INTERNATIONAL STANDARD

ISO/IEC 13156

Second edition
2011-10-01

Information technology — Telecommunications and information exchange between systems — High rate 60 GHz PHY, MAC and PALs

Technologies de l'information — Téléinformatique — PHY, MAC et PALs 60 GHz à haut débit

Reference number
ISO/IEC 13156:2011(E)

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

Page

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 2.

The main task of the joint technical committee is to prepare International Standards. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 13156 was prepared by Ecma International (as ECMA-387) and was adopted, under a special "fast-track procedure", by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, in parallel with its approval by national bodies of ISO and IEC.

This second edition cancels and replaces the first edition (ISO/IEC 13156:2009), which has been technically revised.

Introduction

This International Standard specifies PHY, MAC and PALs for flexible and heterogeneous multi-Gigabit Wireless Personal Area networks. The heterogeneous network consists of two types of devices (Types A and B) that can fully coexist and interoperate but at the same time are able to operate independently. As a result this standard enables a wide range of different implementations and applications ranging from simple and low-power data transfer at short ranges, suitable for handheld devices, to high-rate multimedia streaming at longer distances, when adaptive antenna arrays are employed. Applications include Sync-and-Go, Access points, Wireless desktops and docking stations and uncompressed video streaming.

The Type A device is designed to be the high-end, high-performance device and provides many features including high data rate, longer range, robustness against multipath, support for adaptive antenna arrays and multi-level QoS. On the other hand, Type B devices, designed for handheld devices, are simpler, low power and low cost, while offering high data rates.

Type A and Type B devices offer data rates up to 6,350 Gbps and 3,175 Gbps in a single channel, respectively. This International Standard defines four frequency channels with separation of 2,160 GHz, which may be bonded to each other to increase the data rates by a factor of 2, 3 or 4.

This International Standard defines a single decentralized MAC protocol for both device types, which provides interoperability and coexistence for the device types and features high bandwidth efficiency, QoS provisions, and spatial reuse capability (Figure 1).

Multiple PALs can reside on top of the MAC layer, which interact with the MAC layer through a multiplexing sublayer (MUX). This edition of ISO/IEC 13156 provides an HDMI[1] PAL as well as information regarding IP and USB PALs.

1. HDMI is the registered trademark of the HDMI LLC.

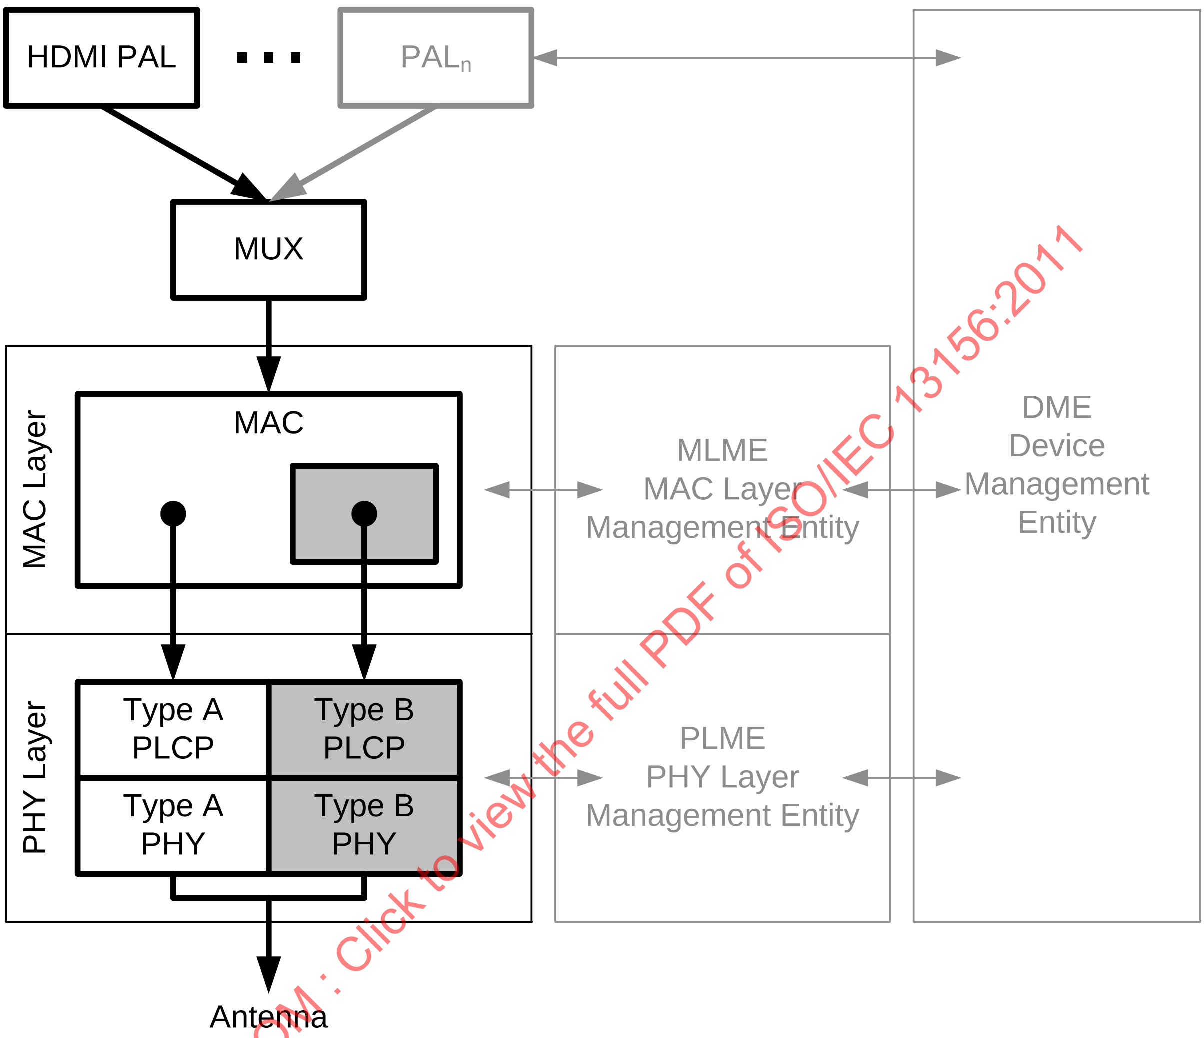

Figure 1 - Protocol structure

NOTE The DME, MLME, PLME, and PALs (except the HDMI PAL) are outside the scope of this International Standard and all references to these are informative.

