case that the certificate was checked and found to be invalid, CERTIFICATE_INVALID.

A frame is considered "invalid" if the above steps result in the frame being discarded or result in the MLME-INCOMING-BEACON.indication primitive being issued with the SecurityStatus parameter set to SIGNATURE_INVALID or CERTIFICATE_INVALID. Otherwise, the frame is considered "valid."

7.4.4 Device initialization procedure¹⁷

Upon initialization, a PD should search the channel on which it intends to operate for the existence of a PPD. The search procedure may be initiated by the NHL of the PD via the MLME-SEARCH.request primitive. The selected channel shall be searched for a period of $5 + m$ superframe periods, where m is an integer selected at random by the NHL from the set $[0, 1, \dots, 89, 90]$. Any beacon frame received during the search shall be passed to the NHL via the MLME-INCOMING-BEACON.indication primitive. A PPD is determined to be present on the channel if the PD receives at least one beacon frame with the Rank subfield equal to one.

At the conclusion of the search, if the PD determines that there is no PPD already present on the TV channel (i.e., no beacon frame was detected), the PD may, at the discretion of an upper layer, promote itself to PPD and begin transmitting periodic beacon frames. Beacon frame transmission is initiated by the NHL via the MLME-START-BEACON.request primitive with the Initialize parameter set to TRUE, which indicates that the PPD is entering the initial transmission period. Upon receipt of this primitive, the MLME shall issue the PLME-INITIALIZE.request primitive such that the PHY shall begin transmitting beacon frames without ICIs. Note that the MAC shall issue a PD-DATA.request primitive in order to forward the actual beacon frame content to the PHY.

During this initial transmission period, the new PPD shall transmit *aInitializationPeriod* superframes, which shall not include ICIs (i.e., Rx periods and ANPs). Each superframe shall be composed of 31 synchronization bursts sent on the synchronization channel. A 120-octet length beacon frame followed by 4 octets of all zeros shall be sent on the beacon channel (5.3). At the conclusion of *aInitializationPeriod* superframes, the PLME shall issue the PLME-INITIALIZE.confirm primitive in order to exit the initial transmission period.

Following the initial transmission period, superframe transmission shall continue; however, the ICI shall be inserted immediately following the beacon frame. Since the ICI has a duration of one slot time, the number of synchronization bursts in a superframe following the initial transmission period shall be 30.

The superframe shall always have a period equal to $(8 \times 124) \text{ bits} / 9609.1 \text{ Hz} = 103.24 \text{ ms}$.

At the conclusion of the search, if the PD determines that there is a PPD on the TV channel (i.e., a PPD's beacon frame was detected), the PD may, at the discretion of an upper layer, send its information to the PPD for inclusion in the PPD's beacon frame rather than begin its own superframe transmissions (i.e., opt to become an SPD). This may be accomplished by the PD synchronizing to the superframe of the PPD, transmitting an RTS burst during the Rx period, receiving an ACK from the PPD, and then transmitting its

¹⁷Figure 32 and Figure 33 illustrate how the primitives may be used to carry out the device initialization procedure.

own beacon frame containing the information it wishes the PPD to include in future beacon frames. Note, however, that the SPD shall wait for a period of *InitializationPeriod* superframes before attempting to communicate with the PPD. For more details, see 7.4.5.

Upon initialization, all MIB and PIB attributes are set to their default values. The NHL of the PD may change the MIB attribute values via the MLME-SET.request primitive. Similarly, the MLME may change the PIB attribute values via the PLME-SET.request primitives.

If a PD does opt to become an SPD, the NHL should write the address of the PPD into the MIB attribute *macPPDAddress* via the MLME-SET.request primitive.

7.4.5 Inter-device communications

This subclause describes the communications between IEEE 802.22.1 PHY devices.

7.4.5.1 Data aggregation

A PPD may aggregate data received from one or more SPDs, as long as these SPDs operate on the same TV channel as the PPD. This process is called LAS channel aggregation. The SPD transfers data to the PPD by sending a beacon frame, which the MLME of the PPD passes to its NHL by issuing the MLME-INCOMING-BEACON.indication primitive. The received data may be aggregated with that of the PPD at the discretion of an upper layer.

Cross-channel aggregation is the process of combining data from two or more PDs that operate on different TV channels. All information necessary to provide cross-channel aggregation is incorporated in this standard; however, it is not currently supported by IEEE P802.22. It may be incorporated in future versions of IEEE P802.22.

7.4.5.2 SPD behavior

An SPD may interrupt the PPD in order to transmit its own beacon frame. To initiate this option, the NHL sends an MLME-START-BEACON.request primitive to the MAC sublayer, causing one beacon frame is to be transmitted. Upon receipt of the primitive, the MAC sublayer shall wait until it receives a PLME-ANP-RESPONSE.indication primitive with the ANPResponse parameter set to NACK from the PLME, indicating that the upcoming beacon frame will be sent by the PPD. Once this happens, the MAC sublayer shall request that the PHY layer transmit a randomly selected RTS codeword (see Table 17) during the Rx period of the upcoming superframe by issuing the PLME-INITIATE-RTS-BURST.request primitive instructing the PHY layer to start the transmission. Note that if the SPD were to transmit an RTS burst during the Rx period following another SPD's beacon frame, there would be no chance of receiving a matching ACK, since the PPD sends at least every other beacon frame. If, in response to the RTS burst, the SPD receives an ACK corresponding to the transmitted RTS codeword (see Table 17) from the PPD during the ANP, the SPD shall transmit its beacon frame in place of the normally-transmitted beacon frame of the PPD during the following superframe. If the SPD received a NACK or an ACK that does not correspond to the transmitted RTS codeword during the ANP, it shall not transmit a beacon frame. Instead, the SPD shall initiate the retry procedure (7.4.2). If the retry procedure fails, the NHL of the SPD is notified of the inability to transmit its beacon frame via the MLME-START-BEACON.confirm primitive with the Status parameter set to TX_FAILURE. The action taken by the NHL on receipt of this primitive is out of the scope of this standard.

In general, the SPD may receive a NACK in the ANP of any superframe. For example, assume two or more SPDs transmit an RTS burst in the Rx period of the same superframe. If either the RTS bursts are not heard by the PPD or the PPD receives at least one RTS burst but elects to transmit its own beacon frame, the PPD shall transmit a NACK. If, instead, the PPD receives at least one RTS burst and does not elect to transmit its own beacon frame, it shall respond with an ACK corresponding to a received RTS burst, which each SPD

shall examine. If the received ACK matches its transmitted RTS codeword, an SPD will determine that its RTS was received by the PPD. Note that it is possible but not likely that more than one SPD will choose the same RTS codeword to be sent in the same Rx period. If this happens, more than one SPD will receive the ACK corresponding to its transmitted codeword, which will cause the SPD's beacon frames to be transmitted simultaneously during the following superframe. If an SPD's beacon frame is not heard by the PPD and the SPD was transmitting data to be aggregated by the PPD, the SPD should notice that its request was not fulfilled by examining the PPD's subsequent beacon frames. In this case, one option for the NHL of the SPD is to use a random backoff time before requesting to send another beacon frame. If an SPD's beacon frame is not heard by the PPD and the only purpose of the SPD's beacon frame was to notify the PPD that it is still active (7.4.6.2), the SPD will be unaware of the collision but should have another chance to send the notification without causing the MLME of the PPD to issue an MLME-SPD-LOST indication to its NHL.

If an SPD wants to send more than one beacon frame, the SPD can send its initial beacon frame with the NST subfields set to one. Upon receipt of this beacon frame, the PPD can grant permission to the SPD to send a second beacon frame by transmitting a Go-On response during an ANP. This process may decrease the RTS collision probability due to the fact that the second beacon frame is sent without transmitting an additional RTS burst during the Rx period. After successfully sending its initial beacon frame (i.e., on receipt of the PD-DATA.confirm primitive), the SPD shall wait for up to *aNSTValidCount* superframes to receive a Go-On response from the PPD. If a Go-On response is received by the MAC sublayer of the SPD via the PLME-ANP-RESPONSE.indication primitive during this time, the MAC sublayer of that SPD shall request that its PHY layer send a second beacon frame during the upcoming superframe by issuing a PD-DATA.request primitive. Note that once this second beacon frame is sent, the SPD shall transmit a new RTS burst before being able to send any further beacon frames. If a Go-On response is not received during this time, the SPD shall terminate the process, and it shall send a new RTS burst before sending another beacon frame. If a Go-On response is received but a new frame is not available from the NHL for transmission, the MAC sublayer shall request that the PHY layer resend the previous beacon frame.

7.4.5.3 PPD behavior

The PPD may send a NACK during the ANP of any superframe regardless of what was heard during the Rx period.

The following text details the policies of the PPD's MAC sublayer on the direction it will give to its PHY layer. At least one PPD beacon frame shall be transmitted between two successive SPD beacon frames. Therefore, the MLME of the PPD shall reserve the appropriate superframes for its own beacon frame transmissions by issuing the PLME-ANP-DECISION.request primitive with the ANPResponse parameter set to NACK whenever necessary, in order to meet this requirement. In addition, if the MLME of the PPD receives a request from the NHL to change the contents of its beacon frame (i.e., the MLME receives an MLME-START-BEACON.request primitive), it shall reserve the upcoming superframe for its own beacon frame transmission by issuing the PLME-ANP-DECISION.request primitive with the ANPResponse parameter set to NACK.

If the PPD receives an SPD's beacon frame with the NST subfield set to one, the PPD may grant the SPD permission to transmit an additional beacon frame without requiring that a second RTS burst be sent. If this beacon frame was received within the last *aNSTValidCount* superframes and the PPD does not have a requirement to send its own beacon frame, the PPD shall reserve the upcoming superframe for that SPD by issuing the PLME-ANP-DECISION.request primitive with the ANPResponse parameter set to GO-ON. If the PPD receives the second beacon frame and that beacon frame has the NST subfield set to one, the PPD shall ignore the SPD's request. Note that once the SPD receives the Go-On and sends the second beacon frame, it is required to transmit a new RTS burst before sending additional beacon frames.

If the PPD does not have a requirement to send its own beacon frame and has not received an SPD's beacon frame with the NST subfield equal to one within the last *aNSTValidCount* superframes, the MLME of the

PPD shall issue the PLME-ANP-DECISION.request primitive with the ANPResponse parameter set to ACK.

On receipt of the PLME-ANP-DECISION.request primitive, the PLME shall consider both the ANPResponse parameter received in PLME-ANP-DECISION.request primitive and what was received over the air during the Rx period before taking any action. The action taken by the PHY layer of the PPD upon receipt of the PLME-ANP-DECISION.request primitive is fully described in 6.5.2. The PLME shall notify the MLME of the final decision via the PLME-ANP-DECISION.confirm primitive.

If the PPD detects an RTS burst from an SPD and has decided to reserve the upcoming superframe for that SPD, the PPD shall transmit the ACK corresponding to the received RTS codeword during the ANP. The PPD shall then enable its receiver for the duration of the SPD's beacon frame. The beacon frame shall be received and passed to the NHL via the MLME-INCOMING-BEACON.indication primitive. Immediately following the beacon frame reception, the receiver shall remain enabled through the Rx period, where the PPD again listens for an RTS burst (or an NPD codeword). No SPD shall transmit an RTS burst during the Rx period following another SPD's beacon frame. However, if the PPD does hear one, it shall be ignored. The PPD shall transmit a NACK during the ANP whether it receives an RTS burst or not, in order to ensure that at least every other beacon frame is transmitted by the PPD.

If the PPD detects more than one RTS burst from different SPDs and has decided to reserve the upcoming superframe for one of these SPDs, the PHY layer of the PPD shall randomly select one of the detected RTS codewords and transmit the corresponding ACK during the ANP.

The PPD may, at the discretion of an upper layer, aggregate the data received by an SPD(s) with its own data and transmit the aggregated data in subsequent beacon frames. The aggregation process could include combining the LAS channels protected by the SPD with those protected by the PPD and transmitting this combined list in the Map field of the PPD's beacon frame.

7.4.5.4 Illustrations

Figure 29 illustrates the scenario of an SPD interrupting a PPD in order to transmit a beacon frame. In (a), the PPD transmits the synchronization bursts and the beacon frame. The PPD then enables its receiver to listen for an RTS burst from an SPD. In this case, the PPD does receive an RTS burst, and consequently, the PPD transmits an ACK. In (b), the SPD transmits the synchronization bursts and the beacon frame. The PPD then enables its receiver to listen for an RTS burst and transmits a NACK whether or not it receives anything. In (c), the PPD transmits the synchronization bursts and the beacon frame. The PPD then enables its receiver to listen for an RTS burst from an SPD. In this case, the PPD does not receive an RTS burst, and consequently, the PPD transmits a NACK. The scenario in (c) repeats until another RTS burst is received, as in (a).

Figure 30 illustrates the Go-On scenario. In (a), the SPD transmits the synchronization bursts and the beacon frame, with the NST subfield set to one, in superframe n . Because the beacon frame was sent by an SPD, no other SPD will transmit an RTS burst during the Rx period, and the PPD transmits a NACK. In (b), the PPD transmits the synchronization bursts and the beacon frame in superframe $n + 1$. In this case, the PPD has decided to allow the SPD from (a) to transmit another beacon frame. Therefore, the PPD ignores any incoming RTS bursts and transmits a Go-On. In (c), the SPD transmits the synchronization bursts and the beacon frame, with the NST subfield set to zero, in superframe $n + 2$. Again, because the beacon frame was sent by an SPD, no other SPD will transmit an RTS burst during the Rx period, and the PPD transmits a NACK.

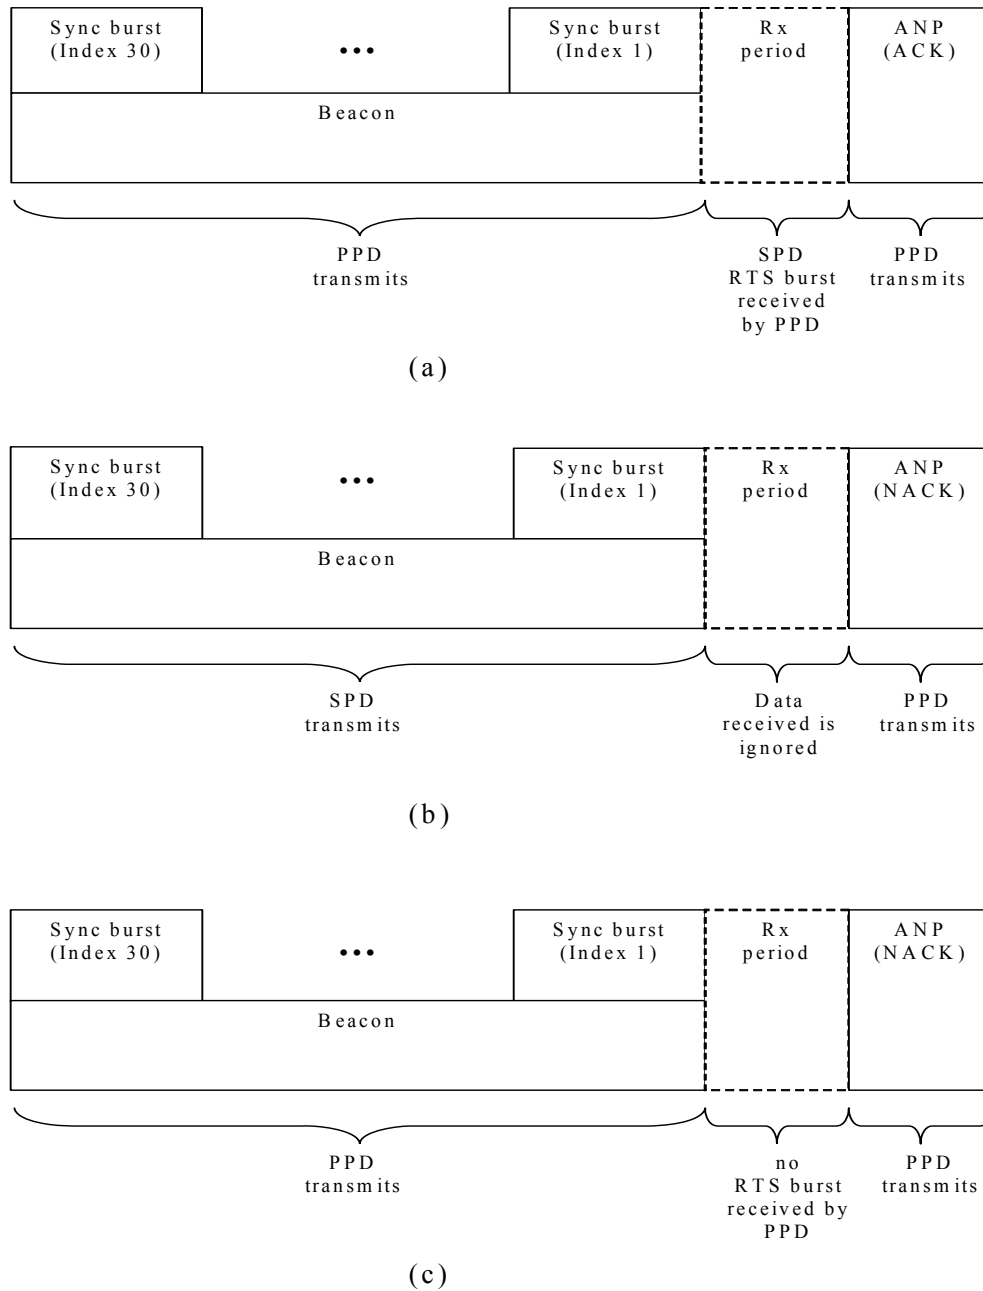
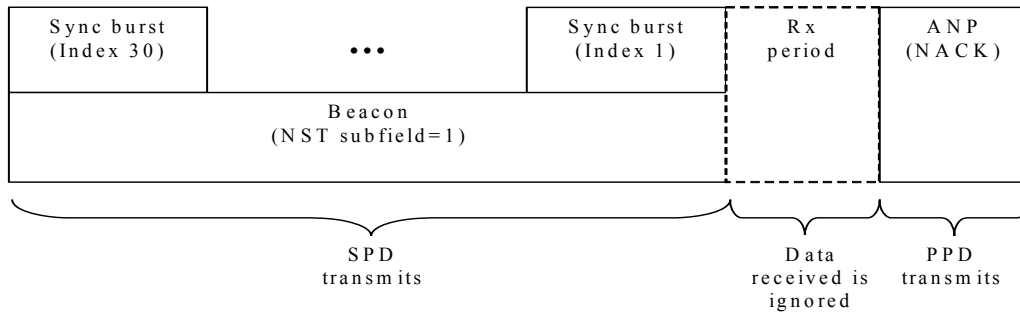


Figure 29—SPD interruption of the PPD in order to transmit a beacon frame


7.4.6 Beacons

7.4.6.1 Primary protecting device (PPD)


The PPD shall transmit a beacon frame in every superframe unless it has been interrupted by an SPD wishing to send its own beacon frame (7.4.5). The NHL of the PPD should issue a new MLME-START-BEACON.request primitive anytime it updates data in the PSDU (required need timer, LAS channel map, etc.).



(a) Superframe n



(b) Superframe $n+1$



(c) Superframe $n+2$

Figure 30—An example of the Go-On process

If the PPD does not transmit a beacon frame due to an interrupt by an SPD, the PPD shall listen for the beacon frame of the SPD and pass it to the NHL via the MLME-INCOMING-BEACON.indication primitive.

7.4.6.2 Secondary protecting device (SPD)

This subclause shall only apply to those SPDs that are not the NPD.

Every SPD shall be responsible for verifying that its protection needs remain satisfied by continuously monitoring the PPD. Every received beacon frame shall be passed to the NHL via the MLME-INCOMING-BEACON.indication primitive. If the SPD does not receive a valid PPD beacon frame within *macMaxMissedPPDBeacons* superframes, the MAC sublayer shall notify the NHL that it no longer detects the PPD's beacon via the MLME-PPD-LOST.indication primitive. The action taken by the NHL on receipt of this primitive is out of the scope of this standard.

The SPD may transmit its own beacon frame by interrupting the PPD (7.4.5). The SPD should send a beacon frame anytime it updates its data (required need timer, LAS channel map, etc.). Every SPD shall send at least one beacon frame every *macActivePeriodSPD* superframes in order to both update its information in the PPD and notify the PPD that it is still active. If the PPD does not receive any beacon frames from a particular SPD within $_macActivePeriodSPD \times _macMissedSPDBeacons$ consecutive superframes, the PPD shall consider this SPD inactive, and the MLME of the PPD shall notify its NHL via the MLME-SPD-LOST.indication primitive.

7.4.6.3 Next-in-line protecting device (NPD)

The PPD should select an SPD to become the NPD in the event that the PPD stops transmitting periodic beacon frames. If the PPD does cease beacon frame transmissions, the NPD shall promote itself and become the new PPD. In this case, the newly promoted PPD should select a new NPD.

7.4.6.3.1 NPD selection

When the PPD wants to select an NPD, it should set the NPD Indication subfield in its beacon frames to 00, indicating that there is no NPD currently in the beaconing network and that one is needed; recall that each time a change is required in the NPD Indication subfield, the NHL will issue a new MLME-START-BEACON.request primitive to the MLME. All SPDs receiving the PPD's beacon frames with the NPD Indication subfield set to 00 shall transmit an RTS burst and then, if permitted by the PPD, a beacon frame. Every time the PPD receives a beacon frame from an SPD, the MLME of the PPD shall pass the information from the beacon frame to its NHL via the MLME-INCOMING-BEACON.indication primitive. The NHL may choose one of the SPDs it has heard to be the NPD. The method for choosing which SPD will be the NPD is out of the scope of this standard. Note that the PPD shall announce its selection of an SPD two superframes following reception of the SPD's beacon frame. In other words if an SPD transmits a beacon frame in superframe n , the PPD shall set the NPD Indication subfield to 01 in the beacon frame sent in superframe $n + 2$ to notify the SPD that it has been chosen as the NPD.

Once an SPD has been selected as the NPD, the NHL of the PPD will inform the MLME of its decision by issuing the MLME-NPD.request primitive with the NPDAddress parameter set to the address of the selected SPD. The PPD shall transmit the next two consecutive beacon frames. The second of these beacon frames shall have the NPD Indication subfield set to 01, indicating that the last SPD that transmitted a beacon frame has been selected as the NPD. The MLME of the PPD shall then respond to the NHL via the MLME-NPD.confirm primitive with the Status parameter equal to COMPLETE.

