

Standard ECMA-392

1st Edition / December 2009

MAC and PHY for Operation in TV White Space

Standard



is the registered trademark of Ecma International



COPYRIGHT PROTECTED DOCUMENT

Contents

Page

1	Scope	1
2	Conformance	1
3	Normative references	1
4	Terms and definitions	1
5	Abbreviations and Acronyms	4
6	General description	7
6.1	Network components	7
6.2	Network formation	7
6.3	Protocol architecture	8
6.4	Addressing	9
6.5	PHY features	9
6.6	Overview of MAC service functionality	10
6.6.1	Logical groups	11
6.6.2	Control algorithms	11
6.6.3	Channel selection	11
6.6.4	The superframe	11
6.6.5	Beaconing	12
6.6.6	Medium access	13
6.6.7	Data communication between devices	13
6.6.8	MAC frame data rates	13
6.6.9	Security	13
6.6.10	Information discovery	14
6.6.11	Support for higher-layer timer synchronization	14
6.6.12	Protection of incumbent users	14
6.6.13	Self-coexistence	15
6.6.14	Rate adaptation	15
6.6.15	Power management	15
7	MAC common part sublayer	15
7.1	MAC Frame Format	16
7.1.1	Frame format conventions	16
7.1.2	General MAC frame format	17
7.1.3	Beacon frames	22
7.1.4	Control frames	25
7.1.5	Command frames	27
7.1.6	Data frames	35
7.1.7	Aggregated data frames	35
7.1.8	Information elements	35
7.2	Frame processing	62
7.2.1	Frame addresses	62
7.2.2	Frame reception	63
7.2.3	Frame transaction	63
7.2.4	Frame transfer	64
7.2.5	Frame retry	64
7.2.6	Inter-frame space (IFS)	65
7.2.7	Duplicate detection	66
7.2.8	RTS/CTS use	66
7.2.9	MAC header fields	66
7.2.10	Information elements	69
7.3	MAC Structure and Beaconing	73

7.3.1	Beacon Period.....	74
7.3.2	Beacon slot state	74
7.3.3	BP length	74
7.3.4	Beacon transmission and reception.....	74
7.3.5	Beacon collision detection	75
7.3.6	BP contraction	75
7.3.7	Merger of multiple beacon groups.....	76
7.3.8	Signalling window.....	79
7.4	Device Synchronization	79
7.4.1	Clock accuracy.....	80
7.4.2	Synchronization for devices in hibernation mode	80
7.4.3	Guard times	80
7.5	Data Transfer Period	81
7.5.1	Prioritized Contention Access (PCA).....	82
7.5.2	Channel Reservation Access (CRA).....	88
7.6	Fragmentation and Aggregation	95
7.6.1	Fragmentation and reassembly.....	95
7.6.2	Aggregation.....	96
7.7	ARQ, Multirate Support and Power Control.....	97
7.7.1	ARQ Policies	97
7.7.2	Multi-rate Support.....	99
7.7.3	Transmit Power Control	100
7.8	Dynamic Channel Selection.....	100
7.9	Power Management Mechanisms	100
7.9.1	Power management modes	100
7.9.2	Device power states	100
7.9.3	Power state transitions	101
7.9.4	Hibernation mode operation.....	102
7.9.5	Hibernation anchor operation	103
7.10	Probe.....	104
7.11	Protection of incumbents	104
7.11.1	Channel Measurement	104
7.11.2	Channel Classification	107
7.11.3	Channel Evacuation	108
7.12	Self-coexistence.....	109
7.12.1	Self-coexistence scenarios.....	109
7.12.2	Distributed self-coexistence mechanisms.....	109
7.12.3	Centralized self-coexistence mechanisms	110
7.13	Network Entry and Initialization	111
7.13.1	Initial Channel SCAN and Device Discovery.....	113
7.13.2	Master-Slave Association	114
7.13.3	Pair discovery	115
7.13.4	Create/join a beacon group	116
7.13.5	Pairing.....	116
7.13.6	Setup connections.....	117
7.14	MAC sublayer parameters	118
8	Security.....	120
8.1	Security mechanisms.....	120
8.1.1	Security operation	120
8.1.2	4-way handshake	120
8.1.3	Key transport.....	121
8.1.4	Freshness protection	121
8.1.5	Data encryption.....	121
8.1.6	Frame integrity protection	121
8.2	Security modes	121
8.2.1	Security mode 0	122
8.2.2	Security mode 1	123
8.2.3	Security mode 2	123
8.3	Temporal keys.....	123

8.3.1	Mutual authentication and PTK derivation	124
8.3.2	GTK exchange	125
8.3.3	Pseudo-random function (PRF) definition	126
8.3.4	PTK and KCK derivation	126
8.3.5	PTK MIC generation	127
8.3.6	Random number generation	127
8.4	Frame reception steps and replay prevention measures.....	128
8.4.1	Frame reception.....	128
8.4.2	Replay prevention	129
8.4.3	Implications on GTKs.....	129
8.5	AES-128 CCM Inputs	129
8.5.1	Overview.....	129
8.5.2	Nonce.....	130
8.5.3	CCM blocks	130
9	PHY	132
9.1	Introduction.....	132
9.2	Symbol description	132
9.2.1	OFDM symbol description	132
9.2.2	Symbol parameters	133
9.3	PPDU.....	134
9.3.1	PLCP preamble	135
9.3.2	PLCP header	137
9.3.3	PSDU.....	142
9.4	Constellation mapping and modulation.....	148
9.4.1	Data modulation	148
9.4.2	Pilot modulation	150
9.5	OFDM modulation.....	150
9.5.1	Data subcarriers	151
9.5.2	Pilot subcarriers	151
9.5.3	Null subcarriers	153
9.5.4	Implementation of Fourier transform	153
9.6	General block diagram for the OFDM PHY	153
9.7	General requirements	154
9.7.1	Operating frequency range	154
9.7.2	Channel bandwidth and numbering	154
9.7.3	PHY layer timing.....	155
9.8	Transmitter requirements.....	155
9.8.1	Transmit center frequency tolerance	155
9.8.2	Symbol clock frequency tolerance	155
9.8.3	Clock synchronization	155
9.8.4	Transmitter constellation error.....	155
9.9	Receiver requirements.....	157
9.9.1	Receiver sensitivity	157
9.9.2	Maximum received signal level.....	157
9.9.3	Center frequency and symbol clock frequency tolerance	158
9.9.4	Link quality estimate	158
9.10	Control mechanisms	158
9.10.1	Device synchronization	158
9.10.2	Transmit power control.....	159
9.11	Multiple antennae (optional).....	159
9.11.1	Multiple antennae normal preamble and burst preamble specification	160
9.11.2	Multiple antennae PLCP header specification	161
9.11.3	Pilot subcarriers for all multiple antennae modes.....	162
9.11.4	Frequency interleaved transmit diversity (FITD).....	162
9.11.5	Alamouti space time block coding (STBC).....	163
9.11.6	Spatial multiplexing (SM) mode	163
Annex A	(normative) MUX sublayer	165
A.1	MUX service	165

A.2	MUX protocol data unit format	165
Annex B (normative)	OFDM parameters for 7 MHz and 8 MHz channel bandwidths	167
Annex C (normative)	Data rates for 7 MHz and 8 MHz channel bandwidths.....	169
Annex D (normative)	MAC policies	171
D.1	Reservation limits	171
D.2	Reservation locations	173
D.3	MAC policies parameters.....	173
Annex E (informative)	FFT-based pilot sensing algorithms	175
E.1	Performance of the algorithms.....	176
Annex F (informative)	An example of TPC algorithm	177
Bibliography.....		181

Introduction

Analogue broadcasting systems have been or are being upgraded to digital technology which frees up channels in the TV frequency bands. Regulation allows using such vacated channels called “TV white spaces”, under certain conditions such as sensing actual vacancy or by querying geo-location databases for the devices wishing to use TV white space. This Standard specifies the physical, and the medium access sub-layer for such wireless devices.

Applications include high speed video streaming and internet access on personal/portable electronics, home electronics equipment, and computers and peripherals.

This Ecma Standard has been adopted by the General Assembly of December 2009.

DISCLAIMER

This document and possible translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to Ecma International, except as needed for the purpose of developing any document or deliverable produced by Ecma International (in which case the rules applied to copyrights must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by Ecma International or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and ECMA INTERNATIONAL DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

MAC and PHY for Operation in TV White Space

1 Scope

This Standard specifies a medium access control (MAC) sub-layer and a physical (PHY) layer for personal/portable cognitive wireless networks operating in TV bands. This Standard also specifies a MUX sublayer for higher layer protocols.

This Standard specifies a number of incumbent protection mechanisms which may be used to meet regulatory requirements, which themselves are outside of the scope of this Standard.

2 Conformance

Conforming devices implement the MUX sub-layer, MAC sub-layer and the PHY layer as specified herein and support at least one of the device types (master, peer, or slave) and at least one of bandwidths (6 MHz, 7 MHz, 8 MHz), and may support multiple antennae modes.

3 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

ISO/IEC 10646:2003, Information technology -- Universal Multiple-Octet Coded Character Set (UCS)

ISO/IEC 18033-3:2005, Information technology -- Security techniques -- Encryption algorithms -- Part 3: Block ciphers

4 Terms and definitions

For the purposes of this document, the following terms and definitions apply.

4.1

alien beacon group

group of devices which BP is not aligned with the BP of the current device

4.2

alien device

member of an alien beacon group

4.3

Access Category

AC

common set of prioritized contention access (PCA) parameters to contend for the medium to transmit MAC protocol data units (MPDUs) with certain priorities

4.4

Beacon Group

BG

set of devices that share the same beacon period start time (BPST)

4.5

Beacon Period

BP

time during which a device sends or listens for beacons

4.6

Beacon Period Start Time

BPST

start of the beacon period

4.7

Channel Reservation Protocol

CRP

protocol to support negotiation and maintenance of channel time reservations

4.8

Contention signalling window

time window for exchanging control or management information in the slotted aloha based manner

4.9

Data integrity

assurance that the data has not been modified from its original form

4.10

Data Transfer Period

DTP

time period within a superframe used mainly for data transfer via prioritized contention access (PCA) or in reservations established using the channel reservation protocol (CRP)

4.11

Device

entity conforming to this Standard

4.12

Extended beacon group

union of a device's beacon group and the beacon groups of all devices in the device's beacon group

4.13

incumbents

regulatory protected transmission systems operating in the TV bands

4.14

incumbent protection mechanisms

mechanisms including DFS, TPC, geo-location/database access, and spectrum sensing

4.15

MPDU

MAC PDU

4.16

MSDU

MAC SDU

4.17**Master****Master device**

device acting as a centralized coordinator of medium access on behalf of at least one slave device

4.18**Master-slave group**

group of devices with a master device and its slave devices

4.19**Message Integrity Code****MIC**

cryptographic checksum generated using a symmetric key

NOTE A MIC is typically appended to data for data integrity and source authentication similar to a digital signature

4.20**Neighbour**

member of a beacon group

4.21**Network Allocation Vector****NAV**

remaining time a neighbour device has indicated it will access the medium

4.22**Outband channel**

channel other than the one being used for data transmission

4.23**Peer****Peer device**

device coordinating medium access with other devices without a centralised coordinator

4.24**Peer-to-peer group**

group of peer devices

4.25**Prioritized Contention Access****PCA**

prioritized CSMA/CA access mechanism

4.26**Proxy**

peer device that coordinates outband channel measurement

4.27**Quiet period**

time period scheduled to detect incumbents

4.28**Reservation**

one or more medium access slots (MASs) within a superframe during which a device has preferential access to the medium

4.29**Reservation signalling window**

time window used for exchanging control or management information in the reservation based manner

4.30

Slave

Slave device

device associated with and coordinated by a master device for medium access

4.31

Stream

logical flow of MSDUs from one device to one or more other devices

4.32

Superframe

periodic time interval to coordinate frame transmissions between devices

4.33

Transmission Opportunity

TXOP

time interval for prioritized contention access (PCA) to initiate transmissions

4.34

TXOP holder

device that has successfully contended for a TXOP

5 Abbreviations and Acronyms

AC	access category
ACK	acknowledgment
A/D	analog-to-digital
AES	advanced encryption standard
AGC	automatic gain control
AIFS	arbitration inter-frame space
ASIE	application-specific information element
AWGN	additive white Gaussian noise
BPOIE	beacon period occupancy information element
BPSK	binary phase-shift keying
BcstAddr	broadcast device address
BP	beacon period
BPST	beacon period start time
B-ACK	block acknowledgment
BW	bandwidth
CBC-MAC	cipher block chaining-message authentication code
CCA	clear channel assessment
CCM	counter mode encryption and cipher block chaining message authentication code
CE	channel estimation
CINR	carrier-to-interference and noise ratio
CP	cyclic prefix
CRC	cyclic redundancy check
CRP	Channel Reservation protocol

CSMA/CA	carrier sense multiple access with collision avoidance
CTS	clear to send
D/A	digital-to-analog
DC	direct current
DestAddr	destination device address
DevAddr	device address
DME	device management entity
DTP	Data transfer period
EO	encryption offset
EUI	extended unique identifier
FCC	Federal Communications Commission
FCS	frame check sequence
FEC	forward error correction
FFT	fast Fourier transform
FITD	frequency interleaved transmit diversity
GF	Galois field
GTK	group temporal key
HDR	header
HEI	header error indicator
I	inphase
ICI	inter-carrier interference
ID	identifier
IE	information element
IFFT	inverse FFT
IFS	inter-frame space
Imm-ACK	immediate acknowledgment
ISI	inter-symbol interference
KCK	key confirmation key
LQE	link quality estimate
LSB	least significant bit
M2S	Master-to-Slave
MAC	medium access control
MAS	medium access slot
MCDU	MAC command data unit
McstAddr	multicast device address
MIB	management information base
MIC	message integrity code
MIFS	minimum inter-frame space
MKID	master key identifier

MLME	MAC sublayer management entity
MPDU	MAC protocol data unit
MSB	most significant bit
MSDU	MAC service data unit
NAV	network allocation vector
No-ACK	no acknowledgement
OFDM	orthogonal frequency division multiplexing
OUI	organizationally unique identifier
P2P	Peer-to-Peer
PCA	prioritized contention access
PER	packet error rate
PHY	physical layer
PLCP	physical layer convergence protocol
PLME	physical layer management entity
PMK	pair-wise master key
PPDU	PHY protocol data unit
ppm	parts per million
PRBS	pseudo-random binary sequence
PRF	pseudo-random function
PSDU	PHY service data unit
PTK	pair-wise temporal key
Q	quadrature
QAM	quadrature amplitude modulation
QP	quiet period
QPSK	quadrature phase-shift keying
RF	radio frequency
RMS	root mean square
RS	Reed-Solomon
RSSI	received signal strength indication
RTG	receive-to-transmit transition gap
RTS	request to send
SAP	service access point
SFC	secure frame counter
SFN	secure frame number
SIFS	short inter-frame space
SM	spatial multiplexing
SNR	signal-to-noise ratio
SrcAddr	source device address
STBC	space time block code

TKID	temporal key identifier
TPC	transmit power control
TTG	transmit-to-receive transition gap
TV	television
TXOP	transmission opportunity
UCA	unused CRP reservation announcement
UCR	unused CRP reservation response
UHF	ultra high frequency
VHF	very high frequency
WM	wireless microphone

6 General description

6.1 Network components

A basic component of a network is a device. Two or more devices communicating on the same physical channel constitute a network. There are three types of devices, master device, slave device, and peer device. The device type of a device is preconfigured. The autonomous transition of device type is not supported in this Standard, although the device type may be reconfigurable by DME which is out of scope of this Standard.

6.2 Network formation

A basic network operates in one of two basic network formation modes: the master-slave mode or the peer-to-peer mode. Both are shown in Figure 1. In the master-slave mode, a device is designated as master and others are associated with the master as slaves. The master coordinates channel access in the master-slave mode. Communication is normally established between slave devices and the master device. A slave device may also directly communicate with another slave device under the coordination of the master.

A peer-to-peer network differs from a master-slave based network mainly in that devices can form a network in the peer-to-peer way and coordinate channel access with distributed beaconing and channel reservation. A peer-to-peer network comprises of peer devices. A peer device is able to access channel via distributed reservation and directly communicate with any other peer device as long as they are in range of one another. In other words, a peer-to-peer network can be ad hoc, self-organizing, and self-healing.

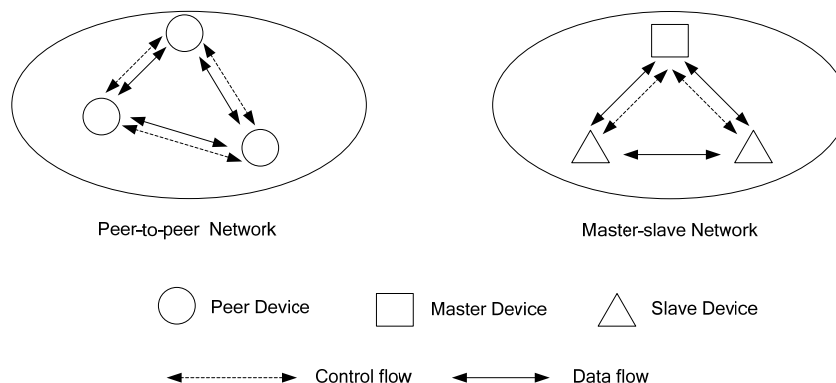


Figure 1 — Basic Network Formation

Two or more networks can share the same channel and may also communicate with each other in a coordinated way.

A number of networks may also form a large-scale network such as a mesh network or a cluster tree network. It allows multiple hops to route messages from any device to any other device in the network. Such functions can be added at the higher layer, but are not part of this Standard.

6.3 Protocol architecture

This Standard specifies a PHY layer and a /MAC sublayer. As shown in Figure 2, the PHY layer and the MAC sublayer correspond to the PHY layer and the MAC sublayer of the OSI basic reference model [5] respectively. In this Standard the MAC entity is represented by a device address.

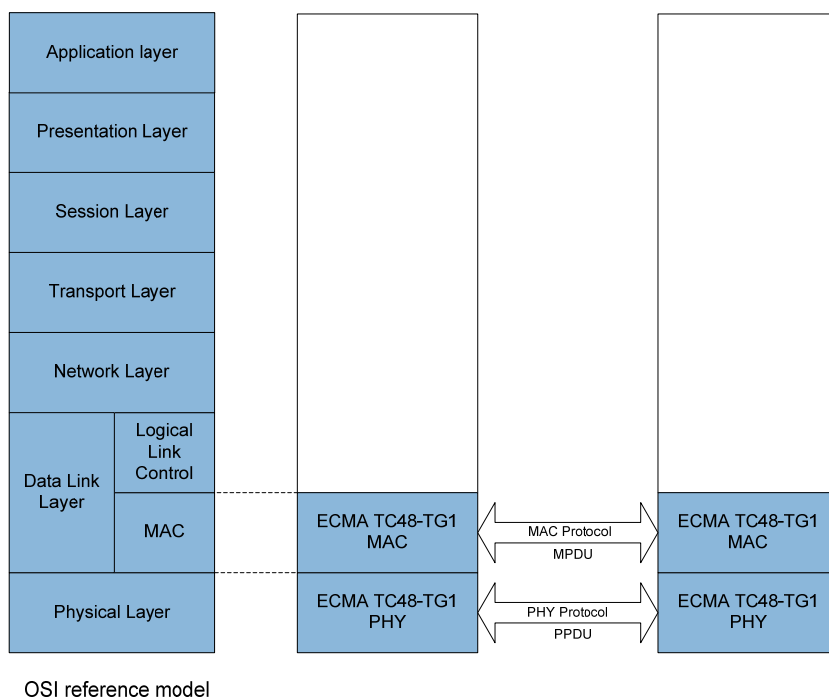


Figure 2 — Architectural reference model

Service access points (SAPs) interaction with PHY and MAC sublayers are illustrated in Figure 3. As a reference, Service access points (SAPs) are provided for both data transfer as well as management of the MAC sublayer. Data transfer for the MAC sublayer is through the MAC SAP. Both the MAC sublayer and the PHY layer conceptually include management entities, called the MAC sublayer management entity (MLME) and physical layer management entity (PLME). These entities provide the layer management service interfaces for the layer management functions. The DME is a layer-independent entity that may be viewed as residing in a separate management plane or as residing “off to the side.” DME may be viewed as being responsible for such functions as the gathering of layer-dependent status from the various layer management entities, and similarly setting the value of layer-specific parameters. The DME typically performs such functions on behalf of the general system management entities and implements standard management protocols. Figure 3 depicts the relationship among the management entities. The specification of SAPs and management entities, as shaded parts of Figure 3, is out of the scope of this Standard.

In order to enable the coexistence of concurrently active higher layer protocols within a single device, a MUX sublayer is specified. This sublayer routes outgoing and incoming MSDUs to and from their corresponding higher layers. The MUX sublayer is described in Annex A.




Figure 3 — The reference model for this Standard

6.4 Addressing

Individual MAC entities are addressed via an EUI-48 [1], and are associated with a volatile abbreviated address called a DevAddr. MAC address is included in beacon and/or control messages for global identification.

Data frames normally use abbreviated DevAddr that identifies a single MAC entity for reducing overhead. DevAddrs are 16-bit values, generated locally within the device. Consequently, it is possible for a single value to ambiguously identify two or more MAC entities. This Standard provides mechanisms for resolving ambiguous DevAddrs.

The MAC addressing scheme includes multicast and broadcast address values. A multicast address identifies a group of MAC entities. The broadcast address identifies all MAC entities.

Device name string may be used for helping user to identify a device, as specified in 7.1.8.16. Device name string can be assigned and changed by DME. Device name string should be included in beacons for assisting device discovery.

A stream ID may be determined locally by device to identify stream originating from itself.

6.5 PHY features

A MAC entity is associated with a single PHY entity.

The MAC sublayer requires the following features provided by the PHY:

- Frame transmission for both normal and burst modes;
- Frame reception for both normal and burst modes;
- Header error indication for PHY and MAC header;
- Clear channel assessment for estimation of medium activity.

Figure 4 shows the structure of a PHY frame.

- There are two types of preamble: normal and burst.
- The PLCP header including MAC and PHY Headers is protected by RS parity.
- The Frame Payload is followed by its frame check sequence (FCS).




Figure 4 — PHY Frame Structure

Frames are transmitted by the PHY from the source device and delivered to the destination device in identical bit order. Throughout this specification reference to the start of a frame refers to the leading edge of the first symbol of the PHY frame at the local antenna and end of a frame refers to the trailing edge of the last symbol of the PHY frame.

Frame transmission and reception are supported by the exchange of parameters between the MAC sublayer and the PHY layer. These parameters allow the MAC entity to control, and be informed of, the frame transmission mode, the frame payload data rate and length, the frame preamble, the PHY channel and other PHY-related parameters.

In normal mode transmission, the MAC entity has full control of frame timing. In burst mode transmission, the MAC entity has control of the first frame timing and the PHY provides accurate timing for the remaining frames in the burst.

6.6 Overview of MAC service functionality

The MAC service specified in this Standard provides:

- Communication between devices within radio range on a single channel using the PHY;
- A reservation-based channel access mechanism;
- A prioritized, contention-based channel access mechanism;
- A synchronization facility for coordinated applications;
- Mechanisms for protection of incumbent user;
- Master-slave operation;
- Peer-to-peer operation;
- Mechanisms for handling mobility and interference situations;
- Device power management by scheduling of frame transmission and reception;
- Secure communication with data authentication and encryption using cryptographic algorithms.

The architecture of this MAC service is either master-coordinated or fully distributed. A device provides required MAC functions and optional functions according to its device type.

Coordination of devices within radio range is achieved by the exchange of beacon frames and/or command frames. Periodic beacon transmission enables device discovery, supports dynamic network organization, and provides support for mobility. Beacons provide the basic timing for the network and carry reservation and scheduling information for accessing the medium.

6.6.1 Logical groups

The MAC protocol is specified with respect to an individual device, which has its own individual neighbourhood. All MAC protocol facilities are expressed with respect to this individual neighbourhood.

In a peer-to-peer network formed with fully distributed medium access coordination, logical groups are formed around each device to facilitate contention-free frame exchanges while exploring medium reuse over different spatial regions. In this Standard, these logical groups are a beacon group and an extended beacon group, both of which are determined with respect to an individual device.

In a master-slave based network formed under the coordination of a master, logical groups are formed with master-slave operation to facilitate contention-free frame exchanges. A master must be a beaconing device. A slave device could be either a beaconing device or nonbeaconing device.

6.6.2 Control algorithms

MAC protocol algorithms attempt to ensure that no member of the extended beacon group transmits a beacon frame at the same time as the device. Information included in beacon frames facilitates contention-free frame exchanges by ensuring that a device does not transmit frames while a neighbour is transmitting or receiving frames.

To permit correct frame reception, MAC protocol algorithms attempt to ensure that a device's DevAddr is unique within the device's extended beacon group.

6.6.3 Channel selection

A device chooses a channel without incumbent presence. If no beacon is detected in the selected channel and no incumbent user is found, a master device or a peer device creates its beacon period (BP) by sending a beacon. If one or more beacons are detected in the selected channel, the device synchronizes its BP to existing beacons in the selected channel. The device exchanges data with members of its beacon group using the same channel the device selected for beacons.

When a slave device is enabled, it scans one or more channels for master beacon. A slave device selects the channel used by its intended master.

Each device operates in a dynamic environment and under unlicensed operation rules. Thus, it is subject to interference from licensed users, other networks, and other unlicensed wireless entities in its channel. To enable the device to continue operation in this type of environment, each device has the capability to dynamically change the channel in which it operates.

If at any time a device determines that the current channel is unsuitable, it uses the dynamic channel selection procedure, as described in 7.8, to move to a new channel.

6.6.4 The superframe

The basic timing structure for frame exchange is a superframe. The superframe duration is specified as mSuperframeLength. The superframe is composed of 256 medium access slots (MASs), where each MAS duration is mMASLength.

Each superframe starts with a BP, which extends over one or more contiguous MASs. The start of the first MAS in the BP, and the superframe, is called the beacon period start time (BPST).

A recurring superframe consists of a beacon period (BP), data transfer period (DTP) and a contention signalling window (CSW). A reservation based signalling window (RSW) could be appended right after BP to support signal exchange between a master and slave devices in the master-slave mode. RSW is not needed for a peer-to-peer network. The signalling windows and beacon period are used for sending and receiving control/management information.

The BP length is adjustable and depends on how many regular beaoning devices participate in the same BP. In the master-slave mode, BP length can be kept as minimum as 1 beacon slot if the master is the only device that does regular beaoning. In the peer-to-peer mode, the number of beacon slots can be as many as the number of peer devices. The beaoning policy is specified in 6.6.5.

All devices which share the same channel shall follow the same superframe structure. Superframe merging is necessary if two networks follow different superframe structures and share the same channel.

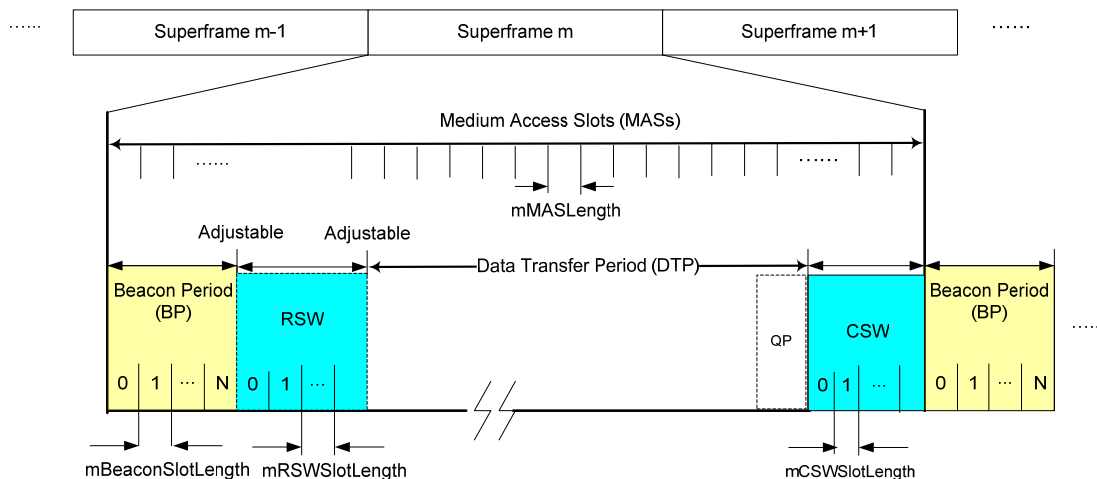


Figure 5 — MAC superframe structure

6.6.5 Beaoning

A device is defined as a beaoning device if it owns a beacon slot in BP and regularly transmits beacons. A peer device or a master device is by default a beaoning device. A slave device is normally a non-beaoning device unless promoted to be a beaoning device. All slave devices must be capable to transmit regular beacons, if the master device so decides.

All the devices shall keep awake during beacon period and CSW in order to capture all the control/management information which might be relevant to every device. A device may exchange data, monitor channel status, or go to sleep mode during data transfer period (DTP).

A beacon packet contains the following important information for network operation, but not limited to:

- Device identification information
 - Important for device discovery and self-coexistence
- Beacon slot occupation information
- Medium reservation information
 - Necessary for QoS provisioning
- TIM (Traffic indication map)
- Quiet period schedule
 - Necessary for in-service channel monitoring (sensing)
- Channel Measurement Report

- Channel classification and change information.

Each device protects its and its neighbours' BPs for exclusive use of the beacon protocol. No transmissions other than beacons are attempted during the BP of any device. Protection of the device's BP is implicit.

A device should protect an alien BP, detected by reception of a beacon frame unaligned with the device's own BP, by announcing a reservation covering the alien BP in its beacon.

6.6.6 Medium access

The medium is accessed in one of three ways:

- During the BP, devices send only beacon frames, according to the rules specified in 7.3.
- During reservations, devices participating in the reservation send frames according to rules specified in 7.5.2.
- During quiet period, all devices keep silent.
- Outside the BP and reservations, devices may send frames using a prioritized contention-based access method, as described in 7.5.1.

6.6.7 Data communication between devices

Data is passed between the MAC entity and its client in MSDUs qualified by certain parameters. MSDUs are transported between devices in data frames. To reduce the frame error rate of a marginal link, data frames can be fragmented and reassembled, as described in 7.6.1. Fragments are numbered with an MSDU sequence number and a fragment number.

If the source device wishes to verify the delivery of a frame, then one of the acknowledgement policies is used, as described in 7.7.1. This Standard provides for three types of acknowledgements to enable different applications. The No-ACK policy, described in 7.7.1.1, is appropriate for frames that do not require guaranteed delivery, or are delay sensitive and a retransmitted frame would arrive too late. The Imm-ACK policy, described in 7.7.1.2, provides an acknowledgement process in which each frame is individually acknowledged following the reception of the frame. The B-ACK policy, described in 7.7.1.3, lets the source send multiple frames without intervening ACK frames. Instead, the acknowledgements of the individual frames are grouped into a single response frame that is sent when requested by the source device. The B-ACK process decreases the overhead of the Imm-ACK process while allowing the source device to verify the delivery of frames to the destination.

If the source device does not receive the requested acknowledgement, then it may retransmit the frame, as described in 7.5.1.7 and 7.5.2.9, or it may discard the frame. The decision to retransmit or discard the frame depends on the type of data or command that is being sent, the number of times that the source device has attempted to send the frame, the length of time it has attempted to send the frame, and other implementation-dependent factors.

6.6.8 MAC frame data rates

MAC beacon frames are intended to be received and interpreted by all devices and hence their frame payloads are transmitted at pBeaconTransmitRate, which can be decoded by all recipients. Other frames are exchanged in a more restricted context and their frame payloads may be transmitted at higher data rates if possible. MAC headers are always transmitted at the most robust data rate supported by the PHY.

6.6.9 Security

Wireless networks present unique security challenges due to the loss of protection provided by wires and shielding. Personal/portable wireless networks present additional challenges due to the wide range of applications and use models that they must support. To name a few, eavesdroppers can overhear data

exchanges not intended for them, whereas imposters can send forged data not using its own identity, can replay previously transmitted data, and can transmit modified data captured from a previous transmission.

This Standard specifies two levels of security: no security and strong security protection. Security protection includes data encryption, message integrity, and replay attack protection. Secure frames are used to provide security protection to data and aggregated data frames as well as selected control and command frames.

Three security modes are defined to control the level of security for devices in their communications. This Standard allows for a device to use one of the two security levels or a combination of them in communicating with other devices by selecting the appropriate security mode.

This Standard further specifies a 4-way handshake mechanism to enable two devices to derive their pair-wise temporal keys (PTKs) while authenticating their identity to each other. A secure relationship is established following a successful 4-way handshake between two devices. A 4-way handshake between two devices is conducted based on a shared master key. How two devices obtain their shared master keys is outside the scope of this Standard.

In addition, this Standard provides means for the solicitation and distribution of group temporal keys (GTKs). While PTKs are used for protecting unicast frames exchanged between two devices, GTKs are employed for protecting multicast and broadcast frames transmitted from a source device to a multicast or broadcast group of recipient devices.

A pseudo-random function is defined based on the MIC generation by CCM using AES-128.

Secure frame counters and replay counters are set up on a per-temporal key basis to guarantee message freshness. No specific mechanisms are created in this Standard to address denial of service attacks given the open nature of the wireless medium.

In this Standard, 128-bit symmetric temporal keys are employed based on AES-128 with CCM to provide payload encryption and message integrity code (MIC) generation.

In general, this Standard specifies security mechanisms, not security policies.

6.6.10 Information discovery

The protocols and facilities of this Standard are supported by the exchange of information between devices. Information may be broadcasted in beacon frames or requested in Probe commands. For each type of information, an Information Element (IE) is defined. IEs may be included by a device in its beacon at any time and may optionally be requested or provided using the Probe command.

A device uses the MAC Capabilities IE and PHY Capabilities IE to announce information about its support of variable or optional facilities. Declaration of capabilities is especially useful when a device detects changes in its immediate neighbourhood.

6.6.11 Support for higher-layer timer synchronization

Some applications, for example, the transport and rendering of audio or video streams, require synchronization of timers located at different devices. Greater accuracy (in terms of jitter bounds) or finer timer granularity than that provided by the synchronization mechanism described in 7.4.

6.6.12 Protection of incumbent users

This Standard provides protocols for detecting incumbent users, protecting incumbent users and enabling seamless device operation, which include coordinated channel measurement, channel management, and TPC as specified in 7.7.3, 7.8 and 7.11.

6.6.13 Self-coexistence

This Standard provides self-coexistence protocol and mechanisms in 7.12 that allow neighbouring networks to share the same channel in the coordinated manner if desired. A master-slave based network can share the same channel with another co-located master-slave network or peer-to-peer network. Similarly, a peer-to-peer network can share the same channel with another peer-to-peer network or master-slave network.

6.6.14 Rate adaptation

A mechanism for data rate adaptation is provided in 7.7.2. A receiver may use this mechanism to inform a transmitter of the optimal data rate to increase throughput and/or reduce the frame error rate (FER).

6.6.15 Power management

A complementary goal of this Standard is to enable long operation time for battery powered devices. An effective method to extend battery life is to enable devices to turn off completely or reduce power for long periods of time, where a long period is relative to the superframe duration.

This Standard provides two power management modes in which a device can operate: active and hibernation. Devices in active mode transmit and receive beacons in every superframe. Devices in hibernation mode hibernate for multiple superframes and do not transmit or receive in those superframes.

In addition, this Standard provides facilities to support devices that sleep for portions of each superframe in order to save power.

To coordinate with neighbours, a device indicates its intention to hibernate by including a Hibernation Mode IE in its beacon. The Hibernation Mode IE specifies the number of superframes in which the device will sleep and will not send or receive beacons or any other frames.

Power management mechanisms are described in 7.9.

7 MAC common part sublayer

This clause specifies the format of MAC frames and MAC sublayer functionality.

The format of MAC frame which includes descriptions of common fields, each frame type and subtype, and information elements is specified in 7.1.

The rules for transmission and reception of MAC frames, including setting and processing MAC header fields and information elements, are specified in 7.2. The MAC superframe structure and beaconing policies are specified in 7.3.

During the data period devices send and receive data using prioritized contention access (PCA) or in reservations established using the channel reservation protocol (CRP), specified in 7.5. PCA permits multiple devices to contend for access to the medium based on traffic priority, and is specified in 7.5.1. The CRP enables a device to gain scheduled access to the medium within a negotiated reservation, and is specified in 7.5.2.

Device synchronization is specified in 7.4. The fragmentation and aggregation of MSDUs is specified in 7.6. 7.7 through 7.10 specify ARQ, probe commands, dynamic channel selection, multi-rate support, transmit power control, power management mechanisms. 7.11 specifies incumbent protections. 7.12 specifies self-coexistence. 7.13 specifies network entry and initialization. 7.14 specifies values for all MAC sublayer parameters.

7.1 MAC Frame Format

This Clause specifies the format of MAC frames. An overview of the MAC frame with descriptions of common fields is followed by Clauses for each frame type and subtype.

7.1.1 Frame format conventions

The following conventions and definitions apply throughout this clause.

7.1.1.1 Figures and Tables

MAC frames are described as a sequence of fields in a specific order. Tables in clause 7 depict fields in the order they are delivered to the PHY SAP from top to bottom, where the top-most field is transmitted first in time, as illustrated in Table 1. In field tables, bits within the field are numbered from the least-significant bit on the top to the most-significant bit on the bottom, as illustrated in Table 2.

Table 1 — Example sequence of fields

Syntax	Size	Notes
Example MAC frame Format {		
First field	2 bytes	First transmitted
Second field	2 bytes	Second transmitted
...
Last field	2 bytes	Last transmitted
}		

Table 2 — Example bitmap of a field

Syntax	Size	Notes
Example Field Format {		
Protocol Version	2 bits	Least significant bits, b1-b0
Secure	1 bit	
ACK Policy	2 bits	
Frame Type	3 bits	
Frame Subtype/Delivery ID	4 bits	
Retry	1 bit	
Reserved	3 bits	Most significant bits, b15-b13
}		

7.1.1.2 Octet order

Unless otherwise noted, fields longer than a single octet are delivered to the PHY SAP in order from the octet containing the least-significant bits to the octet containing the most-significant bits.

7.1.1.3 Encoding

Values specified in decimal are encoded in unsigned binary unless otherwise stated.

A bitmap is a sequence of bits, labeled as bit[0] through bit[N-1]. A bitmap is encoded in a field such that bit[0] corresponds to the least-significant bit of the field and subsequent bitmap elements correspond to subsequent significant bits of the field. Octets of the field are presented to the PHY SAP in order from least-significant octet to most-significant octet.

Reserved fields and subfields are set to ZERO on transmission and ignored on reception.

7.1.2 General MAC frame format

A MAC frame consists of a fixed-length MAC Header and an optional variable-length MAC Frame Body. The MAC Header is specified in Table 3.

Table 3 — MAC header format

Syntax	Size	Notes
MAC_Header_Format {		
Frame Control	16 bits	Refer to 7.1.2.1
DestAddr	16 bits	Destination Address
SrcAddr	16 bits	Source Address
Sequence Control	16 bits	The Sequence Control field identifies the order of MSDUs/MCDUs and their fragments; refer to 7.1.2.4
Access Control	16 bits	Access method, duration, and more frames flag; refer to 7.1.2.5
}		

The MAC frame body, when present, contains a Frame Payload and Frame Check Sequence (FCS) as shown in Table 4.

Table 4 — MAC frame body format

Syntax	Size	Notes
MAC_Frame_Body_Format {		
Frame Payload	<i>variable</i>	Refer to 7.1.2.6
FCS	32 bits	Refer to 7.1.2.7
}		

The Frame Payload length ranges from zero to mMaxFramePayloadSize. If the Frame Payload length is zero, the FCS field is not included, and there is no MAC Frame Body. The Frame Payload length includes the length of the security fields for a secure frame. In secure frames, the Frame Payload includes security fields, as specified in Table 5.

Table 5 — Frame Payload field format for secure frames

Syntax	Size	Notes
Secure_Frame_Payload_Field_Format {		
Temporal Identifier (TKID)	24 bits	Refer to 7.1.2.6.1
Secure Payload	<i>variable</i>	Refer to 7.1.2.6.2
Message Integrity Code (MIC)	64 bits	Refer to 7.1.2.6.3
}		

7.1.2.1 Frame Control

The Frame Control field is specified in Table 6.

Table 6 — Frame Control field format

Syntax	Size	Notes
Frame_Control_Field_Format {		
Protocol Version	2 bits	Refer to 7.1.2.1.1
Secure	1 bit	Refer to 7.1.2.1.2
ACK Policy	2 bits	Refer to 7.1.2.1.3
Frame Type	3 bits	Refer to 7.1.2.1.4
Frame Subtype/Delivery ID	4 bits	Refer to 7.1.2.1.5
Retry	1 bit	Refer to 7.1.2.1.6
Reserved	3 bits	
}		

7.1.2.1.1 Protocol Version

The Protocol Version field is invariant in size and placement across all revisions of this Standard. For this revision of the standard, the Protocol Version is set to zero. All other values are reserved.

7.1.2.1.2 Secure

The Secure bit is set to ONE in a secure frame, ZERO otherwise. Frames with the Secure bit set to ONE use the Frame Payload format for secure frames as specified in Table 5.

7.1.2.1.3 ACK Policy

The ACK Policy field is specified in Table 7.

Table 7 — ACK Policy field encoding

Value	ACK policy type	Description
0	No-ACK	The recipient(s) do not acknowledge the transmission, and the sender treats the transmission as successful without retransmission. The use of this policy is specified in 7.7.1.1.
1	Imm-ACK	The addressed recipient returns an Imm-ACK frame after correct reception, according to the procedures specified in 7.7.1.2.
2	B-ACK	The addressed recipient keeps track of the frames received with this policy until requested to respond with a B-ACK frame, according to the procedures specified in 7.7.1.3.
3	B-ACK Request	The addressed recipient returns a B-ACK frame after reception, according to the procedures specified in 7.7.1.3.

7.1.2.1.4 Frame Type

The Frame Type field is specified in Table 8.

Table 8 — Frame Type field encoding

Value	Frame type	Subclause
0	Beacon frame	Refer to 7.1.3
1	Control frame	Refer to 7.1.4
2	Command frame	Refer to 7.1.5
3	Data frame	Refer to 7.1.6
4	Aggregated data frame	Refer to 7.1.7
5-7	Reserved	

7.1.2.1.5 Frame Subtype/ Delivery ID

The Frame Subtype / Delivery ID field is used to assist a receiver in the proper processing of received frames. In beacon, control or command frames, this field is used as Frame Subtype, as specified in Table 14, Table 21 and Table 25, respectively. In data frames and aggregated data frames, this field is used as Delivery ID as specified in Table 9.

Table 9 — Delivery ID encoding in Frame Control

Bit b12	Bits b11-b9
0	User Priority
1	Stream Index

This field is reserved in all other frame types.

7.1.2.1.6 Retry

The Retry bit is set to ONE in any data, aggregated data, or command frame that is a retransmission of an earlier frame. It is reserved for all other frame types.

7.1.2.2 DestAddr

The DestAddr field is set to the DevAddr of the intended recipient(s) of the frame. The DevAddr specifies a single device for a unicast frame, a group of devices for a multicast frame, or all devices for a broadcast frame.

7.1.2.3 SrcAddr

The SrcAddr field is set to the DevAddr of the transmitter of the frame.

7.1.2.4 Sequence Control

Table 10 — Sequence Control field format

Syntax	Size	Notes
Sequence_Control_Field_Format {		
Fragment Number	3 bits	Refer to 7.1.2.4.1
Sequence Number	11 bits	Refer to 7.1.2.4.3
More Fragments	1 bit	Refer to 7.1.2.4.2
Reserved	1 bit	
}		

The Sequence Control field identifies the order of MSDUs/MCDUs and their fragments. The Sequence Control field is specified in Table 10. The Sequence Control field is reserved in control frames.

7.1.2.4.1 Fragment Number

The Fragment Number field is set to the number of the fragment within the MSDU or MCDU. The fragment number is set zero if the current fragment is the first fragment of an MSDU or MCDU, or if there is no fragmentation applied in the MSDU or MCDU. The fragment number is incremented by one for each successive fragment of that MSDU or MCDU.

7.1.2.4.2 More Fragments

The More Fragments field is set to ZERO if the current fragment is the final fragment of the current MSDU or MCDU, or if there is no fragmentation applied in the MSDU or MCDU. Otherwise, the field is set to ONE.

7.1.2.4.3 Sequence Number

The sequence number is derived from a modulo 2048 counter.

If frame type is beacon frame, the sequence number represents the current superframe number. Superframe number is set zero if the current superframe is the first superframe for the beacon group. The superframe number is incremented by one for each successive superframe. A device shall use a dedicated counter for beacon frame. In addition, all the devices in the same beacon group shall use the same superframe number.

If frame type is command frame, the Sequence Number is set to the sequence number of the MCDU. A device shall increment the sequence number by one for each transmitted MCDU. A device shall use a dedicated counter for MCDUs.

The Sequence Number field for MSDU and MCDU is specified in 7.2.9.3.

The Sequence Number field is reserved in control frames.

The Sequence Number field is used for duplicate frame detection, as described in 7.2.7, and to preserve frame order when using the B-ACK mechanism, as described in 7.7.1.3.

7.1.2.5 Access Control

The Access Control field is specified in Table 11.

Table 11 — Access Control field encoding

Syntax	Size	Notes
Access_Control_Field_Format {		
Duration	14 bits	Refer to 7.1.2.5.1
More frames	1 bit	Refer to 7.1.2.5.2
Access method	1 bit	Refer to 7.1.2.5.3
}		

7.1.2.5.1 Duration

The Duration field is 14 bits in length and is set to an expected medium busy interval after the end of the PLCP header of the current frame. The unit of duration is 4 microseconds. The duration value is set as specified in 7.2.9.1 and used to update the network allocation vector (NAV) according to the procedures specified in 7.5.1.2.

7.1.2.5.2 More frames

In frames sent with the Access Method bit set to ONE, the More Frames bit is set to ZERO if the transmitter will not send further frames to the same recipient during the current reservation block; otherwise it is set to ONE.

In frames sent with the Access Method bit set to ZERO, the More Frames bit is set to ZERO if the transmitter will not send further frames to the same recipient during the current superframe; otherwise it is set to ONE.

The More Frames bit is reserved in beacon and control frames. Additional rules regarding the More Frames field are specified in 7.2.9.2.

7.1.2.5.3 Access method

The Access Method bit shall be set to ONE in all frames transmitted within a hard or private CRP reservation block by the reservation owner or target prior to the release of the reservation block, including the UCA and UCR control frames that release the reservation block.

The Access Method bit may be set to ONE in frames transmitted within a Soft CRP reservation block without backoff by the reservation owner.

The Access Method bit in an Imm-ACK, B-ACK or CTS control frame is set to the same value as the Access Method bit in the corresponding received frame.

The Access Method bit is set to ZERO in frames transmitted at all other times, other than in beacon frames.

The Access Method bit is reserved in beacon frames.

7.1.2.6 Frame Payload

The Frame Payload field is a variable length field that carries the information. In a secure frame, it includes the required security fields as shown in Table 5 and specified below.

7.1.2.6.1 Temporal Key Identifier (TKID)

The TKID field is an identifier for the temporal key used to protect the frame. The TKID uniquely identifies this key from any other temporal keys held by the sender and the recipient(s) of the frame. It does not need to uniquely identify the key for devices not holding the key.

7.1.2.6.2 Secure Payload

The Secure Payload field in the secure frames is the counterpart of the Frame Payload field in non-secure frames. It contains the information specific to individual frame types and protected by the symmetric key identified in the TKID field of the same frame.

7.1.2.6.3 Message Integrity Code (MIC)

The MIC field contains an 8-octet cryptographic checksum used to protect the integrity of the MAC Header and Frame Payload.

7.1.2.7 FCS

The FCS field contains a 32-bit value that represents a CRC polynomial of degree 31.

The CRC is calculated over a calculation field, which is the entire Frame Payload field for this specification. The calculation field is mapped to a message polynomial $M(x)$ of degree $k-1$, where k is the number of bits in the calculation field. The least-significant bit of the first octet presented to the PHY SAP is the coefficient of the x^{k-1} term, and the most-significant bit of the last octet transmitted is the coefficient of the x^0 term.

The CRC is calculated using the following standard generator polynomial of degree 32:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x + 1$$

The CRC polynomial is the one's complement of the modulo 2 sum of the following remainders:

- The remainder resulting from $x^k \times (x^{31} + x^{30} + \dots + x + 1)$ divided (modulo 2) by $G(x)$.
- The remainder resulting from $x^{32} \times M(x)$, divided (modulo 2) by $G(x)$.

The FCS field value is derived from the CRC polynomial such that the least-significant bit is the coefficient of the x^{31} term and the most-significant bit is the coefficient of the x^0 term. Table 12 specifies the encoding of the FCS field for the CRC polynomial:

$$a_{31}x^{31} + a_{30}x^{30} + a_{29}x^{29} + \dots + a_2x^2 + a_1x + a_0$$

Table 12 — FCS field encoding

bits: b_{31}	b_{30}	b_{29}	...	b_2	b_1	b_0
a_0	a_1	a_2	...	a_{29}	a_{30}	a_{31}

In a common implementation, at the transmitter, the initial remainder of the division is preset to all ones and is then modified via division of the calculation field by the generator polynomial $G(x)$. The one's complement of this remainder is the FCS field. At the receiver, the initial remainder is preset to all ones. The serial incoming bits of the calculation field and FCS, when divided by $G(x)$ in the absence of transmission errors, results in a unique non-zero remainder value. The unique remainder value is the polynomial:

$$x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^8 + x^6 + x^5 + x^4 + x^3 + x + 1$$

7.1.3 Beacon frames

MAC Header field settings for beacon frames are described in Table 13. Beacon frames are also referred to as beacons throughout this specification.

Table 13 — MAC header field values for beacon frames

MAC header field	Value
Protocol version	0
Secure	0
ACK Policy	0 (No-ACK)
Frame Type	0 (beacon frame)
Frame Subtype	As specified in Table 14
Retry	Reserved
DestAddr	BcstAddr
SrcAddr	DevAddr of the transmitter
Sequence Control	As specified in 7.1.2.4 and 7.2.9.3. The sequence number contained in sequence control represents the superframe number
Duration	As specified in 7.1.2.5.1 and 7.2.9.1
More Frames	Reserved
Access Method	Reserved

Frame subtype specified in Table 14 indicates four beacon types: regular, signalling, echo beacon.