Information technology — Telecommunications and information exchange between systems — High rate 60 GHz PHY, MAC and PALs

1 Scope

This International Standard specifies a physical layer (PHY), distributed medium access control (MAC) sublayer, and an HDMI protocol adaptation layer (PAL) for 60 GHz wireless networks.

2 Conformance

Conforming devices of Type A, B or C shall implement the MAC sublayer and the PHY layer and may implement the HDMI PAL as specified herein.

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 8802-11:2005/Amd.6, *Information technology — Telecommunications and information exchange between systems — Local and metropolitan area networks — Specific requirements — Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications — Amendment 6: Medium Access Control (MAC) Security Enhancements*

IEEE 100, *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition

4 Terms and definitions

For the purposes of this document, the following terms and definitions apply. For terms and definitions not defined in this Clause, the term and definitions given in IEEE 100, *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition apply.

4.1
Beacon Group
BG
set of devices from which a device receives beacons that identify the same beacon period start time (BPST) as the device

4.2
Beacon Period
BP
period of time declared by a device during which it sends or listens for beacons

4.3
Beacon Period Start Time
BPST
start of the beacon period

4.4
channel
medium over which cooperating entities exchange information

4.5
data integrity
assurance that the data has not been modified from its original form

4.6
device
entity containing an implementation of this International Standard

4.7
Distributed Reservation Protocol
DRP
protocol implemented in each device to support negotiation and maintenance of channel time reservations binding on all neighbours of the reservation participants

4.8
Equivalent Isotropic Radiated Power
EIRP
amount of power that a theoretical isotropic antenna (that evenly distributes power in all directions) would emit to produce the peak power density observed in the direction of maximum antenna gain

4.9
Equivalent Isotropic Received Power
EIRxP
amount of power that a theoretical isotropic antenna (that evenly receives power in all directions) would receive

4.10
extended beacon group
union of a device's beacon group and the beacon groups of all devices in the device's beacon group

4.11
frame
unit of data transmitted by a device

4.12
frame protection
security service provided for a frame, including (but not limited to) payload encryption, message authentication, and replay attack protection

4.13
MAC client
entity above the MAC sublayer that generates MAC service data units for delivery to corresponding entities in other devices, and receives MAC service data units from such entities

4.14
MAC Command Data Unit
MCDU
unit of data exchanged between peer medium access control sublayers in order to manage medium access control functions

4.15
MAC Protocol Data Unit
MPDU
unit of data exchanged between two peer medium access control sublayers using the physical layer

4.16
MAC Service Data Unit
MSDU
information that is delivered as a unit between medium access control service access points (SAPs)

4.17
Master-Slave Pair
MSPr
device-to-device link in which a first device acts as the master (initiates polling) and a second device acts as a slave (responds to a polling inquiry)

4.18
Message Integrity Code
MIC
cryptographic checksum generated using a symmetric key that is typically appended to data in order to provide data integrity and source authentication similar to a digital signature

4.19
neighbour
any device in a device's beacon group

4.20
reservation
named set of one or more medium access slots (MASs) within a superframe during which a device has preferential access to the medium

4.21
reservation block
one or more temporally contiguous medium access slots (MASs) within a reservation not adjacent to other MASs in the reservation

4.22
secure frame
frame in which frame protection is applied

4.23
stream
logical flow of MSDUs from one device to one or more other devices

4.24
superframe
periodic time interval used in this International Standard to coordinate frame transmissions between devices, which contains a beacon period followed by a data period

4.25
symmetric key
secret key shared between two or more parties that may be used for both encryption and decryption as well as for message integrity code computation and verification

4.26
Time Domain Spreading Factor
TDSF
bandwidth expansion ratio due to the application of time domain spreading sequence

5 Notational conventions

The use of the word *shall* is meant to indicate a requirement which is mandated by the Standard, i.e. it is required to implement that particular feature with no deviation in order to conform to the Standard.

The use of the word *should* is meant to recommend one particular course of action over several other possibilities, however without mentioning or excluding these others.

The use of the word *may* is meant to indicate that a particular course of action is permitted.

The use of the word *can* is synonymous with is able to – it is meant to indicate a capability or a possibility.

All floating-point values have been rounded to 3 decimal places.

6 Abbreviated terms

AC	Access Category
ACK	Acknowledgment
AES	Advanced Encryption Standard
ASK	Amplitude Shift Keying
ASIE	Application-Specific Information Element
AWGN	Additive White Gaussian Noise
B-ACK	Block Acknowledgment
BcstAddr	Broadcast Device Address
BP	Beacon Period
BPOIE	Beacon Period Occupancy Information Element
BPSK	Binary Phase Shift Keying
BPST	Beacon Period Start Time
CC	Convolutional Code
CCA	Clear Channel Assessment
CRC	Cyclic Redundancy Check
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance
CTT	Clear To Train
DAC	Digital-to-Analogue Converter
DBPSK	Differential Binary Phase Shift Keying
DBS	Discovery Block Set
DCA	Distributed Contention Access
DestAddr	Destination Device Address
DevAddr	Device Address
DME	Device Management Entity
DQPSK	Differential Quadrature Phase Shift Keying
DRP	Distributed Reservation Protocol
EO	Encryption Offset
EUI	Extended Unique Identifier
FEC	Forward Error Correction
FER	Frame Error Rate
FFT	Fast Fourier Transform
FZ	Frank-Zadoff Sequence
GCM	Galois/Counter Mode

GF	Galois Field
Gbps	Gigabits per second
Gsps	Gigasymbols per second
GTK	Group Temporal Key
HCS	Header Check Sequence
ID	Identifier
IDFT	Inverse Discrete Fourier Transform
IE	Information Element
IFFT	Inverse Fast Fourier Transform
IFS	Inter-Frame Space
Imm-ACK	Immediate Acknowledgment
KCK	Key Confirmation Key
LIFS	Long Inter-frame space
LQE	Link Quality Estimator
LQI	Link Quality Indicator
LSB	Least-Significant Bit
MAC	Medium Access Control
MAS	Medium Access Slot
MCDU	MAC Command Data Unit
Mbps	Megabits per second
McstAddr	Multicast Device Address
MIB	Management Information Base
MIC	Message Integrity Code
MIFS	Minimum Interframe Spacing
MKID	Master Key Identifier
MLME	MAC Sublayer Management Entity
MPDU	MAC Protocol Data Unit
MSB	Most-Significant Bit
MSC	Message Sequence Chart
MSPr	Master-Slave Pair
MSDU	MAC Service Data Unit
No-ACK	No Acknowledgement
OFDM	Orthogonal Frequency Division Modulation
OOB	Out of Band
OOK	On-Off Keying
OUI	Organizationally Unique Identifier
PAA	Phased Array Antenna
PAN	Personal Area Network