If the SPD sees the NPD Indication subfield of the PPD's beacon frame transition from 00 in the first superframe following the transmission of its own beacon frame, to 01 in the second superframe, the SPD shall know that it was selected to be the NPD. The SPD shall then set the MIB attributes *macNPDAddress* and *macNPDPresent* equal to its own address and TRUE, respectively. The SPD shall respond two beacon frames following the reception of the PPD's beacon frame. In other words if the SPD receives the PPD's beacon frame in superframe $n + 2$, the SPD shall respond by sending an NPD codeword during the Rx period in superframe $n + 4$. The process will be complete once the PPD receives the NPD codeword and sets the MIB attributes *macNPDAddress* and *macNPDPresent* equal to the SPD's address and TRUE, respectively.

7.4.6.3.2 Maintaining the NPD's presence

Because the NPD is a special case of an SPD, it shall go through the same procedure as any other SPD to transmit a beacon frame, including the implementation of the retry procedure (7.4.2) when necessary. In addition to occasional beacon transmissions, the NPD shall send an NPD codeword within every *aNPDPeriod* superframes in order to indicate that it is still active. The procedure for sending an NPD codeword, as shown in Figure 31, shall begin once the NPD is within two superframes of reaching the *aNPDPeriod* superframe limit. If the PPD did not send a NACK during the last superframe (i.e., the current beacon frame shall be sent by an SPD), the MLME of the NPD shall generate the PLME-NPD-HEARTBEAT.request primitive, such that the PLME shall send an NPD codeword during the Rx period of the current superframe. Otherwise, if the PPD did send a NACK during the last superframe, the MLME shall wait until to issue the primitive, unless the *aNPDPeriod* superframe limit is reached. Once the limit is reached, the NPD shall send the NPD codeword regardless of what was sent during the last ANP. If the NPD wants to send an RTS burst at the time when the NPD codeword is due to be transmitted, priority shall be given to the NPD codeword and the transmission of the RTS shall be delayed.

When the PHY layer of a PD detects an NPD codeword, it shall notify the MAC sublayer via the PLME-NPD-ACTIVE.indication primitive. The NPD does not require a response from any PD to confirm receipt of the NPD codeword. Since the PPD always transmits something during the ANP, the PPD shall transmit a NACK unless it also received an NST request or an RTS burst, in which case it shall respond appropriately.

If the PPD does not receive either an NPD beacon frame or an NPD codeword within the last ($aMaxMissedNPDCodes \times aNPDPeriod$) superframes, the MLME shall notify the NHL via the MLME-NPD-LOST.indication primitive. The PPD may then decide to select a new NPD by following the procedure in 7.4.6.3.1 again.

Upon receipt of a beacon frame with the NPD Indication subfield set to 01, each SPD shall monitor both the incoming beacon frames and the Rx periods for the presence of the NPD. An SPD can identify an NPD beacon frame by examining the contents of the NPD subfield. An SPD shall listen for an NPD codeword during each Rx period when not sending an RTS burst itself. If an SPD does not receive either an NPD beacon frame or an NPD codeword within the last ($aMaxMissedNPDCodes \times aNPDPeriod$) superframes, the SPD shall conclude that the NPD is no longer active in the beaconing network. The MLME of the SPD shall notify its NHL via the MLME-NPD-LOST.indication primitive.

7.4.6.3.3 De-selecting an NPD

If the PPD changes the NPD Indication subfield in its beacon frame from 01 to 00, indicating that there is no longer an NPD present in the beaconing network and that an NPD is needed, any SPD receiving this beacon frame shall follow the NPD selection process, as described in 7.4.6.3.1. If the PPD changes the NPD Indication subfield in its beacon frame from 01 to 11, no NPD currently exists in the beaconing network, and *macNPDPresent* should be set to FALSE in the respective MIB attributes of all the PDs in the beaconing network.

In either case, the SPD that was previously operating as the NPD shall revert to operating as an ordinary SPD (i.e., an SPD that is not the NPD). Accordingly, the former NPD should remove its address from *macNPDAddress* and set the NPD subfield in all subsequent beacon frames to zero. Similarly, the PPD and any beaconing device that detects the transition in the NPD Indication subfield should also remove the former NPD's address from *macNPDAddress*.

7.4.6.3.4 Promotion of an NPD

If the NPD does not receive a valid beacon frame (7.4.3) from the PPD within *macMaxMissedPPDBeacons* superframes, the MLME of the NPD shall notify its NHL via the MLME-PPD-LOST.indication primitive.

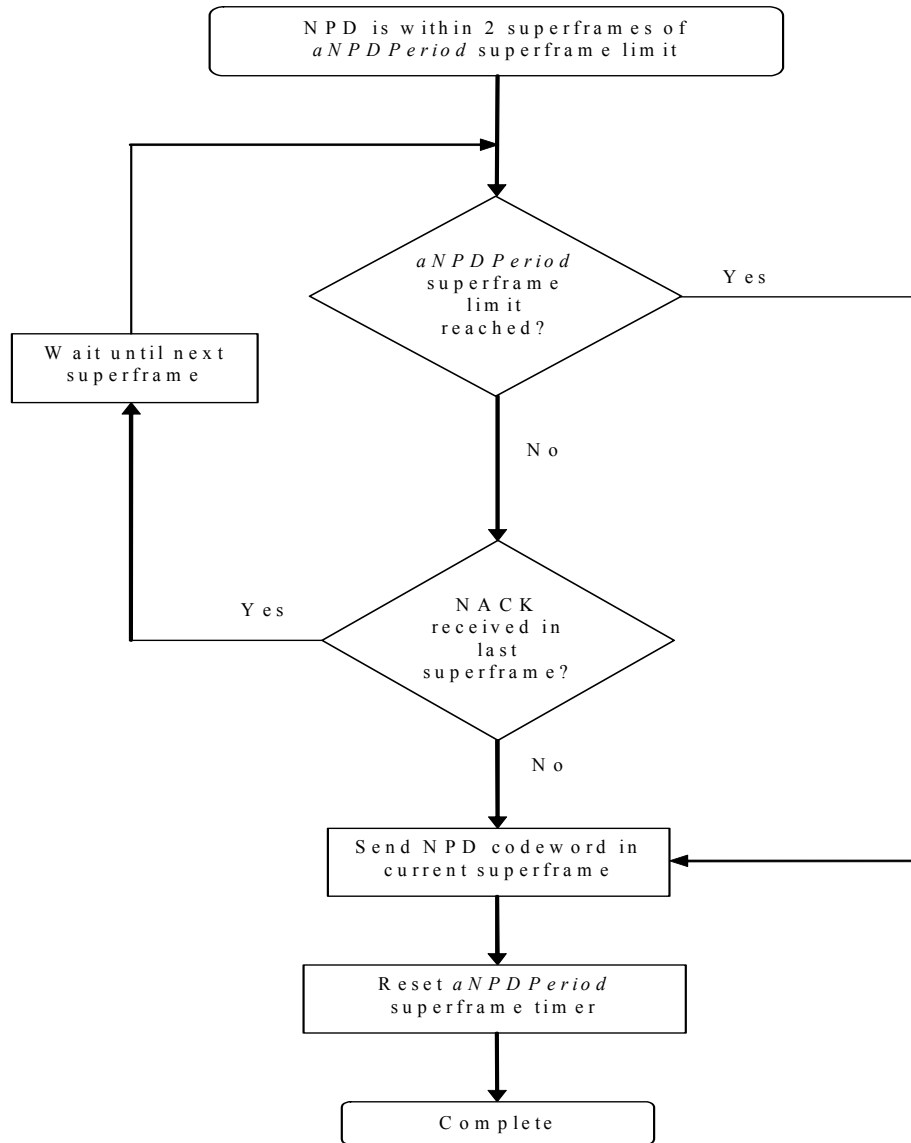


Figure 31—Procedure for sending an NPD codeword

The NHL may then decide to promote the device to PPD (7.4.7.1). In this case, the newly promoted PPD should select a new NPD.

7.4.7 Ceasing transmissions

This subclause describes the mechanism that is used to promote a new NPD or SPD to become the new PPD in case the previous PPD ceases its transmission. Procedure is also explained of what happens when the SPD ceases its transmission or an NPD ceases its transmission and the way this is signaled to the PPD.

7.4.7.1 Primary protecting device (PPD)

A probable scenario is that the PPD will abruptly cease transmitting. If there is an NPD in the beaconing network, the MAC sublayer of the NPD shall issue an MLME-PPD-LOST.indication primitive to the NHL once *macMaxMissedPPDBeacons* consecutive PPD beacon frames are missed. The NHL of the NPD should

then promote the NPD to be the new PPD immediately. Similarly, when an SPD misses *macMaxMissedPPDBeacons* consecutive PPD beacon frames, the MLME of the SPD shall issue an MLME-PPD-LOST.indication primitive to the NHL. The action taken by the NHL of the SPD on receipt of this primitive is out of the scope of this standard. Note, however, that the SPD should know whether there is an NPD present in the beaconing network and should take this into consideration before taking any action.

If the NPD Indication subfield in the PPD's final beacon frame was equal to 01 and the NHL of the SPD did not receive an MLME-NPD-LOST.indication primitive from its MLME, the NHL of the SPD should conclude that an NPD is present in the beaconing network. The SPD should allow the NPD sufficient time to promote itself to PPD before taking any action, assuming that the NPD can provide adequate protection for the SPD. However, if the NPD Indication subfield was set to either 00 or 11 and/or the NHL of the SPD received an MLME-NPD-LOST.indication primitive from its MLME, the NHL of the SPD may conclude that no NPD is present. If this conclusion is reached, one option is for the SPD to promote itself to PPD and begin transmitting periodic beacon frames.

If the PPD is aware that it is about to cease transmission, it shall set the Cease Tx subfield to one upon sending its last beacon frame. If there is an NPD in the beaconing network, the NHL of the NPD shall be notified through the MLME-INCOMING-BEACON.indication primitive. The NHL of the NPD should then promote the NPD to be the new PPD immediately. Similarly, for all SPDs in the beaconing network, the NHL of each SPD shall be notified through the MLME-INCOMING-BEACON.indication primitive. The action taken by the NHL of the SPD on receipt of this primitive is out of the scope of this standard. Again, the SPD should know whether there is an NPD present in the beaconing network and should take this into consideration before taking any action.

Following the PPD's notification that it is ceasing transmission, the NHL of the PPD should initiate the process to stop the PPD's beacon frame transmissions via the MLME-START-BEACON.request primitive with the Start parameter set to FALSE. Upon receipt of this primitive, the MLME shall issue the PLME-SET-TRX-STATE.request primitive with the TRX_STATE1, TRX_STATE2, and TRX_STATE3 parameters all set to TRX_OFF and the Periodic parameter set to TRUE.

If either the NPD or an SPD does promote itself to PPD, this new PPD shall not be required to follow the device initialization procedure (7.4.4). The new PPD shall initially transmit most of the same beacon fields/subfields as the original PPD, in an effort to continue protecting the remaining SPDs. The beacon fields/subfields that shall initially change are the Source Address field, the Location field, the Time Parity subfield, and the NPD Indication subfield.

All SPDs should receive the new PPD's beacon frame, note that the Source Address field in the beacon frame has changed (i.e., the Source Address field is different from the value of *macPPDAddress*) and set the value of *macPPDAddress* to the value of the Source Address field. As usual, every SPD shall send a beacon frame once every *macActivePeriodSPD* superframes in order to both update its information in the PPD and notify the PPD that it is still active. Once every SPD has sent at least one beacon frame, the new PPD shall have the most current information for inclusion in its beacon frames.

7.4.7.2 Secondary protecting device (SPD)

If an SPD is aware that it is about to cease transmission, it should interrupt the PPD to send its own beacon frame (see 7.4.5.2). If the RTS is acknowledged, the SPD shall set the Cease Tx subfield in the beacon frame to one before sending. The NHL of the PPD shall be notified through the MLME-INCOMING-BEACON.indication primitive and should remove all information corresponding to the SPD from its own beacon frames. If the RTS is not acknowledged after several failed attempts (i.e., the retry procedure failed), the NHL of the SPD shall be notified via the MLME-START-BEACON.confirm primitive and should decide how to proceed. The action taken by the NHL of any device is out of the scope of this standard.

Upon receipt of the MLME-START-BEACON.confirm, the NHL of the SPD should initiate the process to stop any further activity on the TV channel by issuing the MLME-START-BEACON.request primitive with the Start parameter set to FALSE. Upon receipt of this primitive, the MLME shall issue the PLME-SET-TRX-STATE.request primitive with the TRX_STATE1, TRX_STATE2, and TRX_STATE3 parameters all set to TRX_OFF and the Periodic parameter set to TRUE.

If the SPD leaves the radio space without notifying the PPD, the MLME of the PPD shall notify its NHL, via the MLME-SPD-LOST.indication primitive, after not receiving any beacon frames from that SPD within $\textit{macActivePeriodSPD} \times \textit{macMissedSPDBeacons}$ consecutive superframes. At that point, the NHL of the PPD should remove all information corresponding to that SPD from its own beacon frames.

7.4.7.3 Next-in-line protecting device (NPD)

If an NPD is aware that it is about to cease transmission, it should interrupt the PPD to send its own beacon frame (see 7.4.5.2). If the RTS is acknowledged, the NPD shall set the Cease Tx subfield in the beacon frame to one before sending. The NHL of the PPD or an SPD shall be notified through the MLME-INCOMING-BEACON.indication primitive, and the PPD should remove all information corresponding to the NPD from its own beacon frames. The NHL of the PPD should then change the NPD Indication subfield in its subsequent beacon frames from 01 to either 00 or 11, depending on whether the PPD wants to assign a new NPD. If the RTS is not acknowledged after several failed attempts (i.e., the retry procedure failed), the NHL of the NPD shall be notified via the MLME-START-BEACON.confirm primitive and will decide how to proceed. The action taken by the NHL of any device is out of the scope of this standard.

Upon receipt of the MLME-START-BEACON.confirm, the NHL of the NPD should initiate the process to stop any further activity on the TV channel by issuing the MLME-START-BEACON.request primitive with the Start parameter set to FALSE. Upon receipt of this primitive, the MLME shall issue the PLME-SET-TRX-STATE.request primitive with the TRX_STATE1, TRX_STATE2, and TRX_STATE3 parameters all set to TRX_OFF and the Periodic parameter set to TRUE.

If the NPD ceases transmitting abruptly, the NHL of the PPD or an SPD shall be notified through the MLME-NPD-LOST.indication primitive after $(\textit{aMaxMissedNPDCodes} \times \textit{aNPDPeriod})$ superframes. The PPD should remove all information corresponding to the NPD from its own beacon frames. The NHL of the PPD should then change the NPD Indication subfield in its subsequent beacon frames from 01 to either 00 or 11, depending on whether the PPD wants to assign a new NPD.

All SPDs should receive the PPD's beacon frames and see that the value of the NPD Indication subfield has changed. If the NPD Indication subfield is equal to 00, indicating that the PPD is requesting a replacement NPD, all SPDs shall follow the NPD selection process, as described in 7.4.6.3.1. If the NPD Indication subfield is equal to 11, indicating that the PPD does not want to replace the NPD, no response shall be sent to the PPD by any SPD.

7.4.8 Setting the internal operating states of the transceiver

The internal operating states of the transceiver are changed via the PLME-SET-TRX-STATE.request primitive for a superframe period or periods, according to the PD's role (i.e., PPD, SPD, or NPD) and current requirement. The transceiver state for the various periods (i.e., beaconing, Rx period, and ANP) of one or more superframes may be set by issuing this primitive one time, in order to reduce the amount of signaling between the MAC sublayer and the PHY. This is possible due to the fact that the PPD transmits most of the beacon frames except on the infrequent occasions that an SPD is granted a chance to send its beacon frame.

The primitive of PLME-SET-TRX-STATE.request primitive has four parameters: TRX_State1 (the transceiver state during the beaconing period), TRX_State2 (the transceiver state during the Rx period),

TRX_State3 (the transceiver state during the ANP) and Periodic, which indicates whether the combination of the TRX_StateN parameters shall be valid for more than one superframe (see 6.2.2.15 for more details). Table 50 shows the relationship between the various PDs and the PLME-SET-TRX-STATE.request primitive parameters.

Table 50—Setting the transceiver states during one or more superframe periods

Type of PD	Requirement	Beaconing period (TRX_State1)	Rx period (TRX_State2)	ANP (TRX_State3)	Valid for more than one superframe (Periodic)
PPD	1 Transmit beacon frames during initial transmission period	TX_ON	TX_ON	TX_ON	TRUE
	2 Transmit a beacon frame	TX_ON	RX_ON	TX_ON	TRUE
	3 Receive a beacon frame	RX_ON	RX_ON	TX_ON	FALSE
SPD	4 Listen to the TV channel	RX_ON	RX_ON	RX_ON	TRUE
	5 Ask permission to transmit a beacon frame (send an RTS burst)	RX_ON	TX_ON	RX_ON	FALSE
	6 Transmit a beacon frame	TX_ON	RX_ON	RX_ON	FALSE
NPD	7 Listen to the TV channel + transmit an NPD codeword	RX_ON	TX_ON	RX_ON	FALSE
	8 Transmit a beacon frame + transmit an NPD codeword	TX_ON	TX_ON	RX_ON	FALSE
All	9 Search the TV channel	RX_ON	RX_ON	RX_ON	TRUE
	10 Cease operation on the TV channel	TRX_OFF	TRX_OFF	TRX_OFF	TRUE

The following explains the conditions leading up to each requirement shown in Table 50 and under which the PLME-SET-TRX-STATE.request primitive is issued:

- 1) Issued following receipt of the MLME-START-BEACON.request primitive with the Initialize parameter set to TRUE.
- 2) Issued following receipt of the PLME-ANP-DECISION.confirm primitive with the ANPResponse parameter equal to NACK.
- 3) Issued following receipt of the PLME-ANP-DECISION.confirm primitive with the ANPResponse parameter equal to ACK or GO-ON.
- 4) Issued following receipt of the PLME-ANP-RESPONSE.indication primitive with the ANPResponse parameter equal to NACK in the absence of the receipt of a MLME_START-BEACON.request primitive.
- 5) Issued following receipt of both the PLME-ANP-DECISION.indication primitive with the ANPResponse parameter equal to NACK and the MLME-START-BEACON.request primitive.

- 6) Issued following receipt of the PLME-INITATE-RTS-BURST.confirm primitive with the Status parameter equal to ACK or GO_ON.
- 7) Issued when the NPD has no requirement to transmit a beacon frame but does have a requirement to transmit an NPD codeword.
- 8) Issued when it becomes necessary to send an NPD codeword and following receipt of the MLME-START-BEACON.request primitive.
- 9) Issued following receipt of the MLME-SEARCH.request primitive.
- 10) Issued following receipt of the MLME-START-BEACON.request primitive with the Start parameter set to FALSE.

7.5 Security suite specifications

The following text specifies the cryptographic mechanisms that are used in this standard.

7.5.1 Notation and representation

This subclause describes the notation and representation of strings and integers for this standard and conforms to conventions described in ANSI X9.63-2001, IEEE Std 1363, and IEEE Std 1363a.

7.5.1.1 Strings and string operations

A string is a sequence of symbols over a specific set (e.g., the binary alphabet $\{0,1\}$ or the set of all octets). The length of a string is the number of symbols it contains. The right-concatenation of two strings x and y of length m and n respectively (notation: $x || y$), is the string z of length $m+n$ that coincides with x on its leftmost m symbols and with y on its rightmost n symbols. An octet is a bit string of length 8.

7.5.1.2 Integers and their representation

Throughout this standard, the representation of integers as bit strings or octet strings shall be fixed. All integers shall be represented as binary strings in most-significant-bit first order and as octet strings in most-significant-octet first order. This representation conforms to the convention in Section 4.3 of ANSI X9.63-2001.

7.5.2 Representation of time

The representation of time for this standard is derived from the \$GPZDA string that provides universal coordinated time (UTC) date and time information, as defined in the NMEA 0183 Interface Standard [B10]. The GPZDA date and time string is formatted as follows:

`$GPZDA,hhtms.ss,dd,aa,yyyy,xx,zz`

where

`hhtms.ss` = UTC Time, where `hh` is hour, `tm` is minutes, and `ss.ss` is seconds

`dd` = Day, 01 to 31

`aa` = Month, 01 to 12

`yyyy` = Year

`xx` = Local zone description, 00 to ± 13 hours

`zz` = Local zone minutes description (same sign as hours)

For example, “\$GPZDA,235958.00,31,12,2005,00,00” is “11:59pm and 58 seconds on Dec. 31, 2005.”

For this standard, the required accuracy of time used when calculating and verifying a signature is plus or minus 5 min, relative to UTC. As such, only the ten's digit t for the minutes is used and the portions of the time string representing the unit's digit for minutes m and seconds s are not used in this standard. Also, the local time zone information is not needed. Therefore, the 11-octet string used for *Time* in the signature generation algorithm, as described in 7.5.4.2, is as follows: *hhtddaayyyy*.

The value for the Time Parity subfield of the Parameter 2 field (7.2.1.4) depends on the value t in this time representation. The value of the Time Parity subfield shall be equal to $(t \bmod 2)$. That is, when t is zero, two, or four, the subfield shall be zero and when t is one, three, or five, the subfield shall be one.

For example, “\$GPZDA,235958.00,31,12,2005,00,00” results in *hhtddaayyyy* = 23531122005 and the Time Parity subfield equals one (since $t = 5$).

7.5.3 Elliptic-curve building blocks

The following elliptic-curve parameters and data elements are defined for use in this standard.

7.5.3.1 Elliptic-curve domain parameters

The elliptic curve domain parameters D used in this standard shall be those for the curve “ansip224k1,” as specified in Appendix J5.3, Example 1, of ANSI X9.63-2001.

All elliptic-curve points (and operations hereon) used in this standard shall be performed on this curve.

7.5.3.2 Elliptic-curve point representation

All elements of the finite field F_p shall be represented as specified in Section 4.1.1 of ANSI X9.63-2001. All elliptic-curve points shall be represented in compressed form, as specified in Section 4.2.1 of ANSI X9.63-2001, and when transmitted, shall be converted to an octet string representation, as specified for compressed points in Section 4.3.6 of ANSI X9.63-2001. Thus, each elliptic-curve point can be represented in 29 octets.