Table 14 — Frame subtype for beacon frames

Value	Beacon frame subtype	Description
0	Regular	Regular beaconing
1	Signalling	Signalling beacon for network entry
2	Echo	Help coverage extension or coexistence operation during master/slave mode
3 – 15	Reserved	Reserved

As specified in 7.1.2.4, the sequence number contained in sequence control represents the superframe number. The superframe number increments once per superframe, following a modulo counter. The device first establishing beacon group initializes the superframe number. Each device includes the superframe number in its beacon frame. Through beacon, a device knows the superframe number of its neighbours. Each device shall align its superframe number to the largest superframe number it observes in the beacon group. Therefore, all the devices in a network can quickly converge to the same superframe number even in the multi-hop case.

7.1.3.1 Regular beacon frame

The beacon frame payload is specified in Table 15.

Table 15 — Regular Beacon frame payload format

Syntax	Size	Notes
Regular_Beacon_Frame_Payload_Format {		
Device Identifier	6 bytes	
Beacon Slot Descriptor	1 byte	Refer to Table 16
Device type	1 byte	Refer to Table 17
For(i=1, i<=N, i++){		
IE _i	<i>variable</i>	Information element, as specified in 7.1.8
}		
}		

The information elements (IEs) are listed in Table 45 in 7.1.8. IEs are included in order of increasing Element ID. CRP IEs that have the same Target DevAddr and Stream Index are adjacent to each other in the beacon. Each beacon frame shall include at least the following IEs: BPOIE, CRP Availability IE, and Regular QP schedule IE.

The Device Identifier field is set to the EUI-48 [1] of the device sending the beacon.

The Beacon Slot Descriptor field format is specified in Table 16.

The Beacon Slot Number field is set to the number of the beacon slot where the beacon is sent within the beacon period (BP), in the range of [0, mMaxBPLength-1].

The Movable bit is set to one if the beacon is movable according to 7.3.6 and is set to ZERO otherwise.

Device type is specified in Table 17.

The operation mode field is set to 00 if the device is peer device, 01 if the device is master device, 10 if the device is nonbeaconing slave device, and 11 if the device is beaconing slave device.

The Security Mode field is set to the security mode at which the device is currently operating.

Table 16 — Beacon Slot Descriptor field format

Syntax	Size	Notes
Beacon_Slot_Descriptor_Field_Format {		
Reserved	2 bits	
Movable	1 bit	
Beacon Slot Num	5 bits	
}		

Table 17 — Device Type

Syntax	Size	Notes
Device_Type_Field_Format {		
Operation mode	2 bits	00: peer 01: master 10: nonbeaconing slave 11: beaconing slave
Security mode	2 bits	00: security mode 0 01: security mode 1 10: security mode 2 11: reserved Specified in 8
Reserved	4 bits	
}		

7.1.3.2 Signalling beacon frame

The signalling beacon frame payload is specified in Table 18. The signalling beacon is transmitted in the contention signalling window (CSW) only when a device performs network entry. More description of the use of the signalling beacon is specified in 7.3.

Table 18 — Signalling Beacon Frame Payload Format

Syntax	Size	Notes
Signalling_Beacon_Frame_Payload_Format {		
Device Identifier	6 bytes	
Device type	1 byte	As specified in Table 17
For(i=1, i<=N, i++){		
IE _i	<i>variable</i>	Information element, as specified in 7.1.8. Information elements are optional.
}		
}		

7.1.3.3 Echo beacon frame

During master/slave operation mode, to help beacon coverage extension or to collect coexistence information among the networks, echo beacons could be transmitted as a relay of the mater beacon. The echo beacon frame payload is specified in Table 19.

Table 19 — Echo Beacon Frame Payload Format

Syntax	Size	Notes
Echo_Beacon_Frame_Payload_Format {		
Device Identifier	6 bytes	Address of echo beacon transmitter
Master Identifier	6 bytes	Address of master device
Device type	1 byte	As defined in Table 17
For(i=1, i<=N, i++){		N is the number of including IEs
IE _i	<i>variable</i>	Information element, as defined in 7.1.8. Information elements are optional.
}		
}		

The Master Identifier field is set to the EUI-48 [1] of the master device sending the regular beacon. Others are the same as regular beacon frame payload format.

The size of echo beacon slot equals to one MAS.

7.1.4 Control frames

Default MAC Header field settings for control frames are listed in Table 20. Specific MAC Header field settings and payload descriptions for each of the control frames are shown in the following subclauses.

Table 20 — MAC Header field values for control frames

Header field	Value
Protocol Version	0
Secure	As specified in 7.1.2.1.2
ACK Policy	0 (No-ACK)
Frame Type	1 (control frame)
Frame Subtype	Value from Table 21
Retry	Reserved
DestAddr	DevAddr of the recipient
SrcAddr	DevAddr of the transmitter
Sequence Control	Reserved
Duration	As described in 7.1.2.5.1 and 7.2.9.1
More Frames	Reserved
Access Method	As described in 7.1.2.5.3

Table 21 lists valid values for the Frame Subtype field for control frames.

Table 21 — Frame Subtype field encoding for control frames

Value	Control frame subtype	Description
0	Imm-ACK	Acknowledges correct receipt of the previously-received frame
1	B-ACK	Acknowledges correct or incorrect receipt of one or more preceding frames
2	RTS	Announces to a recipient device that a frame is ready for transmission and requests confirmation of ability to receive
3	CTS	Responds to an RTS control frame that the recipient is able to receive
4	UCA	Announces to neighbours of the transmitting device that the remainder of a reservation block it owns is available for use by other devices via PCA
5	UCR	Announces to neighbours of the transmitting device that the remainder of a reservation block of which it is the target is available for use by other devices via PCA
6 – 13	Reserved	Reserved
14	Application-specific	At discretion of application owner
15	Reserved	Reserved

7.1.4.1 Immediate Acknowledgement (Imm-ACK)

In Imm-ACK frames, the DestAddr field is set to the SrcAddr of the received frame that is acknowledged. Imm-ACK frames have no frame payload.

7.1.4.2 Block Acknowledgement (B-ACK)

The B-ACK frame acknowledges correct or incorrect receipt of the previous sequence of frames and provides information for the transmission of the next sequence of frames as described in 7.7.1.3. The B-ACK frame payload is specified in Table 22.

In B-ACK frames, the DestAddr field is set to the SrcAddr of the frame that requested the B-ACK. The Buffer Size field specifies the maximum number of octets in the sum of the frame payloads of all frames in the next B-ACK sequence.

The Frame Count Field specifies the maximum number of frames in the next B-ACK sequence.

The Sequence Control and Frame Bitmap fields together specify an acknowledgement window of MSDU fragments and their reception status. The Sequence Control field specifies the Sequence Number and Fragment Number that start the acknowledgement window.

Table 22 — B-ACK frame payload

Syntax	Size	Notes
B-ACK_Frame_Payload_Format {		
Buffer Size	2 bytes	
Frame Count	1 byte	
Reserved	1 byte	
Sequence Control	2 bytes	
Frame Bitmap	variable	
}		

The Frame Bitmap field varies in length. A zero-length Frame Bitmap field indicates an acknowledgement window of length ZERO. Otherwise, the least-significant octet of the Frame Bitmap field corresponds to the MSDU indicated by the Sequence Control field, and each bit of the octet corresponds to a fragment of that MSDU. The least-significant bit in each octet corresponds to the first fragment and successive bits correspond to successive fragments. Successive octets present in the Frame Bitmap field correspond to successive MSDUs, and each bit corresponds to a fragment of the MSDU. The acknowledgement window ends at fragment seven of the MSDU that corresponds to the most-significant octet in the Frame Bitmap.

For all bits within the Frame Bitmap, a value of one indicates that the corresponding fragment was received in either the current sequence or an earlier one. A value of ZERO indicates that the corresponding fragment was not received in the current sequence (although it might have been received in an earlier one). Bits of the least-significant octet of the Frame Bitmap field corresponding to fragments prior to the start of the acknowledgement window are undefined. Frames with a Sequence Number earlier than the Sequence Number indicated in the Sequence Control field were not received in the last B-ACK sequence. Such frames were previously received or are no longer expected.

7.1.4.3 Request To Send (RTS)

In RTS frames, the DestAddr field is set to the DevAddr of the device to receive the following frame from the transmitter. RTS frames have no frame payload.

7.1.4.4 Clear To Send (CTS)

In CTS frames, the DestAddr field is set to the SrcAddr of the received RTS frame. CTS frames have no frame payload.

7.1.4.5 Unused CRP Reservation Announcement (UCA)

The UCA frame is used to explicitly release the remaining time of the current Hard or Private CRP reservation block. The DestAddr field is set to BcstAddr.

The UCA frame payload includes a list of DevAddrs of the devices that are expected to respond with a UCR frame, as shown in Table 23.

Table 23 — Payload format for UCA frames

Syntax	Size	Notes
UCA_Frame_Payload_Format {		
For (i=0; i<N; i++){		
DevAddr _i	2 bytes	
}		
}		

7.1.4.6 Unused CRP Reservation Response (UCR)

The UCR frame is used to respond to UCA frames to explicitly release the remaining time of the current Hard or Private CRP reservation block. The DestAddr field is set to the SrcAddr of the received UCA frame. UCR frames have no frame payload.

7.1.4.7 Application-specific

The payload format for Application-specific control frames is specified in Table 24.

Table 24 — Payload format for Application-specific control frames

Syntax	Size	Notes
Application_Specific_Control_Frame_Payload_Format {		
Specifier ID	2 bytes	Identification of a company or organization
Data		Format and use of the Data field defined by the owner of the Specifier ID
}		

The Specifier ID field is set to a 16-bit value that identifies a company or organization, as listed in [4]. The owner of the Specifier ID defines the format and use of the Data field.

7.1.5 Command frames

Default MAC Header settings for command frames are shown in Table 25.

Table 25 — Default MAC Header field values for command frames

Header field	Value
Protocol Version	0
Secure	As specified in 7.1.2.1.2
ACK Policy	0 (No-ACK) or 1 (Imm-ACK)
Frame Type	2 (command frame)
Frame Subtype	Value from Table 26
Retry	As specified in 7.1.2.1.6
DestAddr	DevAddr of the recipient
SrcAddr	DevAddr of the transmitter
Sequence Control	As specified in 7.1.2.4 and 7.2.9.3
Duration	As described in 7.1.2.5.1 and 7.2.9.1

Table 26 contains a list of valid values for the Frame Subtype field for command frames.

Table 26 — Frame Subtype field encoding for Command frames

Value	Command frame subtype	Description
0	CRP Reservation Request	Used to request creation or modification of a CRP reservation
1	CRP Reservation Response	Used to respond to a CRP reservation request command
2	Probe	Used to request for, or respond with, information elements
3	Pair-wise Temporal Key (PTK)	Used to derive a PTK via a 4-way handshake between two devices
4	Group Temporal Key (GTK)	Used to solicit or distribute a GTK within a secure relationship
5	Beaconing Promotion Request	Used from a nonbeaconing slave device to master to request being promoted as a regular beaconing slave device.
6	Channel Measurement Request	Used to request channel measurement by other devices
7	Channel Measurement Response	Used to respond to channel measurement request
8	Channel Measurement Report	Used to report channel measurement results
9	Channel Measurement Report Acknowledgement	Used to confirm the reception of channel measurement reports
10	Channel Switch Command	Used to request channel switch
11	Channel Switch Response	Used to respond to channel switch request
12	RSW Slot Request	Used by slave devices to request an RSW slot allocated by Master device. The command frame is transmitted in CSW.
13	Application-specific	At discretion of application owner
14	Association Request	Used by a slave device to request association with a master device or used by a peer device to request association with another peer device
15	Reserved	Reserved

7.1.5.1 CRP Reservation Request

The CRP Reservation Request command frame is used to create or modify a CRP reservation. The CRP Reservation Request command frame payload is specified in Table 27.

Table 27 — Payload format for CRP Reservation Request command frames

Syntax	Size	Notes
CRP_Reservation_Request_Frame_Payload_Format		
{		
For (i=0; i<N; i++){		
CRP IE _i	variable	
}		
}		

Each CRP IE field included in the command frame corresponds to a reservation request identified by the Owner DevAddr, Target DevAddr, Stream Index, and Reservation Type in the IE. The CRP IE is defined in Table 48.

7.1.5.2 CRP Reservation Response

The CRP Reservation Response command frame is used to respond to a CRP Reservation Request command frame. The CRP Reservation Response command frame payload is specified in Table 28.

Table 28 — Payload format for CRP Reservation Response command frames

Syntax	Size	Notes
CRP_Reservation_Response_Frame_Payload_Format {		
For (i=0; i<N; i++){		
CRP IE _i	variable	
}		
CRP Availability IE	2-34 bytes	
}		

The CRP Reservation Response command frame includes all the CRP IEs from the reservation request. The CRP Availability IE is included according to the rules specified in 7.5.2.

7.1.5.3 Probe

The Probe command frame is used to request information from a device or respond to a Probe request. The payload format is specified in Table 29.

Table 29 — Payload format for Probe command frames

Syntax	Size	Notes
Probe_Command_frame_Format {		
For(i=0; i < N; i++){		
IE _i	<i>variable</i>	
}		
}		

If the payload includes a Probe IE, the command requests information from the recipient. Each Information Element field contains one information element.

7.1.5.4 Channel Measurement Request

The Channel Measurement Request frame is used to request information from a device of channel measurement, including the presence of incumbents and alien devices. The payload format is specified in Table 30.

Table 30 — Format for Channel Measurement Request Frame Payload

Syntax	Size	Notes
Channel_Measurement_Request_Frame_Payload_Format {		
For(i=0; i < N; i++){		
IE _i	Variable	Channel Measurement Request IEs only, as defined in Table 59
}		
}		

7.1.5.5 Channel Measurement Response

Channel Measurement Response frame is used to confirm the reception of channel measurement request. The payload format is specified in Table 31.

Table 31 — Format for Channel Measurement Response Frame Payload

Syntax	Size	Notes
Channel_Measurement_Response_Frame_Payload_Format {		
For(i=0; i < N; i++){		
IE _i	Variable	Channel Measurement Response IEs only, as defined in Table 61
}		
}		

7.1.5.6 Channel Measurement Report

The Channel Measurement Report frame is used to report channel measurement results. The payload format is specified in Table 32.

Table 32 — Format for Channel Measurement Report Frame Payload

Syntax	Size	Notes
Channel_Measurement_Report_Frame_Payload_Format {		
For(i=0; i < N; i++){		
IE _i	Variable	Channel Measurement Report IE, as defined Table 62
}		
}		

7.1.5.7 Channel Measurement Report Acknowledgement

The Channel Measurement Report Acknowledgement frame is used to confirm the reception of channel measurement reports. The payload format is specified in Table 33.

Table 33 — Format for Channel Measurement Report Acknowledgement Frame Payload

Syntax	Size	Notes
Channel_Measurement_Report_Acknowledgement_Frame_Payload_Format {		
For(i=0; i < N; i++){		
IE _i	Variable	Channel Measurement Report Acknowledgement IE, as defined in Table 65
}		
}		

7.1.5.8 Channel Switch Command

The Channel Switch Command frame is used to request channel switch of other devices. The payload format is specified in Table 34.

Table 34 — Format for Channel Switch Command Frame Payload

Syntax	Size	Notes
Channel_Switch_Command_Payload_Format {		
For(i=0; i < N; i++){		
IE _i	Variable	Channel Change IE, as defined in Table 66
}		
}		

7.1.5.9 Channel Switch Response

The Channel Switch Response frame is used to respond to the Channel Switch Request command. The payload format is specified in Table 35.

Table 35 — Payload format for Channel Switch Response command frame

Syntax	Size	Notes
Channel_Switch_Command_Payload_Format {		
Channel Change Response IE	4 bytes	Refer to Table 68
}		

7.1.5.10 RSW Slot Request command frame

The command frame is used by nonbeaconing slave devices and sent in CSW slot. The payload of the RSW Slot Request command frame is specified in Table 36.

Table 36 — Payload format for RSW Slot Request command frame

Syntax	Size	Notes
RSW_Slot_Request_Payload_Format {		
Device Identifier	6 bytes	MAC address of current device
Reservation Target DevAddr	2 bytes	
}		

The Reservation Target DevAddr is set to the DevAddr of the reservation target of the requested RSW slot(s) which will be used for channel reservation negotiation. If the reservation target is Master device, it shall be set to Master DevAddr. If the reservation target is another slave device, it shall be set to the DevAddr of the slave device.

7.1.5.11 Beaconing Promotion Request

The Beaconing Promotion Request command frame is sent from a nonbeaconing slave device to its master to request to be promoted as a regular beaconing slave device. The format of payload is specified in Table 37.

Table 37 — Payload format for Beaconing Promotion Request frame

Syntax	Size	Notes
Beaconing_Promotion_Request_Payload_Format {		
Reason code	1 byte	0 – CRA by a slave device 1 – Self-coexistence 2 – unspecified 3 – 255 - reserved
}		

7.1.5.12 Association Request

The Association Request command frame is sent from a slave device to a master device, or from a peer device to another peer device, to request association. The format of payload is specified in Table 38.

Table 38 — Payload format for Association Request

Syntax	Size	Notes
Association_Request_Payload_Format {		
MAC address	6 bytes	The MAC address of the device requesting association
}		

7.1.5.13 Pairwise Temporal Key (PTK)

The PTK command frame is used in a 4-way handshake by a pair of devices, as described in 8.3.1, to authenticate each other and to derive a shared symmetric PTK for securing certain unicast traffic between the two devices. The PTK command frame is specified in Table 39.

Table 39 — Payload format for PTK command frame

Syntax	Size	Notes
PTK_Payload_Format {		
Message Number	1 byte	
Status Code	1 byte	
PTKID	3 bytes	
Reserved	11 bytes	
MKID	16 bytes	
I-Nonce / R-Nonce	16 bytes	
PTK MIC	8 bytes	
}		

The Message Number is set to 1, 2, 3, or 4, respectively, in the PTK command containing the first, second, third, or fourth message of the 4-way handshake. The other values of this field are reserved.

The Status Code in a PTK command indicates the current status of the 4-way handshake at the device sending this command. It is encoded as shown in Table 40.

Table 40 — Status Code field encoding in PTK commands

Value	Meaning
0	Normal—the 4-way handshake proceeds.
1	Aborted—the 4-way handshake is aborted per security policy.
2	Aborted—the 4-way handshake is aborted in order to yield to a concurrent 4-way handshake using the same master key.
3	PTKID not accepted—it is the TKID of a PTK or GTK being possessed by this device.
4 – 255	Reserved

The PTKID is set to a non-zero number as the TKID of the PTK to be derived from this 4-way handshake procedure. The initiator of the 4-way handshake chooses this value after determining that this value is different from the TKID of the PTK, if any, that is to be replaced by the new PTK, and the TKID of any PTK or GTK it currently possesses.

The MKID identifies the master key used in this 4-way handshake as described 8.3.1.

The I-Nonce/R-Nonce is a random number generated by the initiator or responder for this 4-way handshake. This field is set to I-Nonce, the random number generated by the initiator in the command containing a Message Number of 1 or 3, and is set to R-Nonce, the random number generated by the responder in the command containing a Message Number of 2 or 4.

The PTK MIC in the PTK command containing a Message Number of 1 is set to zero on transmission and is ignored on reception.

The PTK MIC in the PTK command containing a Message Number of 2, 3, or 4 is set to the MIC that protects the fields in the Frame Payload of this command using the PTK MIC key generated from the first two messages of the 4-way handshake as specified in 8.3.1.

The MAC Header for the PTK command frame is set as indicated in Table 25, with the ACK Policy set to Imm-ACK.

7.1.5.14 Group Temporal Key (GTK)

The GTK command frame is used to solicit or distribute a GTK following a PTK update. The GTK is used to secure certain multicast traffic from a sending device to a group of recipient devices, and is chosen by the sending device. The GTK command frame is always in secure form, and the Secure Payload field is specified in Table 41.

Table 41 — Payload format for GTK command frame

Syntax	Size	Notes
GTK Payload Format {		
Message Number	1 byte	
Status Code	1 byte	
GTKID	3 bytes	
Reserved	3 bytes	
GroupAddr	2 bytes	
GTK SFC	6 bytes	
GTK	16 bytes	
}		

The TKID identifies the PTK used to secure this frame and to generate the GTK being solicited or distributed.

The EO is set to zero, indicating that the secure payload which starts with the Message Number field is encrypted with the PTK indicated in the TKID field.

The SFN is set in the same way as for any secure frame. It is one plus the SFN value used in the last secure frame transmitted with the protection of the PTK indicated in the TKID field.

The Message Number is set to 0 in the GTK command transmitted by a multicast recipient device to solicit a new GTK from a multicast sender. The Message Number is set to 1 in the GTK command transmitted by a multicast sender to distribute a new GTK to a multicast recipient. The Message Number is set to 2 in the GTK command transmitted by a multicast recipient device to respond to the distribution of a new GTK command.

The Status Code in a GTK command indicates the current status of the GTK solicitation or distribution at the device sending this command. It is encoded as shown in Table 42.

Table 42 — Status Code field encoding in GTK commands

Value	Meaning
0	Normal—GTK solicitation or distribution proceeds.
1	Rejected—GTK solicitation or distribution is rejected per security policy.
2	GTKID not accepted—it is the TKID of a PTK or GTK being possessed by this device.
3 – 255	Reserved

The GTKID in the GTK command containing a Message Number of 0 is set to the TKID of the GTK being solicited. It is set to zero if the soliciting device does not know the TKID of the GTK it is soliciting.

The GTKID in the GTK command containing a Message Number of 1 is set to a non-zero number as the TKID of the GTK being distributed. The distributor chooses this value after determining that this value is different from the TKID of the GTK, if any, that is to be replaced by the new GTK, and the TKID of any PTK or GTK the distributor or recipient currently possesses.

The GTKID in the GTK command containing a Message Number of 2 is set to the GTKID in the last received GTK command containing a Message Number of 1.

The GroupAddr is set to the McstAddr or BcstAddr for which the GTK is being solicited or distributed. It is set to 0x0001 if the GTK is applied to all broadcast and multicast traffic from the device distributing this GTK.

The GTK SFC in the GTK command containing a Message Number of 0 is set to zero on transmission and ignored on reception.

The GTK SFC in the GTK command containing a Message Number of 1 is set to the current value of the secure frame counter set up for the GTK being distributed.

The GTK SFC in the GTK command containing a Message Number of 2 is set to the GTK SFC in the last received GTK command containing a Message Number of 1.

The GTK is the GTK distributed by the multicast sender for the McstAddr. In a GTK command soliciting a GTK, the GTK is set to zero prior to encryption.

The MIC is calculated as for any secure frame, using the PTK indicated in the TKID field.

The MAC Header for the GTK command frame is set as indicated in Table 25, with the ACK Policy set to Imm-ACK.

7.1.5.15 Application-specific

The payload format for Application-specific command frames is specified in Table 43.

Table 43 — Payload format for Application-specific command frame

Syntax	Size	Notes
Application-specific_Command_Payload_Format {		
Specifier ID	2 bytes	
Data	variable	
}		

The Specifier ID field is set to a 16-bit value that identifies a company or organization, as listed in [4]. The owner of the Specifier ID defines the format and use of the Data field.

7.1.6 Data frames

MAC Header and Frame Payload fields for data frames are set as described in 7.1.2.

7.1.7 Aggregated data frames

MAC header and frame payload field format for aggregated data frames is defined in 7.6.2.

7.1.8 Information elements

This subclause defines the information elements (IEs) that may appear in beacons and certain command frames.

The general format of all IEs is specified in Table 44.

Table 44 — General IE format

Syntax	Size	Notes
General IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length	1 byte	Set to the length, in bytes, of the IE-specific fields
IE-specific fields	As indicated in Length field	Contain information specific to the IE
}		

Table Information elements contains a list of IEs defined in this draft.

Table 45 — Information elements

Element ID	Information element	Description
0	Traffic Indication Map (TIM) IE	Indicates that a device has data buffered for transmission via PCA
1	Beacon Period Occupancy IE (BPOIE)	Provides information on neighbours' BP occupancy in the previous superframe
2	PCA Availability IE	Indicates the MASSs that a device is available to receive PCA frames and transmit the required response
3 – 7	Reserved	Reserved
8	CRP Availability IE	Indicates a device's availability for new CRP reservations
9	Channel Reservation Protocol (CRP) IE	Indicates a reservation with another device
10	Hibernation Mode IE	Indicates the device will go to hibernation mode for one or more superframes but intends to wake at a specified time in the future
11	BP Switch IE	Indicates the device will change its BPST at a specified future time
12	MAC Capabilities IE	Indicates which MAC capabilities a device supports
13	PHY Capabilities IE	Indicates which PHY Capabilities a device supports
14	Probe IE	Indicates a device is requesting one or more IEs from another device or/and responding with requested IEs
15	Application-specific Probe IE	Indicates a device is requesting an Application-specific IE from another device
16	Link Feedback IE	Provides data rate and power control feedback
17	Hibernation Anchor IE	Provides information on devices in hibernation mode

18	Identification IE	Provides identifying information about the device, including a name string.
19	Master Key Identifier (MKID) IE	Identifies some or all of the master keys held by the transmitting device
20	Relinquish Request IE	Indicates that a neighbour requests that a device release one or more MASs from its reservations.
21	Multicast Address Binding (MAB) IE	Indicates an address binding between a multicast EUI-48 and a McstAddr
22	Regular QP Schedule IE	Indicates the regular quiet period schedule
23	On-demand QP Schedule IE	Indicates the on-demand quiet period schedule
24	Channel Measurement Request IE	Asks a device to measure an outband channel
25	Channel Measurement Response IE	Confirms receiving Channel Measurement Request IE
26	Channel Measurement Report IE	Report Channel Measurement Report to other devices
27	Channel Measurement Report Acknowledgement IE	Acknowledge the reception of Channel Measurement Report IE
28	On-leave IE	Announces the temporary on-leave schedule
29	Proxy Assignment IE	Assigns the next proxy
30	Channel Change IE	Advise a group of device to move to a new available channel
31	Slave device list IE	Used by a master to indicate the list of slave devices associated with itself
32	Channel change response IE	Respond to channel switch request
33	Channel Classification IE	
34	RSW Schedule IE	Indicates the RSW schedule
35	Echo Beacon Position IE	Indicates both MAS position of echo beacon and transmission owner of echo beacon
36	Beaconing Promotion Indication IE	Indicate nonbeaconing slave devices to become regular beaconing slave devices
37	Association Response IE	Used by a device to confirm the association request by another device
38	On-demand QP Negotiation Request IE	To request On-demand QP
39	On-demand QP Negotiation Response IE	To respond to On-demand QP request
40	Device Feature IE	Provides information of device's feature
41	Link Quality Estimate IE	Report link quality estimate
42	Transmit Power Control IE	Provide the recommended transmit power for other devices
43	Disassociation IE	Used by a master or peer device to disassociate a group of slave devices or peer devices, respectively
44 – 254	Reserved	Reserved
255	Application-Specific IE (ASIE)	Use varies depending on the application

7.1.8.1 Beacon Period Occupancy IE (BPOIE)

A device shall always include a BPOIE in its beacon. In the BPOIE the device shall reflect beacons received from neighbours in the recent superframes, as well as information retained based on hibernation mode rules.

The BPOIE provides information on the BP observed by the device sending the IE. The BPOIE is specified in Table 46.

Table 46 — BPOIE format

Syntax	Size	Notes
BPOIE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+K+2×N)	1 byte	N is the total number of devices that occupy beacon slot(s).
BP Length	1 byte	For a master device, BP length ≥ 1; For a peer device, BP length ≥ 2.
Beacon Slot Info Bitmap	K bytes	K = Ceiling (BP_Length/4)
For(i=1; i≤N; i++){		
DevAddr _i	2 bytes	The device address of the slot owner
}		
}		

The BP Length field is set to the length of the BP, measured in beacon slots, as defined in 7.3.3.

The Beacon Slot Info Bitmap field consists of K octets of 2-bit elements to indicate the beacon slot occupancy and movability in the BP, where K = Ceiling (BP_Length/4). Each element n, numbered from 0 to 4×K-1, corresponds to beacon slot n and is encoded as shown in Table 47. Element zero is the least-significant two bits of the field. Unused elements, if any, are set to zero.

Table 47 — Beacon Slot Info Bitmap element encoding

Element value	Beacon slot status
0	Unoccupied (non-movable) No PHY indication of medium activity was received in the corresponding beacon slot in the last superframe, or any frame header received with a valid HEI was not a beacon frame.
1	Occupied & non-movable A beacon frame was received with a valid HEI and FCS in the corresponding beacon slot in the last superframe, and the Movable bit in that beacon was set to zero, or a beacon frame was received in the corresponding beacon slot in a previous superframe that indicated a hibernation period that has not expired, as described in 7.9.
2	Occupied & movable A PHY indication of medium activity was received in the corresponding beacon slot in the last superframe, but did not result in reception of a frame with valid HEI and FCS.
3	Occupied & movable A beacon frame was received with a valid HEI and FCS in the corresponding beacon slot in the last superframe, and the Movable bit in that beacon was set to one.

The DevAddr fields correspond to beacon slots encoded as occupied in the Beacon Slot Info Bitmap. They are included in ascending beacon slot order. If a beacon was received with a valid HEI at a beacon slot in the last superframe, the corresponding DevAddr field is set to the SrcAddr in the MAC header of that received beacon. If a frame was received with an invalid HEI from a beacon slot in the last superframe, the corresponding DevAddr field is set to BcstAddr. If a neighbour of the device is in hibernation mode, the DevAddr field that corresponds to the hibernating neighbour's beacon slot is set to the DevAddr of that neighbour.

7.1.8.2 Channel Reservation Protocol (CRP) IE

A CRP IE is used to negotiate a reservation or part of a reservation for certain MASs and to announce the reserved MASs. The CRP IE is specified in Table 48.

Table 48 — CRP IE format

Syntax	Size	Notes
CRP_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 6+4×N)	1 byte	
CRP Control	2 bytes	Defined in Table 49
Owner DevAddr	2 bytes	DevAddr of reservation owner
Target DevAddr	2 bytes	DevAddr of reservation target
For(i=0; i < N; i++){		
CRP Allocation	4 bytes	Defined in Table 52
}		
}		

The CRP Control field is specified in Table 49.

Table 49 — CRP Control field format

Syntax	Size	Notes
CRP_Control_field {		
Reservation Type	3 bits	
Stream Index	3 bits	
Reason Code	3 bits	
Reservation Status	1 bit	
Conflict Tie-breaker	1 bit	
Unsafe	1 bit	
Reserved	4 bits	
}		

The Reservation Type field is set to the type of the reservation and is encoded as shown in Table 50.

Table 50 — Reservation Type field encoding

Value	Reservation Type
0	Alien BP
1	Hard
2	Soft
3	Private
4	PCA
5 – 7	Reserved

The Stream Index field identifies the stream of data to be sent in the reservation. This field is reserved if the Reservation Type is PCA and Alien BP.

The Reason Code is used by a reservation target to indicate whether a CRP reservation request was successful and is encoded as shown in Table 51. The Reason Code is set to ZERO in a CRP IE sent during negotiation by a reservation owner and by a device maintaining an established reservation. The Reason Code is set to Modified by a device if some of the MASs claimed in the reservation have been removed or if CRP IEs have been combined. The field is reserved if the Reservation Type is PCA.

Table 51 — Reason Code field encoding

Value	Code	Meaning
0	Accepted	The CRP reservation request is granted
1	Conflict	The CRP reservation request or existing reservation is in conflict with one or more existing CRP reservations
2	Pending	The CRP reservation request is being processed
3	Denied	The CRP reservation request is rejected or existing CRP reservation can no longer be accepted
4	Modified	The CRP reservation is still maintained but has been reduced in size or multiple CRP IEs for the same reservation have been combined
5 – 7	Reserved	Reserved

The Reservation Status bit indicates the status of the CRP negotiation process. The Reservation Status bit is set to ZERO in a CRP IE for a reservation that is under negotiation or in conflict. It is set to ONE by a device granting or maintaining a reservation, which is then referred to as an established reservation. The bit is set to ONE if Reservation Type is Alien BP or PCA.

The Conflict Tie-breaker bit is set to a random value of ZERO or ONE when a reservation request is made. The same value selected is used as long as the reservation is in effect. For all CRP IEs that represent the same reservation, the Conflict Tie-breaker bit is set to the same value.

The Owner DevAddr field and Target DevAddr field are set to the DevAddr of the reservation owner and the DevAddr of the reservation target, respectively. These two fields are reserved if the Reservation Type is PCA.

The Unsafe bit is set to ONE if any of the MASs identified in the CRP Allocation fields is considered in excess of reservation limits.

A CRP IE contains one or more CRP Allocation fields. Each CRP Allocation field is encoded using a zone structure. The superframe is split into 16 zones numbered from 0 to 15 starting from the BPST. Each zone contains 16 consecutive MASs, which are numbered from 0 to 15 within the zone.

The format of a CRP Allocation field is specified in Table 52.

Table 52 — CRP Allocation field format

Syntax	Size	Notes
CRP_Allocation_field {		
Zone Bitmap	2 bytes	
MAS Bitmap	2 bytes	
}		

The Zone Bitmap field identifies the zones that contain reserved MASs. If a bit in the field is set to ONE, the corresponding zone contains reserved MASs, where bit zero corresponds to zone zero.

The MAS Bitmap specifies which MASs in the zones identified by the Zone Bitmap field are part of the reservation. If a bit in the field is set to ONE, the corresponding MAS within each zone identified by the Zone Bitmap is included in the reservation, where bit zero corresponds to MAS zero within the zone.

7.1.8.3 CRP Availability IE

The CRP Availability IE is used by a device to indicate its view of the current utilization of MASs in the current superframe.

The CRP Availability IE is specified in Table 53.

Table 53 — CRP Availability IE format

Syntax	Size	Notes
CRP_Availability_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= N)	1 byte	
CRP Availability Bitmap	N (0-32)	variable
}		

The CRP Availability Bitmap field is up to 256 bits long, one bit for each MAS in the superframe, where the least-significant bit of the field corresponds to the first MAS in the superframe and successive bits correspond to successive MASs. Each bit is set to ONE if the device is available for a CRP reservation in the corresponding MAS, or is set to ZERO otherwise. If the CRP Availability Bitmap field is smaller than 32 octets, the bits in octets not included at the end of the bitmap are treated as zero.

7.1.8.4 Link Feedback IE

The Link Feedback IE contains information on the recommended change to the data rate and transmit power level by a recipient device for one or more source devices. The Link Feedback IE is specified in Table 54.

Table 54 — Link Feedback IE format

Syntax	Size	Notes
Link_Feedback_IE_Format {		
Element ID	1 byte	
Length (= 1+3xN)	1 byte	
N	1 byte	Num of Link fields
For (i=1, i <=N, i++) {		
Link _i	3 bytes	Defined in Table 55
}		
}		

The Link field is specified in Table 55.

Table 55 — Link field format

Syntax	Size	Notes
Link_Field_Format {		
DevAddr	16 bits	
Transmit Power Level Change	4 bits	Specified in Table 56
Data rate	4 bits	Specified in Table 137
}		

The DevAddr field is set to the DevAddr of the source device for which the feedback is provided.

The Transmit Power Level Change field is set to the change in transmit power level that the recipient recommends to the source device. The Transmit Power Level Change field encoding is shown in Table 56.

Table 56 —Transmit Power Level Change field encoding

Value	dB
1000 – 1101	Reserved
1110	-2
1111	-1
0000	no change
0001	+1
0010	+2
0011 – 0111	Reserved

The Data Rate field is set to the data rate that the recipient device recommends that the source device use. The Data Rate field is encoded as shown in Table 137.

7.1.8.5 Regular QP Schedule IE

Regular QP Schedule IE, defined in Table 57, shall be included in each beacon frame in order to make sure every device in the beacon group knows the sensing schedule. All the devices in the beacon group should participate in service monitoring.

Table 57 — Regular QP Schedule IE format

Syntax	Size	Notes
Operating_Channel_Monitor_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (=3)	1 byte	
Countdown	1 byte	In the number of superframe. If it equals zero, the regular QP is scheduled in the current superframe.
Sensing cycle	1 byte	In the unit of superframe. Set as mQPfrequency
QP duration	1 byte	In the unit of MAS. Set as mQPD
}		

NOTE if sensing cycle is set as zero, no regular QP is scheduled.

7.1.8.6 On-demand QP Schedule IE

A device may schedule QP on demand as described in 7.11.1.1.2. An on-demand QP schedule IE should be included in beacons. Other devices hearing on-demand QP schedule IE shall keep quiet during QP and may participate in sensing. The format of on-demand QP schedule IE is defined in Table 58.

Table 58 — On-demand QP Schedule IE format

Syntax	Size	Notes
Operating_Channel_Monitor_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3)	1 byte	
Superframe count down	1 byte	The number of superframe left before On-demand QP is enabled. When the value is set to 0, On-demand QP is enabled in current superframe.
QP starting time	1 byte	Starting MAS to perform monitoring
Duration	1 byte	Channel monitoring period, in the unit of MAS
}		

7.1.8.7 Channel Measurement Request IE

A peer device or a master device may ask another device to measure a channel for the detection of incumbents or alien devices. The device shall include channel measurement request IE, defined in Table 59, in its beacon frame or corresponding command frame for such request.

In addition, a peer device or a master device may ask another device to check the availability of a outband channel for backup purpose by using this IE.

Table 59 — Channel Measurement Request IE

Syntax	Size	Notes
Channel_Measurement_Request_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4+2xN)	1 byte	N is the total number of devices invited to monitor outband channel
Channel number	1 byte	Target Channel for measurement. Refer to 9.7.2.
Monitor countdown	1 byte	The starting MAS to perform monitoring
Action code	1 byte	Measurement requirements, as defined in Table 60
Duration	1 byte	The channel monitoring period, in the unit of MAS
For(i=0; i<N; i++){		
DevAddr _i	2 bytes	Invited device
}		
}		

Table 60 — Field Format for Action Code of Channel Measurement

Syntax	Size	Notes
Action_Code_Field_Format {	bit	
TV	1 bit	1 – TV detection required; 0 otherwise
WM	1 bit	1 – WM detection required; 0 otherwise
Alien device	1 bit	1 – Alien device detection required; 0 otherwise
Measured signal strength	1 bit	1 – Need to report measured signal strength; 0 otherwise
Reserved	4 bits	
}		

7.1.8.8 Channel Measurement Response IE

Channel measurement response IE, defined in Table 61, is used to confirm the reception of channel measurement request. Channel measurement response IE may be included in beacon or channel measurement response command frame.

Table 61 — Channel Measurement Response IE

Syntax	Size	Notes
Channel_Measurement_Response_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 5)	1 byte	
DevAddr	2 bytes	The address of the target receiving device of this IE; in other words, the address of the device which requests channel measurement.
Confirmation Code	1 byte	0 – OK; others, reserved
}		

7.1.8.9 Channel Measurement Report IE

Channel measurement report IE, defined in Table 62, is used to report channel measurement results. A device shall report channel measurement results upon receiving channel measurement request. A device may also report channel measurement results once detecting incumbent signal or other interfering sources. Channel measurement report IE may be included in beacon or channel measurement report command frame.

Table 62 — Channel Measurement Report IE

Syntax	Size	Notes
Channel_Measurement_Report_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length	1 byte	$10 + \sum_{i=1}^N (m_i + 1)$ if source type is Alien device; 3 otherwise
Channel number	1 byte	Refer to 9.7.2
Occupancy	1 byte	0: no, 1: TV, 2: WM, 3: Alien device, 4: Unknown, 5-255: reserved
Measured signal strength	1 byte	Defined in Table 63
If (Occupancy = Alien device){		
BPST offset	2 bytes	In the unit of microsecond
Regular QP description	4 bytes	Refer to Table 64
Channel business ratio	1 byte	0-100, in percentage
For(i=0; i<N; i++){		N is the number of alien devices reported
Length (= m _i)	1 byte	
Name String	m _i bytes	Device name string, encoded in Unicode UTF-16LE format, and is specified in Table 77
}/(* for N*)/		
}/(* if Alien device*)/		
}		

The encoding of the received signal strength is specified in Table 63.

Table 63 — Encoding of the received signal strength

Value	Description
0000 0000	-130 dBm or lower
0000 0001	-129 dBm
0000 0010	-128 dBm
...	...
0110 1101	-21 dBm
0110 1110	-20 dBm
0110 1111	-19dBm or higher
0111 0000-11111111	Reserved

The format of Regular QP description field is specified in Table 64.

Table 64 — Regular QP description field format

Syntax	Size	Notes
Regular_QP_description_field_Format {		
QP offset	1 byte	In the unit of superframe. BPST offset + QP offset equals the time difference between the next QP of current beacon group and the next QP of alien devices being reported
QP duration	1 byte	In the unit of MAS
Sensing cycle	1 byte	In the unit of superframe
}		

7.1.8.10 Channel Measurement Report Acknowledgement IE

Channel measurement report acknowledgement IE, defined in Table 65, is used to confirm the reception of channel measurement reports. Channel measurement acknowledgement IE may be included in beacon or channel measurement report acknowledgement command frame.

Table 65 — Channel Measurement Report Acknowledgement IE

Syntax	Size	Notes
Channel_Measurement_Report_Acknowledgement_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4)	1 byte	
DevAddr	2 bytes	The address of the target receiving device of this IE; in other words, the address of the device which reports channel measurement
}		

7.1.8.11 Channel Change IE

Upon detecting incumbent or for the purpose of load balancing, a device is advised to switch to a new available channel. The device shall include Channel Change IE in its beacon before Channel Change.

The Channel Change IE is specified in Table 66.

Table 66 — Channel Change IE format

Syntax	Size	Notes
Channel_Change_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3+2×N)	1 byte	N is the total number of devices advised to change channel
Channel Change Instruction	1 byte	Defined in Table 67
New channel number	1 byte	Set to the backup channel number if available; 255 (Unavailable) otherwise. Refer to 9.7.2.
Channel change countdown	1 byte	The time, e.g., in the unit of the number of remaining superframes, left before channel change
For(i=0; i<N; i++){		Device list advised to change channel
DevAddr	2 bytes	If set as Broadcast address, it indicates every device to change channel
}		
}		

The Channel Change Countdown field is set to the time, e.g., in the unit of the number of superframes remaining until the device changes to the new channel.

The device list includes all the devices which are advised to change to the same new channel.

The Channel Change Instruction field encoding is specified in Table 67. If reason field is set to 0, namely due to incumbent detected, all the devices belonging to same group shall evacuate the current channel and change to the new channel as suggested. If reason encoding is set 1, namely due to the requirement of load balancing, only the devices in the device list are suggested to change to the new channel; other devices receiving the channel change IE could still stay in the current channel. Moreover, if channel copy to resume field is set to 0, the same channel reservation used in old channel is applied to the new channel. This action is enabled only if a backup channel is available and the new channel is clean.

Table 67 — Channel Change Instruction Encoding for Channel Change

Syntax	Size	Notes
Channel_Change_Instruction_Format {		
Reason	1 bit	0 – Incumbent detected 1 – Load balancing
Channel copy to resume	1 bit	0 – Disabled 1 – Enabled. Apply same channel reservation as that of old channel in new channel
Reserved	6 bits	
}		

7.1.8.12 Channel Change Response IE

Table 68 — Channel Change Response IE format

Syntax	Size	Notes
Channel_Change_Response_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 2)	1 byte	
New channel number	1 byte	Set to the backup channel number if available; 255 (Unavailable) otherwise. Refer to 9.7.2.
Channel change countdown	1 byte	The time, e.g., in the unit of the number of remaining superframes, left before channel change
}		

7.1.8.13 On-leave IE

A device needs to visit other channel(s) or enter power saving mode for certain period for some reasons. A device shall include On-leave schedule IE, defined in Table 69, in its beacon before leaving the channel.

Table 69 — On-leave IE format

Syntax	Size	Notes
On_Leave_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4)	1 byte	
On-leave countdown	2 bytes	Time to leave in the unit of 10*microseconds
Duration	2 bytes	On-leave duration in unit of 10*microseconds
}		

7.1.8.14 Proxy Assignment IE

A peer device which is serving as current proxy may appoint the next proxy to coordinate outband channel measurement, by including Proxy Assignment IE, defined in Table 70, in its beacon.

Table 70 — Proxy Assignment IE format

Syntax	Size	Notes
Proxy_Assignment_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4)	1 byte	
DevAddr	2 bytes	Address of the designated device
Proxy Assignment Countdown	1 byte	The number of superframes remaining before the designated device takes over the proxy role
Duration	1 byte	The number of superframes for acting as proxy
}		

The Proxy Assignment Countdown field is set to the number of superframes remaining until the device takes over the proxy role. If this field is zero, the device performs the proxy role immediately.

7.1.8.15 BP switch IE

The BP Switch IE indicates a device will change its BPST to align with an alien BP. It is specified in Table 71.

Table 71 — BP Switch IE format

Syntax	Size	Notes
BP_Switch_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4)	1 byte	
BP Move Countdown	1 byte	
Beacon Slot Offset	1 byte	
BPST Offset	2 bytes	
}		

The BP Move Countdown field is set to the number of superframes after which the device will adjust its BPST. If BP Move Countdown is zero, the next beacon frame transmitted will be at the time specified by this IE.

The Beacon Slot Offset field is set to a positive number by which the device will adjust its beacon slot number when changing its BPST or is set to zero to indicate the device will join the alien BP using normal BP join rules.

The BPST Offset field is set to the positive amount of time the device will delay its BPST, in microseconds.

7.1.8.16 Identification IE

The Identification IE provides identifying information about the device, including a name string. The Identification IE is specified in Table 72.

Table 72 — Identification IE format

Syntax	Size	Notes
Identification_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (=)	variable	
For (i=0;i<N;i++){		N is the number of Identification fields
Identification _i	variable	
}		
}		

The general format of the Identification field is specified in Table 73.

Table 73 — Identification Field Format

Syntax	Size	Notes
Identification_Field_Format {		
Identification Information Type	1 byte	
Identification Information Length (=M)	1 byte	M equals the size of Identification Information Data
Identification Information Data	variable	
}		

The encoding for the Identification Information Type field is shown in Table 74.

Table 74 — Identification Information Type field encoding

Value	Identification Information Data field contents
0	Vendor ID
1	Vendor Type
2	Name String
3 – 255	Reserved

The Identification Information Length field indicates the length, in octets, of the Identification Information Data Field that follows.

The Identification Information Data field, if Identification Information Type is Vendor ID, is specified in Table 75.

Table 75 — Identification Information Data field format for Vendor ID

Syntax	Size	Notes
Vendor_ID_Format {		
Vendor ID	3 bytes	
}		

The Vendor ID is set to an OUI that indicates the vendor of the device. The OUI is a sequence of 3 octets, labeled as oui[0] through oui[2]. Octets of the OUI are passed to the PHY SAP in ascending index-value order.

The Identification Information Data field, if Identification Information Type is Vendor Type, is specified in Table 76.

Table 76 — Identification Information Data field format for Vendor Type

Syntax	Size	Notes
Vendor_Type_Format {		
Vendor ID	3 bytes	
Device Type ID	3 bytes	
}		

The Vendor ID field is set to an OUI that indicates the entity that assigns the values used in the Device Type ID field. The Device Type ID field indicates the type of device.

The Identification Information Data field, if Identification Information Type is Name String, contains the name of the device encoded in Unicode UTF-16LE format, and is specified in Table 77.

Table 77 — Identification Information Data field format for Name String

Syntax	Size	Notes
Name_String_Format {		
For (i=0; i<N; i++){		N is the number of Name String Unicode Char
Name String Unicode Char _i	2 bytes	
}		
}		

7.1.8.17 Channel Classification IE

In a master-slave network, the master device is responsible for selecting the operating channel and assigning it to the MAC/PHY modules. In a peer-to-peer network, a peer device is responsible for selecting its own operating channel and assigning it to the MAC/PHY modules. Channels are classified as six types -

Disallowed, Operating, Backup, Candidate, Protected and Unclassified. Channel Classification IE is used to exchange channel status among devices and specified in Table 78.

Table 78 — Channel Classification IE

Syntax	Size	Notes
Channel_Measurement_Acknowledgement_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+2xN+M)	1 byte	N is the total number of channels classified in this IE
M	1 byte	M is the total number of backup channels plus the operating channel
For (i=1; i<=N, i++){		
Channel Num	1 byte	
Channel description	1 byte	Defined in Table 79
}		
For (i=1; i<=M, i++){		The order of Transmit Power Limit fields is in the order of descending preference. First field of Transmit Power Limit corresponds to the operating channel; the last field of Transmit Power Limit corresponds to the least preferred backup channel.
Transmit Power Limit	1 byte	Defined in Table 80
}		
}		

The Channel Description Field is specified in Table 79.

Table 79 — Channel Description Field Format

Syntax	Size	Notes
Channel_Description_Format {		
Channel type	3 bits	000- Operating 001- Backup 010- Candidate 011- Protected – TV 100- Protected - WM 101- Disallowed 110- Unclassified 111- reserved
Channel preference	5 bits	00000 represents highest preference; 11111 represents least preferred
}		

The Transmit Power Limit field encoding is specified in Table 80. The Transmit Power Limit might be determined by regulation.

Table 80 — Transmit Power Limit field encoding

Value	mW
0	1
1	2
2	3
...	...
98	99
99	100
100-255	reserved

7.1.8.18 RSW Schedule IE

A device may schedule RSW (Reservation based Signalling Windows) as described in 7.3.8.2. An RSW schedule IE only indicates the location of RSW. The usage of each RSW MAS is indicated by CRP IE. The format of RSW schedule IE is defined in Table 81.

Table 81 — RSW Schedule IE format

Syntax	Size	Notes
RSW_Schedule_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 2)	1 byte	
RSW starting time	1 byte	Starting MAS of RSW
Duration	1 byte	in the unit of MAS
}		

7.1.8.19 Echo Beacon Position IE

Echo beacon Position IE shall be included in the Regular Beacon Frame of a master. Echo beacon position IE indicates both MAS position of echo beacon and transmission owner of echo beacon. The Echo Beacon Position IE is specified in Table 82.

Table 82 — Echo Beacon Position IE format

Syntax	Size	Notes
Echo_Beacon_Position_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+3×N)	1 byte	N is the total number of echo beacon slot
N	1 byte	Number of echo beacon sin a super-frame
For(i=0; i<N; i++){		
DevAddr	2 bytes	The device address of the slot owner to transmit echo beacons
MAS position	1byte	Position of Echo beacon in MAS unit. If this field is encoded by '4', then 4-th MAS is allocated for transmitting of echo beacon by slot owner
}		
}		

7.1.8.20 PHY capabilities IE

The PHY Capabilities IE is specified in Table 83.