PAL	Protocol Adaptation Layer
PDU	Protocol Data Unit
PER	Packet Error Rate
PHY	Physical (layer)
PHY-SAP	Physical Layer Service Access Point
PLCP	Physical Layer Convergence Protocol
PLME	Physical Layer Management Entity
PMK	Pair-wise Master Key
PPDU	PLCP Protocol Data Unit
PPM	Parts Per Million
PRBS	Pseudo-Random Binary Sequence
PRF	Pseudo-Random Function
PSD	Power Spectral Density
PSDU	PHY Service Data Unit
PT	Preamble Type
PTK	Pair-wise Temporal Key
QPSK	Quadrature Phase Shift Keying
RF	Radio Frequency
RMS	Root Mean Square
RS	Reed-Solomon
RTT	Request To Train
RX	Receive or Receiver
SAP	Service Access Point
SBA	Switched Beam Antenna
SCS	Segment Check Sequence
SFC	Secure Frame Counter
SFN	Secure Frame Number
SIFS	Short Interframe Spacing
SrcAddr	Source Device Address
TCM	Trellis Coded Modulation
TDSF	Time Domain Spreading Factor
TKID	Temporal Key Identifier
TX	Transmit or Transmitter
TXOP	Transmission Opportunity
WPAN	Wireless Personal Area Network

7 General description (informative)

7.1 PHY general description

This International Standard specifies a physical layer (PHY) for a Wireless Personal Area Network (WPAN), utilizing the unlicensed 60 GHz frequency band. Two types of devices are defined: Type A and Type B. both device types coexist and interoperate with other device types. Furthermore, all device types can operate independently. That is, neither device type requires the presence of another type for operation.

Type A devices operate at an SCBT mandatory mode (A0) at 0.397 Gbps with other optional SCBT modes at data rates 0.794 Gbps to 6.350 Gbps (without channel bonding) and optional OFDM modes at data rates 1.008 Gbps to 4.032 Gbps. Type B devices operate using DBPSK at data rates of 0.794 Gbps to 1.588 Gbps (without channel bonding); with optional modes of DQPSK and UEP-QPSK at data rates of 3.175 Gbps.

Type A devices also support directional antennas via sector antennas or adaptive arrays. This International Standard specifies the necessary training and tracking waveforms and protocols.

There are multiple channels specified in this International Standard. Multiple adjacent channels may be bonded together for increased data rate for Type A (SCBT) and Type B devices. With bonded channels the data rates for these device types increase by a factor proportional to the number of bonded channels. One channel has been designated as the discovery channel. The discovery channel has a lower quality of service but supports the interference prone and time consuming antenna training process. The data channels are optimized for high throughput and spatial reuse.

7.2 MAC general description

7.2.1 General description of the architecture

As illustrated in Figure 2, the MAC is a sublayer of the Data Link Layer defined in the OSI basic reference model [1]. The MAC service is provided by means of the MAC service access point (MAC SAP) to a single MAC service client, usually a higher layer protocol or adaptation layer. In this International Standard the MAC sublayer is represented by a device address.

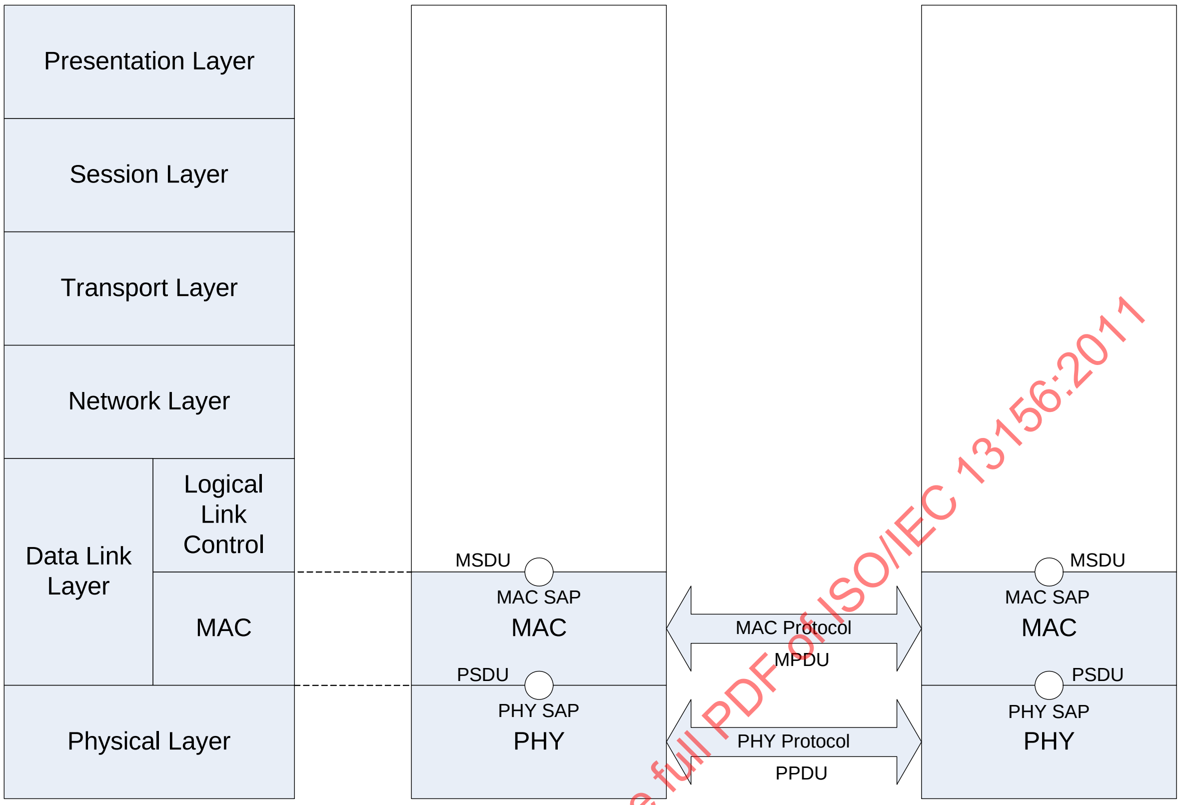

Figure 2 - Architectural reference model

The MAC sublayer in turn relies on the service provided by the PHY layer via the PHY service access point (PHY SAP). The MAC protocol applies between MAC sublayer peers.