7.5.3.3 Elliptic-curve public-key pair

An elliptic-curve-key pair consists of an integer q and a point Q on the curve determined by multiplying the generating point G of the curve by this integer (i.e., $Q = qG$), as specified in ANSI X9.63-2001. Here, Q is called the public key, whereas q is called the private key; the pair (q, Q) is called the public-key pair. Each private key shall be represented as specified in Section 4.3.1 of ANSI X9.63-2001. Each public key shall be represented as defined in 7.5.3.2.

7.5.3.4 Elliptic-curve implicit signature

The signature field of a beacon frame shall be generated and verified using the scheme and procedures described in 7.5.4. The exact format of the 44-octet signature field shall be specified as follows:

KeyID || *RecoverableMessage* || *Signature*

where

- *KeyID* is the 1-octet sequential identifier for the public key corresponding to the signer (identified by the MAC address) that should be used to verify this signature
- *RecoverableMessage* is the 14-octet recoverable message portion of the signature referred to as the c output at step 8 of the signature generation procedure described in 10.5.2 of IEEE Std 1363a
- *Signature* is the 29-octet signature portion of the signature referred to as the d output at step 8 of the signature generation procedure described in 10.5.2 of IEEE Std 1363a

A convention for assigning *KeyIDs* should be established by an industry consortium. For example, the first public key issued by a certificate authority (*CA*) to a particular subject (i.e., MAC address) could start with this identifier set to zero and each subsequently issued public key to this subject could have its associated *KeyID* value incremented by one.

7.5.3.5 Elliptic-curve implicit certificate

The certificate field of a beacon frame shall be generated and processed using the scheme and procedures described in 7.5.5. The exact format of the 31-octet implicit certificate field of the beacon frame shall be specified as follows:

$$KeyIssuerID \parallel ExpirationDate \parallel PublicReconstrKey$$

where

- *KeyIssuerID* is the 1-octet identifier of the *CA* that created the implicit certificate during the execution of the implicit certificate generation protocol
- *ExpirationDate* is the 1-octet identifier used to indicate the date that this certificate expires. The expiration date is October 1, 2007 + *x*, where *x* is the integer value in years represented by this octet. A value of 255 for *x* indicates an infinite expiration date
- *PublicReconstrKey* is the 29-octet representation of the public-key reconstruction data *BEU*, as specified in the implicit certificate generation protocol (7.5.5.2), which is an elliptic-curve point, as specified in 7.5.3.2

7.5.4 Signature scheme

The signature scheme used in this standard is an instance of the Elliptic Curve Signature Scheme with Recovery, Pintsov-Vanstone (ECSSR-PV), as described in 10.5 of IEEE Std 1363a.

7.5.4.1 Signature scheme setup

This subclause establishes the options, as described in 10.5.1 of IEEE Std 1363a, chosen for the signature scheme. The elliptic curve domain parameters *D* specified in 7.5.3.1 shall be used for this signature scheme. In addition, the following options are established:

- The pre-signature, signature and verification primitives shall be ECSP-NR2/PV, ECSP-PV, and ECVP-PV.
- The message-encoding method shall be EMSR2. The encoding parameter *padLen* shall be the integer 14 (indicating the number of amount of added redundancy in octets) and *l* (the input to I2OSP at step 2 of EMSR2) shall be one. The method for combining the pre-signature with the (padded) recoverable message part shall be a stream cipher based on KDF2.
- The redundancy criteria necessary for acceptance of the message after it has been recovered and successfully decoded shall be that the *RecoverableMessage* portion of the signature field (i.e., the *c* output at step 8 of the signature generation procedure described in 10.5.2 of IEEE Std 1363a) equals the octet string "0E" (i.e., the octet string representing *padLen*) repeated 14 times (i.e., "0E 0E 0E 0E 0E 0E 0E 0E 0E 0E 0E 0E 0E 0E").
- The hash function *Hash* used whenever a hash is required in this standard shall be SHA-256, as described in 14.1.3 of IEEE Std 1363a.

7.5.4.2 Signature generation

This subclause describes the inputs and outputs to the signature generation scheme, as described in 10.5.2 of IEEE Std 1363a.

There are two inputs to this signature generation scheme, the recoverable message part M_1 and the nonrecoverable message part M_2 . For the purposes of this standard, the recoverable message part M_1 shall be the empty octet string of length zero and the nonrecoverable message part M_2 shall be the length 33 octet string comprised as follows:

$$M_2 = \textit{Parameter 1} \parallel \textit{Source Address} \parallel \textit{Location} \parallel \textit{Parameter 2} \parallel$$

$$\textit{Parameter 3} \parallel \textit{Map} \parallel \textit{Time}$$

where *Parameter 1*, *Source Address*, *Location*, *Parameter 2*, *Parameter 3*, and *Map* are taken from the to-be-sent beacon frame, as defined in 7.2, and *Time* is the octet string representing the time and date, as defined in 7.5.2. The private key selected for use at step 2 of the signature generation scheme, as described in 10.5.2 of IEEE Std 1363a, shall be *macPrivateKey*, as described in 7.3.2.

The output to the signature generation scheme, as described in 10.5.2 of IEEE Std 1363a, consists of two parts: c (the message representative of the recoverable message part M_1) and d (the signature part). As described in 7.5.3.4, the output c shall be the *RecoverableMessage* portion of the to-be-sent Signature field, and the output d shall be the *Signature* portion of the to-be-sent Signature field.

7.5.4.3 Signature verification

This subclause describes the inputs and outputs to the signature verification scheme, as described in 10.5.3 of IEEE Std 1363a.

There are three inputs to this signature verification scheme: c (the message representative of the recoverable message part M_1), d (the signature), and M_2 (the nonrecoverable message). The input c shall be the *RecoverableMessage* portion of the received Signature field, and the input d shall be the Signature portion of the received Signature field. The input M_2 shall be the length 33 octet string comprised as follows:

$$M_2 = \textit{Parameter 1} \parallel \textit{Source Address} \parallel \textit{Location} \parallel \textit{Parameter 2} \parallel$$

$$\textit{Parameter 3} \parallel \textit{Map} \parallel \textit{Time}$$

where *Parameter 1*, *Source Address*, *Location*, *Parameter 2*, *Parameter 3*, and *Map* are taken from the received beacon frame, and *Time* is the octet string representing the receiving device's current time and date accurate to within 10 min.

The value for the Time Parity subfield of the Parameter 2 field (7.2.1.4) is helpful in synchronizing the received time with the transmitted time. The receiving device may calculate its own value for time parity; that is receiving device's time parity $p' = t \bmod 2$, where t is the receiver's ten's digit for minutes in the time representation, as defined in 7.5.2. If $p' \neq p$ (where p is the Time Parity subfield in the received beacon frame), then the receiving device shall adjust the *Time* octet string in M_2 by either incrementing or decrementing t , the receiver's ten's digit for minutes in the time representation, based on the value of the receiver's unit's digit for minutes m (e.g., if $m < 5$ decrement t ; otherwise, increment t).

The signer's purported public key w' shall be either the public key extracted from the Certificate field of the received beacon frame or the public key obtained via an out-of-scope means (e.g., via an entry in the *macPeerPublicKeyTable* of the MIB). Note that the *KeyID* subfield that is part of the received Signature field (7.5.3.4) along with the MAC address of the device that transmitted this received beacon frame can be helpful in identifying this public key when the over-the-air certificate is not used.

Any error conditions that occur during the signature verification scheme, as described in 10.5.3 of IEEE Std 1363a, such as the processing steps resulting in the output “invalid,” shall be reported to the NHL using the appropriate primitive with SecurityStatus parameter set to SIGNATURE_INVALID.

7.5.5 Certificate scheme

The certificate scheme used in this standard is an instance of the ECQV Implicit Certificate Scheme as described in SEC 4.

In general, as described in SEC 4, an implicit certificate scheme is used by three entities: a Certificate Authority *CA*, a certificate requester *U*, and a certificate processor *V*, where *U* wishes to obtain an implicit certificate from *CA* in order to convey *U*'s associated public key to *V*. The scheme consists of a certificate generation method, a certificate validation method, and a certificate processing method. In the certificate generation method, *U* receives a certificate from *CA*. In the certificate validation method, *U* or *V* can verify that a certificate purportedly provided by *CA* to *U* indeed originated from *CA* and was provided to *U*. In the certificate processing method, *V* processes the certificate to yield a static public key (and associated keying information) purportedly bound to *U*; evidence that this public key is genuinely bound to *U* is only corroborated via subsequent use of the corresponding private key (e.g., a signing transformation involving *U*'s public-key pair).

Figure 32 gives a sequence chart showing an example of the certificate scheme for this standard. Note that this diagram has a separate licensing authority (e.g., an industry consortium or an appropriate regulatory body) and *CA*, though in practice these two authorities might be run by the same entity. Prior to deployment, the system entities will agree on the certificate scheme setup (7.5.5.1). Once deployed, the first step will be for *U*'s MAC address to be sent to the licensing authority so that a certificate can be generated. The owner of *U* trusts the licensing authority to execute the certificate generation method (7.5.5.2) with the *CA*.

When the certificate generation method is completed, the private key and certificate are securely sent from the licensing authority to the owner of beaconing device *U* (e.g., using a smart card). Entity *U* is not required to validate the implicit certificate as described in Section 3 of SEC 4, since it trusts the licensing authority to do this following one of the four methods given in SEC 4. Later, when *V* receives a beacon frame from *U*, it has two options for getting the public key needed to check the signature field of the beacon frame. Option 1 is to get this key from the certificate using the certificate processing method (7.5.5.3). Option 2 is to get this key from an alternate means (e.g., via direct and potentially off-line delivery from the licensing authority) and is out-of-scope from this specification. Whenever *V* successfully validates a certificate from *U*, it can store the public key and other information from this certificate as an entry in the *macPeerPublicKeyTable* of the MIB (e.g., to avoid the need to unnecessarily process the certificate multiple times).

7.5.5.1 Certificate scheme setup

This subclause establishes the options, as described in Section 2 of SEC 4, chosen for the certificate scheme. The following options are established:

- An infrastructure should have been established for the operation of the scheme. For this standard, this infrastructure should use the certificate format as specified in 7.5.3.5, and the identifiers and certificate processing rules specified in this standard. The unique assignment of values to identifiers, such as the MAC address, the *CA* identifier, and the *KeyID* (identifying the public key associated with a subject) to subjects and objects should have been accomplished by other industry, government, or standards organizations. Such organizations may also specify further certificate processing rules that additionally apply when this scheme is used in practice.
- The certificate scheme's elliptic curve domain parameters *D* shall be as specified in 7.5.3.1.
- Entity *U* shall be identified by its 6-octet MAC address and the *CA* shall be identified by a 1-octet identifier included in the certificate as the *KeyIssuerID* subfield described in 7.5.3.5.

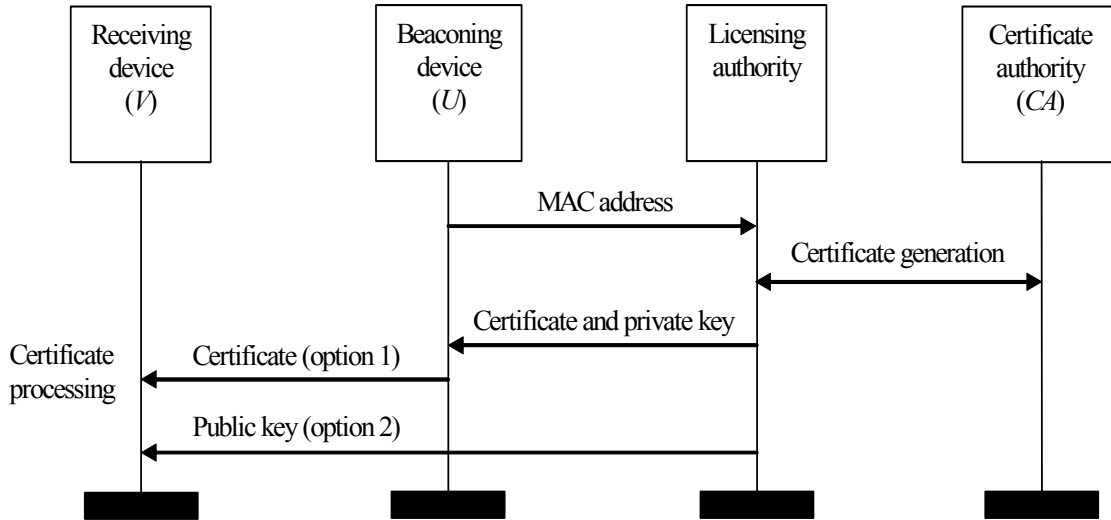


Figure 32—Certificate scheme

- The CA shall be bound to a static public-key pair (wCA , WCA) associated with the system’s elliptic curve domain parameters D . The CA’s public-key WCA and 1-octet identifier shall be stored as an entry in the MIB attribute *macAuthorityPublicKeyTable*.
- The hash function *Hash* chosen for use with the ECQV implicit certificate generation scheme shall be SHA-256, as described in 14.1.3 of IEEE Std 1363a.
- A beaconsing device shall represent elliptic curve points, as described in 7.5.3.2.
- The representation of octets as binary strings shall be as specified in 7.5.1.

7.5.5.2 Certificate generation

This subclause gives a specific instantiation of the certificate generation method, as described in Section 4 of SEC 4, that is relevant to this specification.

The “to-be-signed-certificate data” IU , as specified in Step 6 of the responder transformation in 4.2 of SEC 4, shall be constructed using the following concatenation:

$$IU = KeyID \parallel Subject \parallel KeyIssuerID \parallel ExpirationDate$$

where

- *KeyID* is the public-key identifier, as described in 7.5.3.4
- *Subject* is U ’s 6-octet IEEE address assigned to the device (i.e., the MAC address)
- *KeyIssuerID* is the CA identifier, as described in 7.5.3.5
- *ExpirationDate* is the expiration date of the certificate, as described in 7.5.3.5

U ’s implicit certificate ICU , as specified in Step 7 of the responder transformation in 4.2 of SEC 4, shall be constructed using the following concatenation:

$$ICU = IU \parallel BEU$$

where

- *IU* is the “to-be-signed-certificate data” as specified for Step 6 of the responder transformation in 4.2 of SEC 4
- *BEU* is the public-key reconstruction data *BEU*, as specified for Step 5 of the responder transformation in 4.2 of SEC 4

The certificate sent over-the-air by a beaconing device (i.e., the certificate subfield) is as specified in 7.5.3.5 and is not the output of the certificate generation method *ICU* directly. Instead, the *KeyID* subfield is sent within the Signature field (7.5.3.4) of the beacon frame and the *Subject* subfield is sent in the Source Address field of beacon frame (7.2). The receiving device *V* shall reconstruct *ICU* from these pieces prior to executing the certificate processing method.

7.5.5.3 Certificate processing

This subclause gives a specific instantiation of the certificate processing method, as described in Section 5 of SEC 4, that is relevant to this standard.

The receiving device *V* shall reconstruct *ICU* by pre-pending the received certificate with the *KeyID* subfield sent within the Signature field of the beacon frame (7.5.3.4) and the *Subject* subfield sent in the Source Address field of beacon frame (7.2). That is,

$$ICU = KeyID \parallel Subject \parallel Certificate$$

where

- *KeyID* is the subfield received from the Signature field (7.5.3.4)
- *Subject* is the subfield received from the Source Address field of beacon frame (7.2)
- *Certificate* is the certificate subfield of the received beacon frame (7.5.3.5)

At Step 3 of the implicit certificate processing transformation in Section 5 of SEC 4, the expiration date of the certificate shall be checked. If the certificate is expired, the processing method shall output “invalid” and stop. Although out-of-scope from this specification, the receiving device *V* may also opt to check whether the certificate for this beaconing device has been revoked.

Any error conditions that occur during the implicit certificate processing transformation in Section 5 of SEC 4, such as the processing steps resulting in the output “invalid,” shall be reported to the NHL using the appropriate primitive with the SecurityStatus parameter set to CERTIFICATE_INVALID.

7.6 Message sequence charts (MSC) illustrating MAC-PHY interaction

This subclause uses MSCs to provide example illustrations of how to implement the main tasks specified in this standard. Figure 33 and Figure 34 illustrate the device initialization procedure for the PPD and for an SPD, respectively. Figure 35 and Figure 36 illustrate inter-device communications. Figure 37 and Figure 38 illustrate the process for choosing an NPD.

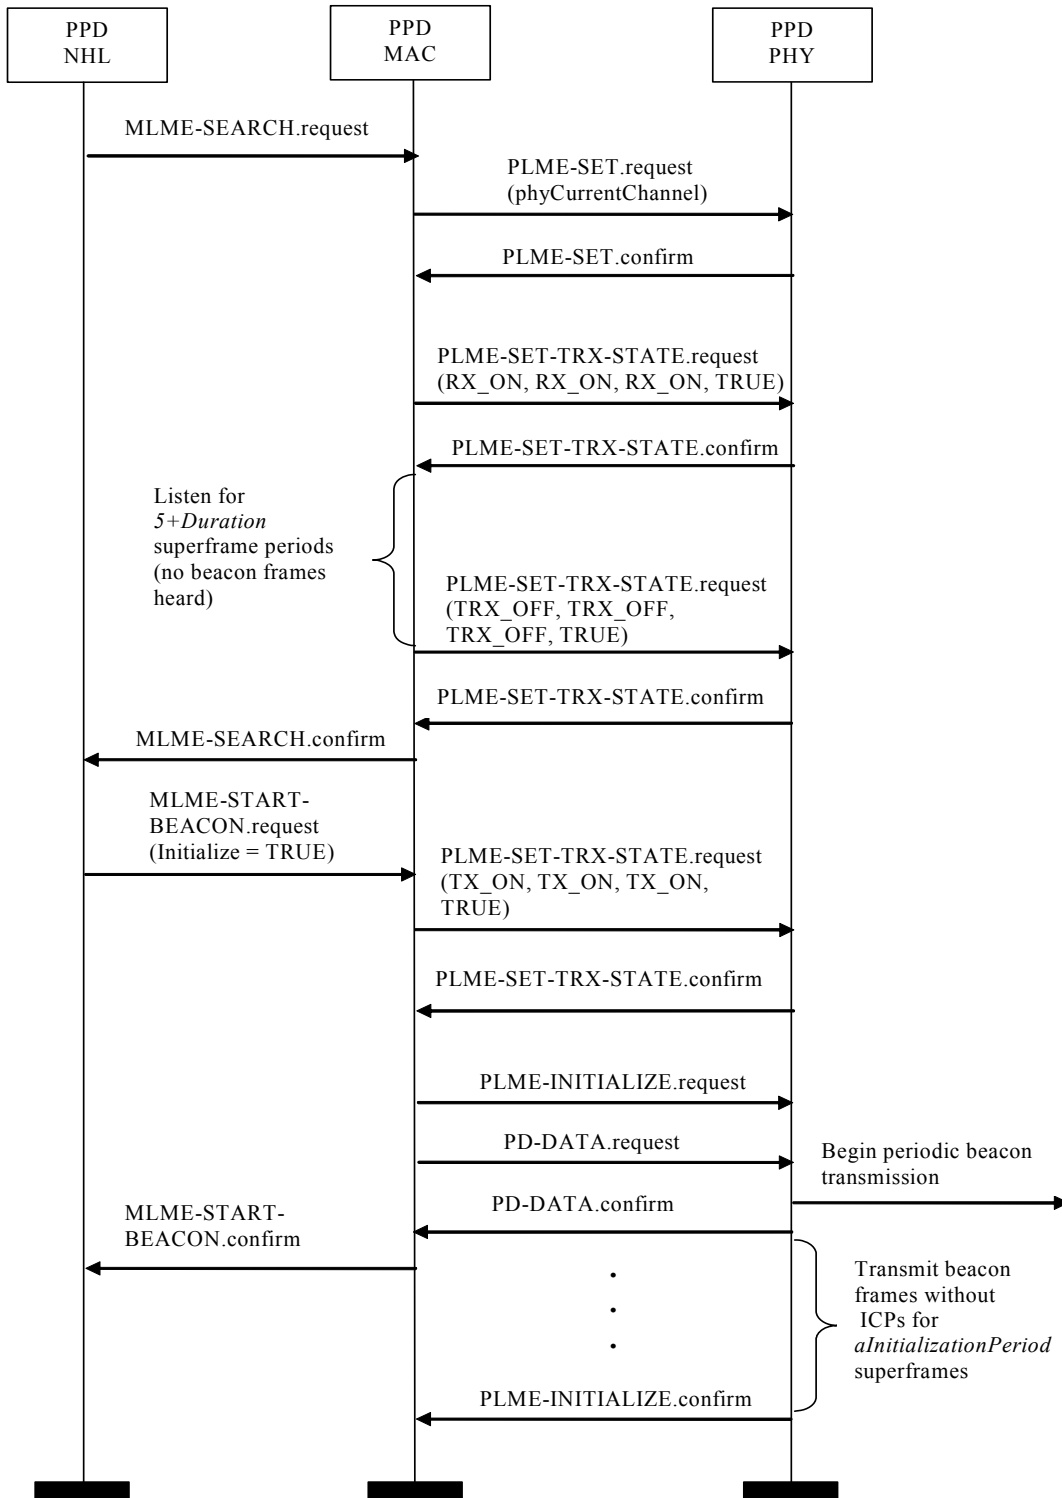
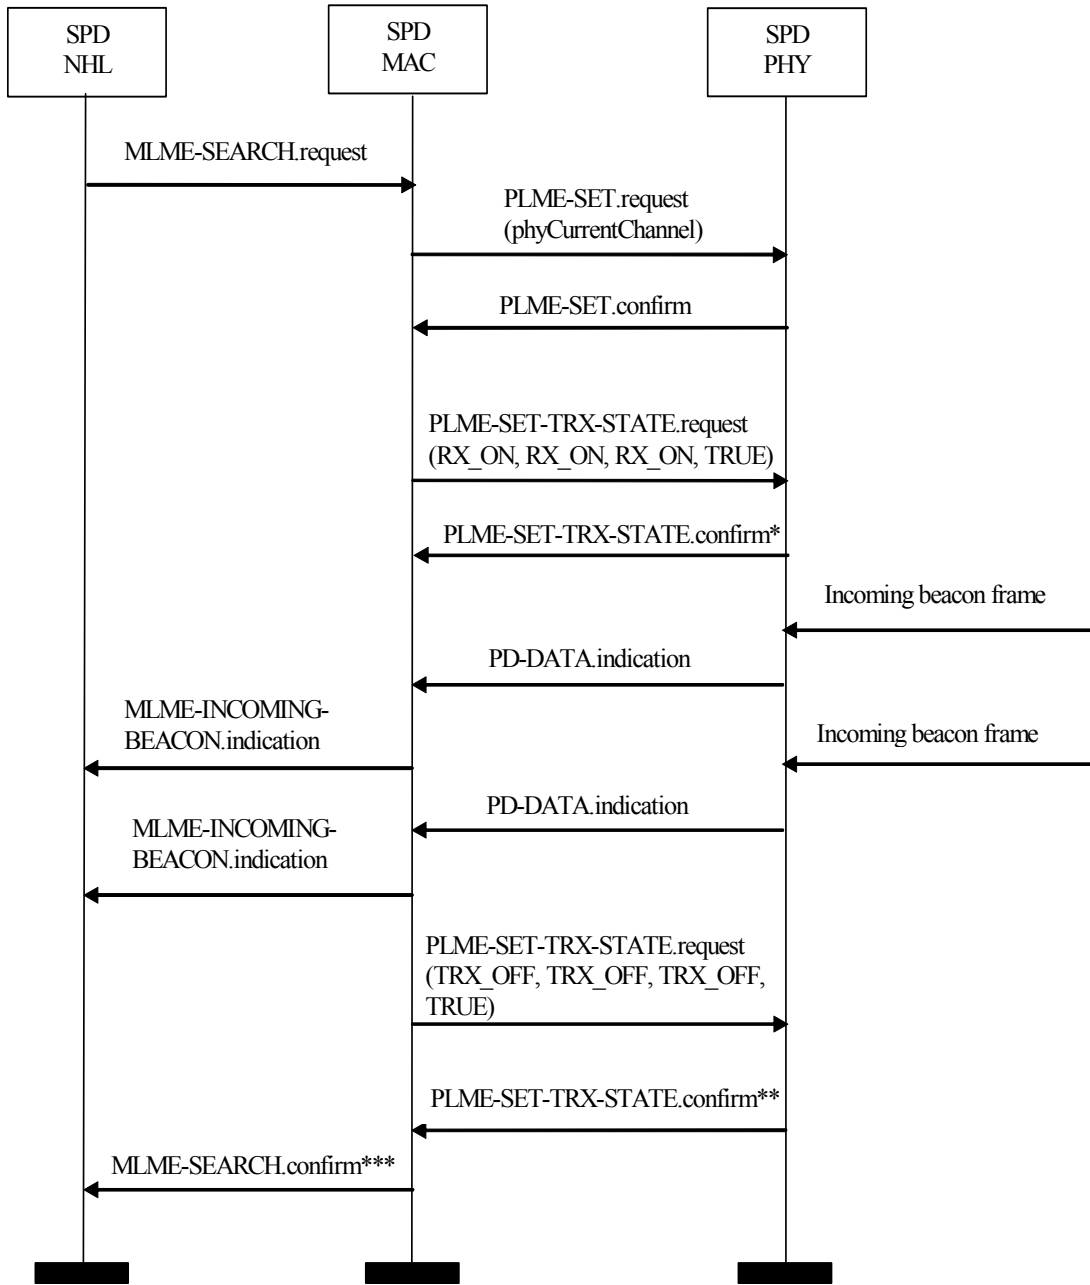


Figure 33—Initialization of the PPD (example)



*Begin the listening period lasting $5 + Duration$ superframe periods.