Table 83 — PHY Capabilities IE format

Syntax	Size	Notes
PHY_Capabilities_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3)	1 byte	
PHY Capability Bitmap	2 bytes	Refer to Table 84
Reserved	1 byte	
}		
}		

The PHY Capability Bitmap field indicates capabilities supported by the PHY, as specified in the Physical Layer specification. A bit is set to ONE if the corresponding attribute is supported, or is set to ZERO otherwise. This field is encoded as described in Table 84. Subsequent octets are reserved.

Table 84 — PHY Capability Bitmap

Octet	Bit	Attribute	Description
0	0	Multiple antennae – FITD TxRx	Support FITD at both transmitter and receiver sides
	1	Multiple antennae – STBC Tx	Support STBC at transmitter side
	2	Multiple antennae – SM Tx	Support SM at transmitter side
	3	Multiple antennae – STBC Rx	Support STBC at receiver side
	4	Multiple antennae – SM Rx	Support SM at receiver side
	5	Bandwidth (6 MHz)	Support 6 MHz bandwidth
	6	Bandwidth (7 MHz)	Support 7 MHz bandwidth
	7	Bandwidth (8 MHz)	Support 8 MHz bandwidth
1	0-7	reserved	

7.1.8.21 MAC Capabilities IE

The MAC Capabilities IE is specified in Table 85.

Table 85 — MAC Capabilities IE format

Syntax	Size	Notes
MAC_Capabilities_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3)	1 byte	
MAC Capability Bitmap	2 bytes	Defined in Table 86
Reserved	1 byte	
}		
}		

The MAC Capability Bitmap field indicates capabilities supported by the MAC entity. A bit is set to ONE if the corresponding attribute is supported, or is set to ZERO otherwise. This field is encoded as described in Table 86. Subsequent octets are reserved.

Table 86 — MAC Capability Bitmap

Octet	Bit	Attribute	Description
0	0	PCA	Capable of transmitting and receiving frames using the PCA mechanism
	1	Soft CRP	Capable of being the owner and target of Soft CRP reservations
	2	Block ACK	Capable of transmitting and acknowledging frames using the B-ACK mechanism
	3	Hibernation anchor	Capable of acting as a hibernation anchor
	4	Centralized self-coexistence	Support centralized self-coexistence mechanisms as specified in 7.12.3
	5	Proxy	Capable of acting as a proxy (as a peer device) to coordinate channel measurement
	6-7	reserved	
1	0 – 7	Reserved	Reserved

7.1.8.22 Slave device list IE

The slave device list IE is used by a master device to indicate the list of slave devices associated with itself. The IE format is specified in Table 87.

Table 87 — Slave Device List IE format

Syntax	Size	Notes
Slave_Device_List_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+2×N)	1 byte	
N	1 byte	Number of slave devices
For(i=0; i<N; i++){		
DevAddr	2 bytes	Slave device address
}		
}		

7.1.8.23 Beaconsing Promotion Indication IE

The Beaconsing Promotion Indication IE is transmitted by a master to indicate a group of nonbeaconsing slave devices to become regular beaconsing slave devices. The IE format is specified in Table 88.

Table 88 — Beaconing Promotion Indication IE format

Syntax	Size	Notes
Beaconing_Promotion_Indication_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+2×N)	1 byte	
N	1 byte	Number of slave devices
For(i=0; i<N; i++){		
DevAddr	2 bytes	Slave device address
}		
}		

7.1.8.24 Association Response IE

The Association Response IE is transmitted by a master device or a peer device to indicate the outcome of an association request. The IE format is specified in Table 89.

Table 89 — Association Response IE format

Syntax	Size	Notes
Association_Response_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1)	1 byte	
Reason code	1 byte	0 – Success 1 – Fail-Invalid Address 2 – Fail-Not enough MAS 3 – Fail-Unspecified 4 – 255 – reserved
}		

7.1.8.25 Hibernation Mode IE

The Hibernation Mode IE is specified in Table 90.

Table 90 — Hibernation Mode IE format

Syntax	Size	Notes
Hibernation_Mode_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 2)	1 byte	
Hibernation Countdown	1 byte	
Hibernation Duration	1 byte	
}		

The Hibernation Countdown field is set to the number of superframes remaining until the device begins hibernation. A value of ZERO indicates that the device will enter hibernation mode at the end of the current superframe.

The Hibernation Duration field is set to the number of superframes for which the device intends to hibernate.

7.1.8.26 Traffic Indication Map (TIM) IE

The TIM IE is used to indicate that an active mode device has data buffered for transmission via PCA. The TIM IE is specified in Table 91.

Table 91 — Traffic Indication Map (TIM) IE format

Syntax	Size	Notes
Traffic_Indication_MAP_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+ 2*N)	1 byte	
N	1 byte	Number of devices for which PCA traffic is buffered
For(i=0; i<N; i++){		
DevAddr _i	2 bytes	DevAddr of a device for which PCA traffic is buffered
}		
}		

Each DevAddr field is set to the DevAddr of a device for which PCA traffic is buffered.

7.1.8.27 On-demand QP Negotiation Request IE

The On-demand QP Negotiation Request IE, as defined in Table 92, is transmitted by a device to request temporary use of certain MASs owned by another device.

Table 92 — On-demand QP Negotiation Request IE

Syntax	Size	Notes
On-demand_QP_Negotiation_Request_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 5)	1 byte	
Device Address	2 bytes	Device address of reservation owner
Superframe countdown	1 byte	The number of superframe left before On-demand QP is intended
QP starting time	1 byte	Starting MAS to intend On-demand QP
Duration	1 byte	Intended QP duration, in the unit of MAS
}		

7.1.8.28 On-demand QP Negotiation Response IE

The On-demand QP Negotiation Response IE, defined in Table 93, is transmitted by a reservation owner to respond to the request to lend certain MASs for On-demand QP. A reservation owner may modify the intended On-demand QP.

Table 93 — On-demand QP Negotiation Response IE

Syntax	Size	Notes
On-demand_QP_Negotiation_Response_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 5)	1 byte	
Device Address	2 bytes	Device address of initiating On-demand QP Negotiation
Superframe countdown	1 byte	The number of superframe left before On-demand QP is agreed to take place. When the value is set to 1, On-demand QP is agreed to take place in next superframe.
QP starting time	1 byte	Starting MAS for agreed On-demand QP
Duration	1 byte	Agreed QP duration, in the unit of MAS
}		

7.1.8.29 Device Features Information element

A device may exchange device features with another device to help decide operation parameter properly. The device features IE is defined in Table 94. The device feature IE may be included in beacon or probe command frame.

Table 94 — Device Features IE format

Syntax	Size	Notes
Device_Features_IE_Format() {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3)	1 byte	
Device Features	3 bytes	Defined in Table 95
}		

The Device attributes field is specified in Table 95.

Table 95 — Device Features field format

Syntax	Size	Notes
Device_Features_field() {		
AC Power support	2 bits	00: AC power support, 01: otherwise 10: not specified, 11: reserved
User-defined Priority of being chosen as master	3 bits	000: highest priority~110: lowest priority 111: this feature is not supported
Security Support	2 bits	00: security mode 0 01: security mode 1 10: security mode 2 11: this feature is not supported Refer to 6.6.9 and 8.2
Maximum Transmission power	8 bits	DEV's Maximum transmission power in dBm
Device Mobility	2 bits	00: fixed location, 01: movable DEV 01: not specified, 11: reserved
Maximum Associated DEVs number	5 bits	Number of maximum associated devices If set as 0, this feature not supported
Reserved	2 bits	
}		

7.1.8.30 Application-specific IE (ASIE)

The ASIE is specified in Table 96.

Table 96 — ASIE format

Syntax	Size	Notes
ASIE_Format() {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 2+N)	1 byte	N equals the length of Application-specific Data
ASIE Specifier ID	2 bytes	
Application-specific Data	Variable	
}		

The ASIE Specifier ID field is set to a 16-bit value that identifies a company or organization, as listed in [4].

The owner of the ASIE Specifier ID defines the format and use of the Application-specific Data field.

7.1.8.31 Application-specific Probe IE

The Application-specific Probe IE is used to request an application-specific IE from a device. It is specified in Table 97.

Table 97 — Application-specific Probe IE format

Syntax	Size	Notes
Application-specific_Probe_Format() {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4+N)	1 byte	N equals the length of Application-specific Request Information
Target DevAddr	2 bytes	
ASIE Specifier ID	2 bytes	
Application-specific Request Information	Variable	
}		

The Target DevAddr field is set to the DevAddr of the device from which an ASIE is requested.

The ASIE Specifier ID is set to a 16-bit value that identifies a company or organization, as listed in [4].

The owner of the ASIE Specifier ID defines the format and use of the Application-specific Request Information field.

7.1.8.32 Hibernation Anchor IE

The Hibernation Anchor IE is specified in Table 98.

Table 98 — Hibernation Anchor IE format

Syntax	Size	Notes
Hibernation Anchor IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3xN)	1 byte	N equals Number of Hibernation Mode Neighbour Devices
For(i=0; i<N; i++){		
Hibernation Mode Neighbour DevAddr	2 bytes	
Wakeup Countdown	1 byte	
}		
}		

The Hibernation Mode Neighbour DevAddr field is set to the DevAddr of the neighbour in hibernation mode.

The Wakeup Countdown field is set to the number of remaining superframes before the device in hibernation mode is expected to wake up. A value of ZERO indicates that the device is scheduled to be in active mode in the next superframe.

7.1.8.33 Master Key Identifier (MKID) IE

The MKID IE is used to identify some or all of the master keys possessed by the device. The MKID IE is specified in Table 99.

Table 99 — MKID IE format

Syntax	Size	Notes
MKID IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 16xN)	1 byte	N equals number of MKID included in this IE
For(i=0; i<N; i++){		
MKID _i	16 bytes	
}		
}		

Each MKID field is set to the identifier of a master key possessed by the device.

7.1.8.34 Multicast Address Binding (MAB) IE

Each device maps multicast EUI-48s to McstAddrs in the 16-bit DevAddr address range. The MAB IE declares the binding between a multicast EUI-48 and the McstAddr that the device will use when transmitting frames destined for that multicast EUI-48.

Table 100 — MAB IE format

Syntax	Size	Notes
MAB IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 8xN)	1 byte	N equals number of Multicast Address Binding Block (MEUI + MDevAddr) included in this IE
For(i=0; i<N; i++){		
MEUI	6 bytes	
MDevAddr	2 bytes	
}		
}		

The format of the MAB IE is shown in Table 100.

The MEUI field is set to the multicast EUI-48 supplied by the MAC client at the MAC SAP.

The MDevAddr field is set to the multicast DevAddr bound to the MEUI field by the MAC entity from the McstAddr address range.

7.1.8.35 PCA Availability IE

The PCA Availability IE identifies the MASSs in which a device will be available to receive PCA traffic and transmit the required response.

The PCA Availability IE is specified in Table 101.

Table 101 — PCA Availability IE format

Syntax	Size	Notes
PCA_Availability_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+N)	1 byte	N equals size of PCA Availability Bitmap
Interpretation	1 byte	
PCA Availability Bitmap	Variable (0 to 32)	
}		

The Interpretation field contains information that specifies the meaning of each bit in the PCA Availability Bitmap field. The Interpretation field is specified in Table 102.

Table 102 — Interpretation field format

Syntax	Size	Notes
Interpretation_Field_Format {		
TIM IE Required	1 bit	
Reserved	7 bits	
}		

The TIM IE Required bit is set to ONE if the device will only be available to receive PCA traffic in the specified MASSs after receiving a TIM IE that addresses it. The bit is set to ZERO if the device will be available to receive PCA traffic in the specified MASSs regardless of TIM IE reception.

The PCA Availability Bitmap field is up to 256 bits long, one bit for each MAS in the superframe, where the least-significant bit of the field corresponds to the first MAS in the superframe and successive bits correspond to successive MASSs. Each bit is set to ONE if the device is available to receive PCA traffic and transmit the required response in the corresponding MAS, or is set to zero otherwise. If the PCA Availability Bitmap field is smaller than 32 octets, the bits in octets not included at the end of the bitmap are treated as zero.

7.1.8.36 Probe IE

The Probe IE is used to request information from a device. It is specified in Table 103.

Table 103 — Probe IE format for standard IEs

Syntax	Size	Notes
Probe IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 2 + N)	1 byte	N equals number of requested IEs included in this IE
Target DevAddr	2 bytes	
For(i=0; i<N; i++){		
Requested Element ID _i	1 byte	
}		
}		

The Target DevAddr field is set to the DevAddr of the device from which IEs are requested or the device that requests IEs.

Each Requested Element ID field is set to the element ID of a requested IE.

7.1.8.37 Relinquish Request IE

The Relinquish Request IE is used to request that a device release one or more MASs from one or more existing reservations. It identifies the target device and the desired MASs, and is specified in Table 104.

Table 104 — Relinquish Request IE format

Syntax	Size	Notes
Relinquish Request IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4 + 4xN)	1 byte	N equals number of Allocations included in this IE
Relinquish Request Control	2 bytes	
Target DevAddr	2 bytes	
For(i=0; i<N; i++){		
Allocation i	4 bytes	
}		
}		

The Relinquish Request Control field is specified in Table 105.

Table 105 — Relinquish Request Control field format

Syntax	Size	Notes
Relinquish Request Control Format {		
Reason Code	4 bits	
Reserved	12 bits	
}		

The Reason Code field indicates the reason for the request, and is encoded as shown in Table 106.

Table 106 — Reason Code field encoding

Value	Code	Meaning
0	Non-specific	No reason specified.
1	Over-allocation	The target device holds more MASs than permitted by policy.
2 – 15	Reserved	Reserved

The Target DevAddr field is set to the DevAddr of the device that is requested to release MASs.

A Relinquish Request IE contains one or more Allocation fields. Each Allocation field is encoded using a zone structure. The superframe is split into 16 zones numbered from 0 to 15 starting from the BPST. Each zone contains 16 consecutive MASs, which are numbered from 0 to 15 within the zone.

The general format of an Allocation field is specified in Table 107.

Table 107 — Allocation field format

Syntax	Size	Notes
Allocation_Field_Format {		
Zone Bitmap	2 bytes	
MAS Bitmap	2 bytes	
}		

The Zone Bitmap field identifies the zones that contain requested MASs. If a bit in the field is set to ONE, the corresponding zone contains requested MASs, where bit zero corresponds to zone zero.

The MAS Bitmap specifies which MASs in the zones identified by the Zone Bitmap field are part of the request. If a bit in the field is set to ONE, the corresponding MAS within each zone identified by the Zone Bitmap is included in the request, where bit zero corresponds to MAS zero within the zone.

7.1.8.38 Link Quality Estimate IE

The Link Quality Estimate IE is used to report link quality estimation based on measurement from a source device to the current device, as specified in Table 108.

Table 108 — Link Quality Estimate IE format

Syntax	Size	Notes
Link_Quality_Estimation_IE_Format() {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 4)	1 byte	
LQE	1 byte	As defined in Table 148
DevAddr	2bytes	DEV address of signal source
}		

7.1.8.39 Transmit Power Control IE (TPC IE)

The Transmit Power Control IE is used by a device to update transmit power of other devices. This IE shall contain information of the recommended transmit power which the destination devices shall keep it for future use. For example, a master device may use Transmit Power Control IE to update the transmit power of a slave device. Similarly, a peer device may use Transmit Power Control IE to update the transmit power of an associated device. The Transmit Power Control IE is defined in Table 109.

Table 109 — Transmit Power Control IE format

Syntax	Size	Notes
Transmit Power Control IE Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 1+3xN)	1 byte	
N	1 byte	Number of TPC fields
For (i=1, i <=N, i++){		
TPC _i	3 bytes	Defined in Table 110
}		
}		

The TPC field is defined in Table 110.

Table 110 — TPC field format

Syntax	Size	Notes
TPC_Field_Format {		
DevAddr	16 bits	DEV address of destination device
Transmit Power	8 bits	Defined in Table 111
}		

The DevAddr field is set to the DevAddr of the destination device. If the DevAddr is set as broadcast address, every device shall update the transmit power as indicated in the TPC field. If the DevAddr is set as multicast address, every device in the multicast group shall update the transmit power as suggested in the TPC field. If the DevAddr is set as Unassociated, the transmit power included in the TPC field applies to new devices to be associated with the master device or the peer device which transmits this TPC IE.

The Transmit Power field encoding is specified in Table 111.

Table 111 — Transmit Power field encoding

Value	mW
0	1
1	2
2	3
...	...
98	99
99	100
100-255	reserved

7.1.8.40 Disassociation IE

To terminate the association of slave devices or peer devices, a master or peer transmits to slave or peer devices by including disassociation IE into its beacon frame. The disassociation IE is specified in Table 112.

Table 112 — Disassociation IE format

Syntax	Size	Notes
Disassociation_IE_Format {		
Element ID	1 byte	Set to the value as listed in Table 45 that identifies the information element
Length (= 3×N)	1 byte	N is the total number of disassociated devices
For(i=0; i<N; i++){		
DevAddr _i	2 bytes	DevAddr of disassociated devices
Reason code	1 byte	Refer to Table 113
}		
}		

The Reason Code is informative for a disassociated device and is encoded as shown in Table 113.

A device receiving the Disassociation IE shall release all reserved MASs and terminate existing connections with the associated device.

Table 113 — Reason Code field encoding

Value	Meaning
0	Lack of resource
2	Self-coexistence problem
3 – 255	Reserved

7.2 Frame processing

This subclause provides rules on preparing MAC frames for transmission and processing them on reception. The rules cover MAC header fields and information elements.

7.2.1 Frame addresses

Frames are addressed using DevAddrs. There are four types of DevAddrs; Private, Generated, Multicast, and Broadcast. Table 114 shows the range for each type of DevAddr.

Table 114 — DevAddr types and ranges

Type	Range
Private	0x0000 – 0x00FF
Generated	0x0100 – 0xFEFF
Unassociated	0xFF00
Multicast (McstAddr)	0xFF01 – 0xFFFE
Broadcast (BcstAddr)	0xFFFF

A device shall associate a DevAddr of either type Private or type Generated with its local MAC entity. A device that uses a NULL EUI-48 shall use a Private DevAddr. If a device uses a Generated DevAddr, it shall select the DevAddr from the Generated DevAddr range at random with equal probability and shall ensure that the generated value is unique among all devices in its extended beacon group. Selection and conflict resolution for Private DevAddrs is out of scope of this Standard.

In all frames transmitted, a device shall set the SrcAddr field to its own DevAddr. In unicast frames, the DestAddr field shall be set to the DevAddr of the recipient. In multicast frames, the DestAddr field shall be set to an address from the Multicast DevAddr range, as specified in 7.2.10.26. In broadcast frames, the DestAddr field shall be set to the Broadcast DevAddr.

A device shall not transmit frames addressed to a recipient with a Private DevAddr at any time outside a Private reservation. A device with a Private DevAddr shall not transmit non-beacon frames outside a Private reservation.

A master or peer device addresses unassociated devices with the DevAddr set as 0xFF00, for example, for transmit power control, specified in 7.7.3.

7.2.1.1 DevAddr Conflicts

A device with a Generated DevAddr shall recognize that its DevAddr is in conflict if any of the following conditions occurs:

- It receives a MAC header in which the SrcAddr is the same as its own DevAddr; or
- It receives a beacon frame in which the BPOIE contains a DevAddr that is the same as its own but corresponds to a beacon slot in which the device did not transmit a beacon and was not in hibernation mode.

A device that recognizes that its DevAddr is in conflict shall generate a new DevAddr to resolve the DevAddr conflict.

7.2.2 Frame reception

Unless otherwise indicated, a frame is considered to be received by the device if it has a valid header error indicator (HEI) and frame check sequence (FCS) and indicates a protocol version that is supported by the device. The HEI is set by the PHY, which indicates whether or not a header error occurred. A valid HEI means the PHY and MAC headers are correctly decoded, invalid otherwise.

A MAC header is considered to be received by the device if it has a valid HEI and indicates a protocol version supported by the device, regardless of the FCS validation.

7.2.3 Frame transaction

A frame transaction consists of an optional RTS/CTS frame exchange, a single frame, and the associated acknowledgement frame if requested by the ACK policy.

Figure 6 shows some frame transaction examples.

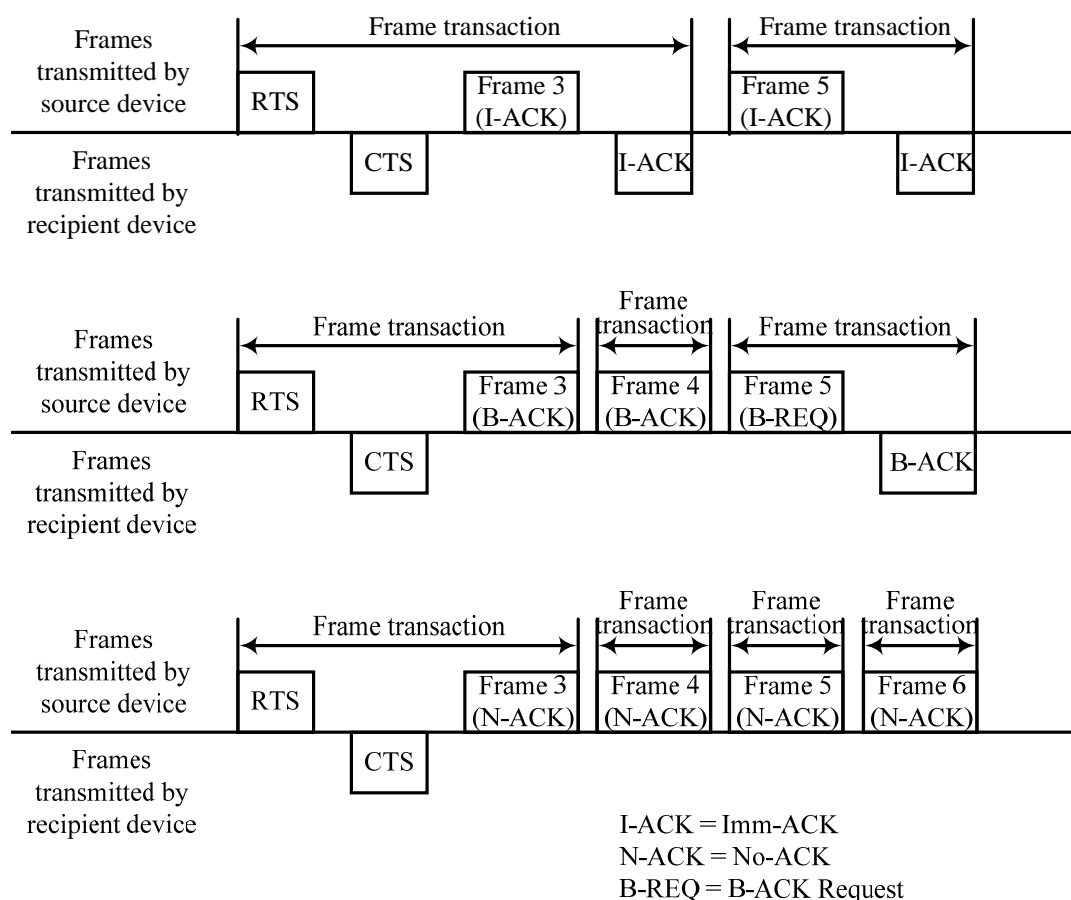


Figure 6 — Frame transaction examples

7.2.4 Frame transfer

A source device shall transmit MSDUs associated with the same Delivery ID and addressed to the same destination EUI-48 in the order in which they arrived at the local MAC SAP. The device shall treat each MSDU of length n as a sequence of octets, labeled MSDU[0] to MSDU[n-1], and shall place these octets in the payload field in ascending index-value order. The device shall transmit fragments of an MSDU or MCDU in order of increasing fragment number.

When using the B-ACK mechanism, a source device may retransmit some previously transmitted frames, causing the sequence numbers and fragment numbers of the retransmitted frames to be out of order with respect to previously transmitted frames.

A source device may reorder MSDUs for transmission if their associated Delivery IDs or destination EUI-48s are different.

A recipient device shall release MSDUs to the MAC client that was transmitted by the same source device with the same Delivery ID in order of increasing sequence number values.

A source device may fragment or aggregate MSDUs for transfer between peer MAC entities, but the recipient device shall deliver whole individual MSDUs through the MAC SAP to the MAC client.

7.2.5 Frame retry

A frame retry is a retransmission of a previously transmitted frame from the same source device to the same recipient device. In a frame that is retransmitted, the source device shall set the Retry bit to ONE.

Unless otherwise stated, in this specification “transmission” means transmission of a new frame or retransmission of a previously transmitted frame.

A device may retransmit a frame as needed, taking into consideration such factors as delay requirements, fairness policies, channel conditions, and medium availability. A device shall apply the medium access rules for new frame transmissions when retransmitting frames, unless stated otherwise.

7.2.6 Inter-frame space (IFS)

Three types of IFS are used in this Standard: the minimum inter-frame space (MIFS), the short inter-frame space (SIFS), and the arbitration inter-frame space (AIFS[i]). There are four values of AIFS depending on the access category of the traffic. The actual values of the MIFS, SIFS, and AIFS are PHY-dependent.

A device shall not start transmission of a frame on the medium with non-zero length payload earlier than MIFS, or with zero length payload earlier than SIFS, after the end of a frame it transmitted previously on the medium. A device shall not start transmission of a frame on the medium earlier than SIFS duration after the end of a previously received frame on the medium.

7.2.6.1 MIFS

Burst frame transmissions are those frames transmitted from the same device where the timing of each frame in the burst after the first is related to the preceding frame through use of the PHY burst mode. In this case a MIFS duration will occur between frames in the burst, as shown in Figure 7. All frames in a burst except the last frame shall be sent with the ACK Policy field set to No-ACK or B-ACK. The last frame in a burst may be sent with any ACK Policy.

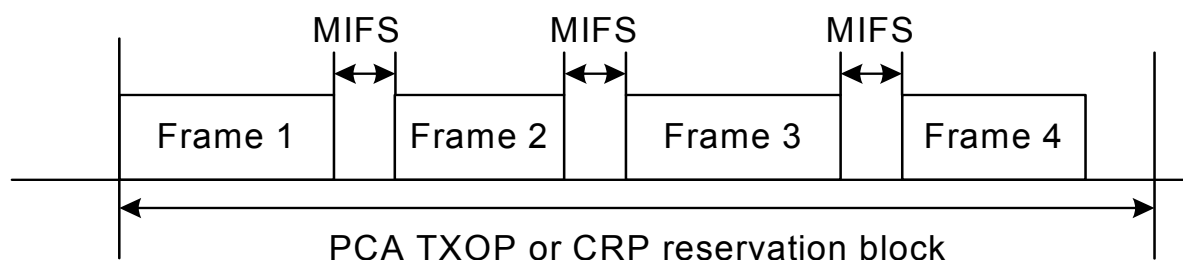


Figure 7 — Use of MIFS

Within a burst, the Duration field shall cover only consecutive frames addressed to the same destination. If the burst continues after the Duration is exhausted, the next frame shall use a standard preamble. The length of MIFS is given by the pMIFS parameter specified in 7.14.

7.2.6.2 SIFS

Within a frame transaction, all frames shall be separated by a SIFS interval.

The length of SIFS is given by the pSIFS parameter specified in 7.14.

7.2.6.3 AIFS

The AIFS is the minimum time that a device using PCA defers access to the medium after it determines the medium to have become idle.

7.2.7 Duplicate detection

Because a device may not receive an Imm-ACK or B-ACK response for a frame it transmitted, it may send duplicate frames even though the intended recipient has already received and acknowledged the frame. A recipient device shall consider a received frame to be a duplicate if the Retry bit is set and the Sequence Control field has the same value as the previous frame received with the same SrcAddr, DestAddr, and Delivery ID field values. A recipient device shall not release a duplicate frame to the MAC client.

7.2.8 RTS/CTS use

An RTS/CTS exchange, when used, precedes data, aggregated data, or command frames to be transferred from a source device to a recipient device. Without a frame body, the RTS frame allows the source device to regain medium access relatively quickly in case of an unsuccessful transmission. With an appropriately set Duration field as specified in 7.2.9.1, the RTS and CTS frames prevent the neighbours of the source and recipient devices from accessing the medium while the source and recipient are exchanging the following frames.

A source device may transmit an RTS frame as part of one or more frame transactions with another device in an obtained PCA TXOP or an established reservation block. In a PCA TXOP, a device should transmit an RTS frame prior to transmitting a sequence of frames using the No-ACK acknowledgment policy or the B-ACK mechanism if those frame transmissions would otherwise not be covered by the Duration field contained in a frame transmitted previously between the same source and recipient devices.

If a reservation target receives an RTS frame addressed to it in the reservation block, from the reservation owner, it shall transmit a CTS frame pSIFS after the end of the received frame, regardless of its NAV setting. If a device receives an RTS frame addressed to it outside a reservation block, it shall transmit a CTS frame pSIFS after the end of the received frame if and only if its NAV is zero and the CTS frame transmission will be completed pSIFS before the start of the next BP or before the start of its own or a neighbour's established reservation block.

On receiving an expected CTS response, the source device shall transmit the frame, or the first of the frames, for which it transmitted the preceding RTS frame pSIFS after the end of the received CTS frame. If the source device does not receive the expected CTS frame pSIFS plus the CTS frame transmission time after the end of the RTS frame transmission, and it transmitted the RTS frame in a PCA TXOP, it shall invoke a backoff as specified in 8.3. If it transmitted the RTS frame in one of its reservation blocks, it shall not retransmit the RTS frame or transmit another frame earlier than pSIFS after the end of the expected CTS frame.

7.2.9 MAC header fields

7.2.9.1 Duration

A device shall set the Duration field in beacon frames to one of the following:

- The time remaining in the BP measured from the end of the PLCP header of the beacon frame, as determined by the largest BP length announced by neighbours of the device in the previous superframe;
- The transmission time of the frame body of the beacon frame; or
- Zero.

A device shall set the Duration field in RTS, command, data, or aggregated data frames to the sum of:

- The transmission time of the frame body of the current frame;
- The transmission time of the expected response frame for the current frame (CTS, Imm-ACK, or B-ACK frame), if any;

A device may estimate the transmission time of a B-ACK frame body based on the expected length and data rate, or may assume a zero-length frame body.

A device shall set the Duration field in CTS, Imm-ACK and B-ACK frames to the larger of zero or a value equal to the duration value contained in the previous frame minus pSIFS, minus the transmission time of the frame body of the received frame to which the CTS, Imm-ACK or B-ACK is responding, minus the transmission time up to the end of the PLCP header of this CTS, Imm-ACK or B-ACK frame.

The following exceptions to previous rules are allowed:

- For frames with ACK Policy set to B-ACK Request, a device shall set the Duration to the sum of the transmission time of the frame body of the B-ACK Request frame plus a SIFS plus the estimated transmission time of the expected B-ACK response frame.
- A device may set the Duration for any frame sent in a Hard or Private reservation block other than UCA or UCR frames to ZERO.

A device shall set the Duration field in UCA and UCR frames to a time interval extending from the end of the PLCP header of the current frame to the time when the remaining CRP reservation block is to be released.

Examples of Duration field values are specified in Figure 8.



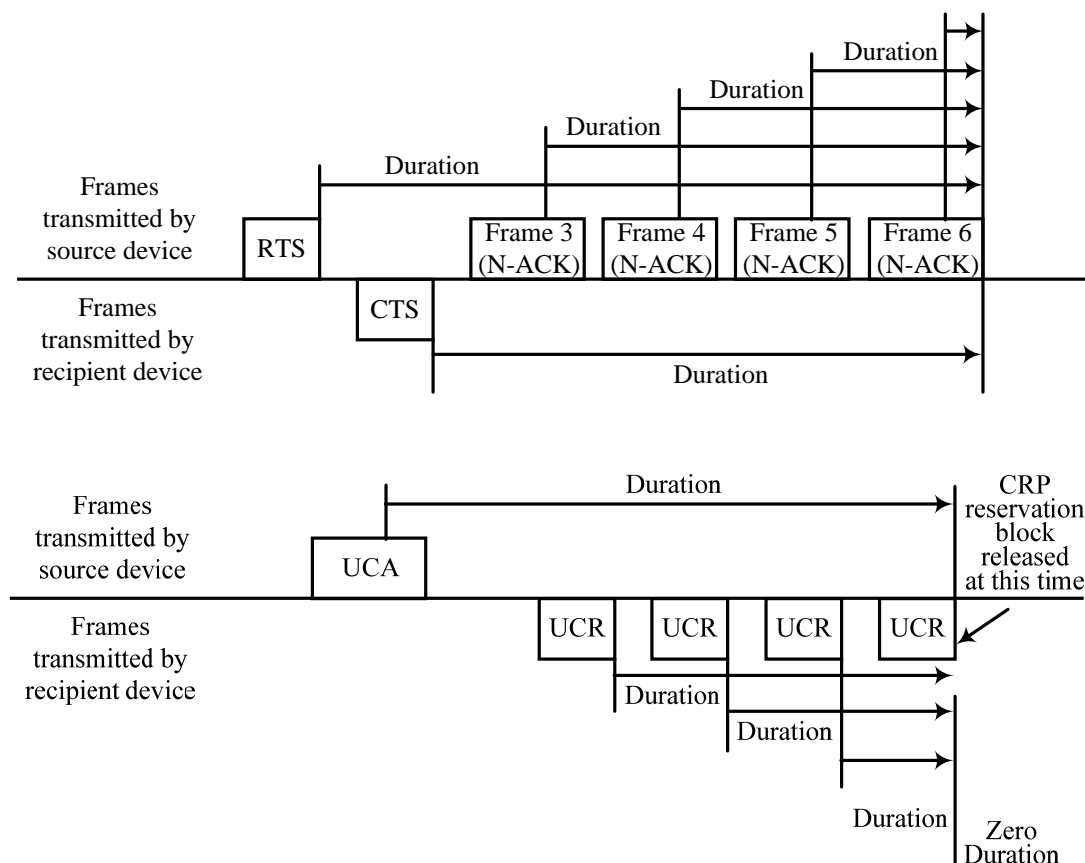


Figure 8 — Duration examples

7.2.9.2 More Frames

If a device sets the More Frames bit to ZERO in a frame sent with Access Method set to ONE, it shall not transmit additional frames to the same recipient(s) within the reservation block.

If a device sets the More Frames bit to ZERO in a frame sent with Access Method set to zero, it shall not transmit additional frames using PCA to the same recipient(s) within the current superframe unless the recipient did not include a PCA Availability IE in its beacon or included a PCA Availability IE in its beacon with the TIM IE Required bit set to ZERO.

7.2.9.3 Sequence Number

The Sequence Number field value is used for duplicate detection for frames sent using the Imm-ACK acknowledgement policy. It is used for both duplicate detection and reordering for frames sent using the B-ACK mechanism.

A device shall assign each MSDU or MCDU transmitted a sequence number from a modulo 2048 counter.

A device shall assign the same sequence number to each fragment of an MSDU or MCDU.

A single sequence number applies to all MSDUs contained in an aggregated data frame. A device shall increment the sequence number counter by one for each transmitted aggregated data frame.

A device shall use a dedicated counter for MCDUs.

A device shall use a dedicated counter for each sequence of MSDUs addressed to the same DestAddr with the same Delivery ID using B-ACK acknowledgement policy.

A device may use one counter for all other MSDUs, or may use a dedicated counter for MSDUs with the same Delivery ID field value addressed to the same DestAddr.

In each beacon frame transmitted in a superframe, a device shall set the Sequence Number field from a dedicated counter that increments once per superframe, modulo 2048, or shall set it to zero.

7.2.10 Information elements

IEs are contained in beacon and command frames. They convey certain control and management information. IEs may be explicitly requested using Probe command frames.

A device shall include IEs in its beacon frame such that they apply to the superframe in which the beacon is transmitted. A device shall interpret IEs contained in beacons received in the current superframe to apply to that superframe.

The remainder of this subclause describes when each IE is generated.

7.2.10.1 Application-Specific IE (ASIE)

A device may include an ASIE in its beacon or command frame for each of its applications which have made the request. The scope of the ASIE is dependent on the application that requested the inclusion of the ASIE.

7.2.10.2 Application-Specific Probe IE

A device may include an Application-specific Probe IE in its beacon or command frame. The scope of the Application-specific Probe IE is dependent on the application that requested the inclusion of the IE.

7.2.10.3 Association Response IE

A device shall include an Association Response IE in its beacon in response to an association request.

7.2.10.4 Beacon Period Occupancy IE (BPOIE)

A regular beaconsting device shall always include a BPOIE in its beacon. In the BPOIE the device shall reflect beacons received from neighbours in the previous superframe, as well as information retained based on hibernation mode rules.

7.2.10.5 Beaconsting Promotion Indication IE

A master device shall include a Beaconsting Promotion Indication IE in its beacon in a superframe when it has received beaconsting promotion request from a slave device or it intends to promote a slave device to be a regular beaconsting slave device.

7.2.10.6 BP Switch IE

A device shall include a BP Switch IE in its beacon prior to changing its BPST, as specified in 7.3.7.

7.2.10.7 Channel Change IE

A device shall include a Channel Change IE in its beacon or Channel Switch Command frame to request other devices to switch channel.

7.2.10.8 Channel Change Response IE

A device shall include a Channel Change Response IE in its beacon or in its Channel Switch Response command frame as response to receiving channel change request from either beacon or Channel Switch

Command frame. A device that includes a Channel Change Response IE shall change channels as indicated in the IE.

7.2.10.9 Channel Classification IE

A master or peer device shall include Channel Classification IE in its beacon if channel classification is initialized or updated.

7.2.10.10 Channel Measurement Report IE

Channel measurement report IE is used to report channel measurement results. A device shall report channel measurement results upon receiving channel measurement request. A device may also report channel measurement results once detecting incumbent signal or other interfering sources. A device shall include Channel measurement report IE in its beacon or channel measurement report command frame.

7.2.10.11 Channel Measurement Report Acknowledgement IE

Channel Measurement Report Acknowledgement IE is used to acknowledge the reception of channel measurement report. A device shall include Channel measurement report acknowledgement IE in beacon or channel measurement report acknowledgement command frame.

7.2.10.12 Channel Measurement Request IE

A device may include channel measurement request IE in its beacon frame or corresponding command frame for channel measurement request.

7.2.10.13 Channel Measurement Response IE

A device shall include Channel Measurement Response IE in its beacon frame or corresponding command frame as response to channel measurement request.

7.2.10.14 Channel Reservation Protocol (CRP) IE

A device shall include CRP IEs in its beacon or command frame for all reservations in which it participates as a reservation owner or target, as described in 7.5.2. A master device shall also include CRP IEs in its beacon for all reservations in which its slave device participates as a reservation owner or target.

7.2.10.15 CRP Availability IE

A device shall include a CRP Availability IE in its beacon or command frame as required to support CRP reservation negotiation, as described in 7.5.2.

7.2.10.16 Device Features IE

A device may include a Device Features IE in its beacon or probe command frame for self-coexistence.

7.2.10.17 Disassociation IE

A device shall include a Disassociation IE in its beacon frame to terminate the association of other devices.

7.2.10.18 Echo beacon Position IE

Echo beacon Position IE shall be included in the Regular Beacon Frame of a master if needed. Echo beacon position IE indicates both MAS position of echo beacon and transmission owner of echo beacon.

7.2.10.19 Hibernation Anchor IE

A device that indicates it is capable of acting as a hibernation anchor shall include a Hibernation Anchor IE in its beacon to provide information on neighbours that are currently in hibernation mode as described in 7.9.

7.2.10.20 Hibernation Mode IE

A device shall include a Hibernation Mode IE in its beacon or command frame before entering hibernation mode, as specified in 7.9. A device that receives a Hibernation Mode IE shall report the beacon slot of the transmitter as occupied and non-movable in the BPOIE included in its beacons during the reported hibernation duration.

7.2.10.21 Identification IE

To provide its own identifying information to neighbours, a device includes an Identification IE in its beacon or command frame.

7.2.10.22 Link Feedback IE

A device may include a Link Feedback IE in its beacon or command frame to provide feedback on a link with a specific neighbour.

7.2.10.23 Link Quality Estimate IE

A device may include a Link Quality Estimate IE in its beacon or command frame to provide link quality estimation on a link with a specific neighbour.

7.2.10.24 MAC Capabilities IE

To negotiate MAC capabilities, a device includes a MAC Capabilities IE in its beacon or command frame.

7.2.10.25 Master Key Identifier (MKID) IE

A device includes a MKID IE in its beacon or command frame to identify some or all of the master keys it possesses.

7.2.10.26 Multicast Address Binding (MAB) IE

A device shall include a MAB IE in its beacon or command frame for at least $mMaxLostBeacons+1$ superframes on registering a multicast address binding for transmission and upon detection of a change in the beacon group.

The MAC entity shall translate the multicast EUI-48 provided by the MAC client along with an MSDU to the bound multicast DevAddr for use in the transmission of the MSDU over the medium.

A device shall not transmit frames with a McstAddr destination address unless a binding to a multicast EUI-48 has been declared by inclusion of a corresponding MAB IE in its beacon.

On receipt of a MAB IE the MAC entity shall establish an association between the source of the MAB IE and the multicast DevAddr and multicast EUI-48 in each Multicast Address Binding Block, to be used in address translations for the bound multicast addresses.

The MAC entity shall deliver received MSDUs addressed to an activated multicast DevAddr to the MAC client on the multicast EUI-48 bound to that multicast DevAddr by the source device of the MSDU.

7.2.10.27 PCA Availability IE

A device includes a PCA Availability IE in its beacon or command frame as needed to facilitate PCA in the presence of reservations or power constraints.

7.2.10.28 PHY Capabilities IE

To exchange PHY capabilities, a device includes a PHY Capabilities IE in its beacon or command frame.

7.2.10.29 Probe IE

A device includes a Probe IE in its beacon or command frame to request certain IEs from another device.

7.2.10.30 On-demand QP Negotiation Request IE

A device includes an On-demand QP Negotiation Request IE in its beacon or probe command frame to initiate the negotiation request.

7.2.10.31 On-demand QP Negotiation Response IE

A device shall include an On-demand QP Negotiation Response IE in its beacon or probe command frame to respond to the negotiation request.

7.2.10.32 On-demand QP Schedule IE

A device shall include on-demand QP schedule IE in its beacon to announce the schedule of On-demand QP.

7.2.10.33 On-leave IE

A device shall include On-leave schedule IE in its beacon before temporally leaving the channel.

7.2.10.34 Proxy Assignment IE

A device shall include Proxy Assignment IE in its beacon to assign a proxy.

7.2.10.35 Regular QP Schedule IE

A regular beaconing device shall include Regular QP Schedule IE in its beacon.

7.2.10.36 Relinquish Request IE

A device includes a Relinquish Request IE in its beacon or command frame to request that a neighbour release one or more MASs from reservations.

If a reservation target receives a request to relinquish MASs included in the reservation, it shall include in its beacon a CRP Availability IE and a Relinquish Request IE identifying those MASs with the Target DevAddr field set to the DevAddr of the reservation owner.

7.2.10.37 RSW Schedule IE

A master device shall include RSW schedule IE in its beacon if RSW is present. An RSW schedule IE only indicates location (starting MAS and duration) of RSW. The usage of each RSW MAS is indicated by CRP IE.

7.2.10.38 Slave device list IE

A master may include slave device list IE in its beacon.

7.2.10.39 Transmit Power Control IE

A master device or a peer device includes Transmit Power Control IE in its beacon to update the transmit power of slave devices or other peer devices.

7.2.10.40 Traffic Indication Map (TIM) IE

A device shall include a TIM IE in its beacon in any superframe when it has frames queued for transmission to one or more recipients that have the TIM IE Required bit set in a PCA Availability IE in the previous superframe. The TIM IE shall include the DevAddr of all such recipients.

7.3 MAC Structure and Beaconsing

The MAC protocol follows a recurring superframe structure, which consists of a beacon period (BP), a data transfer period (DTP) and a contention based signalling window (CSW). A reservation based signalling window (RSW) could be appended right after BP to support signal exchange between a master and slave devices in the master-slave mode. RSW is not needed for pure peer-to-peer based network. The signalling windows and beacon period are used for sending and receiving control/management information.

All devices which share the same channel shall use the same superframe structure. Superframe merging is necessary if two networks follow different superframe structures and share the same channel.

All the devices (except hibernating devices) shall keep awake during beacon period and CSW in order to capture all the control/management information which may be relevant to every device. A device may exchange data, monitor channel, or go to sleep mode during data transfer period (DTP).

A device is defined as a beaconing device if it owns a beacon slot in BP and regularly transmits beacons. A peer device or a master device is by default a beaconing device. A slave device is normally a non-beaconing device unless promoted to be a beaconing device. A beaconing device owns one beacon slot and transmits beacon regularly.

The BP length is adjustable and depends on how many regular beaconing devices participate in the same BP. In the master-slave mode, the BP length should be 1 beacon slot if the master is the only device that does regular beaconing. In the peer-to-peer mode, the number of beacon slots may be as many as the number of peer devices.

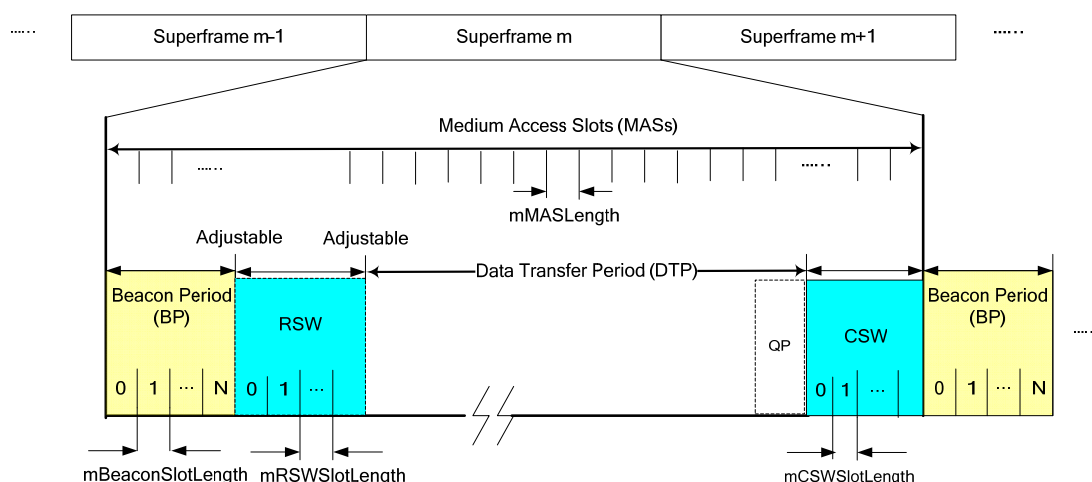


Figure 9 — Example MAC structure

7.3.1 Beacon Period

The maximal length of BP is $mMaxBPLength$ medium access slots. The length of each beacon slot is $mBeaconSlotLength$. The number of beacon slots is adapted to the number of devices in the beacon group. Beacon slots in the BP are numbered in sequence, starting at zero.

A beaconing device shall transmit a beacon frame (defined in 7.1.3 Beacon frames) in its own beacon slot and listen for neighbour's beacons in all beacon slots in BP in each superframe, except as described in 7.3.4.

When transmitting in a beacon slot, a device shall start transmission of the frame on the medium at the beginning of that beacon slot. A device shall transmit beacons at $pBeaconTransmitRate$. The transmission time of beacon frames shall not exceed $mMaxBeaconLength$. This allows for a guard time of at least $mGuardTime$ and $pSIFS$ between the end of a beacon and the start of the next beacon slot.

7.3.2 Beacon slot state

For beacon transmission and BP contraction purposes, a beacon slot should be considered available if in the latest $mMaxLostBeacons+1$ superframes the beacon slot was never encoded as occupied (i.e., no sharing is possible) in the BPOIE of the beacons transmitted or received by the device according to Table 46. A device shall consider a beacon slot unavailable in all other cases.

7.3.3 BP length

A device shall announce its BP length, measured in beacon slots, in its beacon. The announced BP length shall include the device's own beacon slot and all unavailable beacon slots in the BP of the prior superframe. The announced BP length shall not include more than $mBPExtension$ beacon slots after the last unavailable beacon slot in the BP of the prior superframe, unless otherwise indicated in 7.3.7. The announced BP length shall not exceed $mMaxBPLength$. Power-sensitive devices generally should not include any beacon slots after the last unavailable beacon slot in their announced BP length.

The BP length reported by a device varies, as new devices become members of its extended beacon group, and as the device or other devices in its extended beacon group choose a new beacon slot for beacon collision resolution or BP contraction.

7.3.4 Beacon transmission and reception

Before a device transmits any frames, it shall scan for beacons for at least one superframe. If the device receives no beacon frame headers during the scan, it shall create a new BP and send a beacon in the first beacon slot. If the device receives one or more beacon headers, but no beacon frames with a valid FCS during the scan, the device should scan for an additional superframe.

If the device receives one or more beacons during the scan, it shall not create a new BP. Instead, prior to communicating with another device, the device shall transmit a beacon in a beacon slot chosen from up to $mBPExtension$ beacon slots located after the highest-numbered unavailable beacon slot it observed in the last superframe and within $mMaxBPLength$ after the BPST.

If a device detects a beacon collision as described in 7.3.5, it shall choose a different beacon slot for its subsequent beacon transmissions from up to $mBPExtension$ beacon slots located after the highest-numbered unavailable beacon slot it observed in the last superframe and within $mMaxBPLength$ after the BPST.

If the beacon slot chosen for its beacon transmission is located beyond the BP length of any of its neighbours, the device shall also transmit a signalling beacon in the CSW, except as described in 7.3.7.1. The CSW is used under the above conditions regardless of whether a device is sending a beacon for the first time in an existing BP or changing the beacon slot after detecting a beacon collision. A device shall send a beacon in the CSW until its neighbours extend their BP lengths to include its beacon slot but only up to $mMaxLostBeacons$ superframes. After transmitting a beacon in a CSW slot for $mMaxLostBeacons+1$ superframes, a device shall wait for at least $mMaxLostBeacons+1$ superframes before sending a beacon in a CSW slot again.

If two BPs overlap as described in 7.3.7, a device shall wait for a random number of superframes before sending a beacon in a CSW slot to reduce potential collisions.