7.2.2 Device address

Individual MAC sublayers are addressed via an EUI-48 [3], and are associated with a volatile abbreviated address called a DevAddr. Unicast frames carry a destination DevAddr that identifies a single MAC sublayer.

DevAddrs are 16-bit values, generated locally, without central coordination. Consequently, it is possible for a single value to ambiguously identify two or more MAC entities. This International 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.

7.2.3 Features assumed from the PHY

A MAC entity is associated with a single PHY entity via that entity's PHY SAP. The MAC sublayer requires the following features provided by the PHY:

- Frame transmission and reception;
- PLCP header error indication for both PHY and MAC header structures;
- Clear channel assessment for estimation of medium activity.

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 are included in the PLCP header and allow the MAC entity to control, and be informed of, the MAC and PHY related parameters.

Depending on antennas used by the PHY, the MAC sublayer may use the following features provided by the PHY for directional frame transmission and reception:

- Antenna beam switching among different sectors
- Antenna beam steering toward a desired direction

7.2.4 Overview of MAC service functionality

The MAC service defined in this International Standard provides:

- A reservation-based channel access mechanism;
- A contention-based channel access mechanism for antenna training in the discovery channel;
- A synchronization facility among cooperating MAC entities;
- Coexistence and interoperability among Type A and B devices;
- Device power management by scheduling of frame transmission and reception;
- Secure communication with data authentication and encryption using cryptographic algorithms;

Each device provides required MAC functions based on the device type, and optional functions as determined by the application.

Coordination of devices within radio range is achieved by the transmission and reception of beacon and control frames. Periodic beacon transmission provides the basic timing for the network, supports dynamic network organization, and carries reservation and scheduling information for accessing the medium. Exchange of control frames enables antenna training among cooperating devices, and device discovery of Type A and B devices.

Coordination among devices that send periodic beacon frames (referred to as beaconing devices) is fully distributed. Coordination among beaconing devices and devices that do not send beacon frames is achieved by the beaconing devices acting as controllers. Coordination among devices that do not send beacon frames is not specified in this International Standard.

7.2.4.1 Logical groups

Logical groups are formed around each beaconing device to facilitate contention-free frame exchanges while exploring medium reuse over different spatial regions. In this International Standard, these logical groups are a beacon group and an extended beacon group. Both groups are determined with respect to an individual beaconing device, which has its own individual neighbourhood.

7.2.4.2 Device discovery

The MAC sublayer defined in this International Standard enables device discovery through one or more of the following mechanisms:

- Transmission and reception of discovery frames in the discovery channel;
- Exchange of antenna training control frames in the discovery channel;
- Exchange of interoperability control frames between Type A and Type B devices.

7.2.4.3 Channel selection

Once a device discovers another device with which it intends to communicate with, the pair of the devices use explicit channel selection procedure, as described in 16.4.1 to scan one or more channels and to select a channel for frame exchange in a coordinated manner.

If no beacons are detected in the selected channel, the 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.

Each device operates in a dynamic environment and under unlicensed operation rules. Thus, it is subject to interference from 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 without requiring disruption of links with its peers.

If at any time a device determines that the current channel is unsuitable, it uses the implicit channel selection procedure, as described in 16.4.2, to move to a new channel.

7.2.4.4 The superframe

Once a device finds its communication partner and selects a channel, 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).

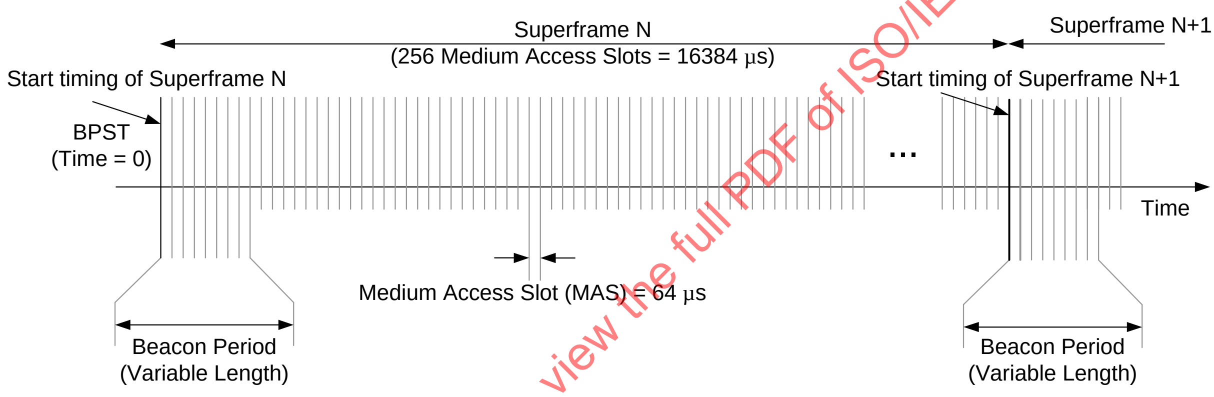

Figure 3 - MAC superframe structure

7.2.4.5 Beacon period protection

No transmissions other than beacons are attempted during the BP of any device.

A device may 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.

7.2.4.6 Medium access

The medium is accessed in one of three ways:

- During device discovery (16.3) and antenna training, devices send beacon and control frames in discovery channel using contention based access according to the rules specified in 16.2.
- During the BP, devices send only beacon frames, according to the rules specified in 16.5.
- During reservations, devices participating in the reservation send frames according to rules specified in 16.6.

7.2.4.7 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 16.9. 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 16.12. This International Standard provides for three types of acknowledgements to enable different applications. The No-ACK policy, described in 16.12.1, 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 16.12.2, provides an acknowledgement process in which each frame is individually acknowledged following the reception of the frame. The B-ACK policy, described in 16.12.3, 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 16.1.5, 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.

7.2.4.8 MAC frame data rates

This International Standard specifies two common PHY modes for the two types of devices respectively. The frame payloads of MAC beacon frames are transmitted using one of the common PHY modes corresponding to its device type, and hence at the rate of the common mode. In device discovery or antenna training, MAC frames are transmitted using one of the discovery modes in the discovery channel. Payloads of other frames may be transmitted at higher data rates if possible.

7.2.4.9 Security

Wireless networks present unique security challenges due to the loss of protection provided by wires and shielding. Distributed 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 International Standard defines two levels of security (Clause 17): 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 International 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 (see 17.2).

This International 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 (see 17.3.1). 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 International Standard.

In addition, this International 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 (see 17.3.2).