**End the listening period.

***Following the reception of the MLME-SEARCH.confirm, the NHL will wait for a period of 100 superframes and then initiate the procedure for the “SPD interrupting the PPD” (Figure 35).

Figure 34—Initialization of an SPD (example)

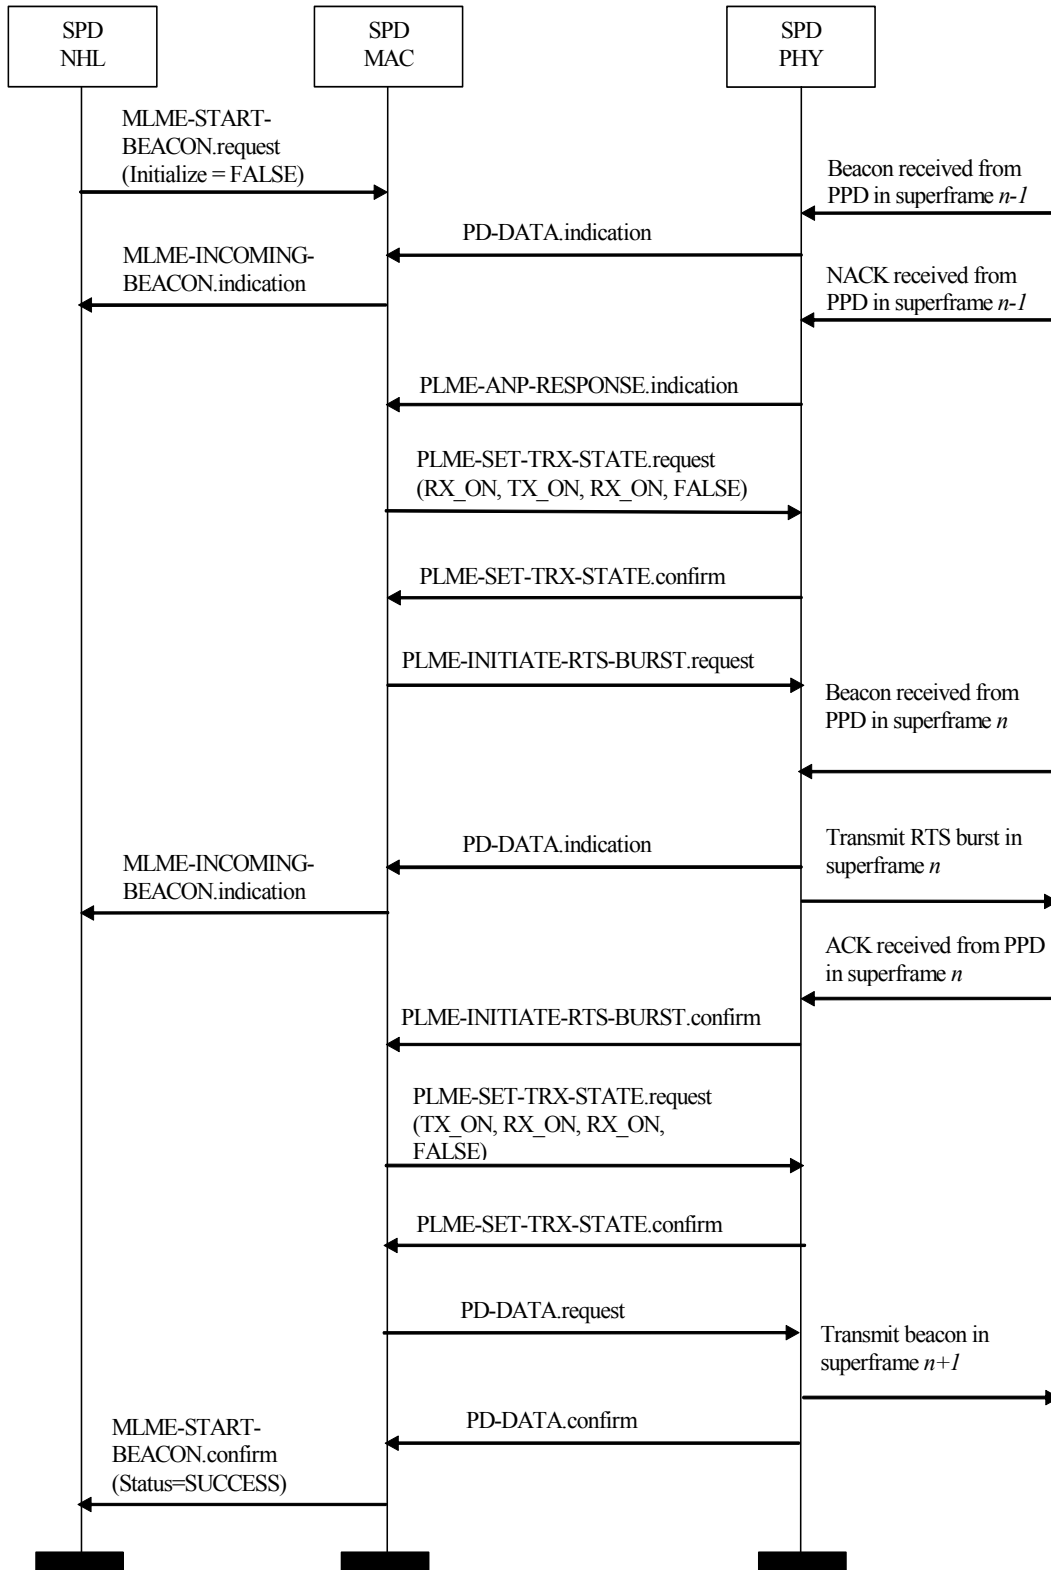
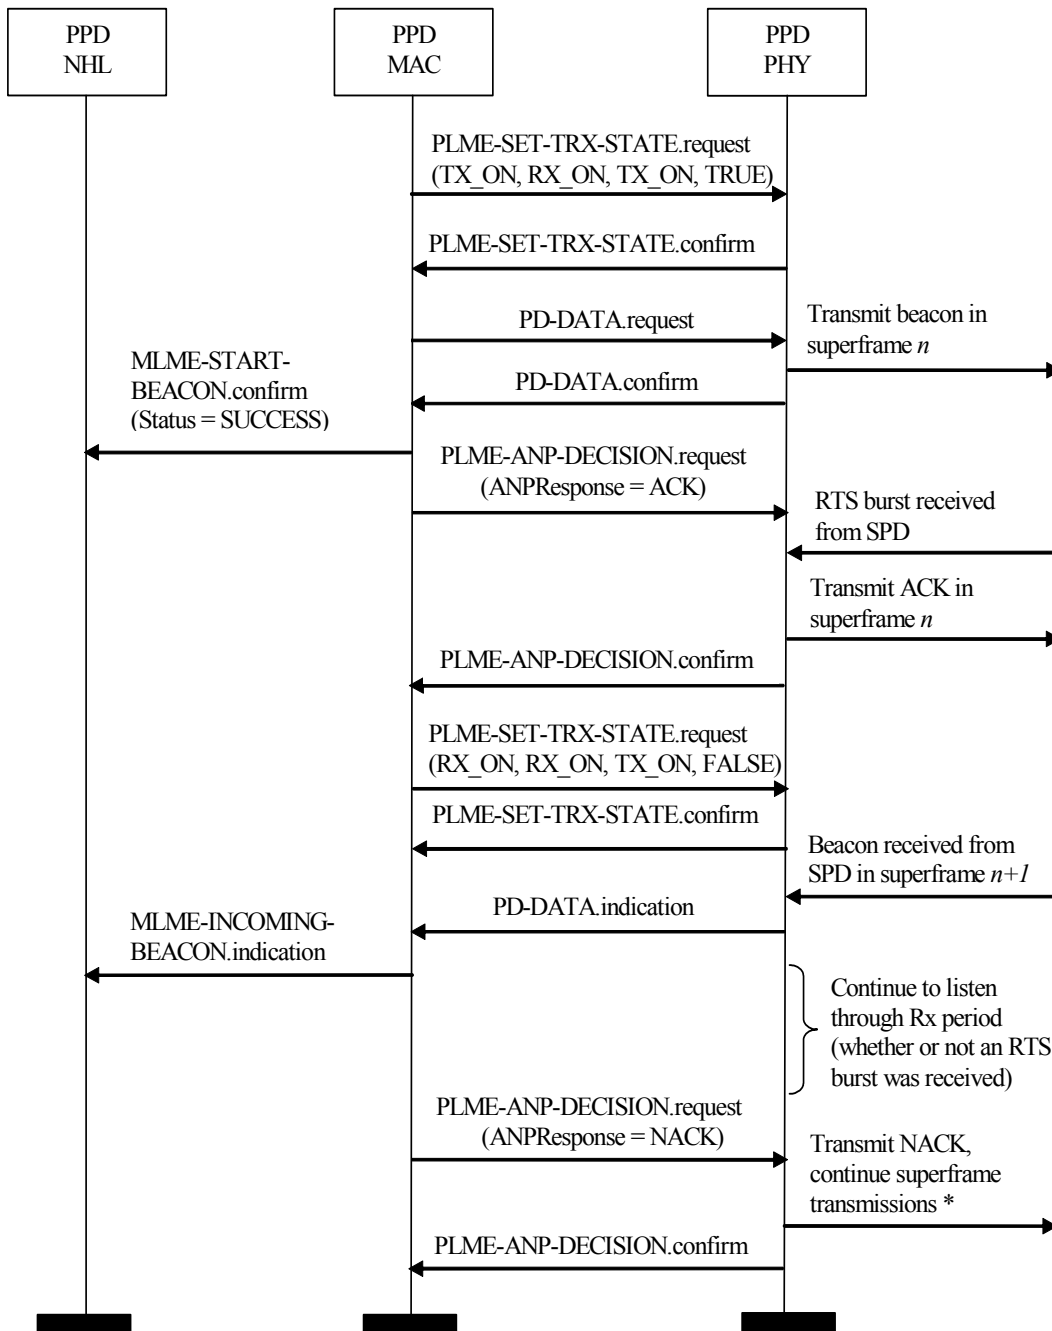
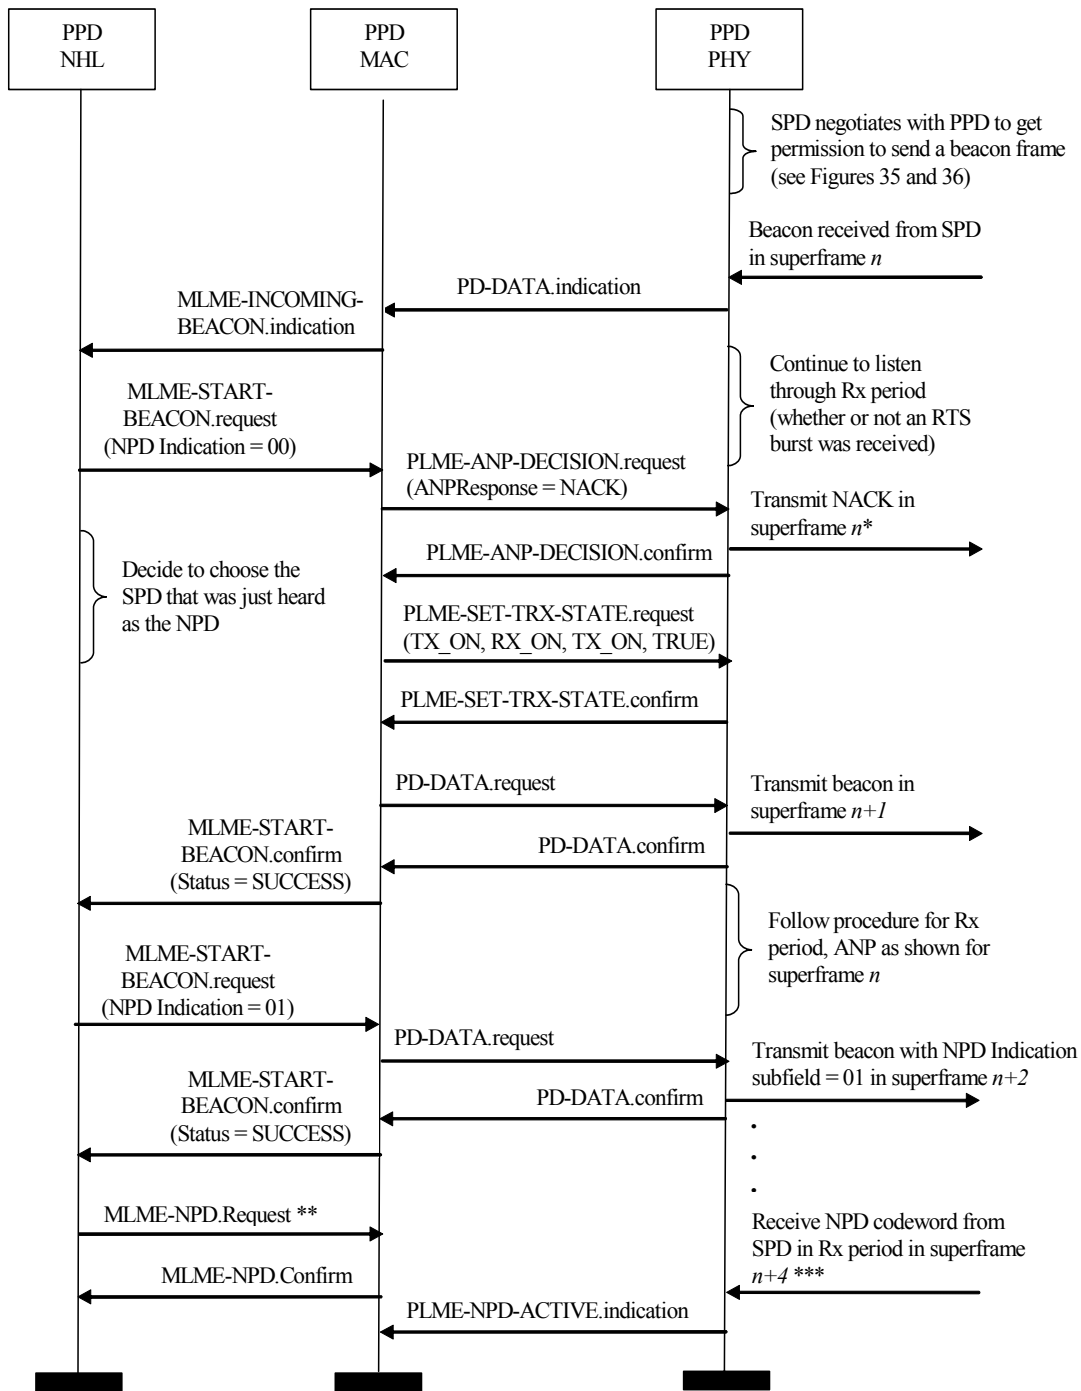


Figure 35—Example illustration of an SPN interrupting the PPD (for inter-device communications)



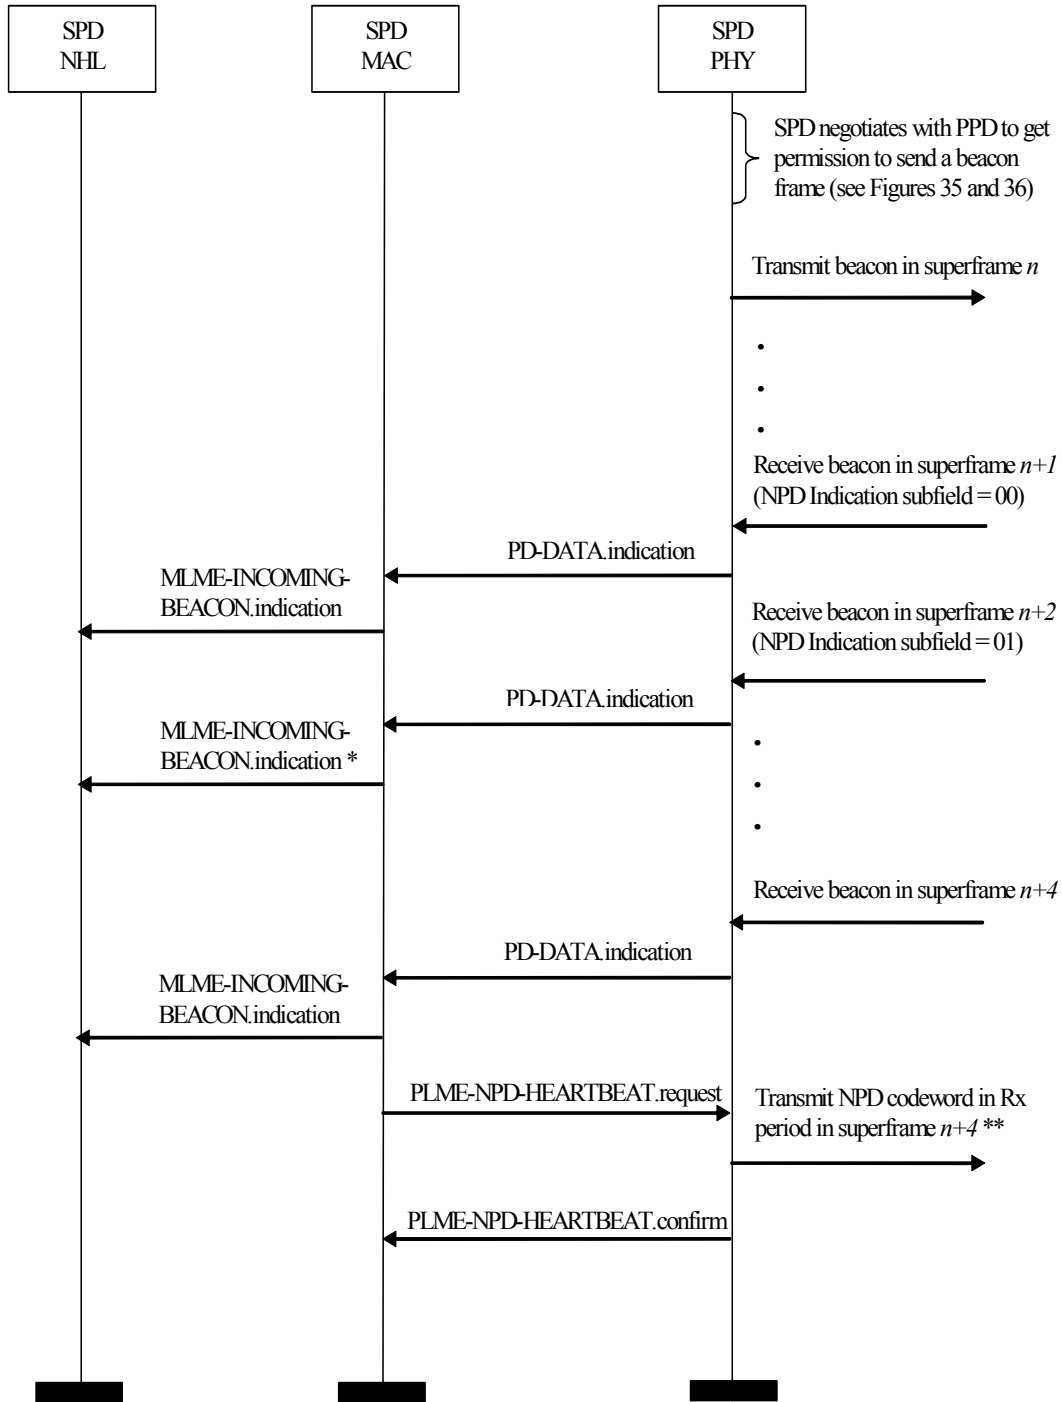
* The NHL may choose to combine the data received in the SPD's beacon with its own and then send the MLME-NEW-BEACON-DATA.request primitive to the MAC to change the beacon content.