A device communicating with other devices shall listen for beacons during the BP length it announced in the last superframe. If a device received a beacon in a CSW slot in the previous superframe, it shall set its BP Length to include the beacon slot indicated in the beacon received in the CSW. If a device received a beacon with invalid FCS, or detected a medium activity that did not result in reception of a frame with valid HEI, in the CSW in the previous superframe, it shall listen for beacons for an additional mBPExtension beacon slots after its last announced BP length, but not more than mMaxBPLength beacon slots.

In order to detect beacon collisions with neighbours, a device shall skip beacon transmission aperiodically, and listen for a potential neighbour in its beacon slot. A device shall skip beacon transmission at least every mMaxNeighbourDetectionInterval.

With the exception of transmitting its own beacon as described in this subclause, a device shall not transmit frames during the announced BP length of any of its neighbours.

If a device does not receive a beacon from a neighbour in the current BP, it shall use information contained in the most-recently received beacon from the neighbour as if the beacon were received in the current superframe, except when determining the contents of the Beacon Slot Info Bitmap and DevAddr fields in its BPOIE. If a device does not receive a beacon from another device for more than mMaxLostBeacons consecutive superframes, it shall not consider the device a neighbour for purposes of this specification.

In a master-slave network, to help extend master beacon coverage, a nonbeaconing slave device selected by its master could transmit echo beacons as a relay of the master's regular beacon. When an echo beacon is scheduled to transmit during a certain superframe, the master device shall include echo beacon position IE (defined in 7.1.8.19) in its regular beacon frame. If a slave device receives echo beacon position IE which indicates itself as the transmission owner of an echo beacon, the slave device shall transmit the echo beacon in the MAS indicated in the echo beacon position IE. The selected slave device generates echo beacon based on echo beacon frame format (defined in 7.1.3.3). The echo beacon shall include CRP availability IE, BPOIE and Regular QP schedule IE.

7.3.5 Beacon collision detection

A device shall consider itself involved in a beacon collision with another device in its extended beacon group if one of the following events occurs:

- Its beacon slot is reported as occupied in the BPOIE in any beacon it receives in the current superframe, but the corresponding DevAddr is neither its own nor BcstAddr.
- Its beacon slot has been reported as occupied and the corresponding DevAddr has been BcstAddr in the BPOIE of a beacon it received in the same beacon slot in each of the latest mMaxLostBeacons superframes.
- After skipping beacon transmission in the previous superframe, its beacon slot is reported as occupied in the BPOIE of any beacon it receives in the current superframe.
- When skipping beacon transmission in the current superframe, it receives in its beacon slot in the current superframe: a MAC header of type beacon frame, or a PHY indication of medium activity that does not result in correct reception of a MAC header.

7.3.6 BP contraction

A device shall consider its beacon to be movable if in the previous superframe it found at least one available beacon slot between the BPST and the beacon slot it indicates in its beacon in the current superframe. However, for purposes of BP contraction, a device may consider an unoccupied beacon slot to be occupied for up to mMaxMovableLatency superframes, if it detects conditions that indicate contraction into that beacon slot might lead to a beacon slot collision, such as a previous beacon slot collision or indication of poor link conditions in that beacon slot.

A device that includes a Hibernation Mode IE in its beacon shall consider its beacon to be non-movable during the announced hibernation period. A device not involved in a beacon slot collision or a BP merge shall shift its beacon into the earliest available beacon slot following the signalling beacon slots in the BP of the next superframe, if in each of the latest $mMaxLostBeacons+1$ superframes:

- a) The device's beacon was movable; and
- b) the device did not receive a beacon from a neighbour that indicated a beacon slot after its own and had the Movable bit set to ONE; and
- c) the device did not receive a beacon from a neighbour that contained a BPOIE that encoded a beacon slot after its own as Movable according to Table 46. However, if in the last $mMaxLostBeacons+1$ superframes the device received a beacon from a neighbour that indicated a BP Length that did not include the device's beacon slot, and that beacon had the Movable bit set to ONE, the device should not change to an earlier beacon slot in the next superframe. Figure 10 shows some examples of BP contraction.




Figure 10 — Illustration for BP contraction by example devices

7.3.7 Merger of multiple beacon groups

Due to changes in the propagation environment, mobility, or other effects, devices using two or more unaligned BPSTs may come into range. This causes overlapping superframes. A received beacon that indicates a BPST that is not aligned with a device's own BPST is referred to as an alien beacon. The BP defined by the BPST and BP length in an alien beacon is referred to as an alien BP.

Synchronization problems could cause the beacon of a fast device to appear to be an alien beacon. A device shall consider a BPST to be aligned with its own if that BPST differs from its own by less than $2 \times mGuardTime$. A device shall consider an alien BP to overlap its own if its BPST falls within the alien BP or if the alien BPST falls within its own BP. A device shall not consider a signalling beacon to be an alien beacon.

If a device does not receive an alien beacon for up to $mMaxLostBeacons$ superframes after receiving one in a previous superframe, it shall use information contained in the most-recently received beacon as if the alien beacon were received at the same offset within the current superframe.

7.3.7.1 Overlapping BPs

If the BPST of a device falls within an alien BP, the device shall relocate its beacon to the alien BP according to the following rules:

1. The device shall change its BPST to the BPST of the alien BP.
2. The device shall adjust its beacon slot number such that its new beacon slot number is its old beacon slot number plus one, plus the number of the highest occupied beacon slot indicated in any beacon received in the alien BP. Alternately, it shall follow normal BP join rules as specified in 7.3.4 to relocate its beacon to the alien BP.
3. The device shall not send further beacons in its previous BP.
4. After changing its BPST, if the device is required to send a signalling beacon in the CSW according to 7.3.4, it shall wait for a random number of superframes before sending a signalling beacon in the CSW.

The device shall choose the random number with equal probability in the range zero to the BP Length declared in its last beacon before relocating to the alien BP.

7.3.7.2 Non-overlapping BPs

If a device detects an alien BP that does not overlap in time with its own BP, it shall merge BPs according to the following rules.

1. The device shall include in its beacon a CRP IE with Reservation Type set to Alien BP for the alien BP. Since the MAS boundaries might not be aligned, the device may need to include an additional MAS in the reservation to completely cover the alien BP. If the device received multiple beacons from the alien BP, it shall include all MASs used by the largest reported BP length in the reservation. If the MASs occupied by the alien BP change over time, the device shall update the CRP IE accordingly.
2. The device shall relocate its beacon to the alien BP, according to 7.3.7.3, within $mBPMergeWaitTime$ if the alien BPST falls within the first half of the superframe, or within $1.5 \times mBPMergeWaitTime$ if the alien BPST falls within the second half of the superframe, but shall not relocate to the alien BP if a beacon received in that alien BP includes a BP Switch IE.

A device that transmits or receives a beacon in its own BP that contains a CRP IE with Reservation Type set to Alien BP shall observe the following rules:

1. The device should not change beacon slots except as required by merge rules in 7.3.7, unless a collision is detected.
2. The device shall listen for beacons during the MASs indicated in the reservation.

7.3.7.3 Beacon relocation

If a device starts or has started the beacon relocation process and receives an alien beacon, it shall follow these rules:

- A. If the device did not include a BP Switch IE in its last beacon, it shall include a BP Switch IE in its beacon in the following superframe with the fields set as follows:
 - A1. The device shall set the BP Move Countdown field to $mInitialMoveCountdown$.
 - A2. The device shall set the BPST Offset field to the positive difference in microseconds between the alien BPST and the device's BPST. That is, the field contains the number of microseconds that the device must delay its own BPST to align with the alien BPST. If multiple alien beacons are received, the device shall set the BPST Offset field to the largest calculated value.
 - A3. The device shall set the Beacon Slot Offset field to:
 - a. one plus the number of the highest occupied beacon slot indicated by any beacon received in the alien BP, based on the Beacon Slot Number field and BPOIE; or
 - b. zero to indicate the device will join the alien BP using normal join rules as specified in 7.3.4.
- B. If the device included a BP Switch IE in its last beacon, it shall modify the BP Switch IE in the following superframe as follows:
 - B1. If the elapsed time between the device's BPST and the following alien BPST is larger than the device's BPST Offset field + $2 \times mGuardTime$, the device shall set the BP Move Countdown field, the BPST Offset field, and the Beacon Slot Offset field as described in A1, A2 and A3 above respectively.
 - B2. If the elapsed time between the device's BPST and the following alien BPST is larger than the device's BPST Offset field - $2 \times mGuardTime$ and smaller than the device's BPST Offset field + $2 \times mGuardTime$, the device shall set the BPST Offset field as described in A2. It shall set the Beacon Slot Offset field as described in A3 if the value in the field is to be increased, or leave it unchanged otherwise. It shall set the BP Move Countdown field to ONE less than the value used in its last beacon if the Beacon Slot Offset field is unchanged, or set it as described in A1 if the Beacon Slot Offset field is changed.

If a device receives a neighbour's beacon that contains a BP Switch IE, it shall follow these rules:

- C. If the device did not include a BP Switch IE in its last beacon, it shall include a BP Switch IE in its beacon in the following superframe with the fields set as follows:
 - C1. The device shall set the BP Move Countdown field to the BP Move Countdown field of the neighbour's BP Switch IE.
 - C2. The device shall set the BPST Offset field to the value of the same field contained in the neighbour's beacon.
 - C3. The device shall set the Beacon Slot Offset field to:
 - a. The larger of: one plus the number of the highest occupied beacon slot indicated by any alien beacon received in the alien BP identified by the neighbour's BP Switch IE, based on the Beacon Slot Number field and BPOIE; or the Beacon Slot Offset field contained in the neighbour's beacon; or
 - b. zero, to indicate the device will join the alien BP using normal join rules.
- D. If the device included a BP Switch IE in its last beacon, it shall modify the BP Switch IE as follows:
 - D1. If the BPST Offset field contained in the neighbour's beacon is larger than the device's BPST Offset field + $2 \times mGuardTime$, the device shall set the BP Move Countdown field, the BPST Offset field, and the Beacon Slot Offset field as described in C1, C2 and C3 above respectively.
 - D2. If the difference between the BPST Offset field contained in the neighbour's beacon and the device's BPST Offset field is smaller than $2 \times mGuardTime$, the device shall modify its BP Switch IE as follows:
 - a. If the Beacon Slot Offset field contained in the neighbour's beacon is larger than the device's Beacon Slot Offset field, the device shall set the BP Move Countdown field, the BPST Offset field, and the Beacon Slot Offset field as described in C1, C2 and C3 above respectively.
 - b. If the Beacon Slot Offset field contained in the neighbour's beacon is equal to or smaller than the device's Beacon Slot Offset field, the device does not receive alien beacons from the alien BP indicated by its current BPST Offset field, and the BP Move Countdown field contained in the neighbour's beacon is less than the device's BP Move Countdown field, then the device shall set the BPST Offset field as described in C2 above. It shall not change the Beacon Slot Offset field. It shall set the BP Move Countdown field to ONE less than the value used in its last beacon.

If a device included a BP Switch IE in its last beacon and none of the conditions within B or D apply, the device shall not change the BPST Offset field or the Beacon Slot Offset field, and shall set the BP Move Countdown field to one less than the value used in its last beacon.

If a device includes a BP Switch IE in its beacon, it shall continue to do so until it completes or halts the relocation process.

- If a device receives an alien beacon that indicates relocation earlier than its planned relocation, the device shall halt the relocation process.
- If a neighbour halts the relocation process, the device shall halt the relocation process.

To halt the relocation process, a device shall include a BP Switch IE in its beacon with BPST Offset field set to 65535, Beacon Slot Offset field set to zero, and BP Move Countdown field set to `mInitialMoveCountdown`. In following superframes, it shall follow the rules above.

At the end of the superframe in which a device includes a BP Switch IE with a BP Move Countdown field equal to zero, the device shall adjust its BPST based on its BPST Offset field. It may transmit a beacon in that superframe, or delay one superframe to begin beacon transmission in its new BP. After relocating its beacon to the alien BP, the device shall include neither the BP Switch IE nor the alien BP CRP IE in its beacon. If the Beacon Slot Offset field was non-zero, the device shall transmit a beacon in the beacon slot with number equal to its prior beacon slot number plus the value from the Beacon Slot Offset field. If this beacon slot number is greater than or equal to `mMaxBPLength`, the device shall follow the normal BP join rules as described in 7.3.4 to relocate its beacon to the alien BP.

7.3.7.4 BP extension

A device that receives an alien beacon with a BP Switch IE with Beacon Slot Offset field greater than zero shall set its BP length to at least the sum of the Beacon Slot Offset field and the BP length reported in the alien beacon, but not greater than mMaxBPLength.

7.3.8 Signalling window

7.3.8.1 Contention based SW (CSW)

A contention signalling window is a time window that is used for exchanging control or management information, for example, network entry messages, channel reservation requests and traffic indication. Any device may use the contention signalling window to send control/management information on demand. Different from beacon period, the whole contention signalling window is shared by all the devices opportunistically; thus it improves channel efficiency for signalling.

The length of contention signalling window is mCSWsize. The channel access method for contention signalling window is contention based. Slotted aloha is used for the contention. Based on the fact that the maximal signalling message length is much less than the maximum length of a regular Medium Access Slot (MAS), the CSW slot length (defined as mCSWSlotLength) is smaller than regular MAS slot length.

7.3.8.2 Reservation based SW (RSW)

In the master-slave mode, a nonbeaconing slave device uses RSW (Reservation-based Signalling Window), a set of reserved MASs, to exchange control or management information with the master device, or with another slave device.

The RSW slot length mRSWSlotLength is the same as that of a regular MAS. If RSW is needed, a master device makes RSW reservation and announce to all slave devices by using RSW Schedule information element in the beacon, as described in 7.1.8.18.

During this procedure, the master device determines the length of RSW by using RSW Schedule IE and allocates each MAS to slaves by using CRP IE (Sec.7.1.8.2). In CRP IE, Owner DevAddr and Target DevAddr are set to the DevAddr of reservation owner and the DevAddr of reservation target, respectively. If the reserved MAS is used for two slave devices to exchange control or management information with each other, Owner DevAddr and Target DevAddr are set to DevAddrs of those two slave devices.

If the Master is the reservation owner of an RSW slot, the Master is the device to initiate transmission of command/control frame in the RSW slot. If a slave device is the reservation owner of an RSW slot, the slave device is the device to initiate transmission of command/control frame in the RSW slot.

A slave device or a device group may request RSW slot allocation by sending an RSW Slot Request command frame (defined in 7.1.5.10) to the master device.

7.4 Device Synchronization

Each device shall maintain a beacon period start time (BPST). The device shall derive all times for communication with its neighbours based on the current BPST. The device shall adjust its BPST in order to maintain superframe synchronization with its slowest neighbour. A device shall synchronize with such a device before it sends its first beacon.

When a device receives a beacon from a neighbour, the device determines the difference between the beacon's actual reception time and the expected reception time. The beacon's actual reception time is an estimate of the time that the start of the beacon preamble arrived at the receiving device's antenna. The expected reception time is determined from the Beacon Slot Number field of the received beacon and the receiving device's BPST. If the difference is positive, then the neighbour is slower. In order to maintain superframe synchronization with a slower neighbour, the device shall delay its BPST by the difference, but limited to a maximum adjustment of mMaxSynchronizationAdjustment per superframe. The adjustment to

BPST may occur at any time following the detection of a slower device, but shall be done before the end of the superframe. Earlier adjustment of the BPST allows a tighter synchronization with the slower neighbour.

A device shall not use signalling beacon or echo beacon for synchronization.

If a device does not receive a beacon from a neighbour, the device may use historical measurements to estimate the impact on superframe synchronization and increment its BPST accordingly. This estimate is applied for up to `mMaxLostBeacons` consecutive superframes.

Beacon transmit time and measured beacon receive time shall be accurate to at least `mClockResolution`.

7.4.1 Clock accuracy

A compliant MAC sublayer implementation shall maintain a clock at least as accurate as `mClockAccuracy`. All time measurements, such as MAS boundary and frame reception time measurements, shall be measured with a minimum resolution of `mClockResolution`.

7.4.2 Synchronization for devices in hibernation mode

Devices in hibernation mode may become unsynchronized beyond the `mGuardTime` value during hibernation. A device in hibernation mode shall wake up at least one superframe before it will send a beacon and shall synchronize to the slowest clock in the beacon group during this superframe.

7.4.3 Guard times

Due to inaccuracies in superframe synchronization and drift between synchronization events, MAS start times for different devices are not synchronized perfectly. To ensure a full SIFS interval between transmissions in adjacent MASs, the devices shall maintain a SIFS interval and guard interval at the end of a reservation block. Guard times apply to all boundaries of CRP reservation blocks and BPs.

Figure 11 is an illustration of how a device uses the guard interval to maintain a SIFS interval between transmissions in adjacent reservation blocks.

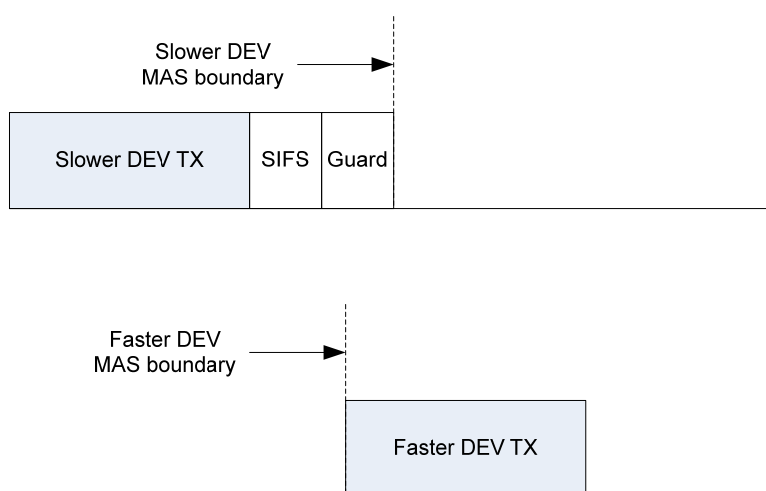


Figure 11 — Guard Time

The length of the guard interval, `mGuardTime`, depends on the maximum difference between devices' MAS boundary times. The difference arises from synchronization error and drift. The guard time is determined as follows:

$$mGuardTime = MaxSynchronizationError + MaxDrift,$$

where MaxSynchronizationError is the worst case error in superframe synchronization and MaxDrift is the worst case drift.

Synchronization is achieved during the BP as described in 7.4. For purposes of determining guard time, MaxSynchronizationError is calculated as twice mClockResolution.

Drift is a function of the clock accuracy and the time elapsed (SynchronizationInterval) since a synchronization event. The maximum drift, MaxDrift, is calculated using the worst case value for clock accuracy, mClockAccuracy, and the longest SynchronizationInterval:

$$MaxDrift = 2 \times mClockAccuracy \text{ (ppm)} \times 1E-6 \times SynchronizationInterval,$$

where

$$SynchronizationInterval = (mMaxLostBeacons+1) \times mSuperframeLength.$$

Propagation delay will also affect timing uncertainty, but in a short-range network propagation delays are small. At 10 m range, the propagation delay is around 33 ns. This is much smaller than mClockResolution and it is ignored in calculating the length of the guard interval.

A device transmitting in a reservation block may start transmission of the preamble for the first frame at the point where it calculates the start of the reservation block to be based on its local clock. For frames that use No-ACK or B-ACK acknowledgement policy, the transmitting device shall ensure that there is enough time remaining in the reservation block to transmit the frame and allow for a SIFS plus mGuardTime before the end of the reservation block as calculated by that device.




Figure 12 — SIFS and guard time in a CRP reservation block – No-ACK

If Imm-ACK is used, or a B-ACK is requested by the last frame, the transmitting device shall also ensure there is enough time for a SIFS interval, the ACK, another SIFS interval, and the guard time, as shown in Figure 13.




Figure 13 — SIFS and Guard Time in a CRP reservation block – Imm-ACK

A device shall be able to receive a frame that is transmitted within the bounds of allowable transmission, accounting for the worst case drift. A device shall begin listening mGuardTime prior to the start of a CRP reservation block, the start of a BP, or the start of a MAS in which the device announced it would be available.

7.5 Data Transfer Period

There are two basic types of channel access mechanisms for data transfer during DTP period. One is prioritized contention access (PCA), specified in 7.5.1. The second one is Channel Reservation access (CRA), specified in 7.5.2.

Data transfer period is shown in Figure 14.

All MASs within DTP are open for access via PCA unless they are reserved.

If QP is scheduled in a superframe, neither PCA nor CRA is allowed during QP.

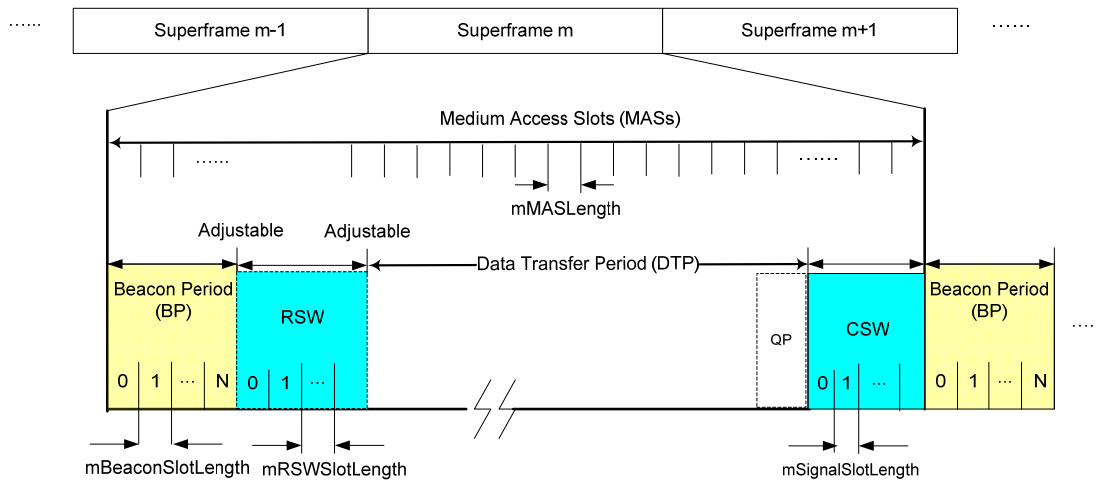


Figure 14 — Data Period

7.5.1 Prioritized Contention Access (PCA)

The PCA mechanism provides differentiated, distributed contention access to the medium for four access categories (ACs) of frames buffered in a device for transmission. A device employs a prioritized contention procedure for each AC to obtain a TXOP for the frames belonging to that AC using the PCA parameters associated with that AC.

For data and aggregated data frames, the four ACs are mapped from eight user priorities as defined in Table 115.

Table 115 — User Priority to access category mappings

Priority	User Priority (Same as 802.1D User Priority)	802.1D Designation	AC	Designation (Informative)
Lowest ↓ Highest	1	BK	AC_BK	Background
	2	-	AC_BK	Background
	0	BE	AC_BE	Best effort
	3	EE	AC_BE	Best effort
	4	CL	AC_VI	Video
	5	VI	AC_VI	Video
	6	VO	AC_VO	Voice
	7	NC	AC_VO	Voice

For command frames, any appropriate AC may be selected.

7.5.1.1 PCA medium availability

A device shall consider the medium to be unavailable for PCA at all of the following times:

- Within the device's BP or neighbours' BPs;
- Within alien BP reservation blocks announced by itself or its neighbours;

- Within RSW or CSW;
- Within regular QP or on-demand QP;
- Within hard and private reservation blocks with Reservation Status set to ONE announced by itself or its neighbours, unless the reservation block has been released;
- Within soft reservation blocks with Reservation Status set to ONE if a neighbour is the reservation target and the reservation owner is not a neighbour, unless the device is the reservation owner; and
- For a zero-length interval at the start of soft or PCA reservation blocks with Reservation Status set to ONE if a neighbour is the reservation owner, for purposes of determining TXOP limits.

At all other times, a device shall consider the medium available for PCA.

7.5.1.2 NAV

A device that transmits or receives frames using PCA shall maintain a network allocation vector (NAV) that contains the remaining time that a neighbour device has indicated it will access the medium. A device that receives a MAC header not addressed to it shall update its NAV with the received Duration field if the new NAV value is greater than the current NAV value. A device shall consider the updated NAV value to start at the end of the PLCP header on the medium.

A device that receives a MAC header with invalid HEI outside its unreleased reservation blocks shall update its NAV as if the frame were correctly received with Duration equal to mAccessDelay.

A device shall reduce its NAV as time elapses until it reaches zero. The NAV shall be maintained to at least mClockResolution.

7.5.1.3 Medium status

For PCA purposes, a device shall consider the medium to be busy for any of the following conditions:

- When its CCA mechanism indicates that the medium is busy;
- When the device's NAV is greater than zero;
- When the device is transmitting or receiving a frame on the medium;
- When the Duration announced in a previously transmitted frame has not yet expired; and
- When the medium is unavailable for PCA.

At all other times a device shall consider the medium to be idle.

7.5.1.4 PCA parameters

A device shall use the set of PCA parameters defined for an AC to obtain a TXOP or perform backoff for this AC. These parameters are summarized below. The parameter values are specified in 7.14.

7.5.1.4.1 AIFS[AC]

A device shall wait for the medium to become idle for AIFS[AC] before obtaining a TXOP or starting/resuming decrementing the backoff counter for the AC. AIFS[AC] is defined below:

$$\text{AIFS[AC]} = \text{pSIFS} + \text{mAIFSN[AC]} \times \text{pSlotTime}$$

AIFS[AC] is related to other timings as specified in Figure 15 and Figure 16.

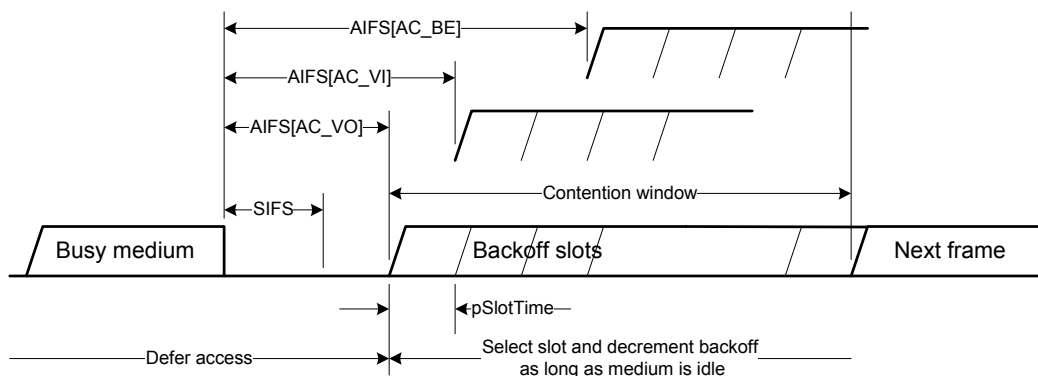


Figure 15 — IFS relationships for PCA

7.5.1.4.2 mCWmin[AC] and mCWmax[AC]

A device shall set $CW[AC]$ to an appropriate integer in the range $[mCWmin, mCWmax]$ after invoking a backoff for the AC, and shall set the backoff counter for the AC to an integer sampled from a random variable uniformly distributed over the interval $[0, CW[AC]]$.

7.5.1.4.3 mTXOPLimit[AC]

A device shall not initiate a frame transaction in a TXOP it obtained for an AC unless the frame transaction can be completed within $mTXOPLimit[AC]$ of the start of the TXOP and $pSIFS$ plus $mGuardTime$ before the medium becomes unavailable for PCA.

7.5.1.5 Obtaining a TXOP

A device shall consider itself to have obtained a TXOP for an AC if it meets the following conditions:

- The device has one or more newly arrived data frames or newly generated command frames belonging to this AC;
- The device had a backoff counter of zero value for this AC and had no frames belonging to this AC prior to the arrival or generation of the new frames;
- The device determines that the medium has been idle for $AIFS[AC]$ or longer; and
- The device has no backoff counters of zero value for other ACs, or has backoff counters of zero value for some other ACs, but such ACs have a lower priority than this AC or the device has no frames belonging to those ACs that are ready for transmission.

The device shall start transmitting a frame belonging to this AC, which may be an RTS frame, as soon as the above conditions are satisfied, subject to the criteria stated in 7.5.1.6. The device shall treat the start of the frame transmission on the wireless medium as the start of the TXOP.

A device shall also consider itself to have obtained a TXOP for an AC if it meets the following conditions:

- The device has one or more frames belonging to this AC buffered for transmission, including retry;

- The device set the backoff counter for this AC to zero in the last backoff for this AC and determines that the medium has been idle for AIFS[AC] since that backoff at the end of the current backoff slot, or the device decrements its backoff counter for this AC from one to zero in the current backoff slot; and
- The device has no backoff counters of zero value for other ACs, or has backoff counters of zero value for some other ACs, but such ACs have a lower priority than this AC or the device has no frames belonging to those ACs that are ready for transmission.

The TXOP shall start at the end of the current backoff slot, i.e., the start of the next backoff slot.

Figure 16 illustrates the timing relationships in obtaining a TXOP.

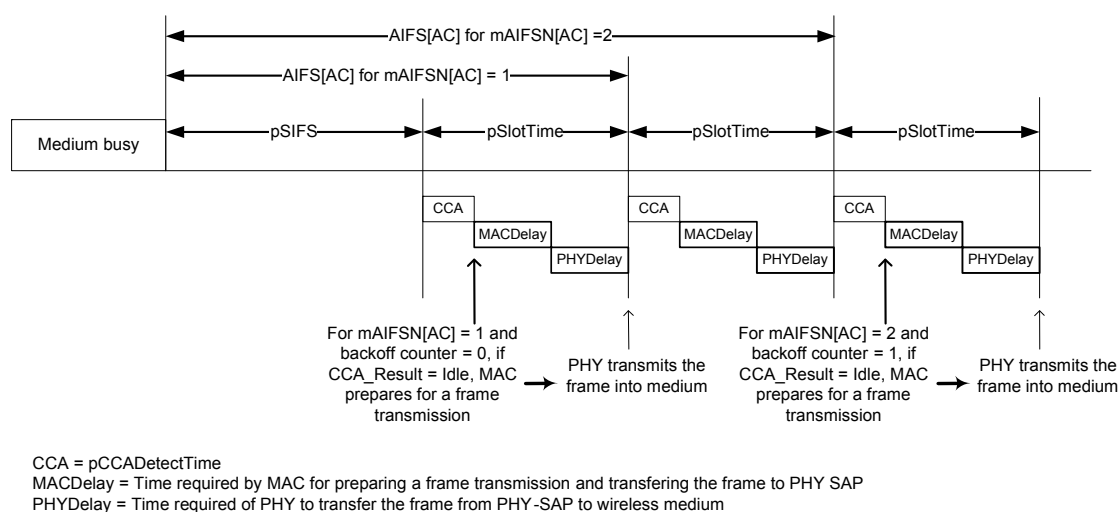


Figure 16 — PCA timing relationships

A device shall ensure that the TXOP it has obtained for an AC shall not be longer than $mTXOPLimit[AC]$ and shall end pSIFS plus mGuardTime before the medium becomes unavailable for PCA.

7.5.1.6 Using a TXOP

A device that has obtained a TXOP is referred to as a TXOP owner. A frame transmission, including a retry, is conducted as part of a frame transaction.

A TXOP owner shall initiate a frame transaction, and continue with one or more frame transactions that belong to the same AC without backoff, in the TXOP it has obtained for this AC, subject to the following criteria:

- Each transaction in the TXOP will be completed within the obtained TXOP; and
- The recipient device will be available to receive and respond during that frame transmission.

A device may retry a frame in a new TXOP that will result in the frame transaction that exceeds the $mTXOPLimit[AC]$ restriction under the following circumstances:

- The frame contains a fragment of an MSDU or an aggregate of MSDUs;
- The frame is the sole frame transmitted by the device in the current TXOP; and
- The frame transaction will be completed pSIFS and mGuardTime before the medium becomes unavailable for PCA.

A recipient device shall not transmit a CTS frame in response to a received RTS frame if its NAV is greater than zero. A recipient device shall not transmit a CTS, Imm-ACK or B-ACK response to a received frame requiring such a response if the response will not be completed pSIFS before the medium becomes unavailable for its PCA.

Under the rules stated above, the following timings apply to transmissions, including responses, in a TXOP (these timings are referenced with respect to transmission to or reception from the wireless medium):

- The TXOP owner shall transmit the first frame of the first or sole frame transaction in the TXOP at the start of the TXOP.
- After transmitting a frame with the ACK Policy set to No-ACK or B-ACK, the TXOP owner shall transmit a subsequent frame pMIFS or pSIFS after the end of that transmitted frame.
- After receiving an RTS frame or a non-RTS frame with the ACK Policy set to Imm-ACK or B-ACK Request, the recipient device shall transmit a CTS frame or an Imm-ACK or B-ACK frame pSIFS after the end of the received frame.
- After receiving an expected CTS, Imm-ACK or B-ACK response to the preceding frame it transmitted, the TXOP owner shall transmit the next frame, or retransmit a frame it transmitted earlier in the case of receiving a B-ACK, pSIFS after the end of the received frame.
- After receiving a requested B-ACK frame with a valid HEI but an invalid FCS, the TXOP owner shall retransmit the last frame it transmitted, or transmit the next frame, pSIFS after the end of the B-ACK frame.

A device shall not transmit frames using PCA to a recipient device within a MAS if the recipient has indicated in a PCA Availability IE in the current superframe that it is unavailable in that MAS. A device that indicates available MASs in a PCA Availability IE in the current superframe and does not set the TIM IE Required bit shall be available to receive frames during those MASs in that superframe if the medium is available for PCA. If the device sets the TIM IE required bit, it shall listen during indicated MASs where the medium is available for PCA in the superframe in which it received a TIM IE containing its DevAddr. If the device does not include a PCA Availability IE in its beacon, it shall be available to receive frames during all MASs available for PCA.

7.5.1.7 Invoking a backoff procedure

A device shall maintain a backoff counter for each AC to transmit frames belonging to the AC using PCA.

A device shall set the backoff counter for an AC to an integer sampled from a random variable uniformly distributed over the range $[0, CW[AC]]$, inclusive, when it invokes a backoff for this AC. The device shall initialize $CW[AC]$ to $mCWmin[AC]$ before invoking any backoff for the AC, adjusting $CW[AC]$ in the range $[mCWmin[AC], mCWmax[AC]]$, inclusive, in the course of performing PCA for the AC as described below.

The device shall set $CW[AC]$ back to $mCWmin[AC]$ after receiving a CTS or Imm-ACK frame or the MAC header of a B-ACK frame expected in response to the last transmitted frame that belonged to the AC, or upon transmitting a frame with ACK Policy set to No-ACK that belongs to the AC. A device shall also set $CW[AC]$ back to $mCWmin[AC]$, but shall not select a new backoff counter value, after discarding a buffered frame belonging to the AC.

A device shall invoke a backoff procedure and draw a new backoff counter value as specified below.

- A. A device shall invoke a backoff for an AC, with $CW[AC]$ set to $mCWmin[AC]$, when it has an MSDU arriving at its MAC SAP, or a command generated at the MAC entity, that belongs to this AC under the following conditions:
 - The device had a backoff counter of zero value for this AC but is not in the middle of a frame transaction belonging to the same AC; and

- The device determines that the medium is busy, or the device has a backoff counter of zero value for another AC, and such an AC has a higher priority than this AC and the device has frames belonging to that AC that are ready for transmission.
- B. A device shall invoke a backoff for an AC, with $CW[AC]$ set to $mCWmin[AC]$, at the end of transmitting a frame with the ACK policy set to No-ACK, or at the end of receiving an expected Imm-ACK or B-ACK response to its last transmitted frame, under the following condition:
- The device has no other frames belonging to this AC for transmission in the current TXOP obtained for this AC.
- C. A device shall invoke a backoff for an AC, with $CW[AC]$ (but not the backoff counter in general) kept to the same value for this AC, at the end of transmitting a frame with the ACK policy set to No-ACK, or at the end of correctly receiving the MAC header of an expected Imm-ACK or B-ACK response frame to its last transmitted frame, under the following conditions:
- The device has one or more frames belonging to this AC that need to be transferred over the wireless medium; and
 - The device finds that there is not enough time remaining in the current TXOP obtained for this AC to complete the pending frame transaction(s) belonging to this AC.
- D. A device shall invoke a backoff for an AC, with $CW[AC]$ (but not the backoff counter in general) kept to the same value for this AC, at the start of a TXOP obtained for the AC under the following condition:
- The device finds that there is not enough time to complete a pending frame transaction belonging to this AC in the obtained TXOP.
- E. A device shall invoke a backoff for an AC, with $CW[AC]$ set to the smaller of $mCWmax[AC]$ or $2 \times CW[AC] + 1$ (the latter $CW[AC]$ being the last CW value for this AC), at the end of the current backoff slot under the following conditions:
- The device has one or more frames belonging to this AC buffered for transmission, including retry;
 - The device set the backoff counter for this AC to zero in the last backoff for this AC and determines that the medium has been idle for $AIFS[AC]$ since that backoff at the end of the current backoff slot, or the device decrements its backoff counter for this AC from one to zero in the current backoff slot; and
 - The device has a backoff counter of ZERO value for another AC, and such an AC has a higher priority than this AC and the device has frames belonging to that AC that are ready for transmission.
- F. A device shall invoke a backoff for an AC, with $CW[AC]$ set to the smaller of $mCWmax[AC]$ or $2 \times CW[AC] + 1$ (the latter $CW[AC]$ being the last CW value for this AC), at pSIFS plus the Imm-ACK frame transmission time after the end of the last frame it transmitted, under the following condition:
- The device does not receive an expected CTS or Imm-ACK frame, or does not correctly receive the MAC header of a requested B-ACK frame by this time.

7.5.1.8 Decrementing a backoff counter

Upon invoking a backoff for an AC, a device shall ensure that the medium is idle for $AIFS[AC]$ before starting to decrement the backoff counter for the AC. To this end, a device shall define the first backoff slot to start at the time when the medium has been idle for pSIFS after the backoff invocation, as illustrated in Figure 16, with subsequent backoff slots following successively until the medium becomes busy. All backoff slots have a length of pSlotTime.

A device shall treat the CCA result pCCADetectTime after the start of a backoff slot to be the CCA result for this backoff slot. After the medium has been idle for $AIFS[AC]$ since it invokes the backoff for the AC, where

AIFS[AC] ends at the end of the current backoff slot, i.e., at the start of the next backoff slot, the device shall decrement the backoff counter for the AC (and its backoff counters for other ACs if appropriate) by one pCCADetectTime after the start of the backoff slot if it finds the CCA result to be idle at this time and determines the medium to be idle for the backoff slot, unless the backoff counter for the AC is already at ZERO value. This procedure is also illustrated in Figure 16.

The device shall freeze the backoff counter for any AC once the medium becomes busy. The device shall treat the residual backoff counter value as if the value were set due to the invocation of a backoff for the AC, following the above procedure to resume decrementing the backoff counter.

7.5.2 Channel Reservation Access (CRA)

The CRA enables a device to reserve one or more MASs that the device can use to communicate with one or more neighbours. All devices that use the CRA for transmission or reception shall announce their reservations by including CRP IEs in their beacons (see 7.1.8.2). A reservation is the set of MASs identified by CRP IEs with the same values in the Owner DevAddr, Target DevAddr, Reservation Type, and Stream Index fields.

All devices shall use CRA for channel reservation.

Reservation negotiation is always initiated by the device that will initiate frame transactions in the reservation, referred to as the reservation owner. The device that will receive information is referred to as the reservation target. This negotiation shall be done by including CRP IE defined in 7.1.8.2 in beacon frame or command frame.

7.5.2.1 Reservation type

Each CRP IE, whether included in a beacon or separately transmitted during explicit CRA negotiation, specifies a reservation type. A device shall decode all CRP IEs in all beacons received from neighbours and shall not transmit frames except as permitted by the reservation type. For all reservation types, a device shall not initiate a frame transaction in a reservation block if that transaction would not complete pSIFS plus mGuardTime before the end of the reservation block.

Reservation types are defined and summarized in Table 116.

Table 116 — Reservation Types

Reservation Type	Description	Reference
Alien BP	Prevents transmission during MASs occupied by an alien BP.	7.5.2.1.1
Hard	Provides exclusive access to the medium for the reservation owner and target; unused time should be released for PCA.	7.5.2.1.2
Soft	Permits PCA, but the reservation owner has preferential access.	7.5.2.1.3
Private	Provides exclusive access to the medium for the reservation owner and target. Channel access methods and frame exchange sequences are out of scope of this specification; unused time should be released for PCA.	7.5.2.1.4
PCA	Reserves time for PCA. No device has preferential access.	7.5.2.1.5

7.5.2.1.1 Alien BP reservations

A device shall announce an alien BP reservation to protect alien BPs as described in 7.3.7. A device shall not transmit frames during an alien BP reservation except possibly to send a beacon in the alien BP.

7.5.2.1.2 Hard reservations

In a hard reservation, devices other than the reservation owner and target(s) shall not transmit frames. Devices other than the reservation owner shall not initiate frame transactions. If there is remaining time in a reservation block that will not be used, the reservation owner and target(s) should release the reservation block by transmitting UCA and UCR frames as described in 7.5.2.8. A device considers the remainder of a reservation block available, subject to other medium access rules, after it has received a UCA or UCR frame that releases the reservation block and the duration indicated in that received frame has expired.

A device shall not transmit a data or aggregated data frame in a hard reservation unless the Delivery ID field is set to a Stream Index that is the same as the Stream Index for the reservation and the DestAddr of the frame is the same as the Target DevAddr for the reservation or the DestAddr of the frame matches the DevAddr of any target of an established multicast reservation. A device may transmit any command or control frame in a hard reservation.

7.5.2.1.3 Soft reservations

In a soft reservation, devices access the medium following PCA rules. The reservation owner may access the medium with the highest priority AIFS and without performing backoff. It may begin transmission at the beginning of each reservation block. It may initiate an additional frame transaction after any transaction it initiated but shall not initiate such a transaction later than pSIFS after the end of the previous frame transaction. The reservation owner shall not transmit a data or aggregated data frame without backoff unless the Delivery ID field is set to a Stream Index that is the same as the Stream Index for the reservation and the DestAddr of the frame is the same as the Target DevAddr for the reservation or the DestAddr of the frame matches the DevAddr of any target of an established multicast reservation. The reservation owner may transmit any command or control frame without backoff. Neighbours of a reservation owner shall follow PCA rules to access the medium. Neighbours of a reservation target that are not neighbours of the reservation owner shall not access the medium.

7.5.2.1.4 Private reservations

The channel access method and frame exchange sequences used during a private reservation are out of the scope of this Standard. Standard frame formats and frame types shall be used during a private reservation. In a private reservation, neighbours of the reservation owner and target(s) shall not transmit frames. If there is remaining time unused during a reservation block, the reservation owner and target(s) should release the reservation block by transmitting UCA and UCR frames. Devices consider the remainder of the reservation block available, subject to other medium access rules, after they have received a UCA or UCR frame that releases the reservation block and the duration indicated in that received frame has expired, as described in 7.5.2.8.

7.5.2.1.5 PCA reservations

During a PCA reservation, any device may access the medium using PCA rules.

7.5.2.2 CRP Availability IE

The CRP Availability IE identifies the MASs where a device is able to establish a new CRP reservation.

The combination of information from CRP Availability IEs and CRP IEs allows an owner to determine an appropriate time for a new CRP reservation. In order to facilitate the CRP negotiation process, devices that are aware of existing neighbours' CRP reservations shall mark the reserved MASs as unavailable.

A device shall mark a MAS unavailable if the device includes it in a CRP IE with the Reservation Status bit set to ONE. It shall mark a MAS unavailable if a neighbour includes it in a CRP IE with a target other than the device, whether the Reservation Status bit is ZERO or one. It shall mark a MAS unavailable if any BP occupies any portion of that MAS, based on information in any beacon received in the latest $mMaxLostBeacons+1$ superframes.

7.5.2.3 CRP reservation negotiation

There are two mechanisms used to negotiate a reservation: explicit and implicit. For explicit negotiation, the reservation owner and target use CRP Reservation Request and CRP Reservation Response command frames to negotiate the desired reservation. For implicit negotiation, the reservation owner and target use CRP IEs transmitted in their beacons. For either negotiation mechanism, the reservation owner completes the negotiation by including an appropriate CRP IE in its beacon.

A device shall not negotiate for MASs that are included in a CRP IE received from a neighbour or any other CRP IE included in the device's beacon, unless the MASs are referenced only in a CRP IE with Reason Code set to Denied.

A device shall announce in the MAC Capabilities IE in its beacon whether it is capable of explicit CRP negotiation. A device shall not initiate an explicit CRP negotiation with devices that do not support it.

A device shall only initiate negotiation for a reservation as the reservation owner.

For reservations of type Alien BP, there is no negotiation with neighbours. A device shall include the appropriate CRP IE with Reservation Status set to ONE on detection of an alien BP, as specified in 7.3.7.

For reservations of type PCA, there is no negotiation with neighbours. A device may select any available MAS to include in a reservation of type PCA. The device may also select MASs included in a neighbour's reservation of type PCA. The device shall not set the Reservation Status bit to ONE in a PCA reservation unless it included a CRP IE in its beacon in the previous superframe that identified the same MASs, with Reservation Type set to PCA and Reservation Status set to ZERO or ONE.

7.5.2.3.1 Negotiation

When negotiating a reservation, the reservation owner shall set the Target DevAddr field of the CRP IE to the DevAddr of the reservation target. It shall set the Reservation Status bit to ZERO and the Reason Code to Accepted in the CRP IE. For new streams, the Stream Index shall be set to a value that is currently not used with this Target DevAddr and has not been used as such for $mMaxLostBeacons+1$ superframes. To negotiate additional MASs for an existing stream, the Stream Index shall be set to the value used for the existing stream.

A reservation owner shall not transmit unicast frames within reserved MASs in a hard, soft, or private reservation unless it and the recipient included CRP IEs with the Reservation Status bit set to ONE in their most-recently transmitted beacons.

When negotiating a reservation, a reservation target shall set the Owner DevAddr field of the CRP IE to the DevAddr of the reservation owner. If a unicast reservation is granted, it shall set the Reservation Status bit to ONE and the Reason Code to Accepted. If a multicast reservation is granted, it shall set the Reservation Status bit to the same value included in the CRP IE by the reservation owner, and shall set the Reason Code to Accepted. If the reservation is not granted, it shall set the Reservation Status bit to ZERO. If the reservation cannot be granted due to a conflict with its own or its neighbours' reservations, the reservation target shall set the Reason Code to Conflict. If the reservation is not granted, it shall set the Reason Code to Denied. If the reservation target cannot grant the reservation immediately, it may set the Reason Code to Pending, and deliver a final response later. For a unicast reservation, the reservation target shall set the CRP Allocation fields to match those in the request. For a multicast reservation, it shall set the CRP Allocation fields to match the request, or to include a subset of the MASs included in the request.

7.5.2.3.2 Explicit negotiation

To start explicit CRP negotiation, the reservation owner shall send a CRP Reservation Request command frame to the target device, as defined in 7.1.5.1.

On reception of a CRP Reservation Request command the reservation target shall send a CRP Reservation Response command, as defined in 7.1.5.2, to the reservation owner. The fields in the CRP IE shall be set according to 7.5.2.3.1. If the reservation cannot be granted due to a conflict with its own or its neighbours'

reservations, the reservation target shall include a CRP Availability IE in the CRP Reservation Response command frame.

In a CRP Reservation Response command frame for a multicast reservation, the reservation target shall include a CRP Availability IE for a Reason Code other than Denied. Final multicast reservations are established implicitly, as described in 7.5.2.3.3.

7.5.2.3.3 Implicit negotiation

Implicit negotiation is carried out by transmitting CRP IE(s) in beacon frames. A device that supports the CRP shall parse all beacons received from neighbours for CRP IE(s) whose Owner DevAddr field and Target DevAddr field match either the device's DevAddr or a multicast DevAddr for which the device has activated multicast reception. From this initial selection, the device shall process the CRP IE(s) that are new with respect to CRP IE(s) included in the most recently received beacon from the same device as a CRP reservation request or a CRP reservation response.

To start implicit negotiation, a reservation owner shall include a CRP IE that describes the proposed reservation in its beacon. The device shall continue to include the CRP IE for at least $mMaxLostBeacons+1$ consecutive superframes or until a response is received.

On reception of a unicast CRP reservation request in a beacon, the reservation target shall include a CRP reservation response in its beacon no later than the next superframe, with fields set as described in 7.5.2.3.1. If the Reason Code indicates Conflict, the reservation target shall include a CRP Availability IE in its beacon.

As long as the reservation owner includes a unicast CRP reservation request in its beacon, the reservation target shall continue to include the CRP reservation response in its beacon. The reservation target shall not change the Reservation Status bit to ONE if there is a reservation conflict with its neighbours.

On reception of a multicast CRP reservation request, a reservation target shall include a reservation response CRP IE in its beacon no later than the next superframe if it is a member of the targeted multicast group. The fields in the CRP IE shall be set according to 7.5.2.3.1. If the Reservation Status bit in the response is ZERO, the reservation target shall include a CRP Availability IE in its beacon unless the Reason Code is set to Denied.

A device that elects to receive traffic in an already established multicast reservation does not negotiate the reservation. To join an established multicast reservation that does not conflict with other existing reservations, a device shall include corresponding CRP IE(s) in its beacon with Reservation Status bit set to ONE and Reason Code set to Accepted.

A device that cannot join an established multicast reservation because of an availability conflict may inform the source by including the corresponding CRP IE(s) in its beacon with Reservation Status bit set to ZERO, and the Reason Code set to Conflict. The device shall also include the CRP Availability IE in the beacon.

7.5.2.3.4 Negotiation conclusion

To conclude negotiation for a unicast reservation, the reservation owner shall set Reservation Status to ONE in the CRP IE in its beacon after receiving a beacon from the reservation target that contains a corresponding CRP IE with Reservation Status set to ONE. To conclude negotiation for a multicast reservation, the reservation owner sets Reservation Status to ONE in a CRP IE in its beacon in the next superframe after transmitting the same CRP IE with Reservation Status set to ZERO, regardless of responses from potential multicast recipients. If a reservation conflict exists, the reservation owner shall not set the Reservation Status bit to ONE except as specified in 7.5.2.5.