A pseudo-random function is defined based on the MIC generation by GCM using AES-128 (see 17.3.3). It can be made available to entities outside the MAC sublayer for random number generation.

Secure frame counters and replay counters are set up on a per-temporal key basis to guarantee message freshness (see 17.4). No specific mechanisms are created in this International Standard to address denial of service attacks given the open nature of the wireless medium.

In this International Standard, 128-bit symmetric temporal keys are employed based on AES-128 with GCM to provide payload encryption and message integrity code (MIC) generation (see 17.5).

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

7.2.4.10 Information discovery

The protocols and facilities of this International Standard are supported by the exchange of information between devices. Information can be broadcast in beacon frames or requested in Probe commands. For each type of information, an Information Element (IE) is defined. IEs can 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.

7.2.4.11 Transmit rate and power adaptation

This International Standard provides transmit rate and power control mechanisms to select the optimal combination of transmit rate and power to increase throughput and/or reduce the frame error rate (FER).

A recipient device may recommend a transmit rate and/or power level change be used by a source device using explicit or implicit link feedback mechanisms. In addition, a source device may request a recipient device provide feedback on the quality of the link, base on which the source determines optimal transmit rate and power.

The transmit power and rate control mechanisms are described in 16.14 and 16.15.

7.2.4.12 Power management

An important goal of this International 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 International 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 International Standard provides facilities to support devices that sleep for portions of each superframe in order to save power.

Power management mechanisms are described in 16.16.

7.2.5 MAC policies

It is desirable to allow and facilitate equitable and efficient coexistence of devices with varying medium access requirements. For this purpose, Annex B specifies policies governing channel selection and sharing of bandwidth. These policies impose, among other things, certain restrictions on the number and configuration of MASs in DRP reservations, on the location of reserved MASs within a superframe, and on channel selection order.

7.2.6 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 16.8 may be an additional requirement. In support of such applications, this International Standard defines an optional MAC facility in Annex C that enables layers above the MAC sublayer to accurately synchronize timers located in different devices. The facility is usable by more than one application at a time.

7.3 MUX general description

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

7.4 HDMI PAL description

This International Standard includes an HDMI pass-through PAL, Clause 18, which preserves the HDMI content protection scheme. The HDMI PAL interfaces with the wired HDMI interface's data channels, clock channel, display data channel and CE control. The HDMI PAL removes the data channels' TMDS encoding prior to wireless transmission and reinstates the TMDS coding prior to forwarding to the HDMI sink.

8 PHY layer (informative)

The PHY contains two functional entities: the PHY convergence function and the PHY layer management function. The PHY service is provided to the MAC through the PHY SAP.

9 Description of signal

9.1 Mathematical framework for SCBT, OFDM, DBPSK, DQPSK, UEP-QPSK, OOK and 4ASK

The transmitted RF signal can be written in terms of the complex baseband signal as follows:

$$s_{RF}(t) = \mathrm{Re}\left\{\sum_{n=0}^{N_{frame}-1} s_n\left(t - nT_{sym}\right)\exp\left(j2\pi f_c t\right)\right\} \quad (1)$$

where Re{.} represents the real part of the signal, T_{sym} is the symbol length, N_{frame} is the number of symbols in the frame, f_c is the centre frequency, and $s_n(t)$ is the complex baseband signal representation for the n^{th} symbol. The exact structure of the n^{th} symbol depends on its location within the frame:

$$s_n(t) = \begin{cases} s_{preamble,n}(t) & 0 \le n < N_{preamble} \\ s_{header,n-N_{preamble}}(t) & N_{preamble} \le n < N_{preamble} + N_{header} \\ s_{payload,n-N_{preamble}-N_{header}}(t) & N_{preamble} + N_{header} \le n < N_{preamble} + N_{header} + N_{payload} \\ s_{ATS,n-N_{preamble}-N_{header}-N_{payload}}(t) & N_{preamble} + N_{header} + N_{payload} \le n < N_{preamble} + N_{header} + N_{payload} + N_{ATS} \end{cases} \quad (2)$$

where $s_{preamble,n}(t)$ describes the n^{th} symbol of the preamble, $s_{header,n}(t)$ describes the n^{th} symbol of the header, $s_{payload}(t)$ describes the n^{th} symbol of the PPDU, $s_{ATS,n}(t)$ describes the n^{th} symbol of the antenna training sequence (ATS), $N_{preamble}$ is the number of symbols in the preamble, N_{header} is the number of symbols contained in the header, $N_{payload}$ is the number of symbols contained in the frame body, N_{ATS} is the number of symbols in the ATS, and $N_{frame} = N_{preamble} + N_{header} + N_{payload} + N_{ATS}$ is the number of symbols in the frame. The exact values of $N_{preamble}$, N_{header}, $N_{payload}$, N_{ATS}, and N_{frame} will be described in more detail in Clause 10.

The potentially complex time-domain signal $s_n(t)$ shall be created by passing the real and imaginary components of the discrete-time signal $s_n[k]$ through digital-to-analogue converters (DACs) and reconstruction filters as shown in Figure 4. When the discrete-time signal $s_n[k]$ is real, only the real digital-to-analogue converter and reconstruction filter need to be used. Clause 10 describes how to generate $s_n[k]$.

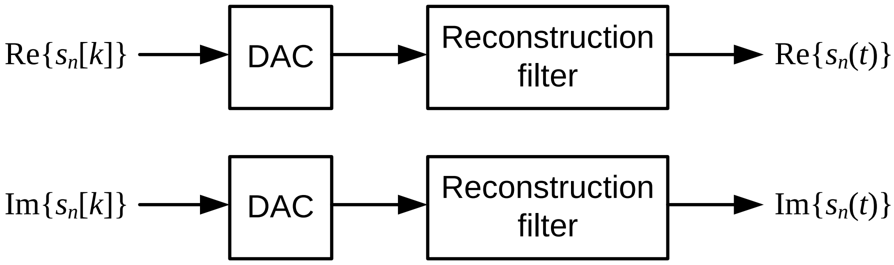

Figure 4 - Conversion from discrete-time signals to continuous-time signals

9.2 Mathematical framework for the narrow band section of the discovery mode preamble

The transmitted RF signal can be written in terms of the complex baseband signal as follows:

$$s_{RF}(t) = \mathrm{Re}\left\{ \sum_{n=0}^{N_{NB}-1} s_{NB,n}(t - nT_{sym}) \left[\exp(j2\pi f_c t) + \exp(j2\pi[f_c - f_0]t) + \exp(j2\pi[f_c + f_0]t) \right] \right\} \quad (3)$$

where $\mathrm{Re}\{.\}$ represents the real part of the signal, T_{sym} is the symbol length, $N_{NB} = 163839$ is the number of symbols in the narrow band preamble, f_c is the centre frequency, f_0=720 MHz is the offset frequency and $s_{NB,n}(t)$ describes the n^{th} symbol of the narrow band preamble.