Figure 36—Example illustration of the PPD being interrupted by an SPD (for inter-device communications)



* A NACK shall always be transmitted in the ANP following the reception of an SPD's beacon frame.
 ** The amount of time required for the NHL to choose an NPD and notify the MAC sublayer is not known. Therefore, exact timing cannot be shown.
 *** The incoming NPD code notifies the PPD that the NPD is active. The PPD must receive an NPD codeword at least once every ($macMaxMissedNPDCodes * macNPDPeriod$) superframes in order to know that the NPD is still active.

Figure 37—PPD selects an SPD as the NPD (example)



* The NHL sees that it was chosen as the NPD.

** The SPD transmits the NPD codeword at least every $macNPDPeriod$ superframes to notify the PPD that it is still active as the NPD.

Figure 38—SPD selected as the NPD (example)

Annex A

(informative)

Example block decoding method

A.1 Introduction

The decoding method described here may be used to recover the Index field of a synchronization burst (6.3).

A.2 Method details

Supposing the received parity bits are denoted by r_{14-r_7} and the received index bits are denoted by r_6-r_{10} , the decoding may be performed as described in the following paragraphs. The \mathfrak{r} vector is flipped left to right to simplify the description of the operations.

First, the syndrome sequence \mathfrak{s} is calculated as shown in Equation (A.1).

$$\mathfrak{s} = (s_0, s_1 \dots s_{14}) = \mathfrak{r}H^T = (r_0, r_1 \dots r_{14}) \times \begin{bmatrix} h_{0,0} & h_{0,1} & \dots & h_{0,14} \\ h_{1,0} & h_{1,1} & \dots & h_{1,14} \\ h_{14,0} & h_{14,1} & \dots & h_{14,14} \end{bmatrix}^T \quad (\text{A.1})$$

where H is the parity check matrix.

The parity check matrix is obtained in the following way. The first row vector of H is $\vec{h}_0 = \{1\ 0\ 0\ 0\ 0\ 0\ 0\ 1\ 0\ 0\ 0\ 1\ 0\ 1\}$. The $(i+1)$ _{th} row vector \vec{h}_i is obtained by performing a cyclic right shift of one element of the i _{th} row vector \vec{h}_{i-1} . Therefore the parity check matrix is as shown in Equation (A.2).

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} \quad (\text{A.2})$$

Next, the error vector, $\hat{e} = (e_0, e_1, \dots, e_{14})$, is calculated. In order to do so, the row index set $A(i) = \{j_1, j_2, \dots, j_4\}$ is defined for the i_{th} element of the error vector, e_i . Each element of $A(i)$ is a row index of H , and the i_{th} element of this row is “1.” If more than half of the members of $s_{A(i)} = \{s_{j_1}, s_{j_2}, \dots, s_{j_4}\}$ equal one, then $e_i = 1$. Otherwise, $e_i = 0$. For example, if $i = 0$, then $A(i) = A(0) = \{0, 1, 3, 7\}$ and $s_{A(i)} = s_{A(0)} = \{s_0, s_1, s_3, s_7\}$. Therefore if more than half of the members of $s_{A(0)}$ equal one, then $e_0 = 1$. Otherwise, $e_0 = 0$.

Then, the corrected sequence \hat{z} is calculated using the results of the two previous steps as shown in Equation (A.3).

$$\hat{z} = (\hat{r} \oplus \hat{e}) \quad (\text{A.3})$$

If $\hat{z}H^T = \vec{0}$, the Index field is $\{z_0, z_1, \dots, z_6\}$. Otherwise, the decoding failed, and the receiver should continue to process the next synchronization burst in order to obtain the Index field.

A.3 Example

Suppose that after coding, the sequence $\{1, 0, 0, 1, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1\}$, which is the coding example given in Table 23 in 6.7.2.1, was transmitted by the PD. Because the Hamming distance of the FEC code is five, up to two erroneously received bits can be tolerated. In the case of this example, bits i_3 and p_2 are erroneously received due to the noise of the wireless channel. The received sequence will then be $\{r_0, r_1, \dots, r_{14}\} = \{1, 0, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, 0, 1, 1\}$.

The following decoding procedure is written using the notation of the C programming language:

- 1) Calculate the syndrome sequence: $s(i)$, where $i = 0, 1 \dots 14$.


```
for (i = 0; i < 15; i++)
{
s(i) = (r[(0+i)%15] + r[(8+i)%15] + r[(12+i)%15] + r[(14+i)%15]) % 2;
}
For this example,  $\{s_0, s_1, \dots, s_{14}\} = \{1, 0, 0, 0, 1, 1, 0, 1, 1, 1, 0, 0, 0, 0, 0\}$ .
```
- 2) Calculate the error vector: $e(i)$, where $i = 0, 1 \dots 14$.


```
for (i = 0; i < 15; i++)
{
e(i) = ((s[(0+i)%15] + s[(1+i)%15] + s[(3+i)%15] + s[(7+i)%15]) > 2 : 1 : 0);
}
For this example,  $\{e_0, e_1, \dots, e_{14}\} = \{0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0\}$ .
```
- 3) Calculate the corrected sequence: $z(i)$, where $i = 0, 1 \dots 14$.


```
for (i = 0; i < 15; i++)
{
z(i) = (r(i) + e(i)) % 2;
}
For this example,  $\{z_0, z_1, \dots, z_{14}\} = \{1, 0, 0, 1, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1\}$ .
```
- 4) Check whether or not the index information is valid following the FEC decoding.


```
for (i = 0; i < 15; i++)
{
s(i) = (z[(0+i)%15] + z[(8+i)%15] + z[(12+i)%15] + z[(14+i)%15]) % 2;
```

```
}
```

For this example, $\{s_0, s_1, \dots, s_{14}\} = \{0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0\}$.

```
SuccessFlg = TRUE;
```

```
for (i = 0; i < 15; i++)
```

```
{
```

```
  if (s(i) == 1)
```

```
    SuccessFlg = FALSE;
```

```
}
```

For this example, (SuccessFlg = TRUE).

- 5) Finally, if (SuccessFlg == TRUE), then the first seven bits of $z(i)$ are the correct index value. Otherwise, the FEC decoding failed.

Annex B

(informative)

Recommended deployment in the United States

B.1 Introduction

Wireless beaconing devices represent an effective alternative for providing protection for Part 74 licensed secondary users of TV channels (e.g., wireless microphones used by TV broadcasters). Data contained in the beacon frame can provide various levels of information that may be decoded depending on the unlicensed device implementation.

A given unlicensed device may be such that it simply detects the presence of a beaconing device. Alternatively, the unlicensed device may choose to decode some or all of the information bits in the beacon frame, which may provide greater assurance of beacon authenticity, or indicate very specific information related to the area being protected, TV channels currently in use, or even TV channel sub-bands (200 kHz segments of a TV channel) that are currently occupied. The beaconing device can be an invaluable piece of equipment both for licensed incumbents (provides protection) and for operators of unlicensed devices (can be an aid for them in their determination of which TV channel is safe to use for whatever service the device is going to provide).

However, care should be taken in the use and deployment of the beaconing device, such that it does not violate Part 74 rules.

It is also important to ensure that the deployment is not such that there is harmful interference or other unwanted effects on the wireless system that is actually meant to be protected. Some of the considerations that may be taken into account in the deployment and operation of beaconing devices are physical proximity of the wireless system, relative to the beaconing device's location, and wireless system operating frequency, relative to the operating frequency of the beaconing device.

This annex provides some recommended practices aimed at ensuring that beaconing device operation will not induce harmful interference on other devices. The practices given here are intended for use when deploying systems in the United States. Each of the topics above is discussed in more detail in subsequent subclauses.

B.2 Physical positioning of the beacon relative to other Part 74 wireless systems

The FCC Part 74 Subpart H [B3] rules specify power limits for devices operating in the TV bands for both UHF and VHF frequencies. The rules stipulate the following:

- The power of the measured unmodulated carrier power at the output of the transmitter power amplifier (antenna input power) may not exceed 250 mW for devices in the frequency ranges 470–608 MHz and 614–806 MHz (i.e., UHF devices).
- The power of the measured unmodulated carrier power at the output of the transmitter power amplifier (antenna input power) may not exceed 50 mW for devices in the frequency ranges 54–72 MHz, 76–88 MHz, and 174–216 MHz (i.e., VHF devices).

It is likely that the wireless beaconing device will utilize the maximum possible transmit power. Maximum transmission power brings with it all of the usual advantages, such as increased detection probability at a given distance from the beaconing device.

However, most FCC Part 74 wireless microphones operate at well below the maximum transmission power allowable by the rules. In fact, FCC Docket Number 04-186 [B4] found that well over 90% of UHF systems made by various wireless microphone manufacturers operate with a transmit power in the 10–50 mW range.

It is important to realize a certain separation distance between the beaconing device and the wireless system receiver. The separation distance to observe depends on several factors, such as the height of the beaconing transmitter and wireless system receiver antennas above ground level (AGL). Figure B.1 illustrates a simplified view of the scenario.

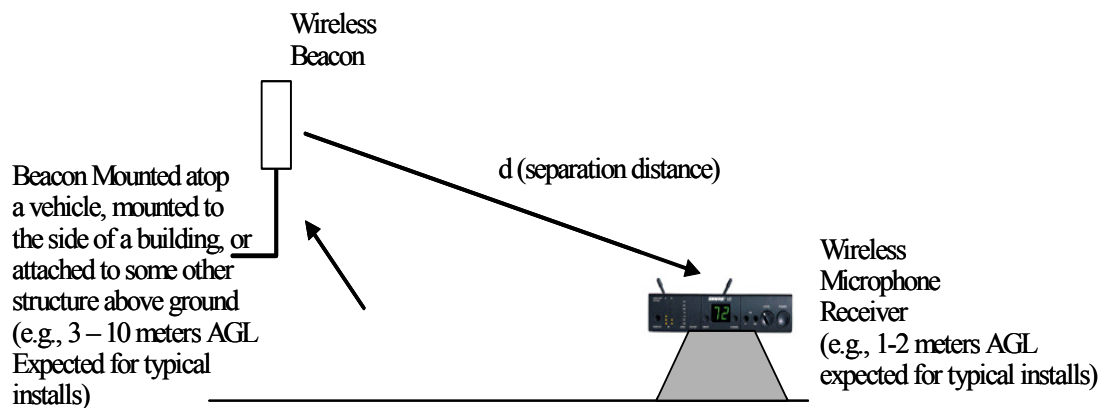


Figure B.1—Beacon/wireless system separation distance

An individual deploying a wireless microphone system with a beaconing device for protection needs to be aware of the approximate distance “d” between the microphone receiver and beacon antennas. A certain minimum separation distance should be maintained in order to ensure proper operation of the wireless microphone system. Some example values for UHF and VHF frequencies, utilizing cross-polarized and co-polarized antennas, respectively, are provided in Figure B.2, Figure B.3, Figure B.4, and Figure B.5 as guidance when setting up these devices. A center frequency of 617 MHz was chosen to calculate the required separation distances for UHF, and a center frequency of 174 MHz was utilized to calculate the required separation distances at VHF. At the lower edge of the UHF band, 470 MHz, approximately 30% more separation is required to protect the wireless microphone receiver, while at the upper limit of 698 MHz, approximately 15% less separation between the beaconing device and the wireless microphone receiver antennas is required.

Note that the distances indicated in Figure B.2, Figure B.3, Figure B.4, and Figure B.5 assume line of sight between the wireless beaconing device and wireless system receiver. This corresponds to the worst case. If there are obstructions in the path, or the installer is experienced and confident about the corresponding impact on signal propagation, these figures could be modified. An example of this might be locating the beaconing device outside on the side of a building or other structure, while locating the wireless system indoors and separated from the beaconing device by, for example, some building material that attenuates the transmitted signal.

Further, Figure B.2, Figure B.3, Figure B.4, and Figure B.5 are conservative, with the aim being that the numbers can be applicable to a wide variety of wireless products with different receiver performance characteristics. Shorter separation distances might prove acceptable, depending on the actual equipment used.

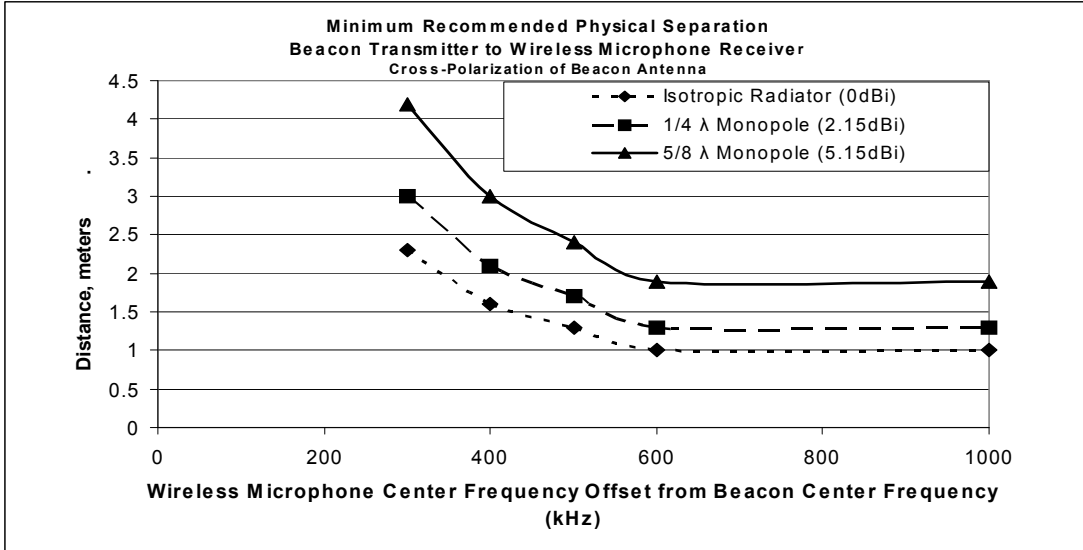


Figure B.2—Minimum recommended physical separation at 617 MHz using a cross-polarized beacon antenna




Figure B.3—Minimum recommended physical separation at 617 MHz using a co-polarized beacon antenna




Figure B.4—Minimum recommended physical separation at 174 MHz using a cross-polarized beacon antenna




Figure B.5—Minimum recommended physical separation at 174 MHz using a co-polarized beacon antenna

When selecting the physical location of the beaconing device, the beacon operator should be aware of the potential for interference to nearby DTV receivers tuned to a lower first adjacent channel at the edge of the DTV service contour. At the edge of the DTV service contour (41 dBu/−84 dBm), predicted interference may be experienced up to a range of 144 m.¹⁸ Many factors affect the radii of potential interference, including DTV receiver front-end design, antenna efficiency, directionality of the antenna, aiming of the

¹⁸Analysis based on two-ray analysis, a 3.5 meter beacon antenna height, a 9 meter TV receiver antenna height, 250 mW conducted beacon power, 2.15 dBi gain beacon antenna, 12dB polarization diversity, and a 40 dB D/U ratio for the adjacent channel. No shadowing has been taken into account; shadowing would reduce the potential interference range, as would other propagation path impairments.

DTV receiving antenna, and the nature of the path encountered between the beaconing device and the DTV receiver antenna. In general, these factors may reduce the region over which potential interference can occur. The potential interference range is, however, reduced as the signal strength of the DTV transmission increases. For example, at the 51 dBu /-74 dBm contour, the potential interference range is reduced to less than 10 m.

It should be noted that potential interference can occur to digital TV reception due to third-order intermodulation and cross-modulation products generated in a DTV receiver. These products are the result of the combination of the beacon signal with a DTV signal located on a channel different from the channel to which the TV is tuned. The effect is documented in the FCC Office of Engineering and Technology Report FCC/OET 07-TR-1003 [B9]. If F_B is the frequency used by the beacon transmitter and F_1 is the frequency of the first DTV signal, and F_D is the frequency of the desired channel tuned by the victim DTV receiver, the third-order intermodulation will occur at frequencies F_{IM3} given by Equation (B.4) and Equation (B.5).

$$F_{IM3} = 2 \times F_B - F_1 \quad (\text{B.4})$$

$$F_{IM3} = 2 \times F_1 - F_B \quad (\text{B.5})$$

Given the possible combinations of F_1 and F_B , there exist conditions where the third-order intermodulation product ($F_{IM3} = F_D$) will interfere with the desired DTV channel in the victim receiver.

If F_2 is the frequency of a second DTV signal, the cross-modulation products from the two DTV signals, F_1 and F_2 , plus the beacon signal, F_B , will occur at frequencies F_{XM} given by Equation (B.6) through Equation (B.8):

$$F_{XM} = F_1 \pm (F_2 - F_B) \quad (\text{B.6})$$

$$F_{XM} = F_B \pm (F_2 - F_1) \quad (\text{B.7})$$

$$F_{XM} = F_2 \pm (F_B - F_1) \quad (\text{B.8})$$

Similarly, given the possible combinations of F_1 , F_2 , and F_B , these exist conditions where the cross-modulation products ($F_{XM} = F_D$) will interfere with the desired DTV channel in the victim receiver.

Martin [B9] demonstrates the complexity of these interference scenarios. Based upon the results of the report, interference may occur if the beacon signal as the undesired signal causes the D/U ratio to go below -21 dB. Beacon users should ensure that the beacon signal does not cause interference to DTV signals if signal levels are sufficiently high enough to induce third-order intermodulation and cross-modulation in the victim DTV receiver.

B.3 Additional beacon antenna deployment considerations

The beacon is expected to be employed to protect operations within a radius of up to 4.5 km. This means that the wireless microphone receivers can be that much closer to the unlicensed interferer than the beaconing device. This impacts the required detection SNR in a negative way. This situation is further exacerbated by the fact that the beacons are expected to be allowed 50 mW for VHF operations and 250 mW for UHF operations, while protecting against 4 W unlicensed devices. Of the parameters that can be adjusted to reconcile these differences, the beacon antenna deployment height offers the greatest advantage with the least impact. Protection of larger radii can be enhanced by using a 9 m antenna height as opposed to a 3 m antenna height. See Table B.1 and Table B.2 for VHF and UHF results, respectively.¹⁹

¹⁹The IEEE P802.22 WRAN-B channel model was used in these calculations. See Kuffner [B7] for a full analysis and assumptions.

Table B.1—Required detection SNRs and expected PERs versus protection radii for VHF

Protected radius (m)	Antenna height (m)	SNR (dB)	MSF1 PER (%)
0	3	6.49	3.3
	9	15.16	< 0.5
500	3	4.83	5.3
	9	13.51	< 0.5
1500	3	1.94	10.6
	9	10.61	0.87
4500	3	-4.64	~ 53
	9	4.03	6.6

Table B.2—Required detection SNRs and expected PERs versus protection radii for UHF

Protected radius (m)	Antenna height (m)	SNR (dB)	MSF1 PER (%)
0	3	13.48	< 0.5
	9	22.15	< 0.5
500	3	10.71	0.83
	9	19.38	< 0.5
1500	3	6.22	3.5
	9	14.90	< 0.5
4500	3	-2.82	~ 33
	9	5.86	3.95

B.4 Recommended wireless system frequencies of operation

The use of a wireless beaconing device will mean that spectrum that was previously available for wireless system use will not be usable. In particular, the wireless beaconing device has been specified to operate in close proximity to the DTV pilot and is centered at 309.4 kHz from the lower edge of the TV channel in use. This number varies somewhat currently, in order to support certain co-existence scenarios with National Television System Committee (NTSC) television, but the value is expected to be fairly constant once the DTV transition is complete.

Operating the beaconing device on the pilot frequency means that other Part 74 wireless systems should avoid operation such that any part of their transmission will overlap the spectrum in this area. In addition, because of the high transmission power associated with the beacon signal, it will be necessary for the wireless system to avoid operation in regions of the TV spectrum that are somewhat close to the beaconing device operating spectrum, as specified in B.2.

Annex C

(informative)

Receiver sensitivity and adjacent channel protection justification

C.1 Introduction

This annex explains the justifications for the receiver sensitivity and adjacent channel protection parameters. The analysis here applies to two IEEE 802.22.1-compliant devices attempting to aggregate. This does not apply to the protection of the IEEE 802.22.1-compliant device from IEEE P802.22-generated interference. In that analysis (Kuffner [B7]), the calculation of link margin is based on the premise that the bulk path loss between the IEEE 802.22.1 transmitter and the IEEE P802.22 sensing receiver is approximately equal to the bulk path loss between the IEEE P802.22 transmitter and the protected device receiver (assuming the IEEE 802.22.1 transmitter is deployed within reasonable proximity to the protected device receiver), with differences in overall path loss being due to differences in antenna heights and frequency selective fading.

The sensitivity analysis begins with the definition of the noise bandwidth. The noise bandwidth of the IEEE 802.22.1-compliant receiver is assumed to be that of the chip matched filter, which is the bandwidth of the spread spectrum signal, $9609.1 \text{ kHz} \times 8x \text{ spreading} = 76.8728 \text{ kHz}$, usually rounded to 77 kHz for calculations. For a nominal IEEE 802.22.1-compliant receiver noise figure around 10 dB, the receiver noise would be as shown in Equation (C.1).

$$-174 \text{ dBm/Hz} + 10 \text{ dB}_{\text{NF}} + 10 \log(77 \text{ kHz}) = -115.1 \text{ dBm} \tag{C.1}$$

In the context of the IEEE 802.22.1-compliant receiver, the adjacent channel is a frequency relative to the IEEE 802.22.1 center frequency, which is 309.4 kHz from the lower edge of the TV channel (see Table 1 6.1.2). The adjacent channel is at 1.075 MHz, as indicated in Annex B. It is noted that this is inside the TV channel and hence co-channel to a TV signal that would be operating on the channel, but it is adjacent channel to the beacon signal. See Figure C.1.