7.5.2.3.5 CRA negotiation between a nonbeaconing device and a master

If a master device is the reservation target and a nonbeaconing slave device is the reservation owner, the channel reservation is described as below. A nonbeaconing device shall scan at least $mMaxLostBeacons+1$ superframes before sending CRP Reservation request command frames to make sure the requested MASS

are not reserved by other devices. The slave device transmits a CRP Reservation Request command frame in an RSW slot. After receiving the CRP request, the master device replies CRP Reservation Response command frames in an RSW slot or include CRP IE in the beacon of next superframe to negotiate the desired reservation (see 7.5.2.3). The Master completes the negotiation by including an appropriate CRP IE in its beacon.

If a master device is the reservation owner and a nonbeaconing slave device is the reservation target, the channel reservation is described as below. The master device initiates channel reservation request either by including CRP IE in its beacon or in a CRP Request command frame in an RSW slot. If the slave device receives CRP IE in the beacon, the slave device replies Reservation Response command frame in an RSW slot. If the slave device receives CRP Request command frame in an RSW slot, the slave device replies CRP Reservation Response command frame in the same or another RSW slot. The Master completes the negotiation by including an appropriate CRP IE in its beacon.

After establishing reservation between a nonbeaconing slave device and its master, the slave device shall either transmit echo beacon regularly, or promote itself as a regular beaconing device, to announce its reservations. If the slave initiates the beaconing promotion process, the slave sends beaconing promotion request command frame to its master, with the reason code set as CRA by a slave device, as defined in 7.1.5.11. After receiving the beaconing promotion request, the master shall send beaconing promotion indication IE in its beacon with device address in the IE set as the requesting slave device to be promoted as regular beaconing slave device, which is defined in 7.1.8.23. Alternatively, the master may initiate beaconing promotion by sending beaconing promotion indication IE in its beacon with device address in the IE set as the target slave device to be promoted as regular beaconing slave device, which is defined in 7.1.8.23. After receiving beaconing promotion indication IE from a master, the slave device starts to join beacon period, following procedures specified in 7.3.

7.5.2.3.6 CRA negotiation between nonbeaconing slave devices

For CRA between nonbeaconing slave devices, to establish reservations, a slave device as reservation owner first sends RSW slot request in the CSW to ask the master to allocate RSW slot. The RSW slot request indicates the reservation target (i.e., the other nonbeaconing slave device) of the CRA. Then the master allocates RSW slots to the requesting slave device to proceed with link measurement with the reservation target. When RSW slots are allocated for link measurement, reservation owner transmits probe command frame without information elements to the reservation target to request link measurement. After link measurement, the reservation target replies with probe command frame with Link Feedback IE or Link Quality Estimate IE to report measurement results. Once link measurement is done, the slave device as reservation owner transmits CRP reservation request command frame to the reservation target. The reservation target replies with CRP Reservation Response. If the reservation is successful, the master will include CRP IEs in its beacon to announce the CRP reservation between the two nonbeaconing slave devices by indicating the reservation owner and reservation target. After receiving CRP reservation announced by the master, the nonbeaconing slave devices can start data communication between each other. If the reservation negotiation between nonbeaconing slave devices fails or the reservation announcement is not heard from master device, the nonbeaconing slave devices may retry the negotiation again after mCRPslaveRetry superframes.

After establishing reservation between nonbeaconing slave devices, both reservation owner and reservation target shall either transmit echo beacon regularly, or promote themselves as regular beaconing devices, to announce their reservations. If a slave initiates the beaconing promotion process, the slave sends beaconing promotion request command frame to its master, with the reason code set as CRA by a slave device, as defined in 7.1.5.11. After receiving the beaconing promotion request, the master shall send beaconing promotion indication IE in its beacon with device address in the IE set as the requesting slave device to be promoted as regular beaconing slave device, which is defined in 7.1.8.23. Alternatively, the master may initiate beaconing promotion by sending beaconing promotion indication IE in its beacon with device address in IE set as the target slave device to be promoted as regular beaconing slave device, which is defined in 7.1.8.23. After receiving beaconing promotion indication IE from a master, the slave device starts to join beacon period, following procedures specified in 7.3.

7.5.2.4 CRP reservation announcements

Once negotiation for a reservation successfully completes, the reservation owner and target shall include CRP IE(s) in their beacons that describe the reservation. Within each CRP IE, the Reason Code shall be set to Accepted and the Reservation Status bit shall be set to ONE. The devices shall include the CRP IEs in each beacon transmitted until the reservation is modified or terminated.

7.5.2.5 Resolution of CRP reservation conflicts

Devices engaged in independent CRP negotiation could attempt to reserve the same MAS, or due to mobility, devices could have reserved the same MAS. A conflict exists between CRP reservations if a MAS is included in both reservations. A device might detect a conflict during a CRP negotiation or after a reservation has been established. Reservations of type Alien BP never conflict with other reservations of type Alien BP. Reservations of type PCA never conflict with other reservations of type PCA.

A device shall apply the following rules to a conflict between a CRP IE included in its beacon and another CRP IE included by a neighbour:

- a) If the device's reservation is of type Alien BP, the device shall maintain the reservation.
- b) If the neighbour's reservation is of type Alien BP, the device shall not transmit frames in conflicting MASs. If the device is the reservation target, it shall also set the Reason Code in its CRP IE to Conflict.
- c) If the device's CRP IE has the Reservation Status bit set to ZERO and the neighbour's CRP IE has the Reservation Status bit set to ONE, the device shall not set the Reservation Status bit to ONE and shall not transmit frames in conflicting MASs. If the device is the reservation target, it shall also set the Reason Code in its CRP IE to Conflict.
- d) If the device's CRP IE has the Reservation Status bit set to ONE and the neighbour's CRP IE has the Reservation Status bit set to ZERO, the device may maintain the reservation.
- e) If the device's CRP IE and neighbour's CRP IE have the Reservation Status bit set to the same value and one of the following conditions is true, the device may maintain the reservation.
 - 1) The device's CRP IE and neighbour's CRP IE have the Conflict Tie-breaker bit set to the same value and the device's occupied beacon slot number is lower than the beacon slot number of the neighbour; or
 - 2) The device's CRP IE and neighbour's CRP IE have the Conflict Tie-breaker bit set to different values and the device's occupied beacon slot number is higher than the beacon slot number of the neighbour.
- f) If the device's CRP IE and neighbour's CRP IE have the Reservation Status bit set to ZERO and ONE of the following conditions is true, the device shall not set the Reservation Status bit to ONE. If the device is the reservation target, it shall set the Reason Code in its CRP IE to Conflict.
 - 1) The device's CRP IE and neighbour's CRP IE have the Conflict Tie-breaker bit set to the same value and the device's occupied beacon slot number is higher than the beacon slot number of the neighbour; or
 - 2) The device's CRP IE and neighbour's CRP IE have the Conflict Tie-breaker bit set to different values and the device's occupied beacon slot number is lower than the beacon slot number of the neighbour.
- g) If the device's CRP IE and neighbour's CRP IE have the Reservation Status bit set to ONE and one of the following conditions is true, the device shall not transmit frames in conflicting MASs. It shall remove the conflicting MASs from the reservation or set the Reservation Status to ZERO. If the device is the reservation target, it shall set the Reason Code in its CRP IE to Conflict.

- 1) The device's CRP IE and neighbour's CRP IE have the Conflict Tie-breaker bit set to the same value and the device's occupied beacon slot number is higher than the beacon slot number of the neighbour; or
- 2) The device's CRP IE and neighbour's CRP IE have the Conflict Tie-breaker bit set to different values and the device's occupied beacon slot number is lower than the beacon slot number of the neighbour.

When a reservation owner withdraws a reservation or part of a reservation due to a conflict, it shall invoke a backoff procedure prior to requesting additional MASs in any reservation. The device shall initialize the backoff window *BackoffWin* to *mCRPBackoffWinMin*. When the backoff algorithm is invoked, the device shall select a random number *N* uniformly from $[0, \text{BackoffWin}-1]$. The device shall not request additional MASs for *N* superframes. If a further negotiation fails due to a conflict, the device shall double *BackoffWin*, up to a maximum of *mCRPBackoffWinMax*. After a negotiation completes, the device shall generate a new backoff *N*. If a device does not request any MASs for $4 \times \text{BackoffWin}$ superframes, the device may terminate this backoff procedure and request MASs at any time unless another conflict occurs.

If a reservation target sets Reason Code to Conflict in any CRP IE in its beacon, it shall include a CRP Availability IE in the same beacon.

7.5.2.6 BPST realignment and existing CRP reservations

A device that realigns its BPST as described in 7.3.7 may assert new CRP reservations with Reservation Status bits set to ONE in the new beacon so long as they are equivalent to its old CRP reservations with the Reservation Status bit set to ONE in the prior BP. For this purpose, two CRP reservations are equivalent if their corresponding Owner DevAddr, Target DevAddr, Stream Index, and Reservation Type fields are the same and the number of MASs claimed by the new reservation is less than or equal to the number claimed by the old reservation.

A device that realigns its BPST shall not assert CRP reservations with MASs that conflict with any BP it announced or detected. The device shall not assert CRP reservations with MASs that conflict with reservations with Reservation Status equal to ONE announced in the new BP unless no other MASs are available. Any conflict with existing reservations shall be resolved according to the procedures specified in 7.5.2.5.

7.5.2.7 Modification and termination of existing CRP reservations

A reservation owner may reserve additional MASs for a stream by negotiating an addition to the reservation using a CRP IE with the same Owner DevAddr, Target DevAddr, Stream Index, and Reservation Type. Once negotiation has completed successfully, the reservation owner shall combine the CRP IEs. When combining CRP IEs, the reservation owner shall set the Reason Code to Modified until a CRP IE is received from the reservation target that describes the combined reservation.

A reservation owner may remove MASs from an established reservation without changing the Reservation Status bit in the CRP IE. If a reservation owner removes some MASs from an established reservation, it shall set the Reason Code in its CRP IE to Modified until the reservation target has changed its CRP IE to match.

A reservation target may remove MASs from an established reservation without changing the Reservation Status bit in the CRP IE due to a conflict, as described in 7.5.2.5, or due to reception of a Relinquish Request IE. If the reservation target is unicast, the reservation owner shall remove the same MASs from the reservation, or terminate the reservation.

To terminate a reservation, the reservation owner shall remove the CRP IE from its beacon.

If a reservation owner changes or removes a CRP IE, the reservation targets shall update or remove corresponding CRP IE from their beacons in the current or following superframe.

To terminate a reservation, a reservation target shall set the Reservation Status bit to ZERO and the Reason Code to an appropriate value, as if responding to an initial reservation request.

If a reservation owner or target does not receive a beacon or any other frame from the other participant in the reservation for more than `mMaxLostBeacons` superframes, it shall consider the reservation terminated, and shall remove the corresponding CRP IE(s) from its beacon.

7.5.2.8 Release of hard or private CRP reservation blocks

If time remains in a hard or private CRP reservation block after a reservation owner completes transmission of associated buffered traffic, it should release the reservation block by sending a UCA frame. If the remaining time in the reservation block is not sufficient for the exchange of UCA and UCR control frames, no action should be taken. The transmitter of the UCA control frame shall include a list of device(s) that shall respond to this announcement. The list should consist of those devices that have previously included the corresponding CRP IE(s) in their beacons. The order in which the DevAddrs of the devices are mentioned in the list is the order in which they shall respond with a UCR control frame. This allows devices around the transmitter as well as the devices in the list to be informed about the early end of the reservation block.

On reception of a UCA control frame, a device shall check whether its DevAddr is included in the device list of the UCA control frame. If its DevAddr is included in the list it shall respond to the UCA control frame with a UCR control frame after a delay given by:

Time to send Response = $pSIFS + pSlotTime + (Position_in_list_in_UCA) \times (UCR_control_frame_duration + pSIFS)$

Time to send Response is calculated from the end of reception of the UCA control frame. Possible values of `Position_in_list_in_UCA` are in the range $[0, N-1]$, inclusive.

UCA and UCR control frames release the time between the end of PLCP header of the last UCR control frame, as indicated by the Duration value in the MAC header of UCA and UCR control frames, and the end of the reservation block. Other MASSs described by the reservation that do not belong to the current reservation block, either in the same superframe or following superframes, are not released.

The Duration value in the UCA control frame shall cover the UCA control as well as all expected UCR control frames. The Duration value in the UCR control frames shall be set to the Duration value in the UCA control frame minus the time between the end of the PLCP header of the UCA control frame and the end of the respective UCR control frame. This value is given by the following equation:

Duration value in UCR = Duration value in UCA – (UCA_frame_body_transmission_time + $pSIFS + pSlotTime + (Position_in_list_in_UCA) \times (UCR_control_frame_transmission_time + pSIFS)$) – UCR_control_frame_transmission_time.

7.5.2.9 Retransmit procedures in CRP reservations

In a hard CRP reservation block, if the reservation owner transmits a frame with ACK Policy set to Imm-ACK or B-ACK, but does not receive the expected acknowledgement frame, it may retransmit the frame within the same reservation block if the reservation block has not been released.

Devices other than the reservation owner that retransmit frames in a soft CRP reservation block shall follow the PCA rules specified in 7.5.1.

A device shall not retransmit a frame earlier than $pSIFS$ after the end of an expected acknowledgement or CTS frame, whether or not it receives the expected frame. A device shall not retransmit a frame in the current reservation block if there is not enough time remaining in the reservation block for the entire frame transaction.

7.6 Fragmentation and Aggregation

7.6.1 Fragmentation and reassembly

A source device may fragment each MSDU/MCDU.

A device shall not fragment any MSDU/MCDU to more than `mMaxFragmentCount` fragments. Fragments may be of varying sizes. Once the MSDU/MCDU is fragmented and a transmission attempted, the device shall not refragment the frame. The device shall not create frame fragments smaller than `mMinFragmentSize`.

The device shall set the Fragment Number field in the first fragment to zero. It shall set each subsequent fragment to the Fragment Number field in the previous fragment plus one. The device shall not increment the Fragment Number field when a fragment is retransmitted.

A device shall assign the same Sequence Number to all fragments of an MSDU/MCDU.

The device shall completely reassemble an MSDU/MCDU in the correct order before delivery to the MAC client. The device shall discard any MSDU/MCDU with missing fragments. If the No-ACK policy is used, the recipient device shall discard an MSDU/MCDU immediately if a fragment is missing. Otherwise, a recipient device shall discard the fragments of an MSDU if the MSDU is not completely received within an implementation-dependent timeout.

If B-ACK is used, unacknowledged fragments from multiple MSDUs belonging to the same stream may be retransmitted in the same sequence. In this case it is the responsibility of the recipient device to deliver the MSDUs in the correct order to the MAC client.

If a source device discards a fragment of an MSDU/MCDU, the device shall discard all fragments of the MSDU/MCDU.

7.6.2 Aggregation

A transmitter may aggregate multiple MSDUs with identical Delivery ID into a single frame payload. A device shall aggregate no more than `mAggregationLimit` MSDUs into an aggregated data frame.

The MAC header of aggregated data frame is specified in Table 117.

Table 117 — MAC header of Aggregated Data Frame

MAC header field	Value
Protocol version	0
Secure	Specified in 7.1.2.1.2
ACK Policy	Specified in 7.1.2.1.3
Frame Type	4
Frame Subtype/Delivery ID	Specified in 7.1.2.1.5
Retry	Specified in 7.1.2.1.6
DestAddr	DevAddr of the receiver
SrcAddr	DevAddr of the transmitter
Sequence Control	Specified in 7.2.9.3
Duration	Specified in 7.2.9.1

The aggregated data frame payload format is defined in Table 118. The frame payload size for an aggregated data frame is subject to the same maximum size as any frame payload.

The aggregation header format is defined in Table 119. The MSDU Count field contains the number of MSDUs included in the aggregated data frame. The Length fields in the Aggregation Header field indicate the length in bytes of the corresponding MSDUs.

Table 118 — Aggregated data frame payload format

Syntax	Size	Notes
Aggregated_Data_Frame_Payload_Format {		
Aggregation header	1+2xN bytes	Defined in Table 119. N equals the number of aggregated MSDUs.
For(i=0; i < N; i++){		
MSDU _i	<i>variable</i>	
}		
}		

Table 119 — Aggregation header format

Syntax	Size	Notes
Aggregation_Header_Format {		
MSDU Count (= N)	1 byte	
For(i=0; i < N; i++){		
Length of MSDU _i	2 bytes	
}		
}		

7.7 ARQ, Multirate Support and Power Control

This subclause define the policy for ARQ, multi-rate and transmit Power control.

7.7.1 ARQ Policies

Three acknowledgement policies are supported by this Standard; no acknowledgement (No-ACK), immediate acknowledgement (Imm-ACK) and block acknowledgement (B-ACK).

A device shall acknowledge all received unicast frames with the ACK Policy field set to either Imm-ACK or B-ACK Request and DestAddr set to the DevAddr of this device. The device shall acknowledge the reception without regard to security validation. A device that receives a frame with a higher Protocol Version than it supports shall discard the frame without acknowledgement.

7.7.1.1 No-ACK

A frame with ACK policy set to No-ACK, shall not be acknowledged by the recipient. The transmitting device MAC entity assumes the frame has been successfully transmitted and proceeds to the next frame upon completion of current frame. All broadcast and multicast frames shall have ACK Policy set to No-ACK.

7.7.1.2 Immediate ACK

On reception of a frame with ACK Policy set to Imm-ACK, a device shall respond with an Imm-ACK frame, transmitted pSIFS after the end of the received frame.

7.7.1.3 Block ACK

The B-ACK mechanism allows a source device to transmit multiple frames and to receive a single acknowledgement frame from the recipient indicating which frames were received and which need to be retransmitted.

A source device initiates the use of the B-ACK mechanism with a recipient device for frames either from the same stream or of the same user priority. If the recipient device accepts use of the B-ACK mechanism, it indicates the maximum number and size of the frames it can buffer. The source device transmits a sequence of frames to the recipient, each from the same stream or of the same user priority, limited by the announced

buffer size and maximum number of frames. The initial frames in the sequence are all transmitted with ACK Policy set to B-ACK. The final frame in the sequence is transmitted with ACK Policy set to B-ACK Request. On receipt of such a frame, the recipient device returns a B-ACK frame giving feedback on the frames received and indicating the buffer space available for the next B-ACK sequence.

A source device may invoke multiple instances of the B-ACK mechanism with the same recipient device, each for a different stream or user priority. A source device may also invoke the B-ACK mechanism with multiple recipient devices.

7.7.1.3.1 Initiation

A source device may activate the B-ACK mechanism independently for any stream or user priority traffic to any potential recipient device advertising B-ACK capability in its MAC Capabilities IE. A source device shall initiate use of the B-ACK mechanism by transmitting a frame with ACK Policy set to B-ACK Request to the recipient device. A source device shall use a dedicated Sequence Number counter for each stream or user priority traffic using the B-ACK mechanism with a recipient. After transmitting the frame, the source device shall follow the rules of operation as described in 7.7.1.3.2.

When receiving a frame with ACK Policy set to B-ACK Request from a source device for a stream or user priority traffic not currently using the B-ACK mechanism, the recipient device shall respond as follows:

- To acknowledge receipt of the frame but reject the request for starting a new instance of B-ACK mechanism, the recipient device shall respond with a B-ACK frame with no frame payload.
- To accept the request for starting a new instance of B-ACK mechanism, the recipient device shall respond with a B-ACK frame with a frame payload indicating the allowed maximum size (in frames and octets) for the next B-ACK sequence. The recipient shall acknowledge the received frame by indicating its reception in the acknowledgement window.

A recipient device may also accept a request to use the B-ACK mechanism even if the request frame has an invalid FCS. To accomplish this, the recipient device shall respond with a B-ACK frame with a frame payload that indicates the allowed maximum size for the next B-ACK sequence, but without acknowledgement of the frame with the invalid FCS.

A recipient device, even though it advertises B-ACK capability in its MAC Capabilities IE, may reject a request to use the B-ACK mechanism for any reason, including a temporary unavailability of resources or a lengthy setup process requiring a delayed start time. Thus, after being rejected, a source device may keep trying to initiate use of the B-ACK mechanism by sending the next frame with ACK Policy set to B-ACK Request.

7.7.1.3.2 Operation

After transmitting a frame with ACK Policy set to B-ACK Request, the source device expects to receive a B-ACK frame in response and takes one of the following actions:

- If the source device does not receive a B-ACK frame, it shall assume that the recipient device did not receive the request frame. To continue B-ACK operation, the source device shall retransmit the request frame with the same ACK Policy using applicable medium access rules as described in 7.5.1 and 7.5.2.
- If the source device receives a B-ACK frame with no frame payload, it shall treat the transmitted frame as received and consider this use of the B-ACK mechanism to be terminated.
- If the source device receives a B-ACK frame with a frame payload and with either Frame Count or Buffer Size set to zero, it shall process the acknowledgement as described below. To continue the B-ACK operation, the source device shall retransmit the requesting frame with the same ACK Policy, independently of whether the frame was indicated as received or not. If the requesting frame was indicated as received, the source device alternatively may transmit a zero-length payload frame with the same Sequence Control and Delivery ID to the recipient device.

- If the source device receives a B-ACK frame with a frame payload containing non-zero values for both Frame Count and Buffer Size, then it shall process the acknowledgement as described below. To continue the B-ACK operation, the source device shall send frames with ACK Policy set to B-ACK or B-ACK Request as described below.

The source device processes the B-ACK frame acknowledgement as follows:

- Frames being held for retransmission with a sequence number earlier than the one indicated by the Sequence Control field were not received in the last B-ACK sequence, but shall not be retransmitted.
- Frames being held for retransmission with sequence and fragment number within the acknowledgement window (specified by the Sequence Control field and the Frame Bitmap field) with corresponding bit set to one were received and shall not be retransmitted.
- Other frames being held for retransmission should be retransmitted in the next sequence, ordered by increasing sequence and fragment numbers.

After receiving a B-ACK frame with non-zero values for Frame Count and Buffer Size, the source device may transmit a sequence of frames. Each sequence of frames shall consist of zero or more frames with ACK Policy set to B-ACK followed by a single frame with ACK Policy set to B-ACK Request. The total number of frames must not exceed the Frame Count value specified in the B-ACK frame and the sum of the lengths of the frame payloads shall not exceed the Buffer Size value specified in the B-ACK frame. The sequence of frames may be transmitted in multiple PCA TXOPs or CRP reservation blocks and may be interleaved with frames to other recipients or of other streams or user priorities, subject to all the medium access rules. Within a sequence, the frames shall be ordered by increasing sequence and fragment numbers. Due to retransmissions, this ordering might not hold from one sequence to the next and frames transmitted within a sequence might not have consecutive sequence and fragment numbers.

When the recipient device receives a frame with ACK Policy set to B-ACK Request, it shall respond using SIFS with a B-ACK frame. To continue operation, the B-ACK frame shall contain a frame payload. If the recipient device receives a frame with a valid HEI but an invalid FCS and with ACK Policy set to B-ACK Request, the device shall also respond with a B-ACK frame with a frame payload. Within the B-ACK frame payload, the recipient device shall set the Frame Count and Buffer Size fields to limit the size of the next sequence of frames. It shall also set the Sequence Control and Frame Bitmap fields to indicate to the source device which frames should be retransmitted.

A recipient device may implement a timeout that indicates when to stop waiting for missing frames, allowing some MSDUs to be released to the MAC client and B-ACK buffer resources to be freed. A recipient device may also implement a timeout to expire an instance of the B-ACK mechanism that appears to be inactive.

7.7.1.3.3 Termination

To terminate use of the B-ACK mechanism, the source device shall transmit a frame from the appropriate stream or of the appropriate user priority to the recipient device with ACK Policy set to anything other than B-ACK or B-ACK Request.

The recipient device may terminate use of the B-ACK mechanism by responding to a frame with ACK Policy set to B-ACK Request with a B-ACK frame with no frame payload.

7.7.2 Multi-rate Support

A device shall transmit beacons at pBeaconTransmitRate.

Devices shall transmit non-beacon frames only at data rates supported by the intended recipient, based on information from the recipient's PHY Capabilities IE.

A recipient device may use the Link Feedback IE to suggest the optimal data rate to be used by a source device, for example, to increase throughput and/or to reduce the frame error rate. The data rate in the Link

Feedback IE is interpreted as the maximum data rate that the source device should use for this particular link, for an acceptable frame error rate.

7.7.3 Transmit Power Control

A device includes Channel Classification IE (defined in 7.1.8.17) to update others regarding the transmit power limit of the operating channel and the backup channels.

A device shall transmit at the lowest power required for reliable communication. A device may recommend a transmit power level change to be used by another device by including a Link Feedback IE (defined in 7.1.8.4) or Transmit Power Control IE (defined in 7.1.8.39) in its beacon. The method to determine transmit power recommendations is out of the scope of this Standard, but the recipient device might use the signal to noise ratio, received signal strength, frame error ratio or other parameters to determine the transmit power change to recommend to the source device.

7.8 Dynamic Channel Selection

Initial channel selection during network entry is described in 7.13. After initial channel selection, a device may change channel due to the detection of incumbents or load balancing. Channel change due to incumbent detected is specified in 7.11.3.

For load balancing and/or better link quality, a device may also change channel. A device may move to a new channel in order to find enough channel resources for its own communication. A device may also volunteer to or be enforced to move to a new channel in order to release channel resources in the old channel to others. A device shall include Channel Change IE in its beacon or channel switch request command frame (as defined in 7.1.5.8) to announce or request channel change. As defined in 7.1.8.11, Channel Change IE includes a list of devices who are invited to change channel. The invited devices are normally the associated devices. Different from the case of channel evacuate which is caused by incumbent detection, a device receiving Channel Change IE is not required to move to new channel unless it is in the device list and it has agreed to move. A device shall include a Channel Change Response IE (as defined in 7.1.8.12) in its beacon or in its Channel Switch Response command frame (as defined in 7.1.5.9) as response to receiving channel change request from either beacon or Channel Switch Command frame. A device that includes a Channel Change Response IE shall change channels as indicated in the IE.

One aspect to load balancing is determining the availability and the load of each individual channel. Nodes may determine the load on a particular channel from the outband channel status report (as defined in 7.1.8.9) sent by a proxy (as defined in 7.1.8.14). Alternatively, a node itself may visit a outband channel and analyze the beacon frames transmitted during a BP.

7.9 Power Management Mechanisms

7.9.1 Power management modes

There are two power management modes:

- Active mode: the device will send or receive beacon(s) or other frames in the current superframe.
- Hibernation mode: the device will not send or receive a beacon or other frames in the current superframe.

Before entering hibernation mode, a device shall announce in previous superframe(s) that it is entering hibernation mode. Note: a master device shall always operate in active mode.

7.9.2 Device power states

Active mode devices are in one of two power states within a superframe:

- Awake: device is able to transmit and receive.

- Sleep: device does not transmit or receive.

A device that is changing from Sleep to Awake state and has a frame in queue to transmit using PCA shall perform CCA until a MAC header is received or up to `mAccessDelay` time, before determining the medium state.

The value of `mAccessDelay` is equal to the time required to transmit one maximum length frame, transmitted at the lowest mandatory data rate, plus the time to transmit an Imm-ACK plus pSIFS.

7.9.3 Power state transitions

Active mode devices shall transition between Awake and Sleep states according to the following rules:

- a) A device shall be in the Awake state `mGuardTime` + `mCSWsize` prior to its BPST in every superframe to participate in the transmission and reception of beacons.
- b) If a device has data traffic pending to be transmitted in CRP reservations in the current superframe, it shall be in Awake state prior to the start of each relevant CRP reservation block to start its transmission. The device may go into Sleep state for the rest of the reservation block if all the pending transmissions completed successfully. The device should release the CRP reservation block before entering Sleep state. A device shall set the More Frames bit to ZERO in the last frame it transmits during a reservation block.
- c) If a device has data traffic pending to be transmitted via PCA in the current superframe, it shall signal its intent with a relevant TIM IE (as defined in 7.1.8.26) in its beacon. Once all of its transmissions are completed, the device may go into Sleep state for the rest of the superframe, subject to other rules in this subclause. A device shall set the More Frames bit to ZERO in the last frame it transmits to a particular recipient using PCA during a superframe.
- d) If a device is expecting to receive transmissions from other devices in a CRP reservation block, as indicated in the beacons of those devices, it shall be in Awake state `mGuardTime` prior to the start of the reservation block for the reception of the planned transmission. It may go into Sleep state either at the end of the reservation block or when all the pending data has been received, as indicated by the More Frames bit. If the device receives a UCA frame, it shall not go into Sleep state until after transmitting a corresponding UCR frame.
- e) If a device is expecting to receive transmissions from other devices via PCA, it may include a PCA Availability IE in its beacon and shall be in Awake state `mGuardTime` prior to the announced MASs. A device that does not include a PCA Availability IE in its beacon shall be in Awake state in all MASs available for PCA in the current superframe. Once all pending data has been received, as indicated by the More Frames bit, the device may go into Sleep state for the rest of the superframe.

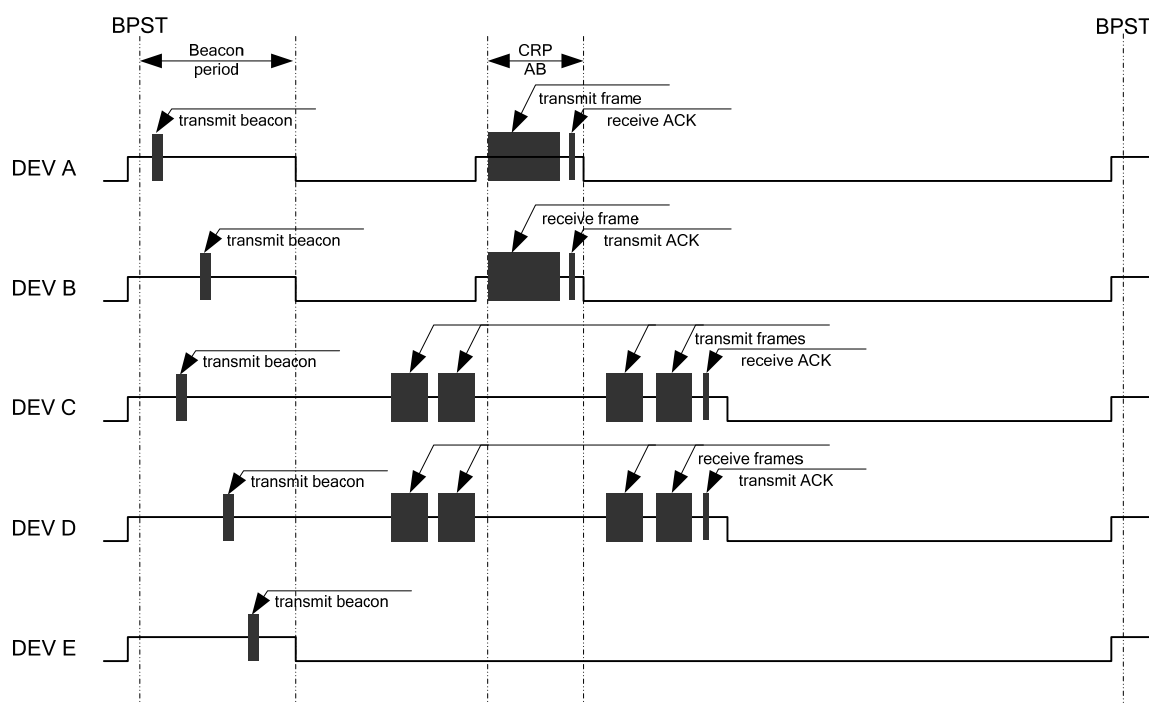


Figure 17 — Power state transition for devices in active mode

Figure 17 illustrates the power state transition for devices in active mode.

- DEV A is a device that has data traffic pending to be transmitted in a CRP reservation block in the current superframe.
- DEV B is a device that is expecting to receive a planned transmission from DEV A in a CRP reservation block in the current superframe.
- DEV C is a device that has data traffic pending to be transmitted via PCA in the current superframe.
- DEV D is a device that is expecting to receive a planned transmission from DEV C via PCA in the current superframe.
- DEV E is a device that does not have any traffic pending in its transmission queues, and is not expecting any planned transmission from other devices.

7.9.4 Hibernation mode operation

A device using hibernation mode shall operate according to the following rules:

- A device shall signal its intent to go into hibernation mode by including a Hibernation Mode IE (as defined in 7.1.8.25) in its beacon or probe command frame. The Hibernation Duration field in the Hibernation Mode IE shall contain a non-zero value that specifies the duration of the hibernation period. A device may signal its intent to go into hibernation mode in several superframes. The value of the Hibernation Countdown field in the Hibernation Mode IE shall be set to indicate the number of remaining superframes before the device enters hibernation mode. In each successive superframe, the device shall reduce the value of the Hibernation Countdown field by one. If this field is set to zero, the device enters hibernation mode at the start of the next superframe.
- When in hibernation mode, the device shall not send a beacon or other traffic. The device should terminate all established CRP reservations before entering hibernation.

- A device may leave hibernation mode prior to the end of its announced hibernation period by sending its beacon.
- A device in hibernation mode shall scan for beacons during the BP for one or more superframes immediately prior to the end of its hibernation period, in order to re-establish synchronization.
- If a device in hibernation mode finds that its former beacon slot is still available in the extended beacon group, the device may transmit a beacon in that beacon slot. Otherwise, the device shall transmit a beacon as if it was doing so for the first time.

Active mode devices in the presence of hibernation mode devices shall operate as follows:

- If an active mode device receives a neighbour's beacon that includes a Hibernation Mode IE, the device shall consider all CRP reservations with that neighbour to be terminated at the start of its hibernation period. An active mode device shall not commence any communication with a hibernation mode device until that device leaves hibernation mode. After receiving a beacon that includes a Hibernation Mode IE with Hibernation Countdown less than or equal to `mMaxLostBeacons`, an active mode device that misses the remaining expected beacons shall consider the device to be in hibernation mode as indicated in the Hibernation Mode IE.
- If an active mode device does not receive an expected beacon from a hibernation mode device, it shall treat the beacon slot of that device as occupied and non-movable, but shall not indicate the beacon slot as occupied by the hibernation mode device in its BPOIE, until the beacon is received, for up to `mMaxHibernationProtection`. During a neighbour's hibernation period an active mode device shall continue to mark the hibernation mode device's beacon slot as occupied and non-movable in its BPOIE. If the active mode device receives another neighbour's beacon in the hibernation mode device's beacon slot, the device shall still advertise the hibernation mode device's `DevAddr` in its BPOIE.
- If an active mode device has unicast traffic for a hibernation mode device, it should buffer its traffic until the hibernation mode device enters active mode.
- If an active mode device has multicast or broadcast traffic it should not delay transmission of the traffic, even if it is aware that some intended recipients are in hibernation mode. It may buffer its multicast traffic for a hibernation mode device until the intended recipient enters active mode, and then deliver the buffered multicast data.

7.9.5 Hibernation anchor operation

Active mode devices that are capable of acting as a hibernation anchor should indicate hibernation anchor capability in its MAC Capabilities IE. In a master-slave network, the master is the default hibernation anchor for slave devices. In a peer-to-peer network, a peer device may act as hibernation anchor for its neighbours.

A device that indicates such capability shall include a Hibernation Anchor IE (as specified in 7.2.10.19) in its beacon to convey information about neighbours in hibernation mode. A device may terminate its role as a hibernation anchor at any time, but at that time it shall remove indication of the capability from its MAC Capabilities IE, as defined in 7.1.8.21.

Devices, such as those that were recently off or in hibernation mode, might not have information about the hibernation state of their neighbours. These devices may use the information provided by Hibernation Anchor IEs for scheduling communication with neighbours in hibernation mode.

Upon reception of a beacon containing a Hibernation Mode IE in which the Hibernation Countdown is set to zero, a hibernation anchor shall include a Hibernation Anchor IE. It shall set the Wakeup Countdown field in the Hibernation Anchor IE based on the Hibernation Duration field in the received Hibernation Mode IE. It shall decrement the Wakeup Countdown field in each successive superframe until the field reaches zero. After it transmits a beacon with a Hibernation Anchor IE that contains a Hibernation Mode Device Information field with Wakeup Countdown set to zero, it shall remove the corresponding Hibernation Mode Device Information field from the Hibernation Anchor IE. It shall not include a Hibernation Anchor IE if there are no Hibernation Mode Device Information fields in the IE.

If the hibernation anchor receives a beacon from a hibernation mode device prior to the end of the announced hibernation duration, the hibernation anchor shall remove the corresponding Hibernation Mode Device Information field from the Hibernation Anchor IE in the next beacon.

After receiving a neighbour's beacon that includes a Hibernation Mode IE with Hibernation Countdown less than or equal to `mMaxLostBeacons`, a hibernation anchor device that misses the remaining beacons from the neighbour shall consider the device to be in hibernation mode as indicated in the Hibernation Mode IE and shall include that device in the Hibernation Anchor IE.

7.10 Probe

The Probe IE and Application-specific Probe IE are used in beacons and probe commands to request one or more IEs from the target device identified in the probe IE. Target devices are not required to respond with all requested IEs. If a target device supports the Probe command frame for one or more IEs, it shall set the Probe bit in its MAC Capabilities IE to ONE, or otherwise it shall set the bit to ZERO. A device shall include a MAC Capabilities IE or a PHY Capabilities IE in its beacon if it is the target of a Probe IE received in a beacon that includes the MAC Capabilities IE Element ID or the PHY Capabilities IE Element ID, respectively. On reception of either probe IE in a beacon, a target device shall include a response in its beacon for the next `mMaxLostBeacons` superframes. On reception of either probe IE in a Probe command frame, a target device shall respond with a Probe command frame addressed to the sender within one superframe or include a response in its beacon for the next `mMaxLostBeacons` superframes. In the Probe command frame or beacon, the target device shall include:

- A Probe IE, with Target DevAddr set to the DevAddr of the requestor, that includes no Requested Element IEs to reject the probe; or
- One or more requested IEs.

7.11 Protection of incumbents

In this clause, we specify necessary functions and procedures to fully protect incumbent services. We specify channel measurement procedures in 7.11.1, channel classification procedures in 7.11.2, and channel evacuation procedures in 7.11.3. Pre-service channel scanning will be discussed in network entry and initialization subclause of the Standard in 7.13.

7.11.1 Channel Measurement

A device needs to measure the operating channel to determine its vacancy. A device should also measure adjacent channel in order to determine the maximal allowed transmission power as specified in 7.7.3. A device may also measure other channels to determine the channel availability for future use.

7.11.1.1 Operating Channel Measurement

In order to detect the presence of an incumbent at very low signal strengths, all the neighbouring devices operating in the same channel shall be quiet at the same time, so that the sensing function is effective. We refer to this as the quiet period (QP). Quiet period (QP) could be scheduled regularly and/or on-demand.

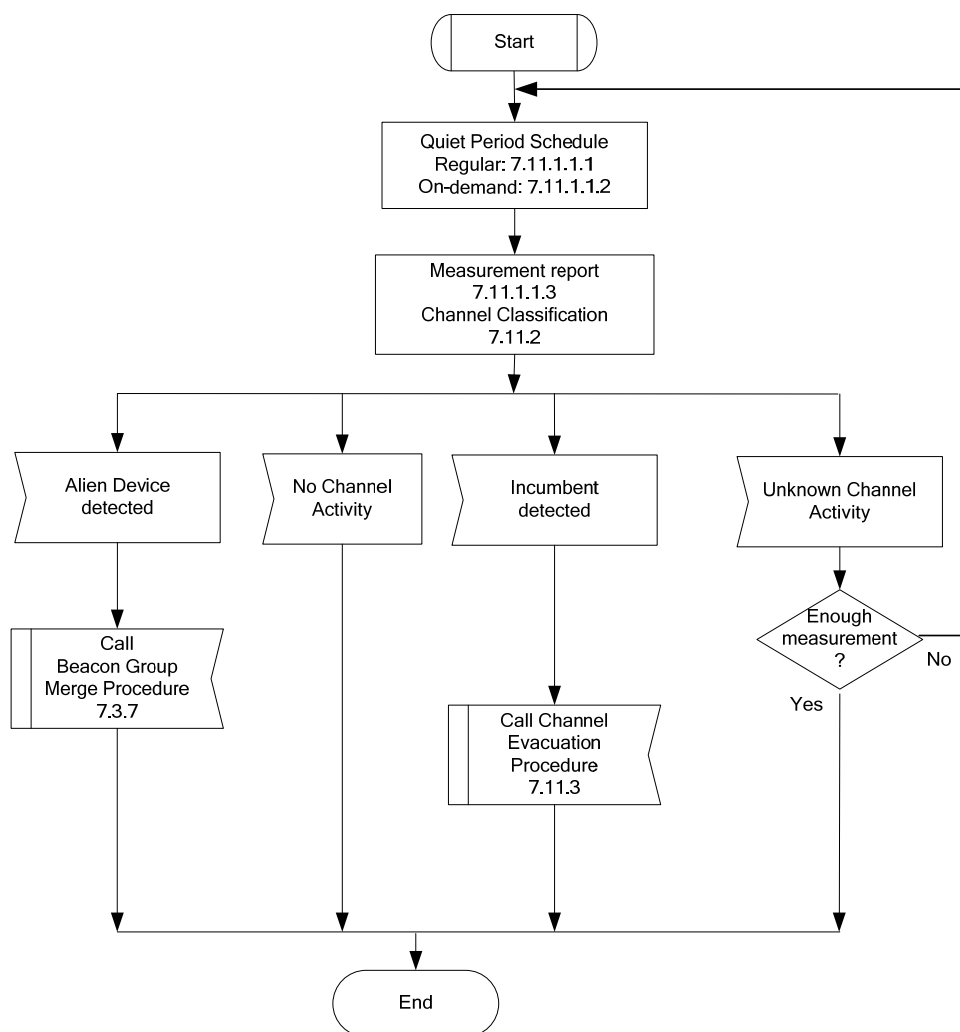


Figure 18 — In-service Channel Monitor

7.11.1.1.1 Regular quiet period

Regular QP is scheduled periodically, once for every mQPfrequency superframes. At a superframe with regular QP, the regular QP precedes the CSW and the duration is mQPD. Devices sharing the same superframe shall follow the same regular QP schedule. A device shall include Regular QP schedule IE (specified in 7.1.8.5) in its beacon if it is a regular beaconing device. When the Countdown field in the regular QP schedule IE is set to 0, any device by default shall keep silent during the QP in current superframe.




Figure 19 — Illustration of quiet period schedule

7.11.1.1.2 On-demand quiet period

A device may activate on-demand sensing based on its own observation, for example, the sensing result of regular QP. In addition, abnormal behaviour driven fine sensing is also defined. Abnormal behaviour is defined as follows; 1) beacon collision, 2) data packet error or 3) substantial SINR change.

This Standard defines three methods to schedule on-demand QP in a timely and network-aware manner with least intervention to other devices.

- a) Use own reservation. If a device has already established a certain portion of reservation (for example, for data communication), the device may utilize its own reservation without spending time to make new reservation for sensing, thus speeding up on-demand sensing. A device may inform others, normally via beacon, about its intention to use its own reservation for sensing. Other devices then may take advantage of the reservation for sensing if they desire to do. A device shall include an on-demand QP schedule IE as defined in Section 7.1.8.6 in its beacon to announce the on-demand sensing schedule.
- b) Make reservation of idle slots. In this method, a device may make reservation of idle slots for on-demand sensing normally via its beacon. Different from normal channel reservation for communication, a device does not need another device to confirm the reservation request with this special reservation type. All the devices receiving the reservation request shall mark the reservation in their reservation information maps. The device making the reservation may refresh such reservation. A device shall include an on-demand QP schedule IE as defined in Section 7.1.8.6 in its beacon to make the on-demand sensing schedule.
- c) Borrow others' reservation. In this method, a device may negotiate with a neighbour to temporally silence the neighbour's reservation for on-demand sensing. To initiate the negotiation, a device sends On-demand QP Negotiation Request IE (as specified in 7.1.8.27) to the device which owns the reservation. On-demand QP Negotiation Request IE shall be transmitted in a beacon or probe command frame. To respond to the negotiation, the reservation owner sends back the On-demand QP Negotiation Response IE (as specified in 7.1.8.28). The On-demand QP Negotiation Response IE shall be transmitted in either beacon or probe command frame. The reservation owner shall give priority to on-demand QP compared with other regular data communications. Once negotiation is successful, the on-demand QP scheduler shall include an on-demand QP schedule IE as defined in Section 7.1.8.6 in its beacon to make the on-demand sensing schedule.

The methods a) and b) will not interrupt others so they are preferred. Method c) is suggested to use only for urgency.

A device other than the on-demand QP scheduler is not required to perform sensing. However, other devices than the scheduler may take such opportunity to perform on-demand sensing to increase its own confidence level of incumbent detection.

7.11.1.1.3 Operating Channel Measurement Report

Once a device detects incumbents, it shall notify others in the following contention signalling window, reservation signalling window, or beacon period. The corresponding information element, Channel Measurement Report IE, is defined in 7.1.8.9. If Channel Measurement Report IE is transmitted in contention signalling window or reservation signalling window, it shall be included in Channel Measurement Report command frame defined in 7.1.5.6. If Channel Measurement Report IE is transmitted in beacon period, it shall be included in regular beacon frame defined in 7.1.3.1. A regular beaconing device should use beacon period to transmit Channel Measurement Report IE. A device other than regular beaconing device should use reservation signalling window to transmit Channel Measurement Report IE. A device other than regular beaconing device may use contention signalling window to transmit Channel Measurement Report IE only in an emergency.

7.11.1.2 Outband channel Measurement

Outband channel measurement is important for channel classification and channel re-selection. Outband channel measurement is also important for adjacent channel operation.

7.11.1.2.1 Outband channel Measurement Request

Outband channel Measurement is coordinated by a master device or a peer device. In a master-slave network, a master device coordinates its slave devices to measure outband channel by using Channel Measurement Request IE, as defined in 7.1.8.7. Channel Measurement Request IE shall be included in a regular beacon frame (defined in 7.1.3.1) or in a Channel Measurement Request command frame (defined in 7.1.5.4). A slave device receiving Channel Measurement Request IE shall confirm the request either by including Channel Measurement Response IE (as defined in 7.1.8.8) in regular beacon or by transmitting Channel Measurement Response command frame, as defined in 7.1.5.5.

In a peer-to-peer network, a peer device may request a group of other peer devices following the same approach as master-slave.

If an outband channel is empty (i.e., not used by other networks) based on previous measurements, a device may visit the empty channel according to its own schedule. To determine whether an outband channel is empty, a device may need to monitor the outband channel for at least one superframe once in a while.

If an outband channel has an existing beacon group, a device should monitor the beacon period of the outband channel to get traffic information and monitor quiet period to detect incumbent.

7.11.1.2.2 Outband channel Measurement Report

A device reports channel measurement to a master device or a peer device by using Channel Measurement Report IE, as defined in 7.1.8.9. Channel Measurement Report IE shall be included in a regular beacon frame (defined in 7.1.3.1) or in a Channel Measurement Report command frame (defined in 7.1.5.6).

A master device or a peer device acknowledges the reception of Channel Measurement Report by using Channel Measurement Acknowledgement IE (as defined in 7.1.8.10) in regular beacon or by using Channel Measurement Acknowledgement command frame, as defined in 7.1.5.7.

7.11.2 Channel Classification

A master device or a peer device classifies channels based on either channel measurement report, or information gathered from geo-location database, or both. Channels are classified as six types - Disallowed, Operating, Backup, Candidate, Protected and Unclassified. Channel classification IE is defined to exchange channel classification status among devices and specified in Table 78, 7.1.8.17. A master or a peer device transmits its generated channel classification IE in its beacon to other devices when needed.

The channels may be classified using the following categories:

- Disallowed: channels that are excluded from use due to operational or regulatory constraints. Disallowed channels need not to be sensed.
- Operating: the current channel used for communication between devices.
- Backup: channels that may become the operating channel immediately in case the current network needs to switch to another channel. The master or peer device may maintain pre-defined number of backup channels at any given time and shall order them according to their relative priorities.
- Candidate: channels that are candidates to become a backup channel. Although backup and candidate channels must be incumbent-free, Backup channels have higher priorities to be used than candidate channels. Candidate channels require less frequent channel information update (including channel measurement) than backup channels.
- Protected: channels in which incumbent or alien network operation has been detected. If an alien network has been detected, the channel could be used as operating channel by self-coexistence mechanisms as specified in 7.12. If incumbent has been detected, the channel shall be protected from being used as operating channel. Once incumbent is gone, the protected channel may be reused as an operating channel.
- Unclassified: channels that have not been classified.

The channel classification algorithm is outside the scope of this Standard.

When a slave device receives channel classification IE from its master, the slave device shall update its channel list according to the classification.

When a peer device receives channel classification IE from another peer device, the peer device shall compare the received channel classification with its own channel classification. If there is discrepancy between these channel classifications, the device shall consolidate in the way that incumbents shall be fully protected.

7.11.3 Channel Evacuation

Upon discovery of an incumbent, a device shall suspend data communication and limit its total transmission time (including for beacons, control messages) to under `mMaxTransTime`. A device needs to evacuate the channel within `mMaxEvacuateCountdown` after detecting incumbents.

Once an incumbent is detected and confirmed, the decision to evacuate this channel may be unilateral. However to maintain communication, a communicating group of which a member detects incumbent shall coordinate to evacuate the channel.

In a master-slave network, the master coordinates the group to move to a backup channel if available. The master transmits the Channel Change IE, as defined in 7.1.8.11, in its beacon. The master may transmit Channel Change IE multiple times to increase its reliability. The Channel Change IE includes the New Channel Number and Channel Change Countdown, which is the remaining time until the device evacuates the current channel. If a backup channel is available, the New Channel Number shall be set to the backup channel number selected by the master. All slave devices which receive the Channel Change IE shall move to the backup channel within the time defined in Channel Change IE. Furthermore, if the backup channel is empty, to speed up recovery, the same channel reservation settings used in the old channel should be used in the new channel. The master enables this by setting Channel copy to resume field defined in Table 67 to 1. Otherwise, after moving to the backup channel, a device shall join a beacon group by following the beacon group joining procedures as specified in 7.3.4 or by following the beacon group merger procedures as specified in 7.3.7. If backup channel is not available, the New Channel Number shall be set to 255 (Unavailable). All devices sending or receiving the Channel Change IE shall leave current channel within the time defined in Channel Change IE and perform channel SCAN as specified in 7.13 to find new channel.