The complex time-domain signal $s_{NB,n}(t)$ shall be created by passing the real and imaginary components of the discrete-time signal $S_{NB}[k]$ through digital-to-analogue converters (DACs) and reconstruction filters as shown in Figure 4. The sequence $S_{NB}[k]$ is described in Clause 10.

10 PLCP sublayer

10.1 General PPDU frame format

This Clause provides a method for converting a PSDU (PHY layer SDU) into a PPDU (PHY layer PDU). Figure 5 shows the general format for the PPDU, which may be composed of four major components: the PLCP preamble, the PLCP header, the PPDU payload, and the Antenna Training Sequence (ATS).

All device types shall follow this general frame format. However, particular parameters or components of the frame format may be different for different device types.

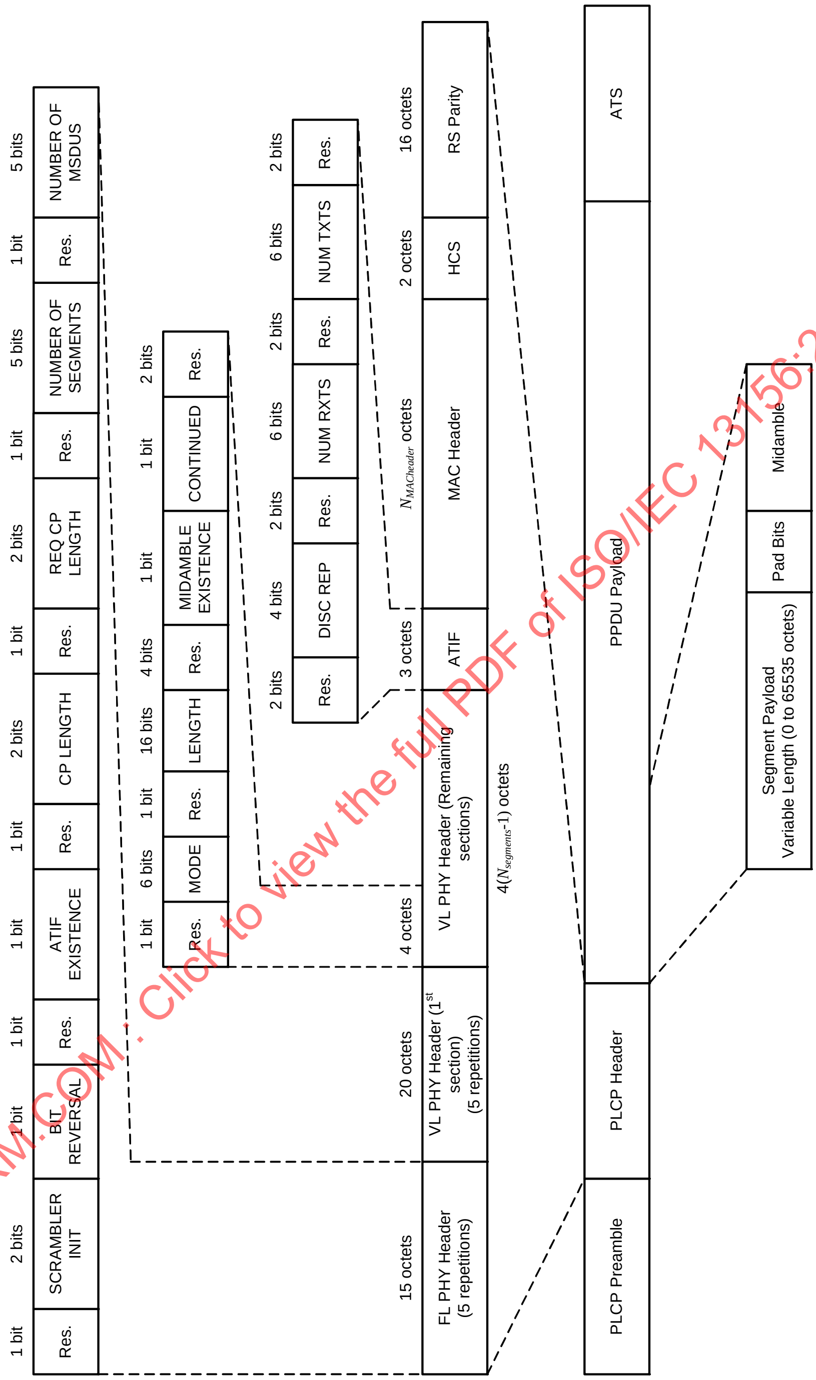

Figure 5 - General PPDU frame format

The components are listed in the order of transmission. The PLCP preamble shall be the first component of the PPDU. It consists of a frame synchronization sequence, and a channel estimation sequence (see 10.1.1). The goal of the PLCP preamble is to aid the receiver in timing synchronization, carrier offset recovery, and channel estimation.

The PLCP header shall be the second major component of the PPDU. The goal of this component is to convey necessary information about both the PHY and the MAC to aid in decoding of the PPDU payload at the receiver. The PLCP header may be further decomposed into PHY header, MAC header, HCS, and Reed-Solomon parity bits (see 10.1.2).

The PPDU payload may be the third major component of the PPDU. The PPDU may contain one or more segments as described in 10.1.3.

The ATS may be the last major component of the PPDU. This sequence is used to train an antenna array (see 10.1.4).

When transmitting the PPDU frame, the PLCP preamble shall be sent first, followed by the PLCP header, potentially the PPDU payload, and potentially the ATS.

10.1.1 PLCP preamble

A PLCP preamble shall be added prior to the PLCP header to aid the receiver in timing synchronization, carrier-offset recovery, and channel estimation.

The PLCP preamble for Type A (SCBT and OFDM) and Type B are different. The details of the PLCP preamble in each case can be found in 10.2.2.3, 10.2.3.3 and 10.3.2.3.

10.1.2 PLCP header

A PLCP header shall be added after the PLCP preamble to convey information about both the PHY and the MAC that is needed at the receiver in order to successfully decode the PPDU payload. The scrambled and Reed-Solomon encoded PLCP header shall be formed as shown in Figure 6:

1.Format the PHY header based on information provided by the MAC (See 10.1.2.1).

2.Calculate the HCS value (2 octets) over the combined PHY header and MAC header (See 10.1.2.2).