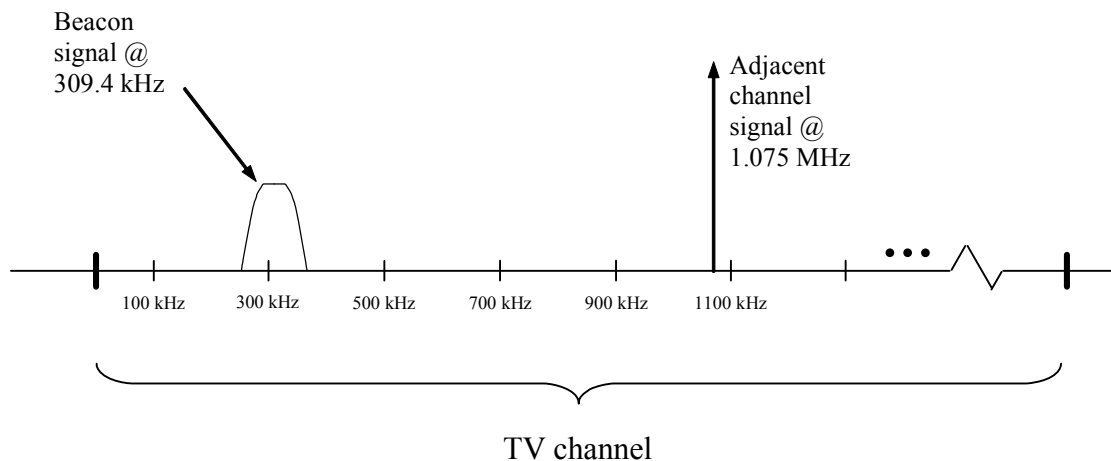


Figure C.1—Graphic definition of IEEE 802.22.1 adjacent channel, which falls co-channel to the would-be TV signal

C.2 Analysis

According to Kuffner [B7], the required E_c/N_0 value for 1% FER in a Gaussian channel is around 4.5 dB (corresponds to a symbol E_s/N_0 of 13.5 dB) without forward error correction. This would put the 1% FER sensitivity in a Gaussian channel at about -110 dBm for -115 dBm receiver noise. An IEEE 802.22.1-compliant receiver with this noise bandwidth and perfect implementation would be expected to demonstrate this sensitivity. Allowing 3 dB of implementation loss due to limited oversampling, frequency offsets, imperfect matched filtering, and the like, the sensitivity could be on the order of -107 dBm. This is the capability of the receiver. A separate question is what is the required sensitivity?

The following analysis considers an example link for determination of both the required sensitivity and the adjacent channel rejection. There are two links to consider in this analysis. See Figure C.2.




Figure C.2—Scenario for link analysis

Link 1 is the desired link between two beacon devices (ENG trucks). This is taken to be 1.5 km maximum. *Link 2* is the interference link between the nearby wireless microphone transmitter and the beacon receiver.

C.2.1 Link 1 Rx signal

A two-ray model is used for the propagation on *Link 1*. The assumptions are shown in Table C.3.

Table C.3—Link 1 assumptions

h_{Tx}	3.5 m
h_{Rx}	3.5 m
P_{Tx}	0.25 W
G_{ant}	7 dBi

The equation used for the two-ray E-field is as follows:

$$E(R, f) = \sqrt{30P_{Tx}} \left(\frac{e^{-j2\pi d_1(R)/\lambda(f)}}{d_1(R)} - \frac{e^{-j2\pi d_2(R)/\lambda(f)}}{d_2(R)} \right) \quad (C.2)$$

where

$$d_1(R) = \sqrt{(h_{Tx} - h_{Rx})^2 + R^2}$$

$$d_2(R) = \sqrt{(h_{Tx} + h_{Rx})^2 + R^2}$$

$$\lambda(f) = 300/f, f \text{ in MHz}$$

The E-field at 1.5 km for 692 MHz, and the assumed parameters are as follows:

$$E(1.5 \text{ km}, 692 \text{ MHz}) = 964 \mu\text{V/m} (59.7 \text{ dBu}) \tag{C.3}$$

$$P_{Rx} = \frac{E^2}{377} \times \frac{\lambda^2}{4\pi} 10^{G_{ant}/10} = -67.3 \text{ dBm} \tag{C.4}$$

Additional path loss assumptions are shown in Table C.4.

Table C.4—Additional path losses

Shadowing	20 dB
Fading (additional rays) ^a	12.5 dB

^aSee, for example, [B8], which shows about 12.5 dB difference in 1% FER Ec/No for Gaussian and WRAN B channels.

This would put the minimum signal at the receiver for 1.5 km range at -99.8 dBm, which is rounded to -100 dBm. This is the sensitivity level (without FEC) quoted in 6.8.6. Simulations of the rate-1/2 convolutional code seemed to indicate about 7 dB of gain, which results in the difference between sensitivity levels specified for MSF 1 and both MSF 2 and MSF 3.

At the other end of the UHF band (470 MHz), the E-field in *Link 1* is actually smaller (656 μV/m vs. 964 μV/m) due to the two-ray model, but this is exactly compensated by the effective aperture of the antenna, putting the nominal (non-faded or shadowed, as in Table C.4) received power at the same level of -67.3 dBm.

For a potentially much shorter 200 m *Link 1* range, the same assumptions lead to a signal strength of -33.5 dBm due to the two-ray model, and 66 dBm with additional shadowing and fading. This interoperation range assumption would relax the IEEE 802.22.1-compliant receiver sensitivity by some 34 dB and greatly simplify adjacent channel selectivity design.

C.2.2 Microphone interfering signal

The assumptions for the interference link (*Link 2*) are shown in Table C.5.

Table C.5—Microphone parameters

EIRP	+10 dBm
h_{Tx}	1.5 m
Emissions	ETSI [B2]
Elevation pattern loss	0.25 dB per end ^a
Polarization loss	3 dB

^aAssumes $G(\theta) = \cos(\frac{\pi}{2} \cos \theta) / (\sin \theta)$ elevation patterns.

The interference signal level due to the microphone at the beacon receiver at 692 MHz is as follows in Equation (C.5):

$$10 \text{ dBm} - 20\log[0.433 \text{ m}/(4\pi \times 10.2 \text{ m})] + 7 \text{ dBi} - 0.25 \text{ dB} - 3 \text{ dB} = -35.9 \text{ dBm} \quad (\text{C.5})$$

This microphone signal is assumed to be located 1.075 MHz above the lower edge of the TV channel, as shown in Figure C.1. With the beacon centered at 309.4 kHz, the microphone signal is about 765 kHz above the beacon receiver center frequency. The power due to the microphone is 3.4 dB larger than Equation (C.5) at -32.5 dBm at 470 MHz using the simple square law model. However, this was reasoned to remain at -36 dBm since the efficiency of the microphone antenna would presumably be poorer at the longer wavelength at the low end of the UHF band. This results in the same adjacent channel rejection numbers for the low and high ends of the band.

In the presence of a large test interferer, convention allows 3 dB of sensitivity degradation due to receiver impairments. With a nominal sensitivity of -100 dBm, this puts the sensitivity in the presence of the -36 dBm interfering signal at -97 dBm without coding (MSF 2, MSF 3) or -104 dBm (MSF 1) with coding.

Note that the out-of-band emissions of the microphone should also be considered. While the sensitivity test would ordinarily use a CW signal, in practice the emissions could be according to the mask in ETSI [B2]. The emissions of the microphone should be -87 dBr/kHz peak/hold around 765 kHz offset according to this reference. With the microphone signal at -36 dBm at the beacon receiver, the out-of-band emissions are at most

$$-36 \text{ dBm} - 87 \text{ dBr/kHz} + 10\log 77 = -104 \text{ dBm} \quad (\text{C.6})$$

at the beacon receiver assuming discrete line spectrum, or more like -114 dBm average if noise-like spectrum is present at this offset, due to the peak/hold nature of the specification and assuming approximately 10 dB peak/average per 1 kHz bin. Thus the out-of-band emissions alone could result in a noise floor at the beacon receiver of -104 dBm, which will be further degraded by the selectivity of the beacon receiver in response to the microphone signal. The 1% FER sensitivity will nominally be about 4–5 dB above the combined out-of-band emissions from the microphone that fall co-channel to the beacon receiver plus the sensitivity degradation due to finite selectivity in the beacon receiver (Kuffner [B7]).

Annex D

(informative)

Next higher layer (NHL) operation

D.1 Introduction

The IEEE 802.22.1 WG was tasked with the development of the protocols, data formats, and other technical details for communication devices used to protect low-power, licensed devices operating in the television broadcast bands from new license-exempt devices that would operate in the same bands. More specifically, this standard details operations of a new class of wireless “beaconing” devices aimed at protecting existing incumbent TV band users (such as licensed wireless microphone operators) from harmful interference that could be generated by new license-exempt devices.

IEEE Std 802.22.1 includes specification of only the MAC and PHY operations of the wireless beaconing device. However, with the ongoing development of the standard, it has become clear that the beaconing device’s operation would be better understood if additional information regarding the functionality of the “next higher” layer (NHL), which would interact with/utilize services of the MAC layer, was provided. Per references made in this draft, the next higher layers would provide services such as selecting a channel, deciding on an operating mode (i.e., PPD or SPD), starting and stopping beacon transmissions, processing incoming beacon frame information, aggregating data, and handling security errors.

This annex provides recommendations for how the NHL functionality should be implemented, e.g., approaches to populating the beacon, and how the various PDs (PPDs, SPDs, and NPDs) should logically behave under some likely operating scenarios.

D.2 Next higher layer (NHL)

Table D.3 illustrates a simplified view of the likely relationship between the MAC sublayer and PHY layer, defined in Clause 6 and Clause 7, respectively, and a “Next Higher Layer.” Note that in general, NHL refers to the layer that would interact directly with the MAC layer; an IEEE 802.22.1 implementation might include also additional “upper layers” that would provide other functions and features.

The MAC sublayer provides an interface between the NHL, PHY, and MLME called the MLME-SAP. The MLME-SAP in turn provides a means of passing information between the NHL and the MAC sublayer via service request, confirm, indication, and response primitives defined in Clause 7.




Figure D.3—NHL relationship to PHY/MAC

D.3 Beaconing device types

This standard identifies several types of Protecting Devices, described as follows:

- *PD (Protecting Device)*: A generic term used to refer to any or all of the devices described next.
- *PPD (Primary Protecting Device)*: The PPD is the main device responsible for providing incumbent protection. The IEEE 802.22.1 draft specification requires that it transmit beacon data at least every other superframe (other transmission opportunities can be dynamically allocated for “inter-device” communication). The protection information within the PPD beacon transmissions may or may not include information aggregated as a result of inter device communication with other PDs.
- *SPD (Secondary Protecting Device)*: An SPD is a PD that has chosen to have another PD (specifically, a PPD) provide protection on its behalf. The protection information is shared with the PPD via inter-device communication, and subsequently broadcasted as part of the PPD’s regular beacon transmissions.
- *NPD (Next-In-Line Protecting Device)*: An NPD is an SPD that has been selected by a PPD to become the new PPD in the event that the current PPD ceases beacon transmission.

The decision of whether to become a particular type of PD is controlled entirely by the NHL, and is dependent on a number of factors. Many of the NHL functions are, in turn, largely a function of the type of PD involved.

In the subclauses that follow, how the NHL could operate is examined focusing on, in particular, the behaviors one would expect to see from the different types of PDs.

D.4 Generic PD behavior

The following subclauses describe some example NHL behaviors that would be performed regardless of the type of PD. One example is a PD’s behavior prior to becoming a PPD, SPD, or NPD.

D.4.1 Searching for existing PDs

Subclause 7.1.1 defines MLME-SEARCH request and confirm primitives, which enable the NHL to initiate scanning of a TV channel in order to determine whether or not there are already PDs operating on the channel. The MLME-SEARCH.request primitive has two parameters, a “ChannelList” bitmap and a “Duration” parameter.

The ChannelList parameter indicates which TV channels should be scanned. Although the 69-bit field supports requesting a scan of many different channels, the current version of this draft standard specifies that the ChannelList parameter should be limited to one TV channel. The TV channel that should be searched is assumed to be known by the operator of the beacon, and could be entered or set by various means. The TV channel used should be one where devices requiring protection are planned to operate.

The Duration parameter indicates how long to search the specified channel as a number of superframe periods, and it is defined such that the minimum time is 5 superframes (Duration = 0) and the maximum time 50 superframes (Duration = 45). It is expected that a new PD device should be able to gather information about its surroundings within 5 superframe times (a single superframe is 103.24 ms). However, in order to avoid conflicts between multiple new PDs, it is recommended that the Duration parameter be set to a random value between 0 and 45. This will decrease the likelihood that these devices will complete search procedures at the same time, which could lead to various problems or performance issues (e.g., increased contention during the Rx period).

The MLME-SEARCH.confirm primitive indicates to the NHL that the scan operation requested has been completed. This primitive could be used to trigger initiation of a decision process in the PD, based on data it will have collected as a result of the search.

D.4.2 Deciding to become a PPD or an SPD after an initial search

A new PD device performs the search procedure described in the previous subclause upon initialization. Specifically, the MLME_SEARCH.request primitive will trigger the possible delivery of many MLME-INCOMING-BEACON.indication primitives from the MLME to the NHL. These primitives originated from one or more existing devices (PPDs, SPDs, or NPDs) operating in the vicinity of the new PD, on the same TV channel.

Following completion of the search procedure (as indicated by receipt of the MLME_SEARCH.confirm primitive at the NHL), the NHL of the new PD should determine if the PD will become an SPD or PPD. The decision should logically be made based on information gathered from the MLME-INCOMING-BEACON.indications received. Note that

- The new PD can distinguish between PPDs and the other types of existing PDs via the device Rank subfield within the Parameter 1 field of each MLME-INCOMING-BEACON.indication primitive it receives. PPDs regularly transmit beacon frames, hence, they can offer the best potential protection for a PD if it determines it wants to become an SPD.
- The new PD can associate MLME-INCOMING-BEACON.indication primitives it receives with particular devices using the 48-bit address found within the indication. This means the new PD can e.g., know how many indications it has received from each PD within its range.

As far as criteria for aiding in the decision

- It is reasonable to assume that other PDs in close proximity to the new PD would potentially represent good candidates to act as a PPD for the new PD. The PD can determine the relative location of each PD from which it receives an MLME-INCOMING-BEACON.indication primitive via the Location parameter.
- Each MLME-INCOMING-BEACON.indication primitive received from the MLME contains a Link Quality Indication (LQI), a representation of the strength and/or quality of the received beacon frame. If multiple MLME-INCOMING-BEACON.indication primitives are received from the same PD, these measurements can be averaged or processed in other ways to get better confidence in the reliability of the measurement.
- Other criteria, such as antenna height and location of the protected receiver equipment are also indicated in the MLME-INCOMING-BEACON.indication primitive and can play a role in the decision-making.

The data just mentioned can be used to develop various schemes for determining whether an “adequate” PPD already exists (the PD can become an SPD). The data can be used to determine some specific details about the PPD’s site, which can then be compared with the potential new PD’s configuration. The results could perhaps be mapped to a simple metric or score that can drive the decision.

It is recommended that a device should employ some mechanism similar to this that allows it to intelligently choose a mode of operation instead of strictly defaulting to operate as a PPD, since operating multiple PPDs in the same vicinity could reduce performance. The random setting of the Duration parameter in the MLME_SEARCH.request primitive actually aims to reduce the likelihood of multiple new devices becoming PPDs at the same time.

If the device does not receive any beacons, it should become a PPD and follow device initialization procedures defined in 7.4.4.

D.4.3 Beacon frame construction and transmission

The NHL can modify the data within a beacon transmission, as well as the transmission mode of the beacon, through the MLME-START-BEACON.request primitive. The meaning of the primitive parameters are documented in 7.1.1.13. The meaning of the corresponding fields and subfields within the beacon frame are well documented in 7.2, and the format varies somewhat depending on the type of PD involved. Table D.6 looks at each primitive parameter and provides guidance on how to set appropriate values in the case of a beacon transmission. It is assumed that a user interface or similar is implemented in the PD to enable the end user to configure the values appropriately.

Table D.6—Recommendations for parameters, fields, and subfields

Parameter, field, or subfield	Recommended value
Parameter 1 Frame Version Number	The NHL should set this value to 0 for this version of the standard.
Parameter 1 Priority Level	The NHL should set this value to 0, unless a hierarchy among beaconing devices has been established. This value can be used by the NHL to arbitrate between them, through assigning each an appropriate priority level.
Parameter 1 Antenna Height	The NHL should set the value to FALSE if all of the devices being protected are 10 m or less AGL. ^a The value is otherwise set to TRUE.
Parameter 1 Rank	The NHL would set this value to TRUE if the PD is a PPD, and to FALSE otherwise.
Source Address	The value is device dependent. Each beacon should have a unique, 48-bit address that can be read and stored in this field by the NHL.
Location	The NHL should obtain the value via geolocation or other means, as described in Annex E.
Parameter 2 Channel Width	The NHL should set the value based on user input.
Parameter 2 Cease Tx	This value could be linked with on/off operation such that it is set to TRUE for some amount of time as part of the normal shutdown sequence of the device.
Parameter 2 Time Parity	The NHL should obtain the data via UTC or other means, as describing in Annex E.
Parameter 2 Keep Out Zone	The NHL should set the value based on the expected operating range of the devices to be protected. The setting likely comes from a user who is familiar with the needs of the device being protected. If there is no idea about the required protection radius, the value should be set to the maximum value of 1 (indicating a protection radius of 4.5 km).
Parameter 2 Sub-group Channels	The NHL should set the value based on user input.
Parameter 2 NPD Indication (PPD only)	The NHL should set the value according to current status or needs of the PPD. In general, it is desirable for the PPD to select an NPD, hence the NPD Indication setting should reflect that state anytime a PPD has no associated NPD. The PPD NHL will know that its selection of an NPD was completed when it receives the MLME-NPD.confirm primitive, and can indicate that state with a different NPD Indication setting.

Table D.6—Recommendations for parameters, fields, and subfields (continued)

Parameter, field, or subfield	Recommended value
Parameter 2 NPD/NST subfields (SPD only)	The NHL of an SPD should set the NPD subfield based on whether or not the PPD has selected the SPD as the NPD or not. The SPD knows it has been selected based on the NPD Indication subfield received in a PPD beacon frame transmission following transmission of one of the SPD’s beacon frames.
Parameter 3 Indoor/Outdoor	The NHL should set the value according to the protected device’s receiver antenna location, as specified by e.g., the end user.
Parameter 3 Required Need Timer	The NHL should set the value according to the amount of time protection is expected to be needed, as specified by e.g., the end user. The NHL logic might include a capability to automatically extend the time, should 6 bits (up to 63 hours of protection) not be enough for a given use scenario.
Map Field Channel Map option	The NHL should set these values based on e.g. end user input. Note that this setting can be useful for large planned events where every slice of spectrum is needed; in those cases it would be undesirable to deploy a beacon on each individual channel requiring protection.
Map Field LAS channel Map option	The NHL should set these values based on e.g., end user input (knowledge of which frequencies the protected devices are operating on within the protected TV channel).
Map Field MSI option	The NHL should set this value based on any manufacturer-specific information needing to be sent.
Channel	The NHL sets the value to the TV channel on which the beacon frame will be transmitted.
Start	The NHL sets the value to TRUE if it wants to start transmitting a beacon frame(s) and FALSE if it wants to stop transmitting.

^aSee E.3 for a recommendation for obtaining altitude information.

If the beaconing device is a PPD, the beacon data transmitted may include a representation of the protection requirements of SPDs that are under the protection of the PPD (aggregated protection requirements). This is described in more detail in D.5.2.

Note that in the case of a PPD, the NHL only needs to send a new MLME-START-BEACON.request primitive to the MLME if one or more of the parameter values change. An SPD is not allowed to transmit more than one frame in a row.

D.4.4 Beacon transmission failure-related enumeration values

Table 37 in Clause 7 defines a number of failure related enumeration values, which are conveyed to the upper layer when an irregular event has happened during an attempt to transmit beacon data. In particular, the NHL behavior for the three following enumerations is considered, in the scenario where the values appear in the security status parameter of the MLME-INCOMING-BEACON.indication primitive:

SIGNATURE_INVALID: Subclause 7.1.2 explains this enumeration as meaning “The signature contained in the received beacon frame is invalid.” In general, an invalid signature could mean that the transmitting device is illegitimate and should not be trusted. Therefore, it is recommended that the receiving beaconing device NHL ignore the content of such transmissions, aside from perhaps logging the event (including the location, time, address, and security credentials received of the potentially rogue device) and reporting it to the trusted authority responsible for managing the beacon security credentials.

CERTIFICATE_INVALID: Subclause 7.1.2 explains this enumeration as meaning “The certificate contained in the received beacon frame is invalid.” A common reason for this would be expiration of the certificate, or perhaps the certificate has been revoked. The NHL should log the event (relevant data including the location, time, address, and security credentials received of the potentially rogue device), and the beacon operator should report it to the trusted certification authority responsible for managing the certificates.

For the two enumeration values just described, care should be taken in the design of the NHL in order to minimize the chance for denial of service attacks by devices transmitting such incorrect security information.

SIGNATURE_NOT_CHECKED: Subclause 7.1.2 explains this enumeration as meaning “The signature contained in the received beacon frame was not checked.” The MAC MIB attribute *macSignatureCheckEnabled* can be used to enable or disable checking. Hence, the NHL probably does not need to take any action, since it can be assumed that the beacon user does not require the signature. The receiver of the beacon data may also resolve to trust the source of the beacon through means outside of the standard (it may chose to use its own security procedure, and would not need the signature scheme described in 7.5.4).

D.5 PPD behaviors

The following subclauses describe some example NHL behaviors of a PPD.

D.5.1 Insertion of Rx period and ANP

Subclause 5.3 specifies that, under control of an upper layer, a receive period (Rx period), and an acknowledgement/non-acknowledgement period (ANP) may be included in the superframe. The resulting interval is called an inter-device communication interval (ICI) (6.5).

The NHL should ensure that the ICI is managed as follows:

- During initial transmissions by a PPD, the PPD should not insert ICIs during the first *aInitialization-Period* superframes. This prevents any transmission by SPDs during that time.
- Subsequent transmissions by the PPD, should include the ICI, enabling SPDs to aggregate with the PPD.

D.5.2 Data aggregation

There are a couple of considerations with respect to data aggregation: whether the PPD should aggregate beacon data that it receives from SPDs, and how the aggregation should be done.

When a PPD’s NHL receives MLME-INCOMING-BEACON.indication primitives from SPDs, it should immediately aggregate the incoming beacon data with its own data. This will prevent unnecessary transmissions over the air by multiple PPDs, since in normal operation, a PPD will transmit much more often than an SPD. Some exceptions where aggregation would not occur include MLME-INCOMING-BEACON.indication primitives with Security Status enumerations, such as SIGNATURE_INVALID or CERTIFICATE_INVALID.