In a peer-to-peer network, the reservation owner shall coordinate the communication group to move to a backup channel if available. The device coordinating channel evacuation transmits the Channel Change IE, as defined in 7.1.8.11, in its beacon. The device coordinating channel evacuation transmits Channel Change IE multiple times to increase its reliability. The Channel Change IE includes the New Channel Number and Channel Change Countdown, which is the remaining time until the device evacuates the current channel. If a backup channel is available, the New Channel Number shall be set to the backup channel number selected by the device coordinating channel evacuation. Other peer devices which receive the Channel Change IE shall move to the backup channel within the time defined in Channel Change IE, if they are in the same communication group as the device sending Channel Change IE. Furthermore, if the backup channel is empty, to speed up recovery, the same channel reservation settings used in the old channel should be used in the new channel. The device coordinating the channel evacuation enables this by setting Channel copy to resume field defined in Table 67 to 1. Otherwise, after moving to the backup channel, a device shall join a beacon group by following the beacon group joining procedures as specified in 7.3.4 or by following the beacon group merger procedures as specified in 7.3.7. If backup channel is not available, the New Channel Number shall be set to 255 (Unavailable). All devices sending or receiving the Channel Change IE shall leave current channel within the time defined in Channel Change IE and perform channel SCAN as specified in 7.13 to find new channel.

In case that the incumbent signal is too strong so that a device cannot exchange beacon/control message, the device shall move to the pre-agreed backup channel (if available) automatically within mMaxEvacuateCountdown duration

7.12 Self-coexistence

7.12.1 Self-coexistence scenarios

This Standard defines three basic self-coexistence scenarios:

- Self-coexistence between two master-slave based networks
- Self-coexistence between two peer-to-peer networks
- Self-coexistence between a peer-to-peer network and a master-slave-based network.

This Standard provides distributed self-coexistence mechanisms in 7.12.2 and centralized self-coexistence mechanisms in 7.12.3.

7.12.2 Distributed self-coexistence mechanisms

The distributed self-coexistence mechanisms include:

- Group Discovery and Notification: this mechanism enables detection of alien beacon group through beacons, and notification through beacons and control/command frames.
- BP merging rules: a set of rules used by devices to decide whether and how they adjust their BP to align with the BP and superframe structure of an alien beacon group.
- Beaconsing Device Promotion (BDP): This mechanism is used to decide which devices are promoted as regular beaconsing devices. Although beacons are important for device discovery, QoS and coexistence, they increase the control overhead. Hence, it is also important to dynamically control the number of beaconsing devices to reach a tradeoff between performance and overhead. Typically, in peer-to-peer operation mode all devices are regular beaconsing devices. However, in master-slave mode, some slave devices may be promoted as regular beaconsing devices by the Master device defines the beaconsing devices according to the BDP mechanism specified in 7.12.2.3.

7.12.2.1 Group Discovery and Notification

The detection is mainly done through regular beacons or echo beacons.

The notification/report process allows devices to announce newly discovered alien beacon group A device shall report alien beacon group by sending Channel Measurement Report IE in its beacons or Channel Measurement Report command frames.

7.12.2.2 Beacon Period Merging Rules

When devices operating on the same channel but with distinct and, possibly, unsynchronized superframe structures (also called Alien BP), discover each other, they must decide whether to merge their BPs. The rules and procedures of BP merging are specified in 7.3.7.

7.12.2.3 Beaconsing Device Promotion (BDP)

In a master-slave network, a slave device may be promoted as a regular beaconsing slave device by its master based on the following rules. The beaconsing promotion process is either initiated by a slave device or by its master. If a slave device observes alien beacons or interference other than incumbents, it shall initiate the beaconsing promotion process. The slave device sends beaconsing promotion request command frame to its master, with the reason code set as self-coexistence, as defined in 7.1.5.11. After receiving the beaconsing promotion request, the master shall send beaconsing promotion indication IE in its beacon with device address in the IE set as the requesting slave device to be promoted as a regular beaconsing slave device, which is defined in 7.1.8.23. Alternatively, a master may initiate beaconsing promotion by sending beaconsing promotion indication IE in its beacon with the address set as the target slave device to be promoted as a regular beaconsing slave device, which is defined in 7.1.8.23. After receiving beaconsing promotion indication IE from a master, the slave device starts to join beacon period, following procedures specified in 7.3.

7.12.3 Centralized self-coexistence mechanisms

7.12.3.1 Self-coexistence between two master-slave based networks

There are two schemes for coexistence between two master-slave-based networks. The first scheme is to merge them into a single master-beacon network. The second scheme is to merge them using two master-beacons (i.e., BP length is greater than 2), based on the distributed self-coexistence mechanism which is specified in 7.12.2. In the first scheme, the DME of one master decides to change itself as slave device (Note: this capability is implementation dependent). As a result, the devices associating with the master changing operation mode as slave (including the master itself) rejoin the other network as slave devices. The remaining master device continues to send its master-beacon which now covers all devices of the original two networks.

7.12.3.2 Self-coexistence between two peer-to-peer networks

There are two basic schemes for coexistence between two peer-to-peer networks. The first scheme is based on the distributed self-coexistence mechanism which is specified in Section 7.12.2. The second scheme is that one peer device elects itself as the master device and other peer devices transition into slave devices and associate themselves with the master device. Note: this operation that a peer device transitions into a master device or a slave device is done at DME level and the capability to support that transition is implementation dependent. The new master device sends its master-beacon which now covers all devices of the original two peer-to-peer networks.

7.12.3.3 Self-coexistence between a master-slave network and a peer-to-peer network

There are two basic schemes for coexistence between a master-slave network and a peer-to-peer network. The first scheme is based on the distributed self-coexistence mechanism which is specified in Section 7.12.2. The second scheme is that all peer devices transition into slave devices and associate themselves with the master device. The master device continues to send its master-beacon which now covers all devices of the original two networks. Note: this operation that a peer device transitions into a slave device is done at DME level and the capability to support that transition is implementation dependent.

7.13 Network Entry and Initialization

Before a device starts e.g. AV streaming or data communication, it needs to scan the channel and associate with the target device(s) and ensure incumbent protection.

This Standard does not presuppose any pre-assigned channel where a device is able to find the target device(s) given the time-varying and unpredictable nature of channel occupancy.

Peer devices should perform the following network entry and initialization as follows:

1. Perform Initial Channel SCAN and Device Discovery, see 7.13.1, if pair device is unknown; otherwise perform pair discovery;
2. Create/join a beacon group, see 7.13.4;
3. Perform pairing, which includes association, authentication and key establishment, see 7.13.5;
4. If indicated as desired by the device during association, perform other initialization procedures such as negotiating basic capabilities;
5. Setup connections, see 7.13.6.

A pair device is defined as the target device with which a device normally communicate.

Figure 20 summarizes the network entry procedure carried out by peer devices.



1. Perform Initial scan channel and device discovery, see 7.13.1;
2. Create/join a beacon group and send master beacon, see 7.13.4;
3. Perform master-slave association, see 7.13.2;
4. If indicated as desired during association, perform other initialization procedures such as negotiating basic capabilities;
5. Setup connections, see 7.13.6.

112

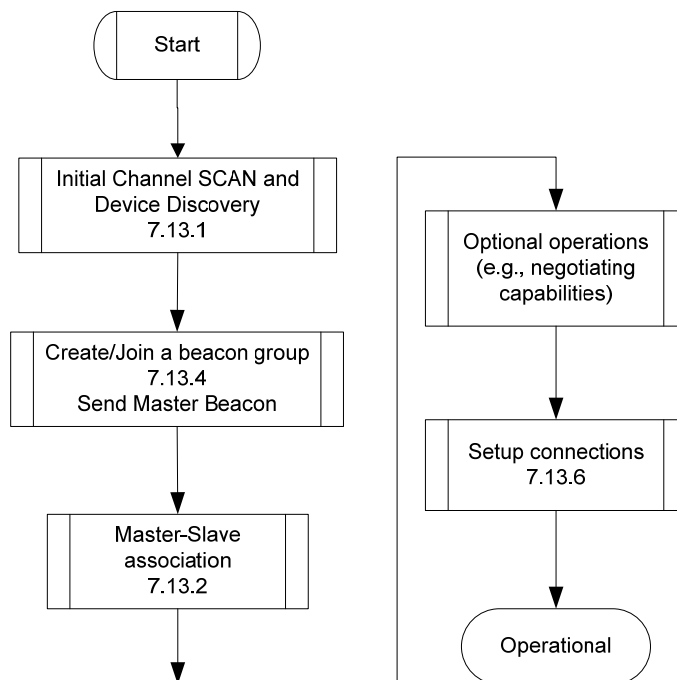


Figure 21 — Network Entry Procedure for A Master Device

For a slave device, the network entry procedure is specified in 7.13.2.

Note that each of these steps taken by a device consists of a set of actions and error verification. The following subclauses specify these steps and their individual responsibilities in detail.

7.13.1 Initial Channel SCAN and Device Discovery

Figure 22 illustrates Initial Channel SCAN and Device Discovery procedure. The scan duration for each channel, i.e., T1, is twice of SuperframeLength.

Energy detection as well as feature detection may be applied to detect channel activities, which could be incumbent presence, device presence, unknown presence, and Clean pending.

Incumbents may be detected by feature detection. If incumbents are present on the current channel, the channel should not be revisited within certain duration, e.g., 10 minutes.

Device presence could be detected by receiving beacons. Once beacons received, the device should listen to and/or probe devices which are sending beacons to get device address, name string, and to derive channel busy ratio.

If the detected energy level is higher than certain level, subject to regulation, but neither an incumbent nor a device can be identified, the channel shall be marked as unknown presence. In case of unknown presence, a device shall not use the channel unless further measurements are taken to make sure if it is incumbent safe.

If energy level is lower than certain level, the channel might be clean. A device may use the channel later but further measurements may be necessary to make sure the judgment is reliable.

If energy level is lower than certain level, the channel may be marked as available.

Upon finishing the initial Channel SCAN and Device Discovery procedure, a device shall classify the channel per 7.11.2.

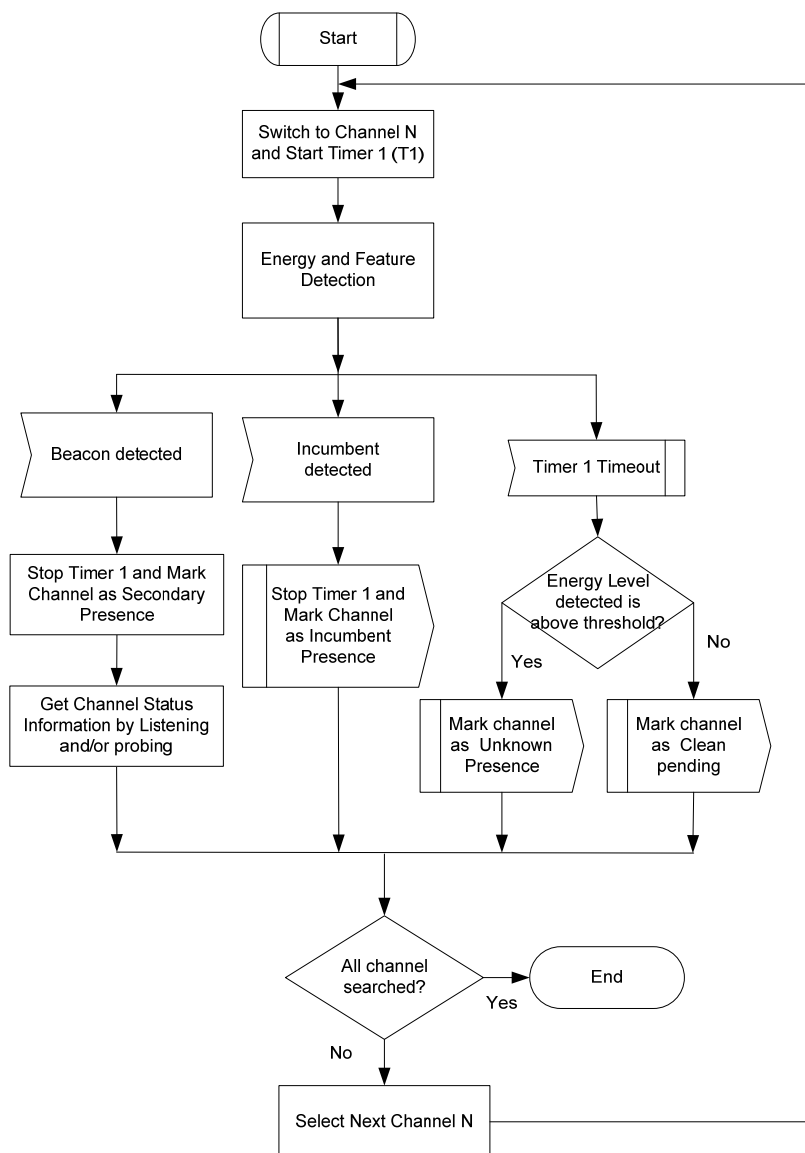


Figure 22 — Initial Channel SCAN and Device Discovery

7.13.2 Master-Slave Association

A un-associated slave device may initiate master-slave association when it detects a master device during network entry. The slave device shall send association request command frame (as specified in 7.1.5.12) to the master device in a contention signalling window. After processing the request, the master device confirms the request implicitly using an association request IE (as specified in 7.1.8.24) included in beacon. A master-slave association procedure is illustrated in Figure 23.

If a slave device sends association request with address that is already used, the master device sends association response IE with “Fail-Invalid Address”. If there are not enough MASs, the master device sends association response IE with “Fail-Not enough MAS”. If the slave device receives association response IE with “Success”, then association is completed.

A slave device that has been associated with a master device is normally not a beaconing device so that BP joining process is not required. However, to assist self-coexistence, a slave device may be promoted by master device to a beaconing device, as specified in 7.12.2.3.

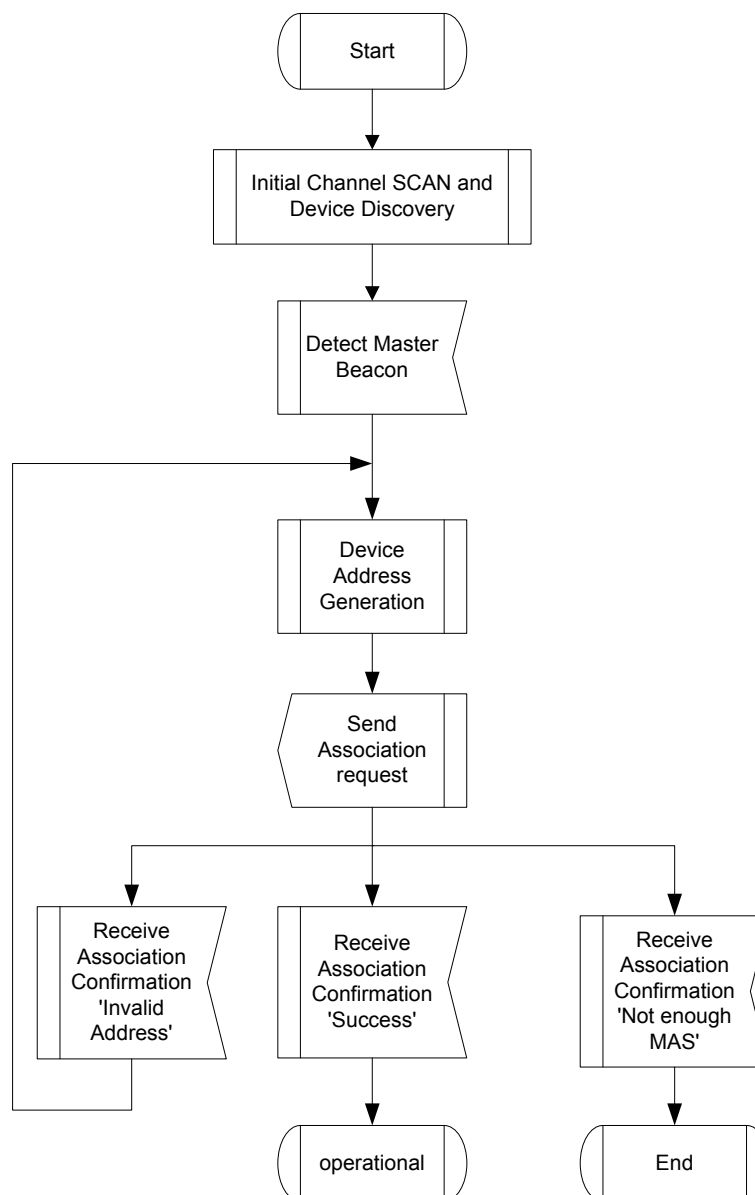


Figure 23 — Master-slave join beacon group

A master device may disassociate a group of slave devices by including a Disassociation IE in its beacon frame. A slave device receiving the Disassociation IE shall release all reserved MASs and terminate existing connections.

7.13.3 Pair discovery

If a device does know who the pair device is after power-up, pair devices should use two interleaving stages, SCAN and STAY for pair discovery. SCAN allows a device to search the pair device on all possible channels. STAY allows a device to stay on a selected channel and advertise itself. The STAY time (T_2) is M times of the maximal SCAN time, where M randomly chosen among (1, 2). The maximal SCAN time equals N times T_1 , N is the total number of channels for scan. Due to the randomization, the SCAN stage of one device has high possibility to fall in the STAY stage of the pair device after limited iteration, thus finding the pair device.

Figure 24 illustrates Pair Discovery procedure. The scan duration for each channel, i.e., T_1 , is twice of SuperframeLength. Once the pair device is found, the pair discovery procedure is completed.

To limit the pair discovery time for the worst case, a device shall stop further discovery after mMaxPairDiscTry iterations and start to do beaconing on a selected channel, see 7.3.4.

After Pair Discovery procedure, a device shall classify the channel per 7.11.2.

As show in Figure 20, further Pair discovery may be activated by a Pairing Command.

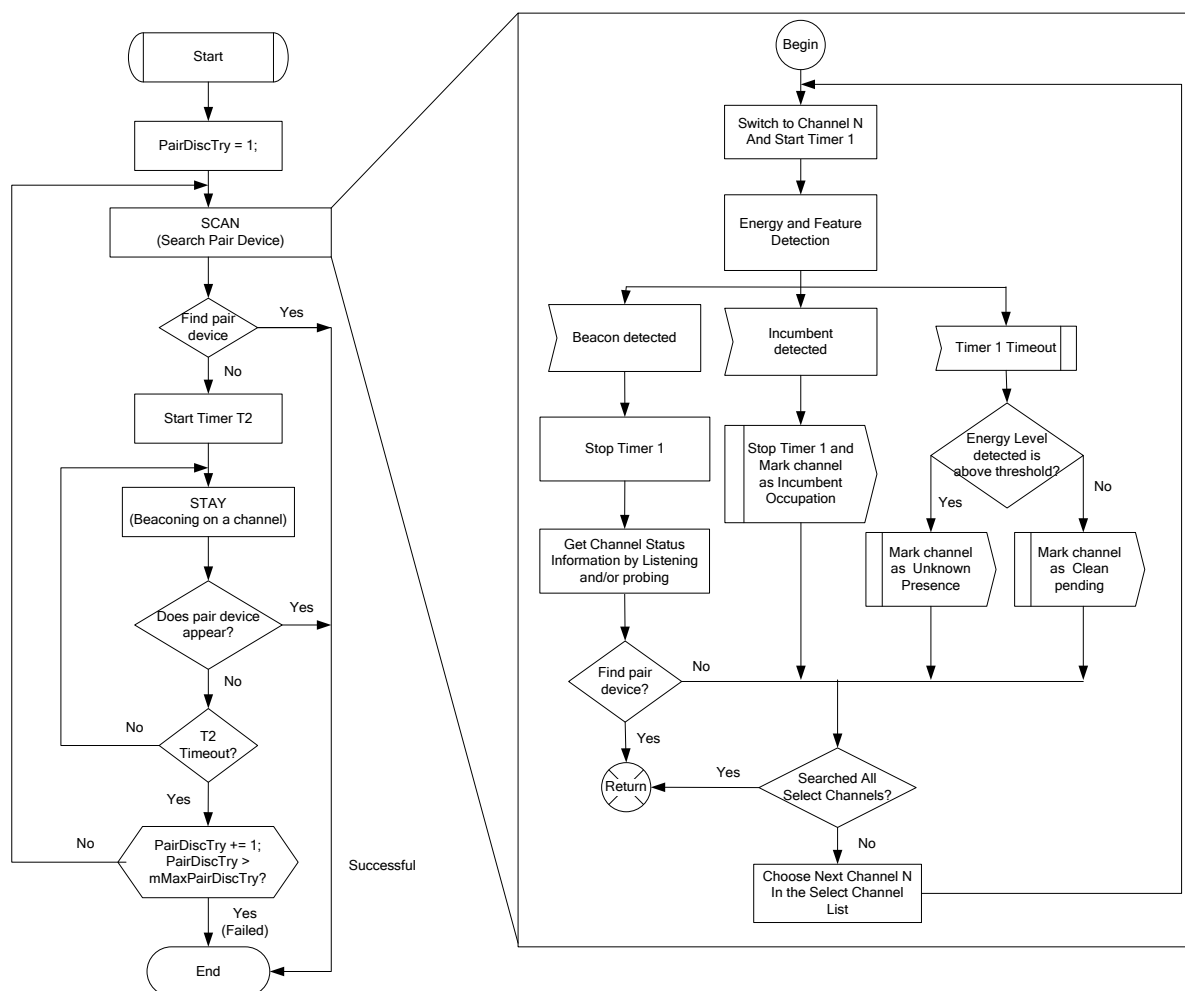


Figure 24 — Pair Discovery

7.13.4 Create/join a beacon group

The device should select an available channel with the least traffic.

Before creating or joining a BG as specified in 7.3.4, the device might be required to ensure incumbent protection using further channel measurements.

7.13.5 Pairing

The pairing procedure enables two devices to register each other, perform authentication, and set up security mechanisms.

Figure 25 illustrates the pairing procedure.

If pair devices have not joined the same BG, they shall join the same BG.

Pair devices shall associate using the Association Request and Response, see 7.1.5.12 and 7.1.8.24 respectively.

If security is enabled, a device shall also perform authentication and TPK establishment, see 8.

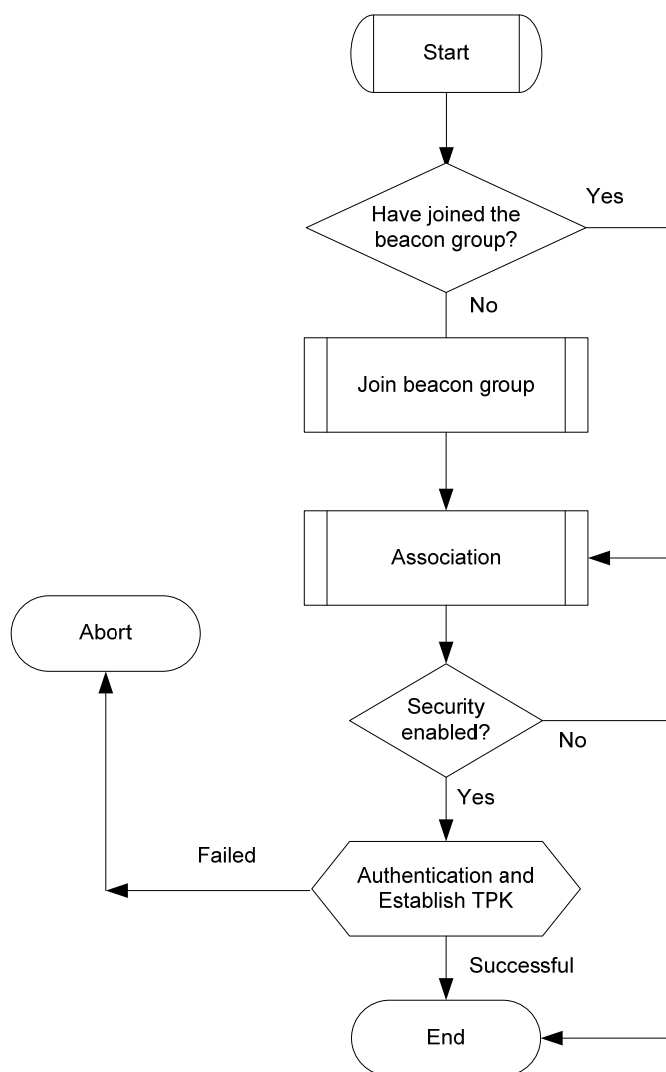


Figure 25 — Pairing Procedure

7.13.6 Setup connections

Figure 26 illustrates connection setup procedure. The source device might need to derive QoS requirement, select an operating channel with enough available bandwidth.

For group communication, the source device needs to bind a multicast address. Moreover, if security is enabled for multicast, group temporal key (GTK) shall be established.

The device shall perform channel reservation, see 7.5.2.

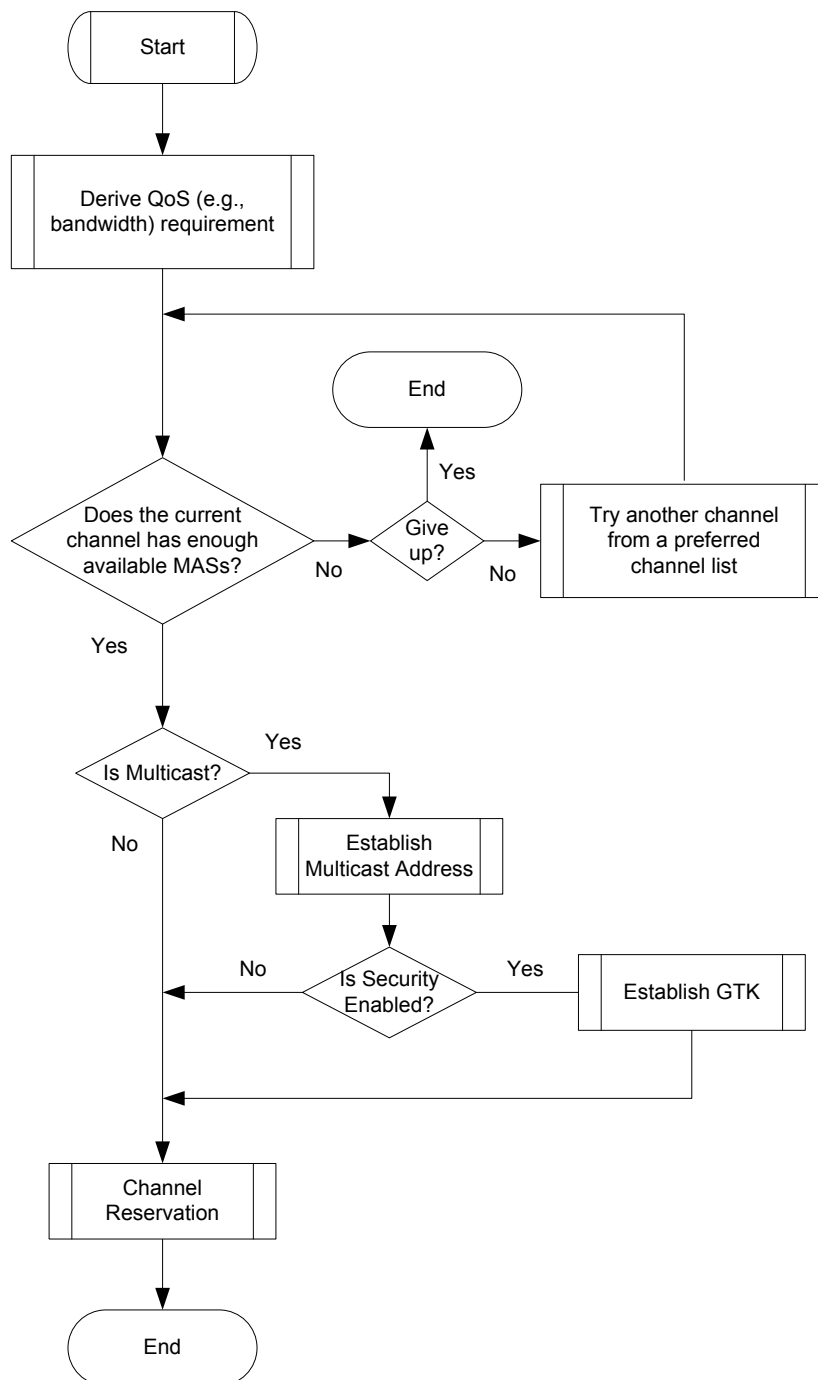


Figure 26 — Connection Setup Procedure

7.14 MAC sublayer parameters

Table 120 contains the values for the MAC sublayer parameters.

Table 120 — MAC sublayer parameters

Parameter	Value
mClockResolution	1 μ s
mMaxSynchronizationAdjustment	4 us
mClockAccuracy	20 ppm
mGuardTime	20 μ s
mMASLength	500 μ s
mMASCount	256
mSuperframeLength	mMASCount \times mMASLength
mCSWSlotLength	mMASLength / 2
mMaxSignalLength	mCSWSlotLength - pSIFS - mGuardTime
mSignalSlotCount	mCSWsize/mCSWSlotLength
mRSWSlotLength	mMASLength
mBeaconSlotLength	2 \times mMASLength
mMaxBeaconLength	mBeaconSlotLength - pSIFS - mGuardTime
mBPExtension	2 beacon slots
mMaxBPLength	12 MAS
mCSWsize	2 MAS
mMaxTransTime	Subject to regulation
mBPMergeWaitTime	128 superframes
mMaxLostBeacons	3
mInitialMoveCountdown	3 \times mMaxLostBeacons
mMaxMovableLatency	32
mAccessDelay	652 μ s
mAggregationLimit	63
mCRPBackoffWinMax	16 superframes
mCRPBackoffWinMin	2 superframes
mMaxFramePayloadSize	pMaxFramePayloadSize
mMaxFragmentCount	8
mCRPslaveRetry	9
mQPD	Duration for regular QP, in the unit of MAS. Minimum is 5 ms / mMASLength. Subject to regulation
mQPFrequency	Regular QP schedule frequency, in the unit of superframes. Subject to regulation. However, if the parameter is set to zero, no regular QP is scheduled.
mMaxPowerUpdateStep	2 dB
mMaxEvacuateCountdown	Subject to regulation.
mMaxHibernationProtection	128 superframes
mMaxPairDiscTry	3
mMaxNeighbourDetectionInterval	128 superframes
mMinFragmentSize	188 bytes
mCWMin[AC_BK]	15
mCWMin[AC_BE]	15
mCWMin[AC_VI]	7
mCWMin[AC_VO]	3
mCWMax[AC_BK]	1023
mCWMax[AC_BE]	1023
mCWMax[AC_VI]	511
mCWMax[AC_VO]	255
mAIFSN[AC_BK]	7
mAIFSN[AC_BE]	4
mAIFSN[AC_VI]	2
mAIFSN[AC_VO]	1
mTXOPLimit[AC_BK]	512 μ s
mTXOPLimit[AC_BE]	512 μ s
mTXOPLimit[AC_VI]	1024 μ s
mTXOPLimit[AC_VO]	256 μ s

Table 121 contains the values of the PHY dependent parameters used by the MAC sublayer for PHY.

Table 121 — PHY-dependent MAC sublayer parameters for the PHY

Parameter	Value
pBeaconTransmitRate	4.75Mbps (QPSK, 1/2), refer to Table 137
pCCADetectTime	10µs
pClockAccuracy	20 ppm
pMaxFramePayloadSize	4095 bytes
pMIFS	2 µs
pSIFS	10 µs
pSlotTime	15 µs

8 Security

This clause specifies the security mechanisms needed to provide the security service introduced in 6.6.9. 8.1 reviews these security mechanisms. 8.2 defines security modes that govern the security operation of devices. 8.3 specifies the 4-way handshake procedure for two devices to establish pair-wise temporal keys (PTKs) and a secure relationship. This subclause also describes how a device solicits or distributes group temporal keys (GTKs) within a secure relationship. 8.4 describes the procedures for frame reception and replay prevention. 8.5 provides the parameters needed in applying the AES-128 CCM cryptography to compute the message integrity code (MIC) and encrypt the secure payload for secure frames.

8.1 Security mechanisms

The security mechanisms specified in this Standard control the security operation of devices by setting appropriate security modes. They allow devices to authenticate each other, to derive PTKs, and to establish secure relationships. They also enable devices to solicit or distribute GTKs within established secure relationships. In addition, the security mechanisms provide replay attack prevention measures through the use of secure frame counters (SFCs) and replay counters. The security mechanisms specify the parameters needed in applying the AES-128 CCM to protect the privacy and integrity of unicast and broadcast/multicast traffic using PTKs and GTKs, respectively. Privacy is protected by encrypting the secure payload, while integrity is protected by including a MIC.

Two devices use a shared master key to establish a secure relationship. The establishment and management of master keys are additional security facilities that need to be provided outside the MAC sublayer.

8.1.1 Security operation

Security modes are defined to control the level of security required of a device in its communications with other devices. Three security modes are provided. Mode 0 allows a device to communicate without security protection. Mode 1 allows a device to use both secure and non-secure frames for data exchanges. Mode 2 restricts a device to use security facilities in transmitting and receiving frames.

A device announces its selected security mode in the Beacon Parameters field in its beacons.

8.1.2 4-way handshake

The 4-way handshake mechanism enables two devices to use a shared master key to authenticate the identity of each other and to establish a new PTK for protecting certain frames exchanged between the two devices. By way of a successful 4-way handshake, the two devices establish a secure relationship with each other.

A device initiates a 4-way handshake with another device only if it has determined that it shares a master key with that device. The master key is not exposed in the 4-way handshake; it is specified by a master key identifier (MKID). A 4-way handshake is affected by use of PTK commands. Following a successful 4-way handshake, each device installs the new PTK into the MAC entity.

8.1.3 Key transport

Two devices establish a new PTK via a 4-way handshake. The PTK is derived from a shared master key and two new random numbers generated by the two devices. A PTK is never transmitted directly in any frame, encrypted or not.

Two devices, after establishing a secure relationship via a successful 4-way handshake, distribute their respective GTKs for protecting their broadcast traffic to each other, if applicable. Additionally, a device may distribute GTKs for protecting certain multicast traffic addressed to those devices with which the device has a valid secure relationship. A device may also request, or solicit, GTKs used to protect multicast traffic from the multicast source devices.

A GTK is solicited or distributed by use of GTK commands. It is sent in encrypted form.

8.1.4 Freshness protection

Freshness protection ensures that no parties can successfully replay previously captured messages as an attack. This Standard defines secure frame counters and replay counters on a per-temporal key basis to provide freshness protection.

8.1.5 Data encryption

Data encryption uses a symmetric cipher to protect data from access by parties not possessing the encryption key. This key is a PTK for unicast traffic transmitted between two devices and a GTK for broadcast/multicast traffic transmitted from a sender to a group of recipients.

Secure frames using a TKID not recognized by the recipient device are reported to the DME.

AES-128 counter mode is used for data encryption in this Standard.

8.1.6 Frame integrity protection

Frames are protected from modification by other parties by message authentication using a MIC. The MIC also provides assurance that the sender of the frame possesses the correct temporal key. This key is shared among a group of devices or only between two devices. The MIC is a cryptographic checksum of the message to be protected.

All secure frames that fail MIC checks are reported to the DME.

AES-128 cipher block chaining – message authentication code (CBC-MAC) is used for MIC calculation in this Standard.

8.2 Security modes

The security mode indicates whether a device is permitted or required to establish a secure relationship with another device for data communications.

Two devices establish a secure relationship by a 4-way handshake based on a shared master key as described in 8.3.

Once two devices establish a secure relationship, they shall use secure frames for frame transfers between them as specified in Table 122 and Table 123. Either device shall discard a received frame from the other device if the frame is required to be a secure frame but was transmitted as a non-secure frame.

Data and aggregated data frames shall be transmitted using the temporal key specified by the TKID passed through the MAC SAP along with the corresponding MSDU. Command and control frames, when transmitted as secure frames in a secure relationship, shall employ a temporal key currently possessed in that secure relationship.

In Table 122, “N” indicates a non-secure frame, and “S” indicates a secure frame. Command frames unlisted in Table 122 are treated as non-secure frames.

Table 122 — Frame protection in a secure relationship

Frame type or subtype	Frame protection	Meaning
Beacon frame	N	Beacon frames shall be sent as non-secure frames.
Imm-ACK control frame	N	Imm-ACK frames shall be sent as non-secure frames.
B-ACK control frame	N	B-ACK frames shall be sent as non-secure frames.
RTS control frame	N	RTS frames shall be sent as non-secure frames.
CTS control frame	N	CTS frames shall be sent as non-secure frames.
UCA control frame	N	UCA frames shall be sent as non-secure frames.
UCR control frame	N	UCR frames shall be sent as non-secure frames.
Application-specific control frame	N, S	Application-specific control frames may be sent as secure or non-secure frames.
CRP Reservation Request command frame	N, S	CRP Reservation Request frames may be sent as secure or non-secure frames.
CRP Reservation Response command frame	N, S	CRP Reservation Response frames may be sent as secure or non-secure frames.
Probe command frame	N, S	Probe frames may be sent as secure or non-secure frames.
PTK command frame	N, S	PTK frames may be sent as secure or non-secure frames.
GTK command frame	S	GTK frames shall be sent as secure frames.
Probe command frame	N, S	Probe frames may be sent as secure or non-secure frames.
Application-specific command frame	N, S	Application-specific command frames may be sent as secure or non-secure frames.
Data frame	S	Data frames shall be sent as secure frames.
Aggregated data frame	S	Aggregated data frames shall be sent as secure frames.

Table 123 specifies the values of the Encryption Offset (EO) field in secure frames.

Table 123 — EO values in secure frames

Frame type or subtype	EO value
Application-specific control frame	Application defined
CRP Reservation Request command frame	Length of Secure Payload
CRP Reservation Response command frame	Length of Secure Payload
PTK command frame	0
GTK command frame	0
Probe command frame	Variable
Application-specific command frame	Application defined
Data frame	Variable
Aggregated data frame	Length of Aggregation Header

8.2.1 Security mode 0

A device operating in security mode 0 shall use non-secure frames to communicate with other devices. Such a device shall not establish a secure relationship with any other device.

If a device operating in this mode receives a secure frame, the MAC entity shall discard the frame.

8.2.2 Security mode 1

A device operating in security mode 1 shall use non-secure frames to communicate with devices operating in security mode 0. The device shall also use non-secure frames to communicate with devices operating in security mode 1 with which it does not have secure relationships. The device shall use secure frames according to Table 122 and Table 123 to communicate with another device operating in security mode 1 with which it has a secure relationship. It shall not establish secure relationships with other devices unless those devices are also operating in security mode 1.

A device operating in security mode 1 may respond to command frames received from other devices with which it does not have a secure relationship.

If a device operating in security mode 1 receives a secure frame from a device with which it does not have a secure relationship, the MAC entity shall discard the frame.

If a device operating in mode 1 receives a non-secure frame from a device with which it has a secure relationship, but the frame is required to be a secure frame per Table 122, the MAC entity shall discard the frame.

A DME that chooses to enable security mode 1 must understand and accept the responsibility that comes with receiving non-secure frames. The DME shall instruct the higher layers to handle the received non-secure frames in a safe and secure manner.

A compliant MAC entity shall never use security mode 1 by default. Security mode 1 shall be entered from either mode 0 or mode 2. Requiring that a DME explicitly select this mode serves as an indication that the DME is aware of the security responsibilities it accepts when enabling security mode 1.

8.2.3 Security mode 2

A device operating in security mode 2 shall not establish a secure relationship with devices operating in either security mode 0 or security mode 1. The device shall use secure frames based on Table 122 and Table 123 to communicate with another device operating in security mode 2 and having a secure relationship with it. A device operating in security mode 2 shall establish a secure relationship with another device operating in the same security mode by a 4-way handshake prior to data exchanges.

If a device operating in mode 2 receives a secure frame from a device with which it does not have a secure relationship, the MAC entity shall discard the frame.

If a device operating in mode 2 receives a non-secure frame from a device with which it has a secure relationship, but the frame is required to be a secure frame per Table 122, the MAC entity shall discard the frame.

8.3 Temporal keys

Two devices establish a secure relationship based on a shared master key by employing a 4-way handshake to derive a PTK as described in this subclause. They may establish a PTK for each master key they share. Two devices have a secure relationship as long as they possess a currently installed PTK. A device's DevAddr is part of the information used in deriving a PTK. Once a PTK is established, it shall not be changed due to a change in the device's DevAddr.

A device solicits a GTK from, or distributes a GTK to, another device sharing a PTK as also described in this subclause.

Master keys are identified by MKIDs. A device is not required to include an MKID IE in its beacon, nor is it required to advertise every MKID it possesses in the MKID IE included in its beacon. They may advertise some or all of the MKIDs they possess in an MKID IE in their beacons. A device may probe another device for the MKIDs possessed by that device by addressing an appropriate Probe IE in a beacon or Probe command to that device. A device shall list all the MKIDs it possesses in the MKID IE in response to a probe request for its MKIDs.

8.3.1 Mutual authentication and PTK derivation

This Standard uses a 4-way handshake to provide mutual authentication and PTK generation for two devices sharing a master key. To perform a 4-way handshake, the two devices assume the roles of “initiator” and “responder”, respectively. A 4-way handshake comprises four messages, called message 1, message 2, message 3, and message 4 in this Standard, that are sent back and forth between the two devices. The device sending message 1 becomes the initiator. The other device becomes the responder.

8.3.1.1 4-way handshake message 1

The initiator shall begin a 4-way handshake by composing and sending message 1 in a PTK command to the responder. In this command, the initiator shall specify the MKID for use in the 4-way handshake, propose a TKID for the PTK to be derived, and include a unique 128-bit cryptographic random number, I-Nonce. The proposed TKID shall be different from any TKID currently installed in the initiator’s local MAC entity or being used in an in-progress 4-way handshake involving this initiator device. The I-Nonce shall be generated anew each time the initiator starts a new 4-way handshake.

On reception of message 1, the responder shall verify that the requested TKID is unique (i.e., not currently installed for an active temporal key or requested by an in-process 4-way handshake exchange). The responder shall perform the following steps:

1. Generate a new 128-bit cryptographic random number, R-Nonce.
2. Derive the PTK and KCK as specified in 8.3.4.
3. Construct and send message 2 in a PTK command.

8.3.1.2 4-way handshake message 2

The responder shall send message 2 to the initiator as specified in 8.3.1.1. In this command, the responder shall include an appropriate Status Code, the newly generated R-Nonce, and the PTK MIC value computed for the message using the newly derived KCK according to 0. If the proposed TKID in message 1 is not unique, the responder shall so indicate in the Status Code.

On reception of message 2, the initiator shall perform the following steps:

1. Derive the PTK and KCK as specified in 8.3.4.
2. Recalculate the PTK MIC for the received message using the KCK according to 0. If the recalculated PTK MIC does not match the PTK MIC field from this message, discard and disregard message 2 and abort the 4-way handshake. Otherwise, consider this message a proof that the responder holds the correct master key, and proceed to the next step.
3. Check the Status Code returned in the received message. If the Status Code indicates an abortion of the 4-way handshake by the responder, stop the 4-way handshake as well. If the Status Code indicates a conflict of the proposed TKID at the responder, restart the 4-way handshake with a different TKID. If the Status Code indicates a normal status, proceed to the next step.
4. Construct and send message 3 in a PTK command.

8.3.1.3 4-way handshake message 3

The initiator shall send message 3 to the responder as specified in 8.3.1.2. In this command, the initiator shall include the same I-Nonce as contained in message 1 and a PTK MIC computed for this message using the newly derived KCK according to 0.

On reception of message 3, the responder shall perform the following steps:

1. Verify the PTK MIC for this message using the KCK according to 0. If the calculated PTK MIC does not match the PTK MIC field from this message, discard and disregard message 3 and abort the 4-way handshake. Otherwise, consider this message a proof that the initiator holds the correct master key, and proceed to the next two steps.
2. Construct and send message 4 in a PTK command.
3. Install the PTK.

8.3.1.4 4-way handshake message 4

The responder shall send message 4 to the initiator as specified in 8.3.1.3. In this command, the responder shall include the same R-Nonce as contained in message 2 and a PTK MIC computed for this message using the KCK according to 0.

On reception of message 4, the initiator shall perform the following step:

1. Verify the PTK MIC for this message using the KCK according to 0. If the calculated PTK MIC does not match the PTK MIC field from this message, discard and disregard message 4 and abort the 4-way handshake.

8.3.2 GTK exchange

Upon successful completion of a 4-way handshake and installation of the resulting PTK, the initiator and responder each shall use GTK command frames (with Message Number set to 1) to distribute their respective GTKs for broadcast traffic to each other. Each may also use a GTK command to distribute a GTK for protecting certain multicast traffic to an intended recipient with which it holds a valid PTK.

On reception of a valid GTK command frame marked as Message Number 1, a device shall verify that the GTKID is a unique TKID. The device shall then respond with a GTK command frame with Message Number set to 2 and Status Code set to the appropriate value.

A recipient may request a GTK for certain multicast traffic in the form of a GTK command (with Message Number set to 0) from the source device if it holds a valid PTK with the source.

On reception of a valid GTK command marked as Message Number 0, the multicast source device shall respond with a GTK command marked as Message Number 1, which may contain the requested GTK. The requesting device, upon receiving this GTK command and verifying the uniqueness of the proposed TKID, shall further return a GTK command with Message Number set to 2 and Status Code set to the appropriate value.

A source device distributing a GTK shall check the Status Code indicated in the returned GTK command (Message Number set to 2). If the Status Code indicates a conflict of the proposed TKID at the recipient device, the source device shall propose a new TKID and re-distribute the GTK to the recipient. After receiving a returned GTK command from the recipient with the Status Code indicating a normal status, the source device shall use the new TKID to re-distribute the GTK to each of the devices to which it has previously distributed the GTK and with which it maintains a secure relationship.

A device installs a newly distributed or received GTK.

A GTK shall be a 128-bit cryptographic-grade random number. A fresh GTK shall be generated when the distributing device establishes a new group relationship. 8.3.6 provides an example means of generating a fresh GTK.

8.3.3 Pseudo-random function (PRF) definition

A PRF is used in several places in the security specification. This subclause defines three PRF variants:

- PRF-64, which outputs 64 bits,
- PRF-128, which outputs 128 bits, and
- PRF-256, which outputs 256 bits.

In the following, K denotes a 128-bit symmetric key, N denotes a 13-octet nonce value, A denotes a unique 14-octet ASCII text label for each different use of the PRF, B denotes the input data stream, $Blen$ specifies the length of this data stream, and \parallel denotes concatenation. Blocks are each 16 octets long, and are defined as inputs to the AES-128 CCM for the MIC generation as specified in 8.5.

CCM-MAC-FUNCTION($K, N, A, B, Blen$)

begin

Form authentication block B_0 from flags = 0x59, N , and $I(m) = 0$

Form authentication block B_1 from $I(a) = 14 + Blen$ and A

Form additional authentication blocks from B

(with last block zero padded as needed)

Form encryption block A_0 from flags = 0x01, N , and Counter_0 = 0

$R \leftarrow \text{MIC}(K, B_0, B_1, \dots, A_0)$

return R

PRF($K, N, A, B, Blen, Len$)

for $i \leftarrow 1$ **to** $(Len + 63)/64$ **do**

$R \leftarrow R \parallel \text{CCM-MAC-FUNCTION}(K, N, A, B, Blen)$

$N \leftarrow N + 1$

return $L(R, 0, Len) = Len$ most-significant bits of R

PRF-64($K, N, A, B, Blen$) = PRF($K, N, A, B, Blen, 64$)

PRF-128($K, N, A, B, Blen$) = PRF($K, N, A, B, Blen, 128$)

PRF-256($K, N, A, B, Blen$) = PRF($K, N, A, B, Blen, 256$)

8.3.4 PTK and KCK derivation

PRF-256 shall be employed to generate the PTK and KCK associated with a 4-way handshake as used in 8.3.1 based on the following parameters as defined in Table 124.

K – The PMK
 N – B12-11= InitiatorDevAddr, B10-9= ResponderDevAddr, B8-6 = PTKID, B5-0 = zero
 A – “Pair-wise keys”
 B – I-Nonce || R-Nonce
 B_{len} – 32

Table 124 — PTK and KCK Generation Parameters

Name	Size (octets)	Description
InitiatorDevAddr	2	DevAddr of device with role of initiator
ResponderDevAddr	2	DevAddr of device with role of responder
I-Nonce	16	Random number selected by initiator (in message 1)
R-Nonce	16	Random number selected by responder (in message 2)
PTKID	3	Negotiated TKID value for the PTK to be derived (in message 1)
PMK	16	A pre-shared pair-wise master key identified by the MKID (in message 1)

The PRF-256 is called with these parameters to compute a 256-bit key stream:

$$\text{KeyStream} \leftarrow \text{PRF-256}(K, N, A, B, B_{len})$$

This key stream is then split to form the desired PTK and KCK. The least-significant 16 octets of KeyStream become the KCK while the most-significant 16 octets become the PTK, as specified in Table 125.

Table 125 — KCK and PTK in KeyStream

Key	Source
KCK	KeyStream octets 0 through 15
PTK	KeyStream octets 16 through 31

8.3.5 PTK MIC generation

The 4-way handshake uses an “out-of-band MIC” calculation for the PTK MIC field in handshake messages 2-4. PRF-64 shall be used to provide the PTK MIC calculation. The PRF-64 parameters shall be defined as follows based on Table 124:

K – The KCK
 N – B12-11 = InitiatorDevAddr, B10-9 = ResponderDevAddr, B8-6 = PTKID, B5-0 = zero
 A – “out-of-bandMIC”
 B – Fields from Message Number to I-Nonce/R-Nonce contained in the PTK command
 B_{len} – Length in octets of B = 48

$$\text{PTK MIC} \leftarrow \text{PRF-64}(K, N, A, B, B_{len})$$

8.3.6 Random number generation

In order to implement the cryptographic mechanisms outlined in this Standard, every platform needs to be able to generate cryptographic grade random numbers. RFC 1750 gives a detailed explanation of the notion of cryptographic grade random numbers and provides guidance for collecting suitable randomness. It

recommends collecting random samples from multiple sources followed by conditioning with PRF. This method provides a means for an implementation to create an unpredictable seed for a pseudo-random generation function. The example below shows how to distill such a seed using random samples and PRF-128.

```
LoopCounter = 0
```

```
Nonce = 0
```

```
while LoopCounter < 32 begin
```

```
    result = PRF-128(0, Nonce, "InitRandomSeed", DevAddr || Time || result || LoopCounter, dataLen)
```

```
    Nonce ← Nonce + 1
```

```
    result ← result || <randomness samples>
```

```
end
```

```
GlobalSeed = PRF-128(0, Nonce, "InitRandomSeed", DevAddr || Time || result || LoopCounter, dataLen)
```

Once the seed has been distilled, it may be used as a key for further random number generation. The 4-way handshake requires each party to supply a 128-bit random number. This number may be generated using the seed and PRF-128.

```
GenerateRandomNonce
```

```
begin
```

```
    N = DevAddr || DevAddr || zero
```

```
    Collect randomness samples
```

```
    result = PRF-128(Global Seed, N, "Random Numbers", <randomness samples>, length of samples)
```

```
return result
```

8.4 Frame reception steps and replay prevention measures

A recipient device shall carry out the reception steps and replay prevention measures as specified in this subclause.