3.The resulting HCS value is appended to the MAC header. The resulting combination (MAC Header + HCS) is scrambled according to 10.2.2.5.1.1.

4.Apply a shortened Reed-Solomon code to the concatenation of variable length PHY header (except for the first section), ATIF, scrambled MAC header and HCS. The shortened Reed-Solomon code shall be as specified in 10.2.2.5.1.2.

5.Prepend five repetitions of the fixed length PHY header and the first section of the variable length PHY header and append the 16 parity octets at the end to form the scrambled and shortened Reed-Solomon encoded PLCP header.

Further encoding and modulation for the header is different for Type A (SCBT and OFDM) and Type B. The details of the further encoding and modulation for the header for each case can be found in 10.2.2.4, 10.2.3.4 and 10.3.2.4.

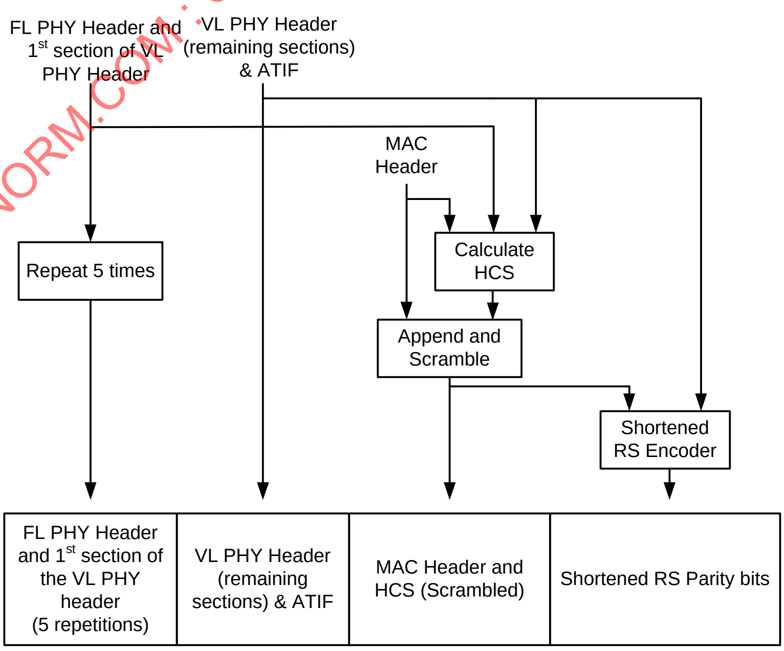

Figure 6- Formation of the header

10.1.2.1 PHY header

The PHY header shall be formed by concatenating five repetitions of the fixed length PHY header and five repetitions of the the first section of the variable length PHY header, followed by the remaining $N_{segments}$-1 sections of the variable length PHY Header and potentially an antenna training indicator field (ATIF). The overall length of the PHY header is therefore equal to $35+4(N_{segments}-1)+3I_{ATIF}$ octets, where I_{ATIF} is equal to zero when the header does not include the ATIF and is equal to one when the header includes the ATIF.

10.1.2.1.1 Fixed length PHY header

The fixed length PHY header contains information about the seed identifier for the data scrambler, the bit reversal state for PPDU payload, the existence of the ATIF, the CP length for the current frame, the requested CP length for the following frame, the number of segments in the frame, and the number of MSDUs in the frame.

The fixed length PHY header field shall be composed of 24 bits, numbered from 0 to 23 as illustrated in Figure 7. Bits 1-2 shall encode the seed value for the initial state of the scrambler, which is used to synchronize the descrambler of the receiver. Bit 3 shall encode whether or not all of the information data bits are inverted. Bit 5 shall encode whether or not the header includes an ATIF. Bits 7-8 shall encode the CP length for the current frame and bits 10-11 shall encode the requested CP length for the following frame. Bits 13-17 shall encode the number of segments in the frame. Bits 19-23 shall encode the number of MSDUs in the frame. All other bits which are not defined in this Clause shall be understood to be reserved for future use and shall be set to zero. The values of the defined fields are described in 10.2.2.4.1.1, 10.2.3.4.1.1 and 10.3.2.4.1.1.

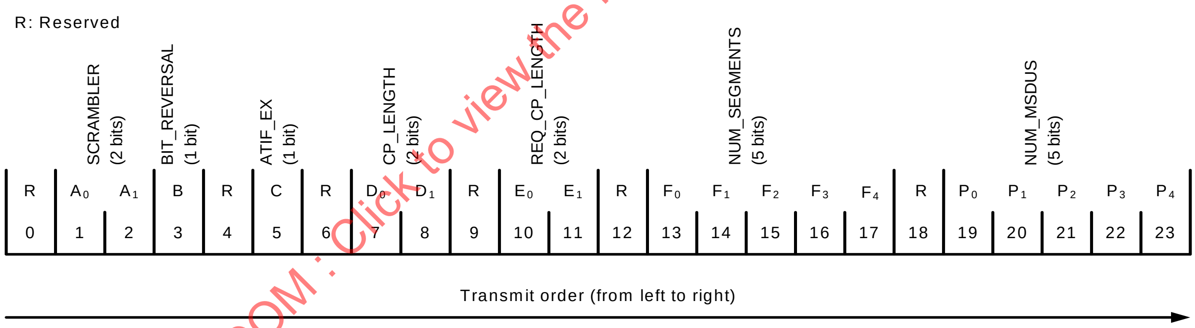

Figure 7 - Fixed length PHY header bit assignment

10.1.2.1.2 Variable length PHY header

Each section of the variable length PHY header contains information about the data rate of the corresponding segment, the length of the payload in the corresponding segment, whether the corresponding segment includes a midamble, and whether the corresponding segment is partitioned at the transmitter and should be reassembled at the receiver.

The variable length PHY header field shall be composed of $32N_{segments}$ bits and shall be constructed by concatenating $N_{segments}$ sections, each for one segment. Each section shall be composed of 32 bits, numbered from 0 to 31, as illustrated in Figure 8. Bits 1-6 shall encode the MODE field, which conveys the information about the type of modulation, the coding rate, and the spreading factor used in the segment. Bits 8-23 shall encode the LENGTH field, with the least-significant bit being transmitted first. Bit 28 shall encode whether the segment is appended by a midamble. Bit 29 shall encode whether the segment has partitioned at the transmitter, and should be reassembled at the receiver. All other bits which are not defined in this Clause shall be understood to be reserved for future use and shall be set to zero. The values of the defined fields are described in 10.2.2.4.1.2, 10.2.3.4.1.2 and 10.3.2.4.1.2.