The actual aggregation of data from an SPD should be done such that both the PPD and the SPD are adequately protecting their associated devices. This means consolidating the same information from one or more sources (PDs) into to one representation within the PPD’s beacon transmission. A PPD’s beacon transmissions should at any time accurately reflect the protection required by it and associated SPDs. This

means that the information should neither fall short or exceed what is required, and that the PPD will need to update the information when the protection required changes. Such updates may occur if, for example, SPDs are added or removed from the list of PDs the PPD protects, or if the protection needed by a given SPD changes. The updating of the aggregation information should be performed independently for each SPD, and should be reflected in subsequent PPD beacon transmissions as soon as possible. Table D.7 describes some simple rules for how data in the beacon transmission of the PPD should be aggregated with that of associated SPDs.

Table D.7—Rules for data aggregation

Parameter, field, or subfield	Recommended value
Parameter 1 Frame Version Number	The NHL should set this value to 0 for this version of the standard.
Parameter 1 Priority Level	The NHL should set this value to the highest value among the PPD and SPD(s) being protected.
Parameter 1 Antenna Height	The NHL should set the value to FALSE if all of the devices being protected by the PPD and SPD(s) are 10 m or less AGL. ^a The value is otherwise set to TRUE.
Parameter 1 Rank	The NHL should set the value to TRUE to indicate a PPD.
Source Address	The NHL should set the value to the address of the PPD.
Location	The NHL should set the value to that of the PPD.
Parameter 2 Channel Width	The NHL should set the value to that corresponding to the PPD. It is not expected that beaconing devices protecting TV channels of different bandwidths will be used in the same area. Should this occur, an acceptable approach would be to not aggregate.
Parameter 2 Cease Tx	The NHL should set this field based on events at the PPD only.
Parameter 2 Time Parity	The NHL should set the value based on PPD timing information only.
Parameter 2 Keep Out Zone	The NHL should set the value to correspond to the largest radius required, among the set of PPD and SPD(s) involved. An alternative approach would be to refine the number based on location information from each PD, collected at the PPD. Together with the PD keep out zone parameter settings from each PD, the PPD could determine the keep out zone with some accuracy.
Parameter 2 Sub-group Channels	Assuming that the TV channel of operation falls within a UHF sub-group (7.2.1.4), the NHL should set the value such that all TV channels requiring protection within that sub-group receive protection.
Parameter 2 NPD Indication (PPD only)	The NHL should set the value according to the current status or needs of the PPD (there is no change due to aggregation).
Parameter 3 Indoor/Outdoor	The NHL should set the value to outdoor (FALSE) if either the PPD or any of the SPDs has a receiver located outside.
Parameter 3 Required Need Timer	The NHL should set the value equal to the amount of time protection is expected to be needed by the PD that is expected to stay active the longest.

Table D.7—Rules for data aggregation (continued)

Parameter, field, or subfield	Recommended value
Map Field - Channel Map option	The NHL should set these values to be the superset of all channels requiring protection by the PPD and SPD(s). If more than five distinct channels appear in the channel lists of the PPD and SPD(s), any five of them can be arbitrarily inserted into the list transmitted by the PPD. Note that it is unnecessary to include the TV channel on which the beacon frames are sent.
Map Field - LAS channel Map option	The NHL should set the bitmap values such that they represent the logical ORing of the bitmaps from the PPD and SPD(s) involved, i.e., the resulting bitmap represents a superset of all PPD and SPD LAS channels that require protection.
Map Field - MSI option	Manufacturer specific.
Channel	The NHL should set the value to the PPD’s channel of operation.
Start	The NHL should set the value based on PPD information only.

^aSee E.3 for a recommendation for obtaining altitude information.

D.5.3 NPD selection

If a PPD is protecting one or more SPDs, it should select one of them as the NPD. Selecting an NPD is desirable because it can shorten the amount of time that the system is operated without protection in the event that the PPD stops transmitting. The NPD can take over immediately if it is given prior warning by the PPD (via transmission of the Cease Tx subfield). Even when the PPD transmission stops irregularly, because there is only one NPD associated with a PPD, there would rarely be issues with determining which SPD should become the new PPD. This means contention delays are eliminated and the overall level of protection for devices is improved.

In order to select an appropriate NPD from the set of SPDs that are aggregated with it, the PPD NHL should determine characteristics of those SPDs. It can collect information on the SPDs in the following two ways:

- Monitor the regular beacon transmissions of the SPDs, which occur at least every *macActivePeriod-SPD* superframes.
- Set the NPD Indication subfield to 00, which indicates to all receiving SPDs that “No NPD currently exists in the beaconing network. All SPDs receiving this beacon frame shall volunteer to be the new NPD,” which will trigger transmissions from all of them.

The PPD captures the incoming beacon transmissions in MLME-INCOMING-BEACON.indication primitives from candidate SPDs. One approach is for the PPD to select one of these SPDs to be the NPD just after receiving its beacon transmission. The selection criteria for choosing an NPD should be driven by the SPD’s ability to provide a level of protection comparable to what the PPD can offer. This implies, for example, that the potential NPD should be selected based on a metric or score similar to that used when an SPD selects a PPD following an initial search (D.4.2). The NPD communicates the choice to the MLME through the MLME-START-BEACON.request primitive.

Alternatively, the PPD may choose to collect information on all of the SPDs before selecting an NPD. This process is begun by capturing incoming beacon transmissions via the MLME-INCOMING-BEACON.indication primitives as in the first approach. The PPD should then set the NPD Indication subfield to 11, indicating that no NPD currently exists and that one is not being requested. This marks the end of the first round of data collection. The PPD should then set the NPD Indication subfield to 00 in order to get all the SPDs to send beacon frames again. Once the NHL of the PPD receives a beacon frame from the

desired SPD (i.e., the SPD is selected based on pre-defined metric), it will issue the MLME-START-BEACON.request primitive to the MLME with the NPD Indication subfield set to 01.

Note that a PPD can “de-select” an NPD by setting the NPD Indication subfield to either 00 or 11 at any time (7.2.1.4). Both of these values indicate to all receiving SPDs that “No NPD currently exists in the beaconing network.” The de-selected NPD recognizes this and subsequently ceases transmission of NPD codewords. It should also no longer set the NPD subfield to TRUE in its beacon transmissions.

D.5.4 SPD lost indication

According to 7.1.1.12, the MLME-SPD-LOST.indication primitive is generated by the MLME of the PPD and issued to its NHL as a notification that the beacon frames from a particular SPD were not heard in the last *macMissedSPDBeacons* superframes.

If the NHL of the PPD receives this indication, it is recommended that the PPD remove all data corresponding to that SPD from its beacon frames. The PPD should ensure that the data is updated such that the rules defined in D.5.2 are followed for the remaining PPD and SPD data.

D.5.5 NPD lost indication

According to 7.1.1.7, the MLME-NPD-LOST.indication primitive is generated by the MLME of a PD other than the NPD, and issued to its NHL as a notification that it has not received either an NPD beacon frame or an NPD code within the last $aMaxMissedNPDCodes \times aNPDPeriod$ superframes.

If the NHL of a PPD receives this indication, it is recommended that the PPD once again execute the procedure described in D.5.3, in order to determine an adequate NPD. Also, if the PPD has not already, it should remove all data corresponding to that NPD from its beacon frames. The PPD should ensure that “de-aggregation” of the NPD data is done such that the rules defined in D.5.2 are followed for the remaining PPD and SPD data.

D.5.6 Planned termination of PPD transmission

Planned termination of PPD beacon transmissions occurs when the PPD stops transmitting in an orderly way, e.g. the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the PPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive with Cease Tx subfield set to TRUE to its MLME. If desired, the turn off sequence of the PPD can be delayed slightly to allow it the possibility of transmitting several such superframes, to improve the chance that SPD will detect the event and have the opportunity to react to it.

D.5.7 Behavior following termination of SPD transmissions without warning

Termination of SPD beacon transmissions without warning refers to instances where the SPD stops transmitting and the PPD does not receive the Cease Tx subfield set to TRUE as part of an MLME-INCOMING-BEACON.indication primitive from that SPD. Examples of such events include someone “pulling the plug” or sudden power outages.

This event will be indicated to the NHL of the PPD via the MLME-SPD-LOST.indication primitive. The corresponding behavior is described in D.5.4.

D.5.8 Behavior following planned termination of SPD transmissions

Planned termination of SPD transmissions occurs when the SPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the SPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive with the Cease Tx subfield set to TRUE to its MLME. If desired, the turn off sequence of the SPD can be delayed slightly to allow it the possibility of transmitting several such superframes, to improve the chance that the PPD will detect the event and have the opportunity to react to it. Note that because there is no assurance that the PPD will grant the SPD even one opportunity to transmit a beacon frame, the approach is strictly “best effort.”

The PPD’s reaction to this event is the same as if the event happened in an unplanned way. The difference is that it can react faster, i.e., it does not have to wait for the MLME-SPD-LOST.indication primitive.

D.5.9 Behavior following termination of NPD transmissions without warning

Termination of NPD beacon transmission without warning refers to instances where the NPD stops transmitting and the PPD does not receive the Cease Tx subfield equal to TRUE as part of an MLME-INCOMING-BEACON.indication primitive from it. Examples of such events include someone “pulling the plug” or sudden power outages.

This event will be indicated to the NHL of the PPD via the MLME-NPD-LOST.indication primitive. The corresponding behavior is described in D.5.5.

D.5.10 Behavior following planned termination of NPD transmissions

Planned termination of NPD beacon transmission occurs when the NPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the PPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive with the Cease Tx subfield set to TRUE to its MLME. If desired, the turn off sequence of the PPD can be delayed slightly to allow the SPD the possibility of transmitting several such superframes, to improve the chance that SPD will detect the event and have the opportunity to react to it.

The PPD’s reaction to this event is the same as if the event happened in an unplanned way. The difference is that it can react faster, i.e., it does not have to wait for the MLME-NPD-LOST.indication primitive.

D.5.11 PPD-specific MIB attributes

The MLME-SET.request primitive can be used by the PPD NHL to modify the value of the MIB attributes. Else, the attributes will revert to the default values, as specified in 7.3.2. For example, if a PPD decides to select an NPD, the PPD NHL should write the address of the NPD it selected into the MIB attribute *macNPDAddress*.

D.6 SPD behaviors

The following subclauses describe suggested NHL behaviors of an SPD.

D.6.1 PPD beacon reception

The NHL of the SPD should regularly monitor and store the PPD beacon data received via MLME-INCOMING-BEACON.indication primitives, such that the SPD can react to events that could compromise the protections it is providing.

- Source Address field/Rank subfield: The SPD NHL should use these fields to confirm that the received beacon data was from its PPD.
- Cease Tx subfield: The SPD NHL should use this field as a warning that the PPD is about to stop transmitting. SPD functionality related to this is described in D.6.13.
- Next-In-Line Protecting Device Indication subfield: The SPD NHL can use this field to determine whether there is an NPD in the system (or not), and if the PPD is actively seeking an NPD.
- SecurityStatus parameter: The SPD NHL can use this field to determine the status of security operation performed on the incoming beacon (if any).

The SPD is responsible for ensuring that its protection needs are being consistently met by the PPD that it has selected. This would include consistently monitoring information in the PPD MLME-INCOMING-BEACON.indication primitives, such as the antenna height and keep out zone data contained within the Parameter1 and Parameter2 parameters, respectively, and the Map and Channel parameters, all of which should have been aggregated to the PPD transmission to reflect the SPD's protection-related information. Note that this data can change dynamically as other SPDs enter and exit the system, but it is each SPD's responsibility to ensure the data continues to represent its protection needs.

D.6.2 Inter-device communication

SPDs may request to interrupt a PPD in order to transmit their own beacon data. Some examples of why the transmission would occur include the following:

- Informing the PPD that the PPD beacon transmissions should be updated to reflect the set of protections the SPD requires (i.e., requesting that the PPD aggregate the SPD's data).
- Informing the PPD periodically that a given SPD still requires protection (else the PPD can remove the SPDs information from its beacon transmissions).

Both of these transmissions are at the discretion of the NHL. It should generate MLME-START-BEACON request primitives to start the process, as per the inter-device communications procedure described in 7.4.5. The MLME-START-BEACON.request primitive will eventually trigger the SPD to randomly select and transmit an RTS codeword during a PPD Rx period. The SPD will know if it has been granted access, based on the subsequent response in the PPD's ANP.

Hence, the SPD NHL should be aware of the following:

- The PD's decision to aggregate and become an SPD; this occurs as a result of the procedures described in D.4.1 and D.4.2.
- The rate with which the SPD periodically updates the PPD, which is every *macActivePeriod* SPD superframes. The NHL can retrieve this information via the MLME-GET.request primitive.

D.6.3 Inter-device communication—retry failure handling

As mentioned previously, the SPD will know if it has been granted an opportunity to transmit its beacon data, based on the subsequent response in the PPD's ANP, which is indicated to the SPD's MLME via the PLME-INITIATE-RTS-BURST.confirm primitive. If that primitive indicates that the SPD was not granted an opportunity to transmit its beacon data, the SPD reverts to the retry procedure mechanism, as described in 7.4.2. The mechanism is designed to allow an SPD to try and get permission to send its beacon data up to

three more times after its initial attempt. If all three retry attempts fail, the SPD NHL will be notified of that event through a TX_FAILURE enumeration in the MLME-START_BEACON.confirm primitive from the MLME.

The SPD NHL could take various approaches to addressing the TX_FAILURE enumeration:

- The SPD could initiate inter-device communication again following a random delay. This will give the SPD four additional opportunities (one initial try and three more retries) to send its beacon data. This solution assumes that the reason for the previous TX_FAILURE was related to a high number of SPDs all wanting to transmit at the same time. The SPD could revert to becoming a PPD if it still cannot complete the communication.
- The SPD could choose to become a PPD and start its own beacon transmissions.

D.6.4 Inter-device communication—selection of the same RTS codeword

There is a small chance that inter-device communication could fail due to selection of the same RTS codeword by multiple SPDs. If the PPD correctly receives the codeword, it will send the corresponding ACK during the ANP, and the SPDs that selected the same codeword will all think that it is safe to transmit a beacon frame. However, the next beacon transmission from the PPD will likely not contain the correct protection data for one or more of the SPDs. The corresponding NHL(s) will realize something is wrong when the PPDs MLME-INCOMING-BEACON.indication primitive is read and found not to contain the relevant data.

To remedy this, NHL of each SPD should trigger another transmission of the beacon data. Although there is a good chance each SPD will send an RTS during the same PPD Rx period, it is unlikely that the SPDs will select the same RTS codewords again, since these are selected at random with each SPD transmission. Nevertheless, it can be a good idea to also implement a backoff mechanism, whereby on subsequent transmission attempts, the SPD waits a random number of superframes (e.g., from one to 16) before it attempts to transmit its beacon frame again.

D.6.5 PPD protection of an SPD

The SPD's protection needs should be adequately captured in PPD beacon transmissions via beacon data aggregation. Hence, if the SPD observes data in the PPD's beacon frame that does not completely represent its protection needs (e.g., it observes that its LAS channels requiring protection are not included in the PPD's beacon), the NHL of the SPD should take some action. Note that even upon initial selection of a PPD and transmission of SPD beacon data through the inter-device communications procedures specified in 7.4.5, the SPD's data should start to appear within the next three PPD beacon frame transmissions. If the SPD receives an MLME-INCOMING-BEACON.indication primitive from the PPD and its data is not present, the SPD can assume that it was not received. In such events, there are a number of actions that could be triggered by the SPD's NHL as follows:

- The SPD can trigger another transmission of its beacon data and wait again to see if it gets aggregated. The NHL may decide to allow several such transmissions, based on some criteria (e.g., a timer or counter), and attempt to promote itself to a PPD only if all such attempts fail.
- The SPD can promote itself to PPD immediately.
- The SPD can leverage MLME-INCOMING-BEACON.indication primitive data it has received from other PDs; e.g., if there were for some reason other PPDs in the area that can meet its protection needs, it could choose one of those as its new PPD.

D.6.6 Transmission of more than one beacon frame

If an SPD wants to send more than one beacon frame in a short amount of time, the SPD NHL can indicate this by setting the NST subfield to one in the first beacon frame to be transmitted. If the PPD wishes to do so, it can issue a special acknowledgement message called a Go-On in response to the SPD's transmission. This will allow the SPD to transmit its second frame without needing to issue another RTS. Transmission in this manner is useful if, for example, the SPD wants to transmit both TV channel information and LAS channel information, since an SPD can only send one or the other in any given beacon frame.

Note that if the SPD sent its first beacon frame during superframe n , the SPD NHL should ensure that it provides the MAC with the second beacon frame prior to superframe $n+2$.

See 6.5.2 and 7.4.5 for more information.

D.6.7 Behavior on SPD receiving a PPD lost indication

The MLME-PPD-LOST.indication primitive is issued to the NHL of an SPD as a notification that communications from the PPD have either stopped (the PPD has not been heard from for $macMaxMissedPPDBeacons$ beacon frames) or have become invalid (the received beacon has an invalid signature or certificate).

The SPD knows whether there was an NPD present in the system before the old PPD stopped transmitting based on receiving the NPD Indication subfield in past PPD beacon transmissions. The SPD will also receive an occasional NPD codeword during the Rx period if it listens during those times. Hence, if an NPD is present when it receives the MLME-PPD-LOST.indication primitive, it can take the following well-defined actions:

- Initiate the previously mentioned search procedure, per the description in D.4.1. The SPD should find a PPD (i.e., the recently “promoted” NPD).
- Execute the logic defined in D.4.2 to determine if the new PPD can provide adequate protection or not. If the new PPD suffices, aggregate with it via the inter-device communications procedures. If the new PPD does not suffice, the SPD should become a PPD and began its own regular beacon transmissions.

An SPD will issue the MLME-NPD-LOST.indication primitive to the NHL if it has not received the NPD's beacon frame or the NPD codeword in the last ($aMaxMissedNPDCodes \times aNPDPeriod$) superframes. After receiving both the MLME-PPD-LOST.indication and MLME-NPD-LOST.indication primitives, the SPD will start the promotion procedure after waiting for a random number of superframes.

D.6.8 Behavior on SPD receiving an NPD lost indication

The MLME-NPD-LOST.indication primitive is issued to the NHL of an SPD as a notification that the communications from the NPD have stopped, i.e., no NPD beacon frame or NPD codeword was heard within the last $aMaxMissedNPDCodes \times aNPDPeriod$ superframes. The SPD also receives information from the PPD's beacon frame regarding the existence of the NPD in the form of the NPD Indication subfield, and it is possible for these two sources of information to conflict. Table D.8 explains what conclusions the SPD should draw based on both the information contained in the PPD's beacon frame and the absence or presence of the NPD transmissions.

If the SPD is outside the coverage area of the NPD, the SPD will not be able to hear the NPD beacon frames, even though the PPD indicates there is an NPD present (i.e., the NPD Indication subfield is set to 01). In this case, the SPD should act as though there is no NPD in the beaconing network.

Table D.8—SPD behavior based on NPD information

Did SPD receive MLME-NPD-LOST.indication?	Value of PPD’s NPD Indication subfield (in MSF 1, Parameter 2)	SPD assumption	SPD recommended action
No (SPD hears an NPD)	00 = No NPD currently exists in the beaconing network, and the PPD is requesting one.	The PPD is in the process of finding another NPD. The NPD transmissions that the SPD was hearing were either from the “old” NPD or from an NPD belonging to another beaconing network that happens to be in range, but to which the SPD does not belong.	The SPD should follow the directions of the PPD, meaning it should attempt to become the NPD and proceed as described in 7.4.6.3.1.
	11 = No NPD currently exists in the beaconing network, and the PPD is not requesting one.	There is an NPD from a different beaconing network close by that is being received, so that no MLME-NPD-LOST indication is generated.	No action required. The PPD does not require an NPD, in any case.
	01 = An NPD currently exists in the beaconing network.	There is an NPD in the network.	No action required.
	10 = reserved	N/A	N/A
Yes (SPD no longer hears an NPD)	00 = No NPD currently exists in the beaconing network, and the PPD is requesting one.	There is no NPD in the network, but the PPD is seeking one.	The SPD should behave as described in 7.4.6.3.1.
	11 = No NPD currently exists in the beaconing network, and the PPD is not requesting one.	There is no NPD in the network.	No action required.
	01 = An NPD currently exists in the beaconing network.	There is an NPD in the network, but the SPD cannot hear it. If the SPD cannot hear the NPD for a long period of time, then the NPD may not be capable of protecting it if the NPD were to become the PPD.	When the PPD disappears, the SPD should act as if there was no NPD in the beaconing network. Therefore, the SPD should consider promoting itself to PPD, as described in D.6.9.
	10 = reserved	N/A	N/A

D.6.9 Behavior if both PPD and NPD are lost

If both the PPD and NPD are not heard by the SPDs: Some actions need to be taken to ensure that all of the SPDs do not attempt to become PPDs at the same time. There are some particular scenarios, depending on whether the PPD and NPD cease transmission with or without warning.

If the PPD and NPD both cease transmission with warning: In this case, if there is neither a PPD nor NPD, each SPD should wait for a random back-off time. If an SPD finds a new PPD before this limit is reached, it can start communication with the new PPD. If an SPD does not find a PPD before the limit is reached, it should promote itself as the new PPD.

If the PPD ceases transmission without warning, but the NPD ceases transmission with warning: If an SPD cannot hear the PPD for a period of $macMaxMissedPPDBeacons$ superframes, at least one of the SPDs should become a PPD. In such cases, a random back-off time should be employed at each SPD prior to it becoming a PPD; the SPD only completes the process of becoming a PPD if no other SPD becomes a PPD first.

If the PPD ceases transmission with warning, but the NPD ceases transmission without warning: In this case, the NHL of the SPD is notified through the MLME-NPD-LOST.indication primitive after $aMaxMissedNPDCodes \times aNPDPeiod$ superframes. In such cases, a random back-off time should be employed at each SPD prior to becoming a PPD; the SPD only completes the process of becoming a PPD if no other SPD becomes a PPD first.