8.4.1 Frame reception

The MAC entity shall perform the following validation steps in sequence when receiving frames:

1. Validate the FCS. If this validation fails, discard the frame. Otherwise, acknowledge the received frame using the appropriate acknowledgment rules, and proceed to the next step.
2. Validate the Secure bit setting in the MAC Header and take the appropriate actions according to its security mode as specified in 8.2. If the frame is not discarded and the Secure bit is set to ONE, proceed to the next step.
3. Validate the TKID. If the TKID does not identify a currently installed PTK or GTK, discard the frame. Otherwise, proceed to the next step.
4. Validate the MIC using the identified PTK or GTK as specified in 8.5. If this validation fails, discard the frame. Otherwise, proceed to the next step.

5. Detect frame replay as specified in 8.4.2. If replay is detected, discard the frame. Otherwise, update the replay counter that was set up for the PTK or GTK used for this frame as also specified in 8.4.2, and proceed to the next step.
6. Process the frame as specified in clause 7, including duplicate frame filtering. If the frame was already received, discard it. Otherwise, proceed to the next step.
7. Decrypt the frame. This step may be taken in parallel with the MIC validation step.

8.4.2 Replay prevention

Each transmitting MAC entity shall set up a 48-bit SFC and initialize it to zero when a temporal key, PTK or GTK, is installed to it. The MAC entity shall increment the SFC by one before transmitting a secure frame—whether a new frame or a retry—that uses the temporal key, and shall set the SFN in that secure frame to the value of the SFC after the increment.

Each recipient MAC entity shall set up a 48-bit replay counter when a temporal key, PTK or GTK, is installed to it. The MAC entity shall initialize the replay counter to zero for an installed PTK, and to the GTK SFC for an installed GTK which was contained in the GTK command distributing the GTK.

Upon receipt of a secure frame with valid FCS and MIC, the recipient shall perform replay attack detection and protection as follows:

The recipient shall compare the SFN extracted from the received frame with the reading of the replay counter for the temporal key used by the frame. If the extracted SFN is smaller than or equal to the replay counter reading, the recipient MAC entity shall discard the frame. Otherwise, the recipient shall set the corresponding replay counter to the received SFN.

The recipient shall insure that the frame passes FCS validation, replay prevention, and MIC verification before using the SFN to update its replay counter.

8.4.3 Implications on GTKs

Because a recipient maintains only one replay counter per installed temporal key, that recipient can receive traffic from only one source using a given temporal key. A scheme that allows multiple source devices to use the same GTK will result in frames sent from some of those sources being seen as replay attacks. To avoid this problem, each source device in a group is required to distribute a unique GTK to the recipients in the group.

8.5 AES-128 CCM Inputs

AES-128 CCM provides confidentiality, authentication, and integrity for secure frames defined in this proposal. This subclause specifies the various fields required for AES-128 CCM operation.

8.5.1 Overview

AES, the Advanced Encryption Standard, is specified in FIPS PUB 197. AES-128 defines a symmetric block cipher that processes 128-bit data blocks using 128-bit cipher keys. CCM, counter with CBC-MAC, is specified in RFC 3610. CCM employs counter mode for encryption and cipher block chaining for authentication. AES-128 CCM combines AES-128 with CCM to encrypt and authenticate messages.

Encryption is done on part or all of the Secure Payload, while authentication is provided by a message integrity code (MIC) that is included in each secure frame. MIC also protects the integrity of the MAC Header and Frame Payload in a secure frame.

CCM has two input parameters – M (number of octets in authentication field) and L (number of octets in length field). For this Standard, M = 8 and L = 2.

CCM requires the use of a temporal key and a unique Nonce for each transmitted frame to be protected. The SFN is combined with frame addressing and temporal key identification information to provide a unique Nonce for every secure frame. Since every frame protection with a key requires a unique Nonce, temporal keys have a known lifetime. Each temporal key may be used to protect up to n frames, where n is the maximum value of the SFN. All security guarantees are void if a nonce value is used more than once with the same temporal key.

In the following figures in this subclause showing the format of Nonce and CCM blocks, the most-significant octet is represented to the left of the other octets.

8.5.2 Nonce

The CCM Nonce is a 13-octet field, consisting of the 2-octet SrcAddr, 2-octet DestAddr, 3-octet TKID, and 6-octet SFN for the current frame. The Nonce is used as a component of authentication block B_0, an input to CBC-MAC. It is also used as a component of input block A_i for CCM encryption. It provides the uniqueness that CCM requires for each instance of authentication/encryption. The CCM Nonce shall be formatted as shown in Table 126. In this Figure, each component of the Nonce is represented with the most-significant octet on the left and the least-significant octet on the right.

Table 126 — Nonce input to the CCM algorithm

Syntax	Size	Notes
Nonce_Format {		
SrcAddr	2 bytes	
DestAddr	2 bytes	
TKID	3 bytes	
SFN	6 bytes	
}		

8.5.3 CCM blocks

The CCM authentication blocks shall be formatted as shown in Table 127 and further described below.

Table 127 — Input to CCM authentication blocks

Syntax	Size	Notes
Input_to_CCM_Authentication_Format {		
Flags (= 0x59)	1 byte	B_0
Nonce	13 bytes	
Encrypted data length $l(m) = P - EO$	2 bytes	
Additional authenticated data length $l(a) = 14 + EO$	2 bytes	B_1
MAC header	10 bytes	
Encryption Offset (EO)	2 bytes	
Security Reserved	1 byte	
0	1 byte	
Secure Payload portion not to be encrypted	EO bytes	B_2, ..., B_(M-1)
Zero padding	0-15 bytes	
Secure Payload portion to be encrypted	$P - EO$ bytes	B_M, ..., B_N
Zero padding	0-15 bytes	
}		

8.5.3.1 Authentication block B_0

Authentication block B_0 is the first input block to the CBC-MAC algorithm. It shall be formatted as shown in Table 128. The component $l(m)$ is represented with the most-significant octet on the left and the least-significant octet on the right. The Nonce component is represented with the least-significant octet on the left and the most-significant octet on the right.

Table 128 — Format of authentication block B_0

Syntax	Size	Notes
Authentication_Block_B_0__Format {		
Flags (= 0x59)	1 byte	
Nonce	13 bytes	
I(m)	2 bytes	
}		

8.5.3.2 Authentication block B_1

Authentication block B_1 is the second input block to the CBC-MAC algorithm. It shall be formatted as shown in Table 129. In this block, the I(a) component is represented with the most-significant octet on the left and the least-significant octet on the right. The EO and MAC Header components are represented with the first octet transmitted into the wireless medium on the left and the last transmitted octet on the right.

Table 129 — Format of authentication block B_1

Syntax	Size	Notes
Authentication_Block_B_1__Format {		
I(a)	2 bytes	
MAC Header	10 bytes	
EO	2 bytes	
Security Reserved	1 byte	
0	1 byte	
}		

8.5.3.3 Authentication blocks B_2, ..., B_n

Authentication blocks B_2, ..., B_(M-1) and B_M, ..., B_N, if any, are additional input blocks to the CBC-MAC algorithm. They shall be formatted as shown in Table 130. They are formed by breaking the Secure Payload portion not to be encrypted into 16-octet blocks and the Secure Payload portion to be encrypted into 16-octet blocks. The last block constructed from the Secure Payload portion not to be encrypted is padded with ZERO values as needed to insure 16-octet block length. Likewise, the last block constructed from the Secure Payload portion to be encrypted is padded with zero values as needed to insure 16-octet block length. The padding octets are not transmitted onto the wireless medium.

Table 130 — Format of authentication blocks beginning from B_2

Syntax	Size	Notes
Authentication_Block_B_2_and_above_Format {		
Secure Payload portion not to be encrypted	EO bytes	B_2, ..., B_(M-1)
Zero padding	0-15 bytes	
Secure Payload portion to be encrypted	P-EO bytes	B_M, ..., B_N
Zero padding	0-15 byte	
}		

In each of the blocks B_2, ..., B_(M-1) or B_M, ..., B_N, the Secure Payload portion not to be, or to be, encrypted shall be represented with the earliest octet transmitted into the wireless medium on the left and the latest transmitted octet on the right. When needed, B_(M-1) and B_N are padded with zeros to the right.

8.5.3.4 Encryption blocks A_0, A_1, ..., A_m

CCM uses encryption blocks A_0, A_1, ..., A_m to generate key stream blocks that are used to encrypt the CBC-MAC and the Secure Payload portion to be encrypted. These blocks shall be formed as shown in Table 131. In this Figure, Counter i is a 2-octet monotonically incrementing counter that shall be initialized to 0 for each secure frame. It shall be incremented by one for each successive encryption block. The Counter i component of A_i shall be represented with the most-significant octet on the left and the least-significant octet on the right. The Nonce component shall be represented with the least-significant octet on the left and the most-significant octet on the right.

Table 131 — Format of A_i blocks

Syntax	Size	Notes
Authentication Block B_1 Format {		
Flags = 0x01	1 byte	
Nonce	13 bytes	
Counter i	2 bytes	
}		

9 PHY

9.1 Introduction

This clause defines the physical (PHY) layer standard for personal/portable applications using TV White Spaces. The specification provides a flexible system that uses a vacant TV channel or a multiple of vacant TV channels to provide the wireless communications, for example, in home distribution of audio and video, wireless internet access, etc. The following clauses of the document provide details on the various aspects of the PHY specifications.

9.2 Symbol description

9.2.1 OFDM symbol description

The transmitted RF signal can be represented mathematically as

$$s_{RF}(t) = \text{Re} \left\{ \sum_{n=0}^{N-1} s_n(t - nT_{SYM}) \exp(j2\pi f_c t) \right\} \quad (1)$$

where $\text{Re}(\cdot)$ represents the real part of the signal, N is the number of OFDM symbols in the PPDU, T_{SYM} is the OFDM symbol duration, f_c is the carrier centre frequency and $s_n(t)$ is the complex base-band representation of the n^{th} OFDM symbol.

$$s_n(t) = 0 \quad 0 > t \geq T_{SYM} \quad (2)$$

The exact form of $s_n(t)$ is determined by the symbol number n . The $s_n(t)$ is composed of three components, the PLCP preamble, PLCP header, and the payload.

$$s_n(t) = \begin{cases} s_{\text{preamble}, n}(t) & 0 \leq n < N_{\text{preamble}} \\ s_{\text{header}, n - N_{\text{preamble}}}(t) & N_{\text{preamble}} \leq n < N_{\text{preamble}} + N_{\text{header}} \\ s_{\text{payload}, n - N_{\text{preamble}} - N_{\text{header}}}(t) & N_{\text{preamble}} + N_{\text{header}} \leq n < N_{\text{frame}} \end{cases} \quad (3)$$

where $s_{\text{preamble},n}(t)$, $s_{\text{header},n}(t)$, and $s_{\text{payload},n}(t)$ denote the n th symbol of the PLCP preamble, the PLCP header, and the payload, respectively. The payload contains the data and the FCS (frame check sequence), the tail bits, and the pad bits, if needed. N_{preamble} , N_{header} , and N_{frame} denote the number of symbols in the PLCP preamble, the PLCP header, and the PLCP frame, respectively.

The complex base-band representation of $s_n(t)$ is defined by Equation (19) in Section 9.5. PLCP preamble is described in 9.3.1. The OFDM symbol of the PLCP header and the payload is defined in 9.5.

9.2.1.1 Time domain description

The time-domain signal is generated by taking the inverse Fourier transform of the length N_{FFT} vector. The vector is formed by taking the constellation mapper output and inserting pilot and guard tones. At the receiver, the time domain signal is transformed to the frequency domain representation by using a Fourier transform. Fast Fourier Transform (FFT) algorithm is usually used to implement Fourier transform and its inverse.

Let T_{FFT} represent the time duration of the IFFT output signal. The OFDM symbol is formed by inserting a cyclic prefix of time duration T_{CP} (shown in Figure 27), resulting in a symbol duration of $T_{\text{SYM}} = T_{\text{FFT}} + T_{\text{CP}}$




Figure 27 — OFDM symbol format

The specific values for T_{FFT} , T_{CP} and T_{SYM} are given in 9.2.2.

9.2.1.2 Frequency domain description

In the frequency domain, an OFDM symbol is defined in terms of its subcarriers. The subcarriers are classified as: 1) data subcarriers, 2) pilot subcarriers, 3) guard subcarriers and 4) Null (includes DC) subcarriers. The classification is based on the functionality of the subcarriers. The total number of subcarriers is determined by the FFT/IFFT size. Except for the DC subcarrier, all the remaining guard/Null subcarriers are placed at the band-edges. The guard subcarriers do not carry any energy. The pilot subcarriers are distributed across the bandwidth. The exact location of the pilot and the data subcarriers are described in 9.5 and are shown in the Figure 28 for one particular symbol.




Figure 28 — OFDM symbol structure

9.2.2 Symbol parameters

For a 6 MHz channel bandwidth, the subcarrier spacing $\Delta F = (BW \cdot 8/7)/128 \approx 53.571$ KHz

$$T_{FFT} = \frac{1}{\Delta F} = 18.667 \mu s \quad (4)$$

where BW represents the channel bandwidth.

The cyclic prefix duration T_{CP} could be one of the following derived values: $T_{FFT}/32$, $T_{FFT}/16$ and $T_{FFT}/8$. Beacon shall be transmitted with CP of $T_{FFT}/8$.

The OFDM symbol duration for different values of cyclic prefix is given in Table 132.

Table 132 — Symbol duration for different cyclic prefixes

	CP = $T_{FFT}/32$	CP = $T_{FFT}/16$	CP = $T_{FFT}/8$
$T_{SYM} = T_{FFT} + I_{CP}$	19.25 μs	19.833 μs	21.0 μs

Table 133 shows the different OFDM parameters and their values for a 6 MHz channel bandwidth. The corresponding numbers for other bandwidths is provided in Annex B.

Table 133 — OFDM Parameters

Parameter	Value
Subcarrier spacing, ΔF (KHz)	53.571
FFT period, T_{FFT} (μs)	18.667
Total number of subcarriers, N_{FFT}	128
Number of guard subcarriers, N_G (L, DC, R)	26(13, 1, 12)
Number of used subcarriers, $N_T = N_D + N_P$	102
Number of data subcarriers, N_D	98
Number of pilot subcarriers, N_P	4
Signal bandwidth (MHz)	5.518

9.3 PPDU

The PPDU is shown in Figure 29. The format for the PPDU includes the PLCP preamble, PLCP Header (HDR), PSDU, tail bits, and pad bits.




Figure 29 — PPDU frame format

9.3.1 PLCP preamble

The PLCP preamble is used by the receiver for frequency and time synchronization and channel estimation.

Two types of PLCP preambles are defined:

- Normal PLCP preamble: Used for all the packets in normal mode and for the first packet in streaming mode.
- Burst PLCP preamble: Used for the second and the subsequent packets in the streaming mode.

9.3.1.1 Normal PLCP preamble

The format of the Normal PLCP preamble is shown in Figure 30. The Normal PLCP preamble is three symbols in duration and consists of a short preamble and a long preamble. The short preamble may be used for initial burst detection; AGC tuning, coarse frequency offset estimation and timing synchronization. The long preamble may be used for channel estimation and for fine frequency offset estimation. The short preamble consists of nine repetitions of a short training sequence while the long preamble consists of two repetitions of a long training sequence.

The length of the cyclic prefix for the Normal PLCP preamble is given as

$$T_{GI} = \frac{1}{8}T_{FFT} \quad (5)$$

and the duration of Normal PLCP preamble is

$$T_{NormalPLCPpreamble} = 3T_{SYM} \quad (6)$$

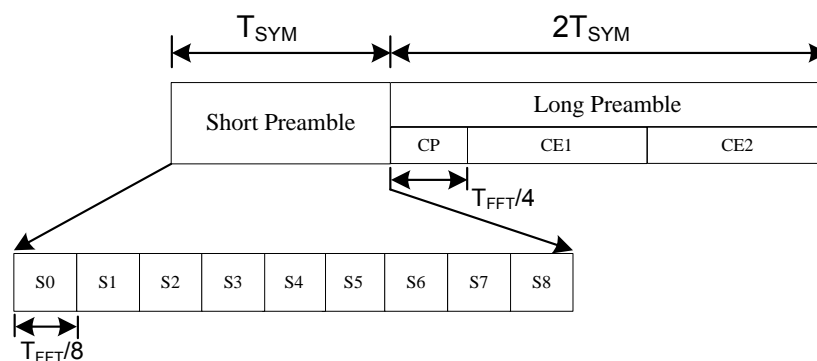


Figure 30 — Normal PLCP preamble format. S0 – S8: short training sequence; CE1, CE2: long training sequence

Short preamble is generated using the following procedure:

- $$\text{PST}(-64:63) = \sqrt{\frac{102}{12}} \{0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, \\ 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, \\ 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, 0, \\ 0, 0, 0, -1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, \\ 0\}$$
- (7)

- CP (S0). The factor $\sqrt{\frac{102}{12}}$ is used to normalize the signal energy.

1. A 128 length frequency domain sequence is defined as shown below

$$\text{PLT}(-64:63) = \{0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, -1, 1, 1, -1, -1, 1, 1, -1, -1, -1, -1, -1, 1, 1, -1, -1, -1, 1, 1, 1, -1, 1, 1, -1, 1, 1, 1, 1, 0, 1, 1, -1, -1, -1, -1, 1, -1, 1, -1, 1, 1, -1, -1, -1, 1, -1, -1, 1, 1, 1, -1, -1, 1, -1, 1, -1, -1, 1, -1, -1, 1, 1, -1, 1, -1, 1, -1, 1, -1, 1, -1, 1, -1, 1, -1, 1, -1, 1, -1, -1, -1, -1, -1, 1, 1, -1, -1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0\}$$

(8)

2. Taking IFFT of the above sequence will generate the time domain representation of the long training sequence.
3. A copy of the time domain sequence is also transmitted in the following symbol.
4. The CP of the two symbols are combined and transmitted at the beginning of the long preamble. Thus the CP of the long preamble is twice the CP of the short preamble.

The format of the burst PLCP preamble is shown in Figure 31. The burst PLCP preamble is one symbol in duration and consists of two repetitions of a burst training sequence. The burst PLCP preamble may be used for channel estimation and for fine frequency offset estimation.

The length of the cyclic prefix for the burst PLCP preamble is given as

$$T_{GI} = \frac{1}{8}T_{FFT} \quad (9)$$

and the duration of the burst PLCP preamble is

$$T_{BurstPLCPPreamble} = T_{SYM} \quad (10)$$



1. A 128 length frequency domain sequence is defined as shown below

(11)

- ### 9.3.2 PLCP header

PHY Hdr (40)	MAC HDR (80)	RS Parity (64)	Tail bits (12)
-----------------	-----------------	----------------------	----------------------

The numbers in parenthesis represent the number of bits allocated for that field.

The format of the PHY header is shown in Figure 33. The PHY header includes data rate, length, transmission mode, scrambler initialization seed, interleaver option, multiple antenna mode, CP mode and transmits power fields. It also includes a number of reserved bits that may be used to define additional modes in future revisions. The reserved bits are set to 0.




Figure 33 — PHY header format

9.3.2.1.1 Cyclic prefix mode (CP)

The CP mode bits are used to determine the CP duration used for the payload symbols. This field is 8 bits long and occupies the positions 0 to 7 in the PHY header. The first two bits of this field i.e. C0 and C1 determine the CP duration as defined in the Table below.

Table 134 — CP mode bits C0 and C1 definition

CP Mode bits C1 – C0	CP Duration
00	1/32
01	1/16
10	1/8
11	Reserved

The remaining CP mode bits (C2 – C7) are a repetition of C1 and C0 bits. The bits C2, C4 and C6 are a copy of bit C0, while the bits C3, C5 and C7 are a copy of bit C1.

9.3.2.1.2 Transmission mode (MODE)

The transmission mode field is used to represent the stream type to which the next packet/frame belongs and also to indicate the preamble type for the next packet. The MODE bit M0 is located at position 8 in the PHY header and indicates the stream type as shown in Table 135 below.

Table 135 — MODE bit M0 definition

MODE Bit M0	Stream Type
0	Normal
1	Burst/Streaming

The MODE bit M1 is located at position 9 in the PHY header and indicates the preamble type for the next packet as shown in Table 136 below.

Table 136 — MODE bit M1 definition

MODE Bit M1	Preamble Type
0	Normal
1	Burst

9.3.2.1.3 Data rate field (RATE)

The data rate field in the PHY header defines a combination of modulation, coding and spreading schemes. The size of the rate field is 4 bits and occupies the positions 10 to 13 in the PHY header (see Figure 33). Table 137 defines the mapping of the PHY layer transmission parameters to the rate field.

Table 137 — Mapping of the PHY layer parameters to Rate field. The data rates are derived using the parameters defined in Table 133. CP duration = 1.167 μ s (1/16 T_{FFT})

RATE R3 – R0	Modulation	Outer Coding	Inner Coding Rate	Data Rate (Mb/s)	Spectral Efficiency (bit/s/Hz)
0000 (0)	QPSK	(245,255,5)	1/2	4.75	0.79
0001 (1)	QPSK	(245,255,5)	2/3	6.33	1.05
0010 (2)	16-QAM	(245,255,5)	1/2	9.49	1.58
0011 (3)	16-QAM	(245,255,5)	7/12	11.08	1.85
0100 (4)	16-QAM	(245,255,5)	2/3	12.66	2.11
0101 (5)	64-QAM	(245,255,5)	1/2	14.24	2.37
0110 (6)	64-QAM	(245,255,5)	7/12	16.62	2.77
0111 (7)	64-QAM	(245,255,5)	2/3	18.99	3.16
1000 (8)	64-QAM	(245,255,5)	3/4	21.36	3.56
1001 (9)	64-QAM	(245,255,5)	5/6	23.74	3.96
1010 – 1111 (10 – 15) Reserved					

9.3.2.1.4 Multiple antenna field (ANT)

The ANT field bits are used to determine the type of multiple antennae transmission scheme used to transmit the payload symbols. This field is 2 bits long and occupies the positions 14 and 15 in the PHY header. The multiple antennae field bits are defined in Table 138.

Table 138 — ANT field bits A1 and A0 definition

Multiple antennae field bits A1 – A0	Multiple antennae transmission scheme
00	Reserved
01	Frequency Interleaved Transmit Diversity (FITD)
10	Space Time Block Code (STBC)
11	Spatial Multiplexing (SM)

If a transmitting device uses 2 antennae, then the ANT mode bits A1 and A0 in the PHY header denote whether the transmitter is using FITD, STBC or SM for the data payload. The multiple antenna modes are described in 9.11. If a transmitting device uses only a single antenna, then the ANT mode bits shall be set to “00”.

9.3.2.1.5 PLCP length field (LENGTH)

The PLCP length field indicates the length of the PSDU/MPDU (including aggregation) in bytes, which includes the data and FCS, and does not include tail bits, and the pad bits. The LENGTH field is 12 bits in size and represents PSDU/MPDU sizes from 0 to 4095 bytes. The LENGTH field occupies positions 16 to 27 in the PHY header with LSB (L0) bit at position 16 and MSB (L11) bit at position 27.

9.3.2.1.6 Scrambler initialization seed (SCR)

The scrambler initialization seed is used to initialize the pseudo random sequence generator (PRBS) used in the scrambler. The SCR field is located in positions 28 and 29 in the PHY header.

9.3.2.1.7 Interleaver option (INTLVR)

The interleaver option is used to indicate the interleaver parameters used in the encoding of the PSDU. The INTLVR field is located in positions 32 and 33 in the PHY header and is defined as shown in Table 139.

Table 139 — Interleaver option field definition

INTLVR bits I1 – I0	Interleaver column size N_{col}
00	14
01	Not used
10	Not used
11	7

9.3.2.1.8 Transmit power field (TXPWR)

The transmit power field indicates the relative transmit power level that will be used to transmit the current packet/frame. The TXPWR field is located in positions 37 – 39 in the PHY header and is defined as shown in Table 140.

Table 140 — Transmit power field definition

TXPWR (P2 – P0)	Relative Transmit Power Level (dB)
000	0
001	3
010	6
011	9
100	12
101	15
110	18
111	21

9.3.2.1.9 Reserved bits

The bits in positions 30, 31, and 34-36 in the PHY header are reserved.

9.3.2.2 MAC header

The MAC header field received from the MAC is incorporated in to the PLCP header without any modifications.

9.3.2.3 Encoding of PLCP header

In order to enable the receiver to obtain the CP mode information before completing the full decoding of PLCP header, the encoding of the PLCP header is performed differently compared to other packet based standards. The PLCP header is RS encoded as described in 9.3.2.3.1, and then convolutionally encoded with a code rate of $R=1/2$ as described in 9.3.3.2.2, bit interleaved as described in 9.3.3.3 with $N_{\text{col}}=14$, and transmitted using QPSK modulation. Figure 34 shows the different steps in the PLCP header encoding process. The PHY header and MAC header bytes are input to a systematic $(23, 15, 4)$ RS encoder. The resultant 8 parity bytes are appended to the PHY and MAC header bytes to result in a 184 bit vector. Tail bits are then inserted in the middle and at the end of this vector. The resultant 196 bit vector is then split in to two equal parts of 98 bits and then each is independently encoded and mapped in to two OFDM symbols HDR1 and HDR2. The CP ratio of each OFDM symbol is $1/8$.

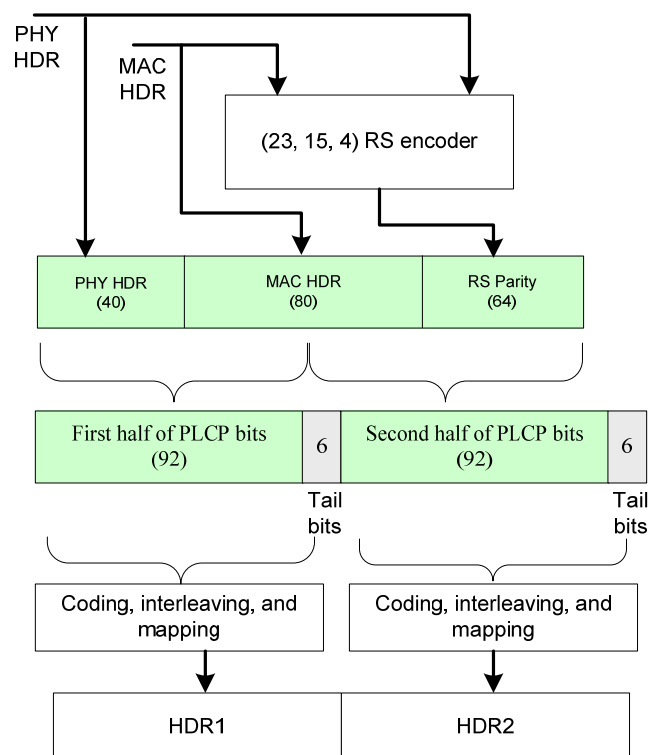


Figure 34 — PLCP header encoding process

9.3.2.3.1 RS coding for PLCP header

A RS encoder specified in 9.3.3.2.1 is also used to encode the PHY and MAC header bytes. Since the number of octets in the header is much smaller than the code length, less number of parity octets is used for the header field. The parity octets for the header are derived as follows:

- 1) The header message octets (size \tilde{K}) are pre-fixed with zero octets, as described in 9.3.3.2.1, such that the length of the message octets equals K .
- 2) The message block is then encoded using a (N, K, T) RS code to generate $2T$ parity octets
- 3) The last $(N - K - 2\tilde{T})$ parity octets are removed from the generated code word
- 4) The $(K - \tilde{K})$ padded octets are removed from the code word to form a $(\tilde{N}, \tilde{K}, \tilde{T})$ code block

Where $N = 255$, $K = 245$, $T = 5$ and $\tilde{N} = 23$, $\tilde{K} = 15$, $\tilde{T} = 4$.

At the receiver, if the RS decoder cannot recover the PHY and MAC headers (i.e. the code block has more than 4 byte errors) then the HEI is set to ZERO (invalid), otherwise the HEI is set to ONE (valid).

9.3.3 PSDU

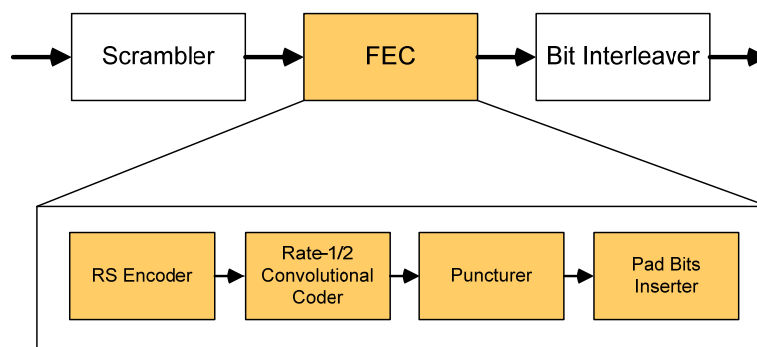


Figure 35 — PSDU encoding process

9.3.3.1 Data scrambling

The scrambler shall be used to scramble the PSDU/MPDU with the pseudo-random binary sequence (PRBS). The polynomial for generating the PRBS, $P(x)$, shall be : $P(x) = x^9 + x^4 + 1$. The scrambler shall be initialized on each PSDU/MPDU by using the seed value, S_0 and S_1 , which are specified by the MAC and defined within the PLCP header. $P(x)$ and Initial vector are shown in Figure 36, and the sequences of the scrambler output (S_{out}) according to each seed value are shown in the Table 141. The 511-bit sequence will be repeatedly generated by scrambler, and the same scrambler is used to scramble transmit data and to descramble the received data.




Figure 36 — Data scrambler

Table 141 — Sequences of scrambler output (S_{out})

[illegible]

(1, 1)	0, 0, 0, 0, 1, 1, 1, 1, 0, 1, 1, 1, 0, 0, 0, 0, 1, 0, 1, 1, 0, 0, 1, 1, 0, 1, 1, 0, 1, 1, 1, 0, 1, 0, 0, 0,
	0, 1, 1, 1, 0, 0, 1, 1, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 1, 0, 1, 0, 1, 1, 1, 0, 1, 0, 1, 1, 1, 0, 0, 1, 0,
	0, 1, 0, 1, 1, 1, 0, 0, 1, 1, 1, 0, 0, 0, 0, 0, 0, 1, 1, 0, 1, 1, 1, 0, 1, 0, 0, 1, 1, 1, 1, 0, 1, 0, 1, 0,
	0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 1, 0, 1, 0, 1, 0, 1, 0, 1, 1, 1, 1, 0, 1, 0, 1, 1, 0, 1, 0, 0, 0, 0, 0, 1, 1,
	0, 1, 1, 1, 0, 1, 1, 0, 1, 1, 0, 1, 0, 1, 1, 0, 0, 0, 0, 0, 1, 0, 1, 1, 1, 0, 1, 1, 1, 1, 0, 0, 0, 1, 1, 1,
	1, 0, 0, 1, 1, 0, 1, 0, 0, 1, 1, 0, 1, 0, 1, 1, 1, 0, 0, 0, 1, 1, 0, 1, 0, 0, 0, 1, 0, 1, 1, 1, 1, 1, 1, 0,
	1, 0, 0, 1, 0, 1, 1, 0, 0, 0, 1, 0, 1, 0, 0, 1, 1, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 1, 1, 0, 0,
	1, 0, 1, 0, 1, 1, 0, 0, 1, 0, 0, 1, 1, 1, 1, 1, 1, 0, 1, 1, 0, 1, 0, 0, 1, 0, 0, 1, 0, 0, 1, 1, 0, 1, 1, 1,
	1, 1, 0, 0, 1, 0, 1, 1, 0, 1, 0, 1, 0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 1, 0, 0, 1, 1, 1, 0, 1, 1, 0, 0, 1, 0, 1, 1,
	1, 1, 0, 1, 1, 0, 0, 0, 0, 1, 1, 0, 1, 0, 1, 0, 1, 0, 0, 1, 1, 1, 0, 0, 1, 0, 0, 0, 0, 1, 1, 0, 0, 0, 1, 0, 0,
	0, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 1, 0, 0, 1, 0, 0, 0, 1, 1, 1, 0, 1, 0, 1, 0, 1, 1,
	0, 1, 1, 0, 0, 0, 1, 1, 1, 0, 0, 0, 1, 0, 0, 1, 0, 1, 0, 0, 0, 1, 1, 0, 1, 1, 0, 0, 1, 1, 1, 1, 0, 0, 0,
	1, 1, 1, 1, 0, 0, 0, 1, 0, 1, 1, 0, 1, 1, 1, 0, 0, 1, 0, 1, 0, 0, 1, 0, 0, 0, 0, 0, 1, 0, 0, 1, 1, 0, 0, 1, 1,
	1, 0, 1, 0, 0, 0, 1, 1, 1, 1, 1, 0, 1, 1, 1, 1, 0, 0, 0, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,

9.3.3.2 Forward error correction (FEC)

9.3.3.2.1 Reed-Solomon outer coding

The RS coder is based on a systematic ($N = 255$, $K = 245$, $T = 5$) RS code using GF(256). Here N represents the total number of octets after encoding, K represents the number of octets before encoding, and T represents the error correcting capability of the coder. The number of parity octets is equal to $2T$.

The code generator polynomial shall be generated by $g(x) = \sum_{i=1}^{2T} (x + \alpha)^i$, where α is the primitive root of the field generator polynomial $p(x) = x^8 + x^4 + x^3 + x^2 + 1$.

Figure 37 describes the RS encoding process. Bit b0 in represents the first output of the scrambler. The output of the scrambler is first arranged into octets (B0 to BM) as shown in (b) which is then grouped into message blocks as shown in (c). The message block is fed to the RS encoder as shown in (d) with V_{244} being the first input octet. The output of the RS encoder is represented as $U = \{U_{254}, U_{253}, \dots, U_2, U_1, U_0\}$ where $\{U_{254}, U_{253}, \dots, U_{12}, U_{11}, U_{10}\} = \{V_{244}, V_{243}, \dots, V_2, V_1, V_0\}$ and $\{U_9, U_8, \dots, U_2, U_1, U_0\}$ represents the parity octets with U_9 being the first parity octet and U_0 the last parity octet. The RS encoder output octets are converted to bits using the inverse operation to the one described in (b) and are then fed to the convolutional encoder.

The last message block could have less than 245 octets in which case the code is shortened to derive RS coder with $K < 245$. In this case, an appropriate number of zero octets are pre-fixed to the encoder message block. For example, if the last message block length is 100 octets then V_{244} to V_{100} in the message block L will be made zero. The zero octets are removed from the encoder output after the encoding process.

b _J	b _{J-1}	---	b _{j+1}	b _j	---	b ₁₆	b ₁₅	b ₁₄	b ₁₃	b ₁₂	b ₁₁	b ₁₀	b ₉	b ₈	b ₇	b ₆	b ₅	b ₄	b ₃	b ₂	b ₁	b ₀
----------------	------------------	-----	------------------	----------------	-----	-----------------	-----------------	-----------------	-----------------	-----------------	-----------------	-----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------	----------------


a) Scrambler output

MSB												LSB MSB												LSB
b _J	b _{J-1}	---		b _{j+1}	b _j	---		b ₁₆	b ₁₅	b ₁₄	b ₁₃	b ₁₂	b ₁₁	b ₁₀	b ₉	b ₈	b ₇	b ₆	b ₅	b ₄	b ₃	b ₂	b ₁	b ₀
B _M		---		B _m		---		B ₁								B ₀								

b) Conversion of scrambler output to octets

Message block L	---	Message block I	---	Message block 1				Message block 0			
				B 489	---	B 246	B 245	B 244	---	B 1	B 0
				V 0	---	V 243	V 244	V 0	---	V 243	V 244

c) Input to RS encoder



d) RS encoding process

Code block L	---	Code block I	---	Code block 1				Code block 0			
				C 509	---	C 256	C 255	C 254	---	C 1	C 0
				U 0	---	U 253	U 254	U 0	---	U 253	U 254

e) Output of RS encoder

Figure 37 — RS encoding process

9.3.3.2.2 Convolutional inner coding

The data burst is encoded using a rate – ½ binary convolutional encoder. The constraint length of this coder is equal to 7 and its generator polynomials are 133_o and 171_o. Figure 38 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 convolutional coder shall be initialized at the beginning of the PLCP header and at the beginning of the PSDU. Tail bits (6 in number) shall be used with the PSDU data to bring back the encoder to zero state.

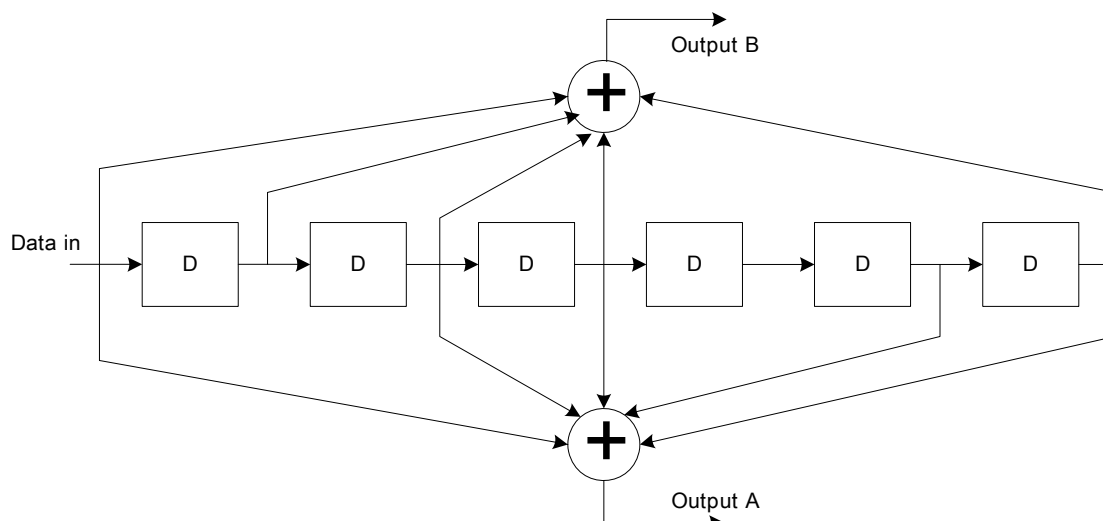


Figure 38 — Rate $\frac{1}{2}$ convolutional coder with generator polynomials 133_o, 171_o.
The delay element represents a delay of 1 bit

9.3.3.2.3 Puncturing

Puncturing is used to derive additional code rates from a base rate $R = 1/2$ convolutional code. Table 142 shows the puncturing process for the different code rates. The second row of this Table shows the convolutional coder output where A and B represent the first and second outputs bits respectively. The subscripts represent the output order. The third row of this Table shows the output of the puncturer, which is also the input to the bit-inserter module at the receiver. The output is derived by omitting some of the encoded bits at the transmitter. At the receiver, an inverse operation is performed by inserting zeroes into the decoder input at the locations corresponding to the omitted bits. The output of the bit-inserter (or the input to the decoder) is shown in the fourth row of this table. For the last block of bits, the puncturing pattern is applied to the partially filled block.

Table 142 — Puncturing and bit-insertion for the different coding rates

Code rate	1/2	2/3	3/4	5/6	7/12
Convolutional coder output	A_1B_1	$A_1B_1A_2B_2$	$A_1B_1A_2B_2A_3B_3$	$A_1B_1A_2B_2A_3B_3A_4B_4A_5B_5$	$A_1B_1A_2B_2A_3B_3A_4B_4A_5B_5A_6B_6A_7B_7$
Puncturer output/bit-inserter input	A_1B_1	$A_1B_1A_2$	$A_1B_1A_2B_3$	$A_1B_1A_2B_3A_4B_5$	$A_1B_1A_2B_2A_3B_3A_4B_4A_5B_5A_6A_7$
Decoder input	A_1B_1	$A_1B_1A_20$	$A_1B_1A_200B_3$	$A_1B_1A_200B_3A_400B_5$	$A_1B_1A_2B_2A_3B_3A_4B_4A_5B_5A_60A_70$

9.3.3.2.4 Pad bits insertion

The total number of coded bits, N_{CB} , shall be the multiple of N_{CBPS} , the number of coded bits (196, 392, or 588 bits) in one OFDM symbol according to the modulation order (QPSK, 16QAM, 64QAM). Therefore, if the N_{CB} is not multiple of N_{CBPS} , the pad bits shall be added at the rear of the encoded message. The number of pad bits, N_{PAD} , is computed with using the number of OFDM symbols, N_{SYM} , as follow:

$$N_{SYM} = \text{Ceiling}(N_{CB} / N_{CBPS}) \quad (12)$$

$$N_{PAD} = (N_{SYM} \times N_{CBPS}) - N_{CB} \quad (13)$$

The function Ceiling (•) is a function that returns the smallest integer value greater than or equal to its argument value. The appended bits are set to the values produced by the same scrambler defined in 9.3.3.1. The all ZERO bits with the length of N_{PAD} will be inserted into scrambler which is initialized with initial vector.

9.3.3.3 Bit interleaving

All encoded data bits shall be interleaved by a block interleaver with a block size corresponding to the number of bits in a single OFDM symbol, N_{CBPS} . The interleaver is defined by a two-step permutation. The first permutation ensures that adjacent coded bits are mapped onto nonadjacent subcarriers. The second ensures that adjacent coded bits are mapped alternately onto less and more significant bits of the constellation and, thereby, long runs of low reliability (LSB) bits are avoided.

We shall denote by k the index of the coded bit before the first permutation; i shall be the index after the first and before the second permutation, and j shall be the index after the second permutation, just prior to modulation mapping.

The first permutation is defined by the rule:

$$i = (N_{CBPS} / N_{col})(k \bmod N_{col}) + \text{floor}(k / N_{col}) \quad k = 0, 1, \dots, N_{CBPS} - 1 \quad (14)$$

The function floor (.) denotes the largest integer not exceeding the parameter.

The second permutation is defined by the rule:

$$i = s \times \text{floor}(i / s) + (i + N_{CBPS} - \text{floor}(N_{col} \times i / N_{CBPS})) \bmod s \quad i = 0, 1, \dots, N_{CBPS} - 1 \quad (15)$$

The value of s is determined by the number of coded bits per subcarrier, N_{CBPC} , according to:

$$s = \max(N_{CBPC}/2, 1) \quad (16)$$

where the parameter, N_{col} , is determined from the INTLVR field.

The deinterleaver, which performs the inverse relation, is also defined by two permutations.

Here we shall denote by j the index of the original received bit before the first permutation; i shall be the index after the first and before the second permutation, and k shall be the index after the second permutation, just prior to delivering the coded bits to the convolutional (Viterbi) decoder.

The first permutation is defined by the rule:

$$i = s \times \text{floor}(j / s) + (j + \text{floor}(N_{col} \times j / N_{CBPS})) \bmod s \quad j = 0, 1, \dots, N_{CBPS} - 1 \quad (17)$$

where s is as defined above.

The second permutation is defined by the rule:

$$k = N_{col} \times i - (N_{CBPS} - 1) \text{floor}(N_{col} \times i / N_{CBPS}) \bmod s \quad i = 0, 1, \dots, N_{CBPS} - 1 \quad (18)$$


9.4 Constellation mapping and modulation

9.4.1 Data modulation

The output of the bit interleaver is entered serially to the constellation mapper. The input data to the mapper is first divided into groups of N_{CBPC} (2, 4 or 6) bits and then converted into complex numbers representing QPSK, 16-QAM or 64-QAM constellation points. The mapping is done according to Gray-coded constellation mapping as shown in Figure 39. The input bit, b_0 , is the first bit among the N_{CBPC} bits. The complex valued number is scaled by a modulation dependent normalization factor K_{MOD} . Table 143 shows the K_{MOD} values for the different modulation types defined in this clause.

Table 143 — Modulation dependent normalization factor

Modulation Type	K_{MOD}
QPSK	$1/\sqrt{2}$
16-QAM	$1/\sqrt{10}$
64-QAM	$1/\sqrt{42}$



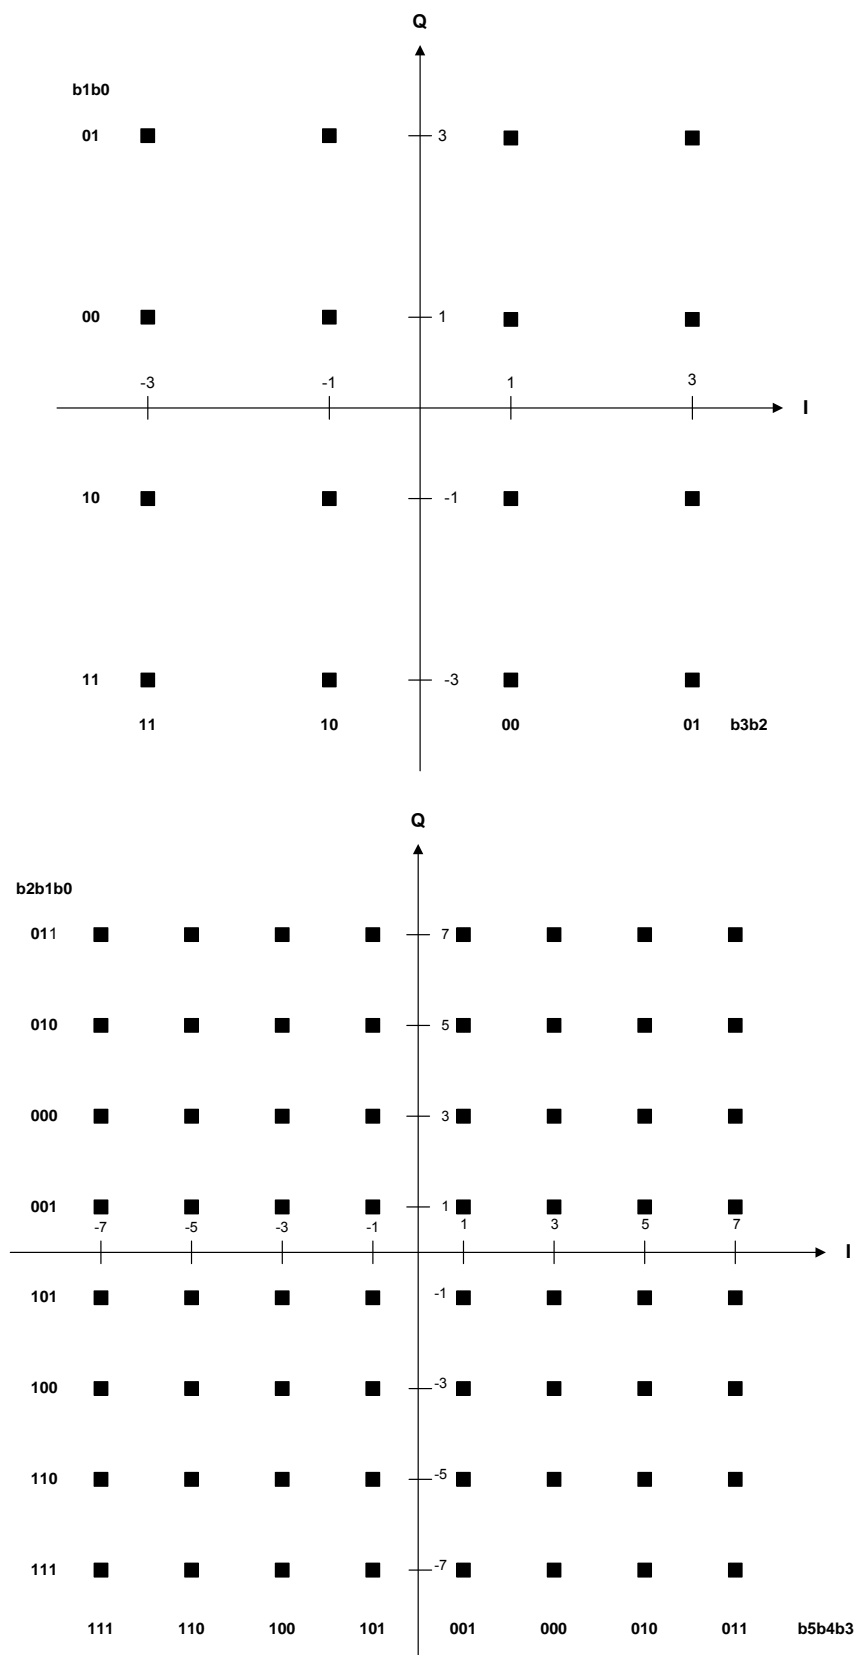


Figure 39 — QPSK, 16-QAM, and 64-QAM constellations

9.4.2 Pilot modulation

The pilot subcarriers shall be modulated by using BPSK constellation mapping as shown in Figure 40. The input bit, b_0 , shall be provided by the PRBS generator shown in 9.3.3.1. The normalization factor K_{MOD} for BPSK constellation is 1.




Figure 40 — BPSK constellation

9.5 OFDM modulation

The transmitted RF signal, $s_{RF}(t)$, was introduced in 9.2.1. The complex base-band representation of the n th OFDM symbol, $s_n(t)$, is defined as

$$s_n(t) = \sum_{k=-\frac{N_{FFT}}{2}}^{\frac{N_{FFT}}{2}-1} C_{k,n} \exp(j2\pi k \Delta f (t - T_{GI})) , \quad 0 \leq t < T_{SYM} \quad (19)$$

where

n	denotes the symbol number
k	denotes the subcarrier number
N_{FFT}	is the number of total subcarriers, i.e. FFT size
$C_{k,n}$	is the complex constellations corresponding to subcarrier k of OFDM symbol n
Δf	is the subcarrier spacing, i.e. the inverse of time duration of useful symbol
T_{GI}	is the time duration of cyclic prefix
T_{SYM}	is the time duration of OFDM symbol

The complex constellations, $C_{k,n}$, are modulated by four types of subcarriers which are data, pilot, guard, and DC. Each subcarrier has a different purpose within OFDM symbol. T_{GI} and T_{SYM} could have different values for the PLCP preamble, PLCP header, and payload.

9.5.1 Data subcarriers

In all OFDM symbol following the PLCP preamble, 98 subcarriers among 102 used subcarriers are used for data transmission. These data subcarriers carry the complex constellations described in 9.5. The stream of complex constellations from mapper is divided into groups of $N_D=98$ complex constellations. A group of complex constellations are sequentially mapped to the IFFT inputs from -51 to 51, excluding the IFFT inputs for pilot and DC subcarriers.