If PPD and NPD cease transmission without warning: If an SPD cannot hear the PPD for a period of $macMaxMissedPPDBeacons$ superframes and the NHL of the SPD receives the MLME-NPD-LOST.indication primitive after $aMaxMissedNPDCodes \times aNPDPeiod$ superframes, a random back-off time should be employed at each SPD prior to becoming a PPD. If before that happens, an SPD finds a PPD, it can start communication with the new PPD. If an SPD does not find a PPD before the limit is reached, it should promote itself as the new PPD.

D.6.10 Behavior when selected as the NPD

The NHL of an SPD realizes that it has been chosen as the NPD if it observes the NPD Indication subfield of the PPD's beacon frames transition from 00 in the first superframe following the transmission of its own beacon frame, to 01 in the second superframe. The newly selected NPD's NHL should set the NPD subfield to TRUE to indicate its new status in all subsequent beacon frame transmissions and should also start to transmit regular NPD codes during the PPD Rx periods at a rate defined by $aNPDPeiod$.

D.6.11 Planned termination of SPD transmission

Planned termination of SPD beacon transmission occurs when the SPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the SPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive to its MLME with the Cease Tx subfield set to TRUE. If desired, the turn off sequence of the SPD can be delayed slightly to allow the SPD the possibility of transmitting several such superframes (i.e., by the NHL sending several MLME-START-BEACON.request primitives), to improve the chance that the PPD will detect the event and have the opportunity to react to it. Note that the PPD still controls beacon transmissions and ultimately determines if the SPD will be able to send anything and if so, how often. The SPD can, however, attempt to leverage the NST feature in order to get multiple transmission opportunities.

D.6.12 Behavior following termination of PPD transmissions without warning

Termination of PPD beacon transmission without warning refers to instances where the PPD stops transmitting and the SPD does not receive the Cease Tx subfield set to TRUE as part of an MLME-INCOMING-BEACON.indication primitive from that PPD. Examples of such events include someone "pulling the plug" or sudden power outages.

This event will be indicated to the NHL of the SPD via the MLME-PPD-LOST.indication primitive. The corresponding behavior is described in D.6.7.

D.6.13 Behavior following planned termination of PPD transmissions

Planned termination of PPD transmissions occurs when the PPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the PPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive with the Cease Tx subfield set to TRUE to its MLME. If desired, the turn off sequence of the PPD can be delayed slightly to allow it the possibility of transmitting several such superframes, to improve the chance that the SPD will detect the event and have the opportunity to react to it.

The SPD's reaction to this event is the same as if the event happened in an unplanned way. The difference is that it can react faster, i.e., it does not have to wait for the MLME-PPD-LOST.indication primitive to take action.

D.6.14 Behavior following termination of NPD transmissions without warning

Termination of NPD beacon transmissions without warning refers to instances where the NPD stops transmitting and the SPD does not receive the Cease Tx subfield set to TRUE as part of an MLME-INCOMING-BEACON.indication primitive from it. Examples of such events include someone "pulling the plug" on the NPD, or sudden power outages.

This event will be indicated to the NHL of the SPD via the MLME-NPD-LOST.indication primitive. The corresponding behavior is described in D.6.8.

D.6.15 Behavior following planned termination of NPD transmissions

Planned termination of NPD beacon transmission occurs when the NPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the NPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive with the Cease Tx subfield set to TRUE to its MLME. If desired, the turn off sequence of the NPD can be delayed slightly to allow it the possibility of transmitting several such superframes (i.e., by the NHL sending several MLME-START-BEACON.request primitives), to improve the chance that SPDs will detect the event and have the opportunity to react to it. Note that the PPD still controls beacon frame transmissions and ultimately determines if the NPD will be able to send anything and if so, how often. The NPD can, however, attempt to leverage the NST feature in order to get multiple transmission opportunities.

An SPD's reaction to this event is the same as if the event happened in an unplanned way. The difference is that it can react faster, i.e., it does not have to wait for the MLME-NPD-LOST.indication primitive to take action.

D.6.16 SPD-specific MIB attributes

The MLME-SET.request primitive can be used by the SPD NHL to modify the value of the MIB attributes. Else, the attributes will revert to the default values, as specified in 7.3.2. For example, if a PD decides to become an SPD, the SPD NHL should write the address of the PPD into the MIB attribute *macPPDAddress*.

D.7 NPD behaviors

The following subclauses describe some example NHL behaviors of an NPD.

D.7.1 Behavior on reception of the PPD NPD Indication subfield

MLME-INCOMING-BEACON.indication primitives from the PPD contain the NPD Indication subfield, which governs the behavior of the SPDs with respect to selection or deselection of them as an NPD. The NHL of an existing NPD should behave as follows when it determines the NPD Indication subfield information:

- A value of 00 indicates that the PPD wants to deselect the NPD and seek a new SPD to behave as the NPD. This behavior may be triggered by an unfavorable change in the parameters that caused the first NPD to be selected. For example, the NPD may change location, making it a less desirable choice. Hence the PPD should constantly monitor the regular beacon transmissions of its NPD, as well as the NPD codewords that should be regularly transmitted by the NPD. Once the PD concludes its role as NPD, it should cease transmission of NPD codewords or beacon frames with the NPD subfield set to TRUE and still attempt to transmit beacon frames as an SPD via the normal inter-device communications procedures.
- A value of 11 indicates that no NPD is required. The NPD should cease transmission of NPD codewords or beacon frames with the NPD subfield set to TRUE. The PPD will remove protection data associated with the NPD.
- A value of 01 indicates that the NPD should remain active (no change to current operation).

D.7.2 Behavior on NPD receiving a PPD lost indication

The MLME-PPD-LOST.indication primitive is issued to the NHL of an NPD as a notification that communications from the PPD have either stopped (the PPD has not been heard from for *macMaxMissedPPDBeacons* beacon frames) or have become invalid (the received beacon has an invalid signature or certificate). Upon receipt of this primitive, the NHL should promote itself to PPD and begin transmitting as the PPD immediately. The NHL of the new PPD should set the NPD Indication subfield accordingly:

- If the new PPD wants to find an NPD, it should set the subfield to 00 to provoke SPDs to send data (this is the preferred operation).
- If the new PPD does not want to find an NPD, it should set the subfield to 11 to indicate that no NPD is required.

D.7.3 Planned termination of NPD transmissions

Planned termination of NPD beacon transmission occurs when the NPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the NPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive to its MLME with the Cease Tx subfield set to TRUE. If desired, the turn off sequence of the NPD can be delayed slightly to allow the NPD the possibility of transmitting several such superframes, to improve the chance that the PPD and SPD will detect the event and have the opportunity to react to it. Note that the PPD still controls beacon transmissions and ultimately determines if the NPD will be able to send anything and how often. The NPD can, however, attempt to leverage the NST feature in order to get multiple transmission opportunities.

D.7.4 Behavior following termination of PPD transmissions without warning

Unplanned termination of PPD beacon transmission refers to instances where the PPD stops transmitting and the PPD does not receive the Cease Tx subfield set to TRUE as part of an MLME-INCOMING-BEACON.indication primitive from that PPD. Examples of such events include someone “pulling the plug” or sudden power outages.

This event will be indicated to the NHL of the NPD via the MLME-PPD-LOST.indication primitive. The corresponding behavior is described in D.7.2.

D.7.5 Behavior following planned termination of PPD transmissions

Planned termination of PPD transmissions occurs when the PPD stops transmitting in an orderly way, e.g., the device on/off switch is pressed or corresponding SW user interface is used. In this scenario, the NHL of the PPD should indicate its intention to stop transmitting by sending an MLME-START-BEACON.request primitive with the Cease Tx subfield set to TRUE to its MLME. If desired, the turn off sequence of the PPD can be delayed slightly to allow it the possibility of transmitting several such superframes, to improve the chance that the NPD will detect the event and have the opportunity to react to it.

The PPD's reaction to this event is the same as if the event happened in an unplanned way. The difference is that it can react faster, i.e., it does not have to wait for the MLME-PPD-LOST.indication primitive.

D.7.6 NPD-specific MIB attributes

The MLME-SET.request primitive can be used by the NPD NHL to modify the value of the MIB attributes. Else, the attributes will revert to the default values, as specified in 7.3.2.

Annex E

(informative)

Acquisition of time and location information

E.1 Introduction

There are several means of acquiring proper time and location information for incorporation into the beacon frame; MSF 1 conveying location parameters and MSF 2 and MSF 3 incorporating time into the generation of security/certificate parameters.

E.2 Time

As stated in this standard, time shall be utilized in universal coordinated time (UTC) standard, reference to non-advanced time along the prime meridian; Greenwich, England. Utilization of the time information should be accurate to ± 10 minutes of UTC. Granularity of utilized time to this minimal accuracy allows the use of several sources. These include the following:

- a) Global Positioning Satellite (GPS) information based on the United States constellation of satellites now in service
- b) Galileo, the European Union improved position/time reference satellite system
- c) WWVB – 60 kHz carrier frequency – transmission of time of day/date data, United States
- d) WWV / WWVH 2.5, 5, 10, 15, and 20 MHz HF transmissions of time of day/date information
- e) CHU – 3330, 7335, 14670 kHz HF transmissions of time of day information
- f) DCF – 77 kHz carrier frequency – transmission of time of day / date data, Germany
- g) MSF – 66 kHz carrier frequency – time of day transmission of time of day/date data, Rugby, England
- h) GLONASS, the Russian Federation satellite-based system providing time of day/date and location information
- i) Various other HF-based time of day/date radio service worldwide
- j) LORAN-C, Time/Location transmissions broadcast at 100 kHz carrier frequency
- k) Accurate Real Time Clock integrated circuits
- l) Web-based time sources, such as the U.S. Naval Observatory website²⁰
- m) Direct entry of time in UTC format

The just-listed HF and VLF radio transmissions time signals (3–9) are all usable as a time reference for IEEE 802.22.1 beacon frame security information generation; all report time in UTC format. These sources would primarily be useful in fixed access locations such as studio operations. Data from common time reference systems commonly found within the studio environment, such as SMPTE time code information, would be satisfactory for use by the beaconing device. Benefits arising from the use of signals such as WWVB are that a clear view of the sky is not needed as would be the case for GPS or other satellite-based systems; VLF signals, in particular, exhibit long-range reception characteristics as well as good building penetration.

²⁰The U.S. Naval Observatory time source can be found at <http://tycho.usno.navy.mil/cgi-bin/timer.pl>.

For mobile applications, the easiest method of acquiring time is to utilize information obtained from GPS, Galileo, or GLONASS; the latter receiving support from the Russian Federation and India. In the case of GPS, common NMEA data streams (NMEA0183 format [B10]) such as \$GPGGA will provide the required time of day and date information required. It should be noted that GPS time differs from UTC by several seconds. GPS time was fixed at 0 hour on 6 January, 1980 and has not been corrected by leap-seconds since commencement of operation. For example, on October 2007, GPS time led UTC by 14 s. Since time accuracy should be reported with a granularity of 10 min, this is of no consequence to the generation of signature and certificate for authentication of the beacon transmission. It is expected that Galileo equipment will provide similar data streams that may be utilized to obtain similar information.

An older, yet fully usable source of time and date information is LORAN-C. LORAN-C time/date information, like GPS, also does not utilize UTC time; rather, time is accurate as initialized on 1 January 1958 at 0 hour. No leap-seconds are incorporated into the signaling of time. LORAN-C, therefore, exhibits a 23 s lead over UTC at the time of writing.

Real Time Clock chips can eliminate the need entirely for external acquisition and reporting of time. Low-cost RTC chips offer 5PPM accuracy while newer designs operate with 2PPM accuracy or better. They have service life usually exceeding 10 years and could be set at the time of manufacture. An error of 5PPM results in 2.63 min/year error while 2PPM results in 1.05 min/year error. In the case of a 5PPM error, simply resetting the RTC at the time of certificate renewal would suffice to maintain the required accuracy needed by the beacon authentication system. Use of a 2PPM device would reduce the need to reset the RTC (using an Internet connection/certificate renewal or external source such as GPS) to once every 5 years. This system would appear to be the most desirable method of retaining accurate time within a beacon device.

Finally, time could be manually entered into the beaconing device manually. Such practice would still require the use of an accurate clock device to maintain proper time after it is set and it poses the risk of improper operation through operator error. While it is recommended that the user be able to input UTC time into the beaconing device, such provision would usually be utilized to correct long-term drift of an RTC device.

E.3 Location information

Location information is directly transmitted in MSF 1 of the IEEE 802.22.1-compliant beacon frame transmission. It provides information necessary for other devices, such as an IEEE P802.22 WRAN, to provide proper protection from interference to protected operations. Accuracy of location is important; the minimum keep-out zone protected by the beacon is 1500 m; accuracy of data sent by the beacon should be bounded by a radius of error no greater than 100 m.

There are several methods that can be utilized to obtain location information. The most common are as follows:

- a) Global Positioning Satellite (GPS) information based on the United States constellation of satellites now in service
- b) Galileo, the European Union improved position/time reference satellite system
- c) GLONASS, the Russian Federation satellite-based system providing time of day/date and location information
- d) LORAN-C, Time/Location transmissions broadcast at 100 kHz carrier frequency
- e) Proprietary methods of location derivation such as time of arrival, etc.
- f) Direct entry of location in WGS84 format as obtained from maps such as those published by the USGS in the United States. Similar maps can be obtained from the proper authorities in most countries.

At the time of writing, the most common source of location information will undoubtedly be GPS. The accuracy of GPS is dependant upon several factors. Initially, GPS was offered to the public in a reduced-accuracy form due to the incorporation of selective availability; a slow dithering of the location information that reduced the accuracy of GPS to 100 m. The primary GPS information is obtained from satellite data received at approximately 1575 MHz and occupies approximately 20 MHz RF bandwidth. A second augmentation signal, located at approximately 1175 MHz was used by authorized parties to increase the accuracy to 1 m. In the 1990s, selective-availability was discontinued and GPS, from that point forward, exhibits a maximum location error of 20 m. Further refinements such as the use of differential GPS (from land-based LW transmitters) and WAAS (from geo-stationary satellites) increase the accuracy to approximately 1 m. All forms of GPS are sufficiently accurate for use by the IEEE 802.22.1 beaconing system.

Galileo will provide accuracies at least one order of magnitude better than GPS. Galileo operates at approximately 1200 MHz with a bandwidth of 14 MHz. The Galileo system will provide the necessary accuracy required by the beaconing system.

GLONASS operates at approximately 1610 MHz and occupies 13 MHz of RF bandwidth. It currently offers accuracies of approximately 70 m; however, improvements are planned and are expected to be launched.

LORAN-C is the oldest of the usable, operating location systems. It operates with a carrier frequency of 100 kHz and, therefore, offers some building penetration. It is being phased out for use by aircraft but it is still widely used for other navigation systems such as navel / ship operations. The raw accuracy of LORAN-C can be as poor as 160–400 m; however, after initial corrections are taken into account, the repeatable accuracy is approximately 90 m. With the incorporation of differential information, the accuracy can be increased to 30 m or better.

Direct entry of map information is also useful for fixed point operations such as at permanent studios. The operator should be careful to input the correct coordinates from a map, in WGS84 DATUM format or extract from one of several internet-based map service providers that can convert a street address to latitude and longitude coordinates. Beaconing devices are expected to allow direct entry of location information, but it may be easier to simply connect a GPS receiver that has taken a reading just outside of, or on the roof of the studio or other fixed point location and report that to the beaconing device via a NMEA0183 string such as \$GPGGA.

It is also anticipated that additional sources of location information data will become available for use by the beaconing device. Care should be taken to ensure that the data is parsed properly and is send in the correct DATUM format; WGS84.

Another important factor related to location is the specified height of the protected device's antenna. The antenna height is specified relative to "ground level," which is defined as the highest point within the radius of protection (i.e., the keep-out zone radius). The recommended source of altitude data is found using a terrain database. There are many free sources of terrain data, such as the U.S. Geological Survey (USGS) National Elevation Dataset (NED) and the Shuttle Radar Topography Mission (SRTM) database.²¹

²¹For free database downloads of both of these, visit <http://seamless.usgs.gov/>.

Annex F

(informative)

Concept of service primitives

This annex, which is largely drawn from IEEE Std 802.15.4™-2006 [B5], provides an overview of the concept of service primitives.

The services of a layer N are the capabilities it offers to the NHL (or sublayer) $N + 1$ by building its functions on the services of the next lower layer (or sublayer) $N - 1$. Figure 38 shows the service hierarchy and the relationship of the $N + 1$, N , and $N - 1$ layer (or sublayer) peer protocol entities. In the case of this standard, the $N + 1$, N , and $N - 1$ layers correspond to the NHL, MAC sublayer, and PHY layer, respectively.

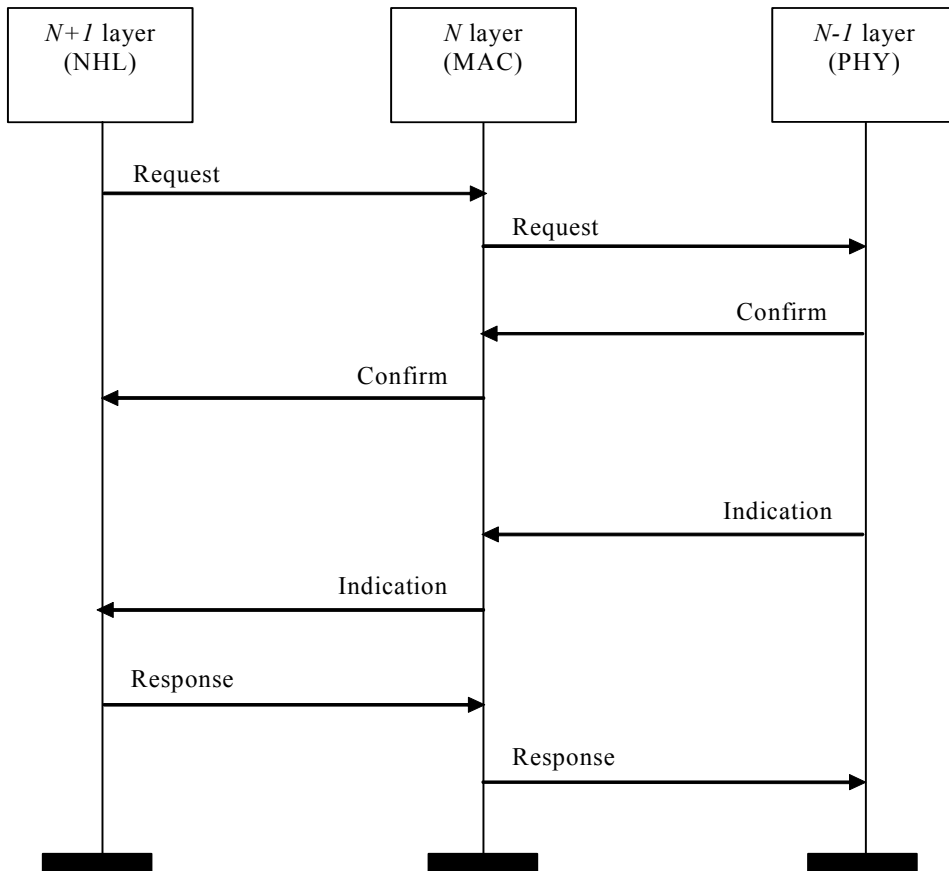


Figure 38—Service primitives

The services are specified by describing the information flow between a layer and the layer beneath it. This information flow is modeled by discrete, instantaneous events, which characterize the provision of a service. Each event consists of passing a service primitive from one layer to the other through a Service Access Point (SAP). A SAP is an interface that supports the transport of information between peer layer (or sublayer) entities. Service primitives convey the required information by providing a particular service. These service

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

A service is specified by describing the service primitives and parameters that characterize it. A service may have one or more related primitives constituting the activity related to that particular service. Each service primitive may have zero or more parameters that convey the information required to provide the service.

A primitive can be one of four generic types:

- *Request*: The request primitive is passed from the $N + 1$ layer to the N layer (or from the N layer to the $N - 1$ layer) to request that a service is initiated.
- *Confirm*: The confirm primitive is passed from the N layer to the $N + 1$ (or from the $N - 1$ layer to the N layer) to convey the results of one or more associated previous service requests.
- *Indication*: The indication primitive is passed from the N layer to the $N + 1$ layer (or from the $N - 1$ layer to the N layer) to indicate an event that is significant to the $N + 1$ layer (or N layer). This event may be logically related to a remote service request, or it may be caused by an N layer (or $N - 1$ layer) internal event.
- *Response*: The response primitive is passed from the $N + 1$ layer to the N layer (or from the N layer to the $N - 1$ layer) to complete a procedure previously invoked by an indication primitive.

Annex G

(informative)

Bibliography

[B1] ANSI C63.17-1998, American National Standard for Methods of Measurement of the Electromagnetic and Operational Compatibility of Unlicensed Personal Communications Services (APSCS) Devices.²²

[B2] ETSI EN 300 422-1 V1.2.2, Electromagnetic compatibility and Radio spectrum Matters; Wireless microphones in the 25 MHz to 3 GHz frequency range; Part 1: Technical characteristics and test methods, 2000.

[B3] FCC 47 CFR Section 74.861, Subpart H—Low Power Auxiliary Stations, May 18, 2007.²³

[B4] FCC Docket Number 04-186, Comments of Shure Incorporated, November 30, 2004.

[B5] IEEE Std 802.15.4-2006, IEEE Standard for Information Technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs).

[B6] ISO/IEC 7498-1:1994, Information technology—Open Systems Interconnection—Basic Reference Model: The Basic Model.²⁴

[B7] Kuffner, S., “TG1 Link Margin Simulation Results Using Static Multipath Magnitudes,” IEEE P802.22-07/0241r2, June 2007.

[B8] Kuffner, S., “TG1 Required Detection SNRs Versus Protected Radii,” IEEE P802.22-08/0219r0, July 2008.

[B9] Martin, Stephen R., “Interference Rejection Thresholds of Consumer Digital Television Receivers Available in 2005 and 2006,” OET Report FCC/OET 07-TR-1003, March 30, 2007.

[B10] NMEA 0183 Interface Standard, Version 3.01, January 2002.²⁵

²²ANSI publications are available from the Sales Department, American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (<http://www.ansi.org/>).

²³Federal Communications Commission publications are available from <http://www.fcc.gov>.

²⁴ISO publications are available from the ISO Central Secretariat, Case Postale 56, 1 rue de Varembe, CH-1211, Genève 20, Switzerland/Suisse (<http://www.iso.ch/>). ISO publications are also available in the United States from the Sales Department, American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (<http://www.ansi.org/>).

²⁵National Marine Electronics Association publications are available from <http://www.nmea.org/>.