9.5.2 Pilot subcarriers

In all OFDM symbol following the PLCP preamble, four of the subcarriers are allocated for pilot signals in order to make the coherent detection and to provide the robustness of the transmission system against the frequency offsets and phase noise. These pilot signals shall be inserted in subcarriers for 13 OFDM symbols, as defined in Table 144. The pilot insertion pattern is repeated per every 13 OFDM symbols.

Table 144 — Pilot subcarrier index during 13 OFDM symbols

Symbol Index modulo 13	0	1	2	3	4	5	6	7	8	9	10	11	12
Subcarrier Index	-51	-39	-31	-45	-35	-27	-49	-41	-33	-47	-29	-37	-43
	-25	-13	-5	-19	-9	-1	-23	-15	-7	-21	-3	-11	-17
	1	13	21	7	17	25	3	11	19	5	23	15	9
	27	39	47	33	43	51	29	37	45	31	49	41	35

NOTE The first OFDM symbol starts after the long preamble from 0.

The pilot insertion pattern is shown in Figure 41.




Figure 41 — Pilot insertion pattern

The pilot signals shall be BPSK modulated by a pseudo random binary sequence to avoid the generation of line spectral frequencies. The polynomial for the PRBS generator shall be the same as that used for the data scrambler, as specified in 9.3.3.1. The PRBS generator shall be initialized by the seed 1 1 1 1 1 1 S1 S0, where S1=1 and S0=0. Thus, the sequence for pilot modulation is the same as S1=1 and S0=0 case, as shown in Table 141. The sequence shall be used for pilot subcarriers from the PLCP header symbol in a successive manner. The first four bits, 0, 0, 0, and 1, are used for pilot modulation in the first PLCP header symbol, while the next four bits, 1, 1, 1, and 0, are used for pilot modulation in the second PLCP header symbol, and so on.

9.5.3 Null subcarriers

Null subcarriers include the DC subcarrier and the guard subcarriers. No power is allocated to the null subcarriers. The DC subcarrier falling at 0th subcarrier is not modulated in order to prevent any saturation effects or excess draw at the amplifier. For each OFDM symbol, 25 subcarriers are allocated as guard subcarriers. These guard subcarriers are located on either edge of the OFDM symbol. The 13 and 12 subcarriers are used as left and right guard subcarriers, respectively. The guard subcarriers are used to fit the spectrum within the allocated bandwidth, thus reduce the interference between adjacent channels and relax the specs on analog tx/rx filters.

9.5.4 Implementation of Fourier transform

The inverse Fast Fourier Transform (IFFT) is a common implementation for performing the inverse Discrete Fourier Transform. The left side of the spectrum from -51 to -1 is mapped into the lower IFFT inputs from 77 to 127, while the right side of the spectrum from 1 to 51 is mapped into the upper IFFT inputs from 1 to 51. And the DC of the spectrum is mapped into 0. Due to the DC subcarrier and the guard band subcarriers on the left side and right side of the spectrum, the IFFT inputs 0 and from 52 to 76 shall be set to null. The subcarrier falling at DC is not used since it may cause the problems in D/A and A/D converter and carrier feed-through in the RF system.

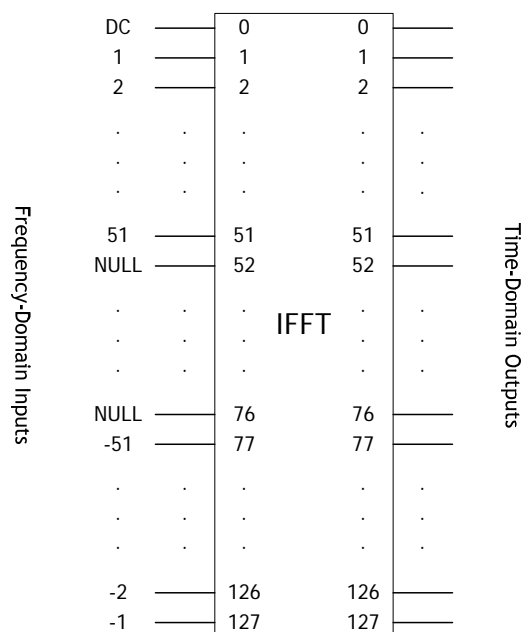


Figure 42 — Inputs and outputs of IFFT

9.6 General block diagram for the OFDM PHY

The general block diagram of the transmitter and receiver applicable to the OFDM PHY layer is shown in Figure 43. This sub-clause specifies the functional processing of the baseband signal for OFDM PHY layer.

The binary sequences for transmission are supplied to the PHY layer from the MAC layer. These are input to a channel coding processor which includes data scrambler, RS encoder, convolutional encoder, puncturer, and bit interleaver. It is specified in 9.3.3 in detail. The interleaved binary sequences shall be mapped to data constellations according to modulation schemes specified in 9.4.1. These data constellations shall be mapped onto the data subcarrier k of the OFDM symbol n using Equation (19). In the frequency-domain, an OFDM symbol contains the data, pilot, and null subcarriers, as described in 9.5.1 through 9.5.3. In order to support the synchronization, channel estimation, and tracking process, the preamble is inserted in the first 1 or 3 OFDM symbols of each frame specified in 9.3.1, and the pilot subcarriers are transmitted at fixed positions in the frequency domain within each OFDM symbol as specified in 9.5.2. The resultant stream of constellations is subsequently input to an inverse Fast Fourier Transform (IFFT). The OFDM symbol is cyclically extended by a cyclic prefix to prevent inter-symbol interference (ISI) and inter-carrier interference (ICI) due to delay spread. Finally, the OFDM signal is delivered to the front-end modules via an AD converter. At the receiver, the operations are the reverse of that for the transmitter. In addition to the functional processing of the data, the synchronization and channel estimation shall be performed at the receiver.

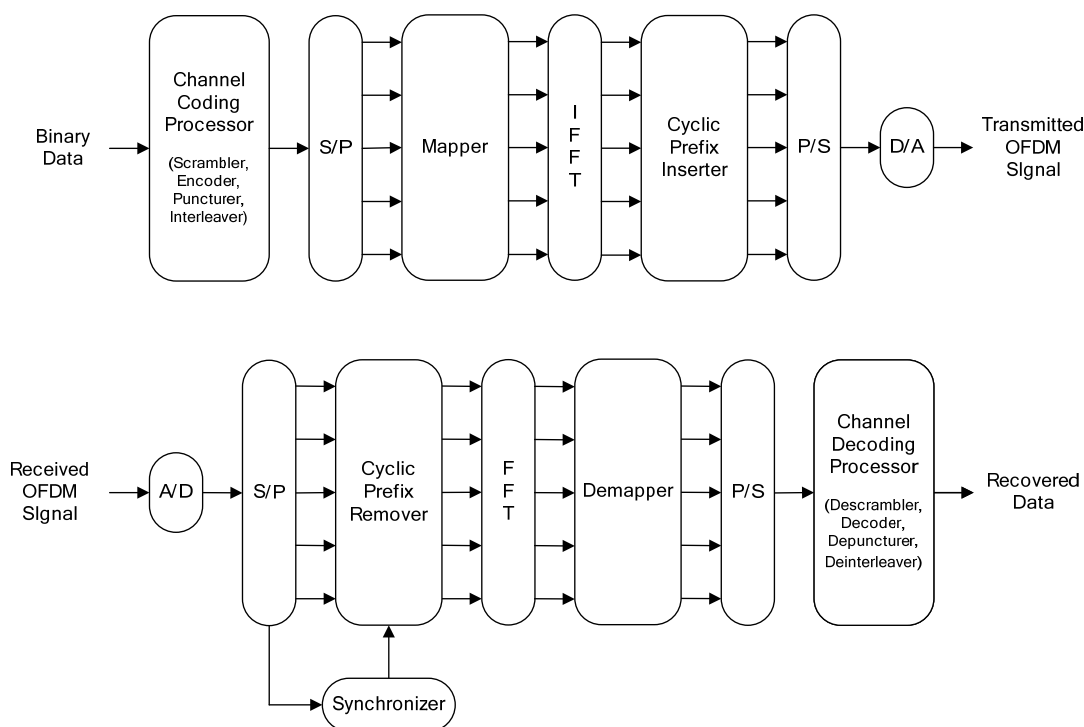


Figure 43 — Transmitter and receiver block diagram for the OFDM PHY

9.7 General requirements

9.7.1 Operating frequency range

This PHY operates in the VHF/UHF TV broadcasting frequencies, subject to regulation. The extremes of the international regulations range of TV broadcast band is from 47 MHz to 910 MHz.

9.7.2 Channel bandwidth and numbering

The channel bandwidth is equivalent to that of one broadcast channel, i.e., 6 MHz, 7 MHz, or 8 MHz. The channel bandwidth and channel numbering might be dependent on the geographic location where the device is intended for operation.

9.7.3 PHY layer timing

The values for the PHY layer timing parameters are defined in Table 145. The interframe spacing parameters are defined by pMIFS or pSIFS.

Table 145 — PHY layer timing parameters

PHY Parameter	Value
pMIFS	2 μ s
pSIFS	10 μ s
pSlotTime	15 μ s
pCCADetectTime	10 μ s

9.7.3.1 Receive-to-transmit transition gap (RTG)

The receive-to-transmit transition gap (RTG) is defined as the time period required for a device to process the last received symbol and then switch to transmit mode. The period starts when the device has received the last sample of the last symbol and ends when the device is ready to send the first sample of the preamble. It should be less than pSIFS.

9.7.3.2 Transmit-to-receive transition gap (TTG)

The transmit-to-receive transition gap (TTG) is defined as the time period required for a device to complete transmission of the last symbol and then switch to receive mode. The period starts when the device has transmitted the last sample of the last symbol and ends when the device is ready to receive next transmission. It should be less than pSIFS.

9.7.3.3 Time between successive transmissions

In normal mode, the time between successive transmissions equals SIFS period. In burst mode, the time between successive transmissions equals MIFS period.

9.8 Transmitter requirements

9.8.1 Transmit center frequency tolerance

The transmitter center frequency tolerance shall be within ± 20 ppm.

9.8.2 Symbol clock frequency tolerance

The symbol clock frequency tolerance shall be within ± 20 ppm.

9.8.3 Clock synchronization

The transmitter center frequency and the symbol clock frequency should be derived from the same reference oscillator.

9.8.4 Transmitter constellation error

The degradation of the receiver SNR shall be no more than 0.5dB due to the relative constellation error in the transmitter. The relative constellation RMS error, averaged over all data and pilot subcarriers of the OFDM

symbols and over all of the frames, shall not exceed the values defined in the Table 146. All data transmission types are considering the (255, 245, 5) RS encoding.

Table 146 — Allowed relative constellation error versus data transmission type

Modulation - code rate	Data rate (Mbits/s)	Relative constellation error (dB)
QPSK - 1/2	4.75	11.7 dB
QPSK - 2/3	6.33	13.6 dB
16-QAM - 1/2	9.49	17.1 dB
16-QAM - 7/12	11.08	18.6 dB
16-QAM - 2/3	12.66	19.7 dB
64-QAM - 1/2	14.24	21.9 dB
64-QAM - 7/12	16.62	23.6 dB
64-QAM - 2/3	18.99	24.8 dB
64-QAM - 3/4	21.36	26.6 dB
64-QAM - 5/6	23.74	28.1 dB

The relative constellation RMS error shall be calculated with using a device capable of converting the transmitted signal into a stream of complex samples at 48/7 Msamples/s or more, with sufficient accuracy in the I/Q imbalance, DC offset, phase noise, etc. The sampled signal shall then be processed in a manner similar to that of an ideal receiver. An example of the minimum steps necessary for receiver processing as following:

1. Detect the start of the frame.
2. Estimate the coarse and fine frequency offset and correct them.
3. Establish the fine timing of OFDM symbol.
4. Estimate the channel frequency response and equalize the channel.
5. For each of the data and pilot subcarriers, find the closest constellation point and compute the Euclidean distance.
6. Compute the RMS error, averaged over all the data and pilot subcarriers in the OFDM symbols for payload and over all frames, as follows:

$$RMS_{error} = \frac{\sum_{i=1}^{N_{Packet}} \sqrt{\sum_{n=1}^{N_{SYM}} \left[\frac{\sum_{k=1}^{N_D} |D[n, k] - D_0[n, k]|^2 + \sum_{k=1}^{N_P} |P[n, k] - P_0[n, k]|^2}{(N_D + N_P) \cdot N_{SYM} \cdot P_0} \right]}}{N_{Packet}} \quad (20)$$

where

- N_{Packet} the number of packet
- N_{SYM} the number of OFDM symbols in the packet, not including the OFDM symbols for PLCP preamble and PLCP header
- N_D the number of data subcarriers in an OFDM symbol
- N_P the number of pilot subcarriers in an OFDM symbol
- P_0 the average power of the data and pilot constellations
- $D_0[n,k]$ the ideal constellation point for k^{th} data subcarrier for the n^{th} OFDM symbol
- $P_0[n,k]$ the ideal constellation point for k^{th} pilot subcarrier for the n^{th} OFDM symbol
- $D[n,k]$ the observed constellation point for k^{th} data subcarrier for the n^{th} OFDM symbol
- $P[n,k]$ the observed constellation point for k^{th} pilot subcarrier for the n^{th} OFDM symbol.

NOTE The test shall be performed over a minimum of $N_{Packet} = 100$ packets, where the PSDU of each packet is at least 30 symbols in length and is generated from random data.

9.9 Receiver requirements

9.9.1 Receiver sensitivity

The minimum receiver sensitivity levels for the different PHY modes in an AWGN channel are listed in Table 147. These numbers correspond to a packet error rate of 1% with a payload size of 1960 bytes and assuming a noise figure of 6.0 dB, and an implementation loss of 6.0 dB. The minimum input levels are measured at the antenna connector.

Table 147 — Minimum receiver sensitivity levels for different PHY mode

PHY Mode	Minimum Sensitivity (dBm)
0	-92.1
1	-90.2
2	-86.7
3	-85.2
4	-84.1
5	-81.9
6	-80.2
7	-79.0
8	-77.2
9	-75.7

9.9.2 Maximum received signal level

The receiver shall provide a maximum PER of 10% for a PSDU length of 2048 bytes for a maximum received input level of -30 dBm for all PHY modes.

9.9.3 Center frequency and symbol clock frequency tolerance

The receiver center frequency and the symbol clock frequency should be derived from the same reference oscillator and shall be within tolerance ± 20 ppm.

9.9.4 Link quality estimate

To communicate reliably, receivers should estimate and report the Link Quality Estimate (LQE) of the received channel. LQE is an implementation specific estimation of the mean SNR on the packet duration available after FFT. Devices shall report either “not supported” or the estimates through the LQE field in the LQE IE as described in 7.1.8.38. The mean SNR shall be quantized in 1 dB increments in the range from -10 dB to +40 dB. The mapping between the estimated value and the link quality estimate (LQE) is defined in Table 148. Values outside those ranges shall be assigned the closest extreme value. The all-zero bits indicate that the device cannot support the reporting of LQE.

Table 148 — Encoding of the link quality estimates

LQE	Description
0000 0000	The report of LQE is not supported
0000 0001	-10 dB or smaller
0000 0010	-9 dB
...	...
0011 0011	+40 dB
0011 0100 - 1111 1111	Reserved

The accuracy of the LQE is defined as the standard deviation of the packet-by-packet SNR estimates for a static AWGN channel and a fixed SNR value. The test for the accuracy of the LQE shall be performed over a minimum of 1000 packets with a PSDU generated from random data with length equal to 1024 bytes. Table 149 shows the allowed standard deviation for the different estimated values.

Table 149 — Allowed standard deviation for the link quality estimates

Link Quality Estimate (LQE)	Allowed standard deviation
-10 dB, ..., -6 dB	1.3 dB
-5 dB, ..., -1 dB	1.1 dB
0 dB, ..., 6 dB	0.9 dB
7 dB, ..., 15 dB	0.7 dB
16 dB, ..., 40 dB	0.6 dB

9.10 Control mechanisms

9.10.1 Device synchronization

All the receiver devices shall be synchronized with the transmitter using the PLCP preamble. The PLCP preamble shall be used by the receiver for timing and frequency synchronization. At the receiver device, the center frequency and clock frequency shall be synchronized and locked to the transmitter with a tolerance of maximum 2% of subcarrier spacing.

9.10.2 Transmit power control

All devices shall provide transmit power control (TPC). The objective of power control is to limit their transmit power to the minimum necessary for reliable communication. It prevents harmful interference to incumbent users and mitigates interference to other unlicensed users. The transmit power, P_{next} , is defined as follows:

$$P_{next} = \begin{cases} P_i, & \text{at the first transmission} \\ P_{prev} + \Delta P, & \text{following the first transmission} \end{cases} \quad (21)$$

At the first transmission, the power level, P_{next} , shall use the initial transmit power, P_i . When the transmit power change for the device is needed, the power level shall be adjusted with the proper amount of power update, ΔP . Thus the key of power control mechanism is to determine the optimum initial transmit power, P_i , and power update, ΔP , considering the neighbouring environments and transmission schemes.

The initial transmit power and power update shall be determined by considering the existence of incumbents in adjacent channels. The initial transmit power and power update should also consider unlicensed signals in neighbouring networks. The existence of incumbents is determined by incumbent protection mechanisms.

For the initial transmit power, P_i , if no signal is detected in adjacent channels and adjacent networks, the initial transmit power shall be allowed up to the maximum transmit power. However, when the incumbent signals appear in the adjacent channels or other unlicensed signals appear in neighbouring networks, the initial transmit power should be reduced if necessary. The differential decrement of initial transmit power compared to the maximum power level is implementation-dependent. For the power update, ΔP , should be limited to $mMaxPowerUpdateStep$ except certain condition that ΔP has to be larger than $mMaxPowerUpdateStep$, e.g., the detection of incumbents in the adjacent channels. When ΔP is no larger than $mMaxPowerUpdateStep$, a device may use either link feedback IE or Transmit Power Control IE to update the transmit power of other devices. When ΔP is larger than $mMaxPowerUpdateStep$, a device shall use Transmit Power Control IE to adjust the transmit power of other devices.

The power control mechanism should also take into account the number and frequency location of the incumbents or unlicensed users. Using the incumbent protection mechanisms, the device might be able to obtain the number and frequency location of the incumbents (e.g., narrowband wireless microphones) in adjacent channels.

For a master device or a peer device, the device may decide the initial transmit power and power update by itself using incumbent protection mechanisms. For a slave device, the initial transmit power and power update shall be determined from the master device using Link Feedback IE or Transmit Power Control IE or Channel Classification IE.

The transmit power control should also be affected by the transmission scheme, such as modulation and coding schemes. For example, the applications using higher-order modulation need higher transmit power to maintain the link quality. And, in order to rapidly adjust the transmit power level, the applications using high-order modulation need larger step of power update than the value for low-order modulation.

The transmit power should be periodically updated by monitoring the signal quality. The decision algorithm to update the transmit power is implementation-dependent. Thus the parameters such as signal quality (i.e. RSSI, CINR, PER, etc.), decision threshold, and decision method are not specified here.

9.11 Multiple antennae (optional)

Devices may support three different types of multiple antenna transmitter schemes: (1) Frequency Interleaved Transmit Diversity (FITD), (2) Alamouti Space Time Block Coding (STBC) and (3) Spatial Multiplexing (SM). In FITD and STBC 2 antennae are required at the transmitting device, however the receiver may have only one antenna. In SM both transmitting and receiving devices need 2 antennae. This specification supports a maximum of 2 transmit antennae. Even though FITD is optional at the transmitter, all receivers shall support FITD reception with 1 receive antenna. This is because, in order to maintain robustness of the header and beacon and compatibility between devices with different capabilities, if a transmitting device uses 2 transmit

Auto-correlation at time instant m is given as

$$\rho(m) = \sum_{p=0}^{P-1} r(m-p)r^*(m-N-p) \quad (26)$$

where p is the correlation window size.

In the case of multiple transmission modes, the synchronization field also includes the information on the transmission mode embedded within. Note that auto-correlation derived from Equation (26) does not provide information about the transmission mode. In order to identify the transmission mode, the receiver performs a parallel correlation of the received signal with different training sequences as shown below.

$$\rho_1(m) = \sum_{n=0}^{N-1} r(m-n)P_{ST1}(n) \quad (27)$$

$$\rho_{2A}(m) = \sum_{n=0}^{N-1} r(m-n)P_{ST2}(n) \quad (28)$$

The maximum of all the correlator outputs is selected and is compared against a threshold to detect the presence of the signal and also to identify the transmission mode.

If the receiver determines that a packet is transmitted using dual antenna mode then it uses FITD decoding process to identify the multiple antenna transmission mode. This information is used to determine the appropriate processing steps to decode the rest of the packet.

9.11.3 Pilot subcarriers for all multiple antennae modes

For all multiple antennae modes, pilot subcarriers will be inserted as follows: in all OFDM symbols following the PLCP preamble, two of the subcarriers in each OFDM symbol will be allocated for pilot subcarriers. Antenna 1 shall use the subcarrier indices specified in Rows 1 and 3 of Table 144 and Antenna 2 shall use the subcarrier indices specified in Rows 2 and 4. The pilot symbols will be generated in the same way as described in 9.5.2.

9.11.4 Frequency interleaved transmit diversity (FITD)

The FITD transmission is done as follows: the bit stream is convolutionally encoded and punctured, interleaved and mapped to symbols in the same way as for single-antenna transmission. After that, symbols with indices $\{-51:2:51\}$, skipping the 4 pilot positions are mapped to the corresponding indices for Antenna 1 (48 symbols) and symbols with indices $[-50:2:50]$ skipping the DC are mapped to corresponding indices in Antenna 2 (50 symbols).




Figure 45 — Block diagram of Frequency Interleaved Transmit Diversity (FITD) transmission scheme

9.11.5 Alamouti space time block coding (STBC)

STBC mode shall NOT be used with rate $\frac{1}{2}$ QPSK, but may be used with all other modes. STBC is performed over the same frequency index over 2 consecutive OFDM symbols. Hence STBC encoding will always have an even number of transmitted OFDM symbols and the data will have to be zero-padded to make it so. In order to keep the total transmitted power the same as for single antenna stream, the data symbol stream on each antenna is normalized by multiplying with $1/\sqrt{2}$. The block diagram for STBC encoding is as follows:



Figure 46 — Block diagram of Space Time Block Code (STBC) transmission scheme

9.11.5.1 Pad bits insertion for STBC

Since the total number of transmitted OFDM symbols have to be even, the number of padded bits shall be calculated as follows:

$$N_{SYM,STBC} = 2 \text{Ceiling}\left(\frac{N_{CB}}{2N_{CBPS}}\right) \quad (29)$$

$$N_{PAD,STBC} = (N_{SYM,STBC} \times N_{CBPS}) - N_{CB} \quad (30)$$

9.11.6 Spatial multiplexing (SM) mode

In this mode, the bit stream is split into two streams after convolutional coding and puncturing. Stream 1 is then bit interleaved with $N_{col} = 14$ and Stream 2 is bit interleaved with $N_{col} = 7$. Each stream is then mapped to symbols and transmitted. When the INTLVR bits are equal to "11" then Stream 1 is bit interleaved with $N_{col} = 7$ and Stream 2 is bit interleaved with $N_{col} = 14$. Each stream will have the same data rate chosen from Table 137 for total data rate twice that of each stream. In order to keep the total transmitted power the same as for single antenna stream, the data symbol stream on each antenna is normalized by multiplying with $1/\sqrt{2}$. Figure 47 shows the Block diagram of Spatial Multiplexing (SM) transmission mode.

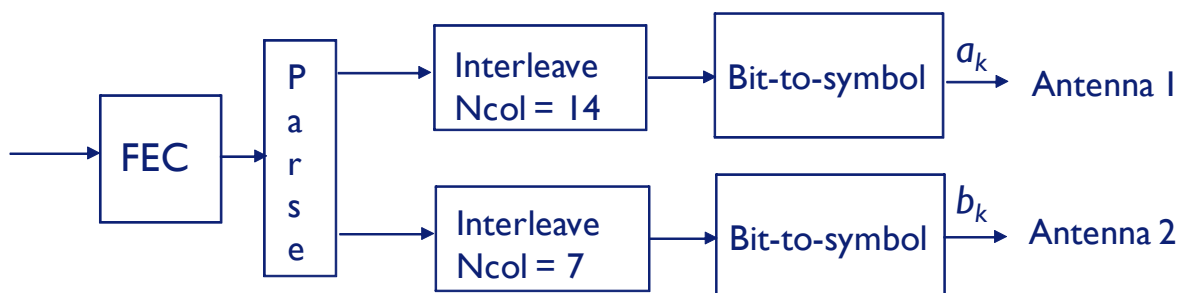


Figure 47 — Block diagram of Spatial Multiplexing (SM) transmission mode

Annex A (normative)

MUX sublayer

The MUX sublayer is the MAC client as depicted in Figure 2 and routes data between the MAC sublayer and MUX clients.

A.1 MUX service

The MUX sublayer is expressed in terms of the MUX SAP, the MUX service, and the MUX client. Each MUX client is associated with a unique protocol. Service data units presented at the MUX SAP by the MUX client are therefore associated with that protocol.

The protocol is encoded in a MUX header as either:

- A protocol identifier and an OUI; or
- A protocol identifier; or
- An IEEE EtherType value [2].

The MUX service adds a MUX header to the MUX service data unit to construct a MUX protocol data unit. The MUX sublayer makes use of the service provided by the MAC sublayer for the transfer of its protocol data units.

On receipt of a MUX protocol data unit from the MAC Sublayer, the MUX service removes the MUX Header and delivers the transported service data unit to the appropriate MUX client based on the identified protocol.

A.2 MUX protocol data unit format

A MUX protocol data unit consists of a MUX Header and a MUX Payload and is shown in Figure A.1.

octets: 2 or 5	N
MUX Header	MUX Payload

Figure A.1 — MUX protocol data unit format

The MUX Payload field contains the MUX service data unit that is a payload data unit of the protocol identified in the MUX Header.

The first two octets of the MUX Header are encoded as unsigned binary values, and are delivered to the MAC sublayer in order from the octet containing the most-significant bits to the octet containing the least-significant bits. The octet order for this field is the reverse of that for most fields in this specification.

The MUX Payload is a sequence of octets labeled as MUX Payload[0] through MUX Payload[N-1]. Octets are passed to the MAC sublayer in ascending index-value order.

The MUX Header and MUX Payload together form the payload of the MAC sublayer.

There are three versions of the MUX Header, which are distinguished based on the value of the first two octets of the header.

The first version has a length of five octets and is specified in Figure A.2.

octets: 2	3
Protocol ID (0x0000 – 0x00FF)	OUI

Figure A.2 — Format of first version of MUX Header

The Protocol ID field is restricted to values from 0 through 255 and is set to a value that identifies a protocol defined by the owner of the OUI specified in the OUI field. The OUI is a sequence of 3 octets, labeled as oui[0] through oui[2]. Octets of the OUI are passed to the PHY SAP in ascending index-value order.

The second version of the MUX Header has a length of two octets and is defined in Figure A.3. The Protocol ID field is restricted to values from 256 through 1535 and is reserved for future use.

octets: 2
EtherType (0x0100 – 0x05FF)

Figure A.3 — Format of second version of MUX Header

The third version of the MUX Header has a length of two octets and is specified in Figure A.4.

octets: 2
EtherType (0x0600 – 0xFFFF)

Figure A.4 — Format of third version of MUX Header

The EtherType field is restricted to values from 1536 through 65535 and is set to the value of an EtherType [2] identifying a protocol.

Annex B (normative)

OFDM parameters for 7 MHz and 8 MHz channel bandwidths

The sampling frequency for the 7 MHz and 8 MHz channel bandwidth is 8 MHz and 64/7 MHz, respectively. According to the channel bandwidth, the OFDM parameters may be changed as shown in Table B.1 and Table B.2. The change of channel bandwidth will affect the subcarrier spacing, the FFT period, the cyclic prefix duration, signal bandwidth, and so on.

Table B.1 — OFDM Parameters for 7 MHz bandwidth

Parameter	Value
Subcarrier spacing, ΔF (KHz)	62.5
FFT period, T_{FFT} (μs)	16.0
Total number of subcarriers, N_{FFT}	128
Number of guard subcarriers, N_G (L, DC, R)	26(13, 1, 12)
Number of used subcarriers, $N_T = N_D + N_P$	102
Number of data subcarriers, N_D	98
Number of pilot subcarriers, N_P	4
Signal bandwidth (MHz)	6.438

Table B.2 — OFDM Parameters for 8 MHz bandwidth

Parameter	Value
<i>Subcarrier spacing, ΔF (KHz)</i>	71.429
<i>FFT period, T_{FFT} (μs)</i>	14.0
<i>Total number of subcarriers, N_{FFT}</i>	128
<i>Number of guard subcarriers, N_G (L, DC, R)</i>	26(13, 1, 12)
<i>Number of used subcarriers, $N_T = N_D + N_P$</i>	102
<i>Number of data subcarriers, N_D</i>	98
<i>Number of pilot subcarriers, N_P</i>	4
<i>Signal bandwidth (MHz)</i>	7.357

Annex C (normative)

Data rates for 7 MHz and 8 MHz channel bandwidths

Table C.1 — Mapping of the PHY layer parameters to Rate field. The data rates are derived using channel bandwidth = 7 MHz and with CP duration = 1.0 μ s (1/16 T_{FFT})

RATE R3 – R0	Modulation	Outer Coding	Inner Coding Rate	Data Rate (Mb/s)	Spectral Efficiency (bit/s/Hz)
0000 (0)	QPSK	(245,255,5)	1/2	5.54	0.79
0001 (1)	QPSK	(245,255,5)	2/3	7.38	1.05
0010 (2)	16-QAM	(245,255,5)	1/2	11.08	1.58
0011 (3)	16-QAM	(245,255,5)	7/12	12.92	1.85
0100 (4)	16-QAM	(245,255,5)	2/3	14.77	2.11
0101 (5)	64-QAM	(245,255,5)	1/2	16.62	2.37
0110 (6)	64-QAM	(245,255,5)	7/12	19.39	2.77
0111 (7)	64-QAM	(245,255,5)	2/3	22.15	3.16
1000 (8)	64-QAM	(245,255,5)	3/4	24.92	3.56
1001 (9)	64-QAM	(245,255,5)	5/6	27.69	3.96
1010 – 1111 (10 – 15) Reserved					

Table C.2 — Mapping of the PHY layer parameters to Rate field. The data rates are derived using channel bandwidth = 8 MHz and with CP duration = 0.875 μ s (1/16 T_{FFT})

RATE R3 – R0	Modulation	Outer Coding	Inner Coding Rate	Data Rate (Mb/s)	Spectral Efficiency (bit/s/Hz)
0000 (0)	QPSK	(245,255,5)	1/2	6.33	0.79
0001 (1)	QPSK	(245,255,5)	2/3	8.44	1.05
0010 (2)	16-QAM	(245,255,5)	1/2	12.66	1.58
0011 (3)	16-QAM	(245,255,5)	7/12	14.77	1.85
0100 (4)	16-QAM	(245,255,5)	2/3	16.88	2.11
0101 (5)	64-QAM	(245,255,5)	1/2	18.99	2.37
0110 (6)	64-QAM	(245,255,5)	7/12	22.15	2.77
0111 (7)	64-QAM	(245,255,5)	2/3	25.32	3.16
1000 (8)	64-QAM	(245,255,5)	3/4	28.48	3.56
1001 (9)	64-QAM	(245,255,5)	5/6	31.65	3.96
1010 – 1111 (10 – 15) Reserved					

Annex D **(normative)**

MAC policies

D.1 Reservation limits

A reservation consists of a row component and a column component.

Row component: A portion of a reservation that includes an equal number of MASs at the same offset(s) within every zone, optionally excluding zone zero, as indicated in the CRP IE(s).

Column component: The portion of the reservation that is not a row component.

Rules stated in this subclause apply independently to a device whether it is a reservation owner or a reservation target.

A device should consider contiguous reservation blocks from multiple column components in the same zone as if they were a single reservation block in a single column component.

A device shall not allocate more channel time than necessary for its optimal operation.

A device shall set the Unsafe bit of the CRP Control field of a CRP IE according to the following rules:

- a) A device shall not identify more than Y consecutive MAS in the same zone within a column component in CRP IEs with the Unsafe bit set to zero, where Y is a function of the MAS number within the zone (counting from zero) of the earliest reserved MAS within the set of consecutive MASs, as shown in Table D.1.
- b) A device shall not set the Unsafe bit to ONE in CRP IEs except as necessary to comply with a).

Table D.1 — Reservation block size limits

First MAS number	Y
0	8
1	7
2	6
3	5
4	4
5	4
6	4
7	4
8	4
9	4
10	4
11	4
12	4
13	3
14	2
15	1

A device may at any time send a Relinquish Request IE in its beacon where the Target DevAddr identifies a device transmitting its beacon with one or more CRP IEs with the Unsafe bit of the CRP IE Control field set to ONE (unsafe CRP IEs). The device shall not set the Target DevAddr field to identify a device if that device does not include any unsafe CRP IEs in its beacon, unless forwarding a received Relinquish Request IE to its reservation owner, as specified in 7.1.8.37.

The Allocation fields of the Relinquish Request IE shall identify MASs in one or more unsafe CRP IEs.

The Reason Code of the Relinquish Request Control field shall be set to a valid Reason Code indicating the reason for requesting the identified MASs.

If a device receives a beacon or command frame that contains a Relinquish Request IE with Target DevAddr set to its own DevAddr that identifies MASs it includes in an unsafe CRP IE, it shall:

- Modify its CRP IEs to remove the identified MASs; or
- Modify its CRP IEs such that the Unsafe bit in any CRP IE that includes one or more identified MASs is set to ZERO per the previous rules in this subclause.

The device shall make this adjustment within mUnsafeReleaseLimit superframes after first receiving the Relinquish Request IE.

If a device requests a neighbour to release MASs in an unsafe CRP IE, the device shall not include a new unsafe CRP IE in its beacon or change a CRP IE to set the Unsafe bit to ONE until mOwnerUnsafeHoldoff has passed.

If a device includes an unsafe CRP IE in its beacon and it receives a Relinquish Request IE that identifies MASs included in the CRP IE, it shall not include a new unsafe CRP IE in its beacon or change a CRP IE to set the Unsafe bit to ONE until the neighbour requesting release establishes a new CRP IE or mTargetUnsafeHoldoff has passed.

D.2 Reservation locations

As referenced in this subclause, a reservation owner shall select MASs based on a minimum latency requirement of not less than 4 milliseconds, or on a medium utilization efficiency or power consumption requirement for a minimum reservation block length.

In the row component of a reservation, the reservation owner shall select reservation blocks such that the lowest MAS number selected within a zone is maximized, except that the reservation owner is not required to use more than one reservation block per zone.

In the column component of a reservation, the reservation owner shall select reservation blocks that meet its requirements such that each block is located within the first eight MASs of its zone, if possible. If not possible, the reservation owner shall select reservation blocks that meet its requirements and that minimize the highest MAS number selected in any zone.

If multiple potential zone locations meet the previous requirements, the reservation owner shall select reservation blocks in zones such that the latest used set in the following ordered list of sets is as early as possible:

[{8}, {4 or 12}, {2, 6, 10, or 14}, {1, 3, 5, 7, 9, 11, 13, or 15}].

If there are multiple possible zone locations that use the same latest set, the reservation owner should minimize the highest MAS number selected within the zones. The reservation owner shall place each reservation block at the earliest available location within its zone.

If the MASs available for reservation by the reservation owner change, it shall determine if its reservations still meet the reservation location rules as listed above. If not, the reservation owner shall change its reservations within mCompactionLimit superframes such that they meet the reservation location rules.

A reservation owner may disregard the reservation location rules for MASs identified in CRP IEs with the Unsafe bit set to ONE according to the Y limit stated in D.1.

D.3 MAC policies parameters

Table D.2 contains the values for the MAC policies parameters.

Table D.2 — MAC policies parameters

Parameter	Value
mCompactionLimit	32 superframes
mOwnerUnsafeHoldoff	32 superframes
mTargetUnsafeHoldoff	2×mMaxLostBeacons
mUnsafeReleaseLimit	4 superframes

Annex E (informative)

FFT-based pilot sensing algorithms

The FFT-based pilot sensing techniques described in this Annex are **non-blind** (ATSC-specific) sensing techniques that meet the sensing sensitivity requirements.

The ATSC VSB signal has a pilot at the lower band-edge in a known location relative to the signal. This description assumes that the signal to be sensed is a band-pass signal at a low-IF of 5.38 MHz with the nominal pilot location at 2.69 MHz and is sampled at 21.52 MHz. However, the basic steps of the sensing algorithm may be implemented with suitable modifications for any IF and sampling rate. The essential features of the method are as follows:

- (1) Demodulate the signal to baseband by the nominal frequency offset of $f_c = 2.69$ MHz. Hence, if $x(t)$ is the real, band-pass signal at low-IF, $y(t) = x(t)e^{-j2\pi f_c t}$ is the complex demodulated signal at baseband.
- (2) Filter $y(t)$ with a low-pass filter of bandwidth, e.g., 40 kHz (± 20 kHz). The filter bandwidth should be large enough to accommodate any unknown frequency offsets.
- (3) Down-sample the filtered signal from 21.52 MHz to 53.8 KHz, to form the signal $z(t)$.
- (4) Take FFT of the down-sampled signal $z(t)$. Depending on the sensing period, the length of the FFT will vary. For example, a 1 ms sensing window will allow a 32-point FFT while a 5 ms window will allow a 256-point FFT.
- (5) Determine the maximum value, and location, of the FFT output squared.

Steps (1) and (2) above are shown in Figure E.1.




Figure E.1 — Frequency Domain Description

Signal detection may then be done either by setting a threshold on the maximum value, or by observing the location of the peak over successive intervals. Instead of the FFT, other well-known spectrum estimation methods, such as the Welch periodogram may also be used in step (4) above.

The basic method described above may be adapted to a variety of scenarios as described below:

- (1) Multiple fine sensing windows, e.g. 5 ms sensing dwells every 10 ms. The 256-point FFT outputs squared from each sensing window may be averaged to form a composite statistic as well as the location information from each measurement may be used to derive a detection metric.

(2) If a single long sensing window, e.g. 10 ms is available, a 512-point FFT or periodogram may be used to obtain better detection performance.

The parameters of the sensor may be chosen depending on the desired sensing time, complexity, probability of missed detection and probability of false alarm. Detection based on location is robust against noise uncertainty since the position of the pilot can be pinpointed with accuracy even if the amplitude is low due to fading. Various combining schemes may be developed for both pilot-energy and pilot-location sensing.

(1) Pilot-energy sensing: For a single sensing window, the FFT output is simply squared and the maximum value is compared to a threshold. For multiple sensing dwells, there are 2 possibilities: (i) the decision from each dwell is saved and a “hard-decision” rule is applied to declare “signal detect” if the number of positives is greater than a certain number, or (ii) the square of the FFT output of all dwells is averaged and the maximum level is compared to a threshold. The choice of threshold in all cases is determined by the desired P_{FA} .

(2) Pilot-location sensing: This is usually used for multiple dwells.

- Let no. of dwells = N .
- Let $f_{\max}^{(1)}$ be the location of the maximum of the FFT-output averaged over the first $N/2$ dwells.
- Let $f_{\max}^{(2)}$ be the location of the maximum of the FFT-output averaged over the second $N/2$ dwells.
- Detection statistic:

$$D = |f_{\max}^{(1)} - f_{\max}^{(2)}|$$

- If $D < N_T$, signal present.
- Other variations:
 - Any PSD algorithm may be used instead of FFT.
 - Other averaging intervals could be used.
- P_{FA} using this method is extremely low and robust. Threshold value N_T will depend on the FFT-size.

E.1 Performance of the algorithms

Both pilot-energy and pilot-location based sensing algorithms were tested with the 12 DTV signals specified. The sensing time was multiples of 5 ms, which allowed the use of a 256-point FFT. Table E.1 shows the required SNR for a $P_{FA} = 0.05$ and $P_{MD} = 0.10$ and no noise uncertainty.

Table E.1 — Required SNR for DTV signal detection (Average over 12 signals)

Method	5 ms (N = 1)	10 ms (N = 2)	30 ms (N = 6)	50 ms (N = 10)
Pilot-Energy	-18 dB	-20.5 dB	-23.5 dB	-24.5 dB
Pilot-Location ($N_T = 2$)	-	-18.5 dB	-22.0 dB	-24.0 dB

Annex F (informative)

An example of TPC algorithm

As an example, the initial transmit power and power update may be determined considering the status of the first adjacent channel and adjacent cell, as well as the modulation scheme. The initial transmit power is composed into two parts as shown in Equation F.1. One of the two parts is the initial transmit power considering the neighbouring signals, $P_{i,NS}$, the other part is the initial transmit power considering the modulation scheme, $P_{i,MS}$.

$$P_i = P_{i,NS} + P_{i,MS} \quad (\text{F.1})$$

First, $P_{i,NS}$ is determined from the status of the first adjacent channel and adjacent cell. Based on the information in the database or the results of spectrum sensing, $P_{i,NS}$ is determined as defined in Table F.1. If no signal is detected in the first adjacent channel and adjacent cell, $P_{i,NS}$ is allowed up to a level $P_{i,MS,64QAM}$ below the maximum transmit power of 100 mW. When the incumbent signals or other unlicensed signals appear in the first adjacent channel for the operating cell and adjacent cell for the operating channel, the first component of initial transmit power, $P_{i,NS}$, is reduced step by step. Where the differential decrement of $P_{i,NS}$ compared to the maximum power level is implementation-dependent. In Table F.1, the values in parentheses correspond to the values for sensing-only devices. It is same in the following Table F.2 through Table F.4. Since the cell radius of TV stations is much larger than that of portable/personable unlicensed devices, and the sensing range of wireless microphone signals is larger than that of data transmission, the incumbent signals in the adjacent cell for the operating channel may be detected using database access and/or spectrum sensing. In these cases, since the operating channel will be switched to another channel, the transmit power control mechanisms does not need to consider the incumbent signals in the adjacent cell for the operating channel.

Table F.1 — Decision of initial transmit power, $P_{i,NS}$, considering the neighbouring signals

The status of the first adjacent channel The status of the adjacent cell		Signal type in the first adjacent channel for the operating cell		
		No signal	Other unlicensed or unknown signals	Incumbent signals (TV, WM, etc.)
Signal type in the adjacent cell for the operating channel	No signal	$P_{i,NS1} = 100 - P_{i,MS,64QAM}$ mW ($50 - P_{i,MS,64QAM}$ mW)	$P_{i,NS2} = P_{i,NS1} \geq P_{i,NS} \geq P_{i,NS3}$ mW ($P_{i,NS1} \geq P_{i,NS} \geq P_{i,NS3}$ mW)	$P_{i,NS3} = 40 - P_{i,MS,64QAM}$ mW ($40 - P_{i,MS,64QAM}$ mW)
	Other unlicensed or unknown signals	$P_{i,NS4} = P_{i,NS1} \geq P_{i,NS} \geq P_{i,NS3}$ mW ($P_{i,NS1} \geq P_{i,NS} \geq P_{i,NS3}$ mW)	$P_{i,NS2} \geq P_{i,NS}, P_{i,NS4} \geq P_{i,NS} \geq P_{i,NS5}$ mW ($P_{i,NS2} \geq P_{i,NS}, P_{i,NS4} \geq P_{i,NS} \geq P_{i,NS5}$ mW)	$P_{i,NS5} = P_{i,NS3} \geq P_{i,NS}$ mW ($P_{i,NS3} \geq P_{i,NS}$ mW)

And second, $P_{i,MS}$ is determined considering the modulation scheme, as defined in Table F.2. The modulation scheme is classified as: 1) 64QAM, 2) 16QAM, 3) QPSK. The differential decrement of $P_{i,MS}$ compared to the maximum power level is implementation-dependent. The ratio of $P_{i,NS}$ to $P_{i,MS}$ is also implementation-dependent.

Table F.2 — Decision of initial transmit power, $P_{i,MS}$, considering the modulation scheme

Modulation scheme	Initial transmit power
64QAM	$P_{i,MS,64QAM} =$ $100 \geq P_{i,MS} \text{ mW}$ $(50 \geq P_{i,MS} \text{ mW})$
16QAM	$P_{i,MS,16QAM} =$ $P_{i,MS,64QAM} \geq P_{i,MS} \text{ mW}$ $(P_{i,MS,64QAM} \geq P_{i,MS} \text{ mW})$
QPSK	$P_{i,MS,QPSK} =$ $P_{i,MS,16QAM} \geq P_{i,MS} \text{ mW}$ $(P_{i,MS,16QAM} \geq P_{i,MS} \text{ mW})$

According to the Equation F.1 and the values from Table F.1 and Table F.2, when the incumbent users present in the first adjacent channel for the operating cell, the initial transmit power does not exceed 40 mW. If no signal is detected in the first adjacent channel and adjacent cell, the initial transmit power shall be allowed up to 100 mW. For the sensing-only devices, the initial transmit power does not exceed 50 mW.

The transmit power is updated using Equation F.2. In the same manner as the decision of initial transmit power, the power update is also composed into two parts. One of the two parts is the power update considering the neighbouring signals, ΔP_{NS} , the other part is the power update considering the modulation scheme, ΔP_{MS} .

$$\Delta P = \Delta P_{NS} + \Delta P_{MS} \quad (\text{F.2})$$

First, ΔP_{NS} is also determined from the status of the first adjacent channel and adjacent cell. Based on the information in the database or the results of spectrum sensing, ΔP_{NS} is determined as defined in Table F.3. The maximum adjustment step of power update is no more than 2 dB. For the sensing-only devices, the maximum adjustment step is reduced to 1 dB, a half of the step size compared to the value for the database-accessible devices. If no signal is detected in the first adjacent channel and adjacent cell, ΔP_{NS} is allowed up to a level $\Delta P_{MS,64QAM}$ below the maximum adjustment step of 2 dB. When the incumbent signals or other unlicensed signals appear in the first adjacent channel for the operating cell and adjacent cell for the operating channel, the first component of power update, ΔP_{NS} , is reduced step by step. The differential decrement of ΔP_{NS} compared to the maximum adjustment step is implementation-dependent.

Table F.3 — Decision of power update, ΔP_{NS} , considering the neighbouring signals

The status of the first adjacent channel The status of the adjacent cell		Signal type in the first adjacent channel for the operating cell		
		No signal	Other unlicensed or unknown signals	Incumbent signals (TV, WM, etc.)
Signal type in the adjacent cell for the operating channel	No signal	$\Delta P_{NS1} = \pm 2 - \Delta P_{MS,64QAM} \text{ dB}$ ($\pm 1 - \Delta P_{MS,64QAM} \text{ dB}$)	$\Delta P_{NS2} = \pm \Delta P_{NS1} \geq \Delta P_{NS} \geq \Delta P_{NS3} \text{ dB}$ ($\pm \Delta P_{NS1} \geq \Delta P_{NS} \geq \Delta P_{NS3} \text{ dB}$)	$\Delta P_{NS3} = \pm \Delta P_{NS2} \geq \Delta P_{NS} \text{ dB}$ ($\pm \Delta P_{NS2} \geq \Delta P_{NS} \text{ dB}$)
	Other unlicensed or unknown signals	$\Delta P_{NS4} = \pm \Delta P_{NS1} \geq \Delta P_{NS} \geq \Delta P_{NS3} \text{ dB}$ ($\pm \Delta P_{NS1} \geq \Delta P_{NS} \geq \Delta P_{NS3} \text{ dB}$)	$\pm \Delta P_{NS2} \geq \Delta P_{NS}, \Delta P_{NS4} \geq \Delta P_{NS} \geq \Delta P_{NS5} \text{ dB}$ ($\pm \Delta P_{NS2} \geq \Delta P_{NS}, \Delta P_{NS4} \geq \Delta P_{NS} \geq \Delta P_{NS5} \text{ dB}$)	$\Delta P_{NS5} = \pm \Delta P_{NS3} \geq \Delta P_{NS} \text{ dB}$ ($\pm \Delta P_{NS3} \geq \Delta P_{NS} \text{ dB}$)

And second, ΔP_{MS} is determined considering the modulation scheme, as defined in Table F.4. The higher the modulation order, the larger the step of power update. The differential decrement of ΔP_{MS} compared to the maximum adjustment step is implementation-dependent. The ratio of ΔP_{NS} to ΔP_{MS} is also implementation-dependent.

Table F.4 — Decision of power update, ΔP_{MS} , considering the modulation scheme

Modulation scheme	Power update
64QAM	$\Delta P_{MS,64QAM} = \pm 2 \geq \Delta P_{MS} \text{ dB}$ ($\pm 1 \geq \Delta P_{MS} \text{ dB}$)
16QAM	$\Delta P_{MS,16QAM} = \pm \Delta P_{MS,64QAM} \geq \Delta P_{MS} \text{ dB}$ ($\pm \Delta P_{MS,64QAM} \geq \Delta P_{MS} \text{ dB}$)
QPSK	$\Delta P_{MS,QPSK} = \pm \Delta P_{MS,16QAM} \geq \Delta P_{MS} \text{ dB}$ ($\pm \Delta P_{MS,16QAM} \geq \Delta P_{MS} \text{ dB}$)

According to the Equation F.2 and the values from Table F.3 and Table F.4, when the incumbent signals or other unlicensed signals appear in the first adjacent channel for the operating cell and adjacent cell for the operating channel, the power update should have a minimum step. If no signal is detected in the first adjacent channel and adjacent cell, the power update is allowed up to $\pm 2 \text{ dB}$ and $\pm 1 \text{ dB}$, for the database-accessible devices and sensing-only devices, respectively.

Bibliography

- [1] "Guidelines for use of a 48-bit Extended Unique Identifier (EUI-48™)",
<http://standards.ieee.org/regauth/oui/tutorials/EUI48.html>
- [2] IEEE STD 802.3 -2005 IEEE Standard for Information technology - Telecommunications and information exchange between systems- Local and metropolitan area networks- Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.
- [3] High Rate Ultra Wideband PHY and MAC Standard, ECMA-368.
- [4] Assigned numbers for Specifier ID in Application-specific control frames, command frames, IEs and Probe IEs. http://www.ecma-international.org/publications/standards/UWB_specifier_id.htm
- [5] ISO/IEC 7498-1:1994, Information Technology-Open Systems Interconnection-Basic Reference Model: The Basic Model.