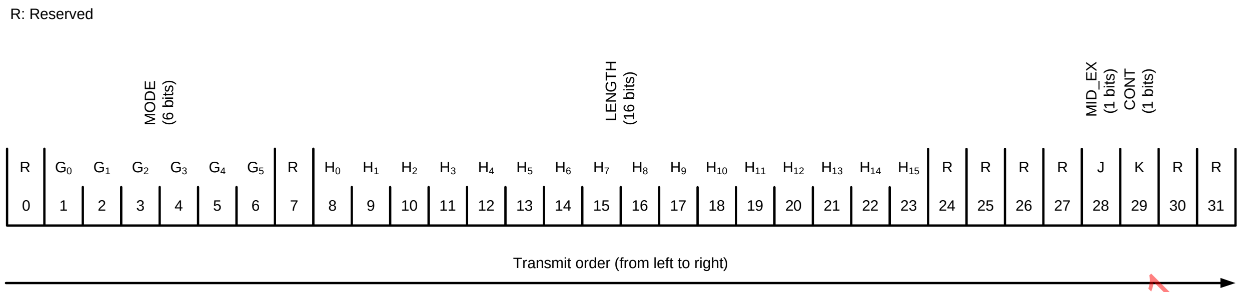

Figure 8 - Variable length PHY header bit assignment

10.1.2.1.3 Antenna training indicator field

The Antenna Training Indicator Field (ATIF) shall follow the variable length PHY header if the ATIF_EXISTENCE bit in the fixed length PHY header is set to 1 (see 10.1.2.1.1). The ATIF contains information about the number of block repetitions in the discovery mode, the number of training symbols for receive antenna training, and the number of training symbols for the transmit antenna training.

The ATIF shall be composed of 24 bits, numbered from 0 to 23, as illustrated in Figure 9. Bits 2-5 shall encode the DISC_REP field which conveys the number of repetitions in the discovery mode. Bits 8-13 shall encode the field NUM_RXTS which conveys the number of training symbols for receive antenna training. Bits 16-21 shall encode the field NUM_TXTS which conveys the number of training symbols for transmit antenna training. All other bits which are not defined in this Clause shall be understood to be reserved for future use and shall be set to zero. The values of the defined fields are described in 10.2.2.4.1.3, 10.2.3.4.1.3 and 10.3.2.4.1.3.

.

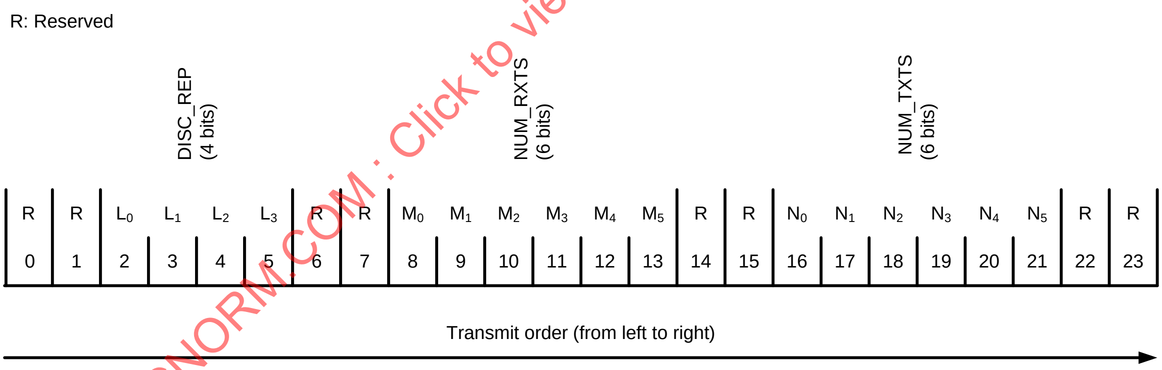

Figure 9 - Antenna Training Indicator Field bit assignment

10.1.2.2 Header check sequence (HCS)

The combination of PHY header and the MAC header shall be protected with a 2 octet CCITT CRC-16 header check sequence (HCS). The CCITT CRC-16 HCS shall be the ones complement of the remainder generated by the modulo-2 division of the combined PHY and MAC headers by the polynomial: $x^{16} + x^{12} + x^5 + 1$. The HCS bits shall be processed in the transmit order. All HCS calculations shall be made prior to data scrambling. A schematic of the processing order is shown in Figure 10. The registers shall be initialized to all ones.

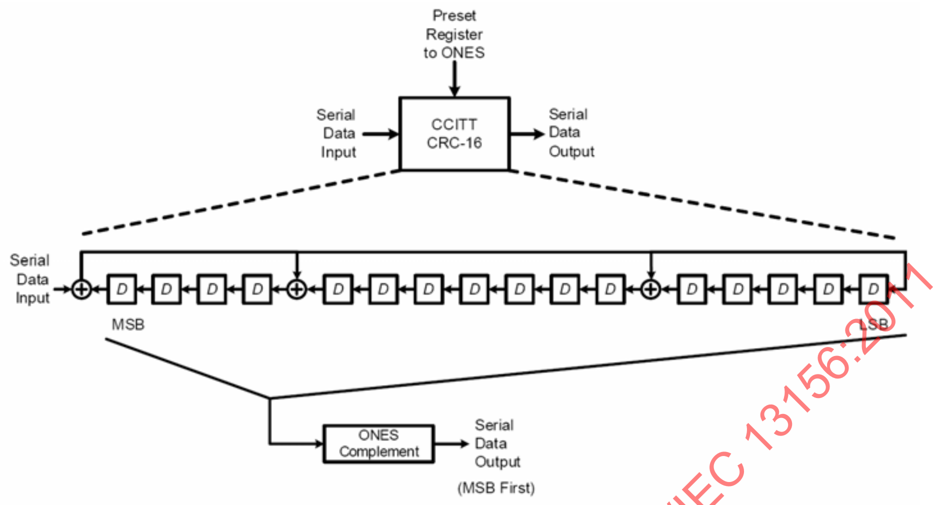

Figure 10 - CCITT CRC-16 block diagram

10.1.3 PPDU payload

The PSDU shall be split into the MAC header and up to 28 PSDU parts, where a PSDU part consists of an MSDU, an FCS and padding. Each PSDU part may be further split into two or more data bit blocks. The total number of PSDU parts and data bit blocks shall not be larger than 57. A number of consecutive PSDU parts that use the same transmission mode may be combined into one data bit block. The resulting data bit blocks shall be encoded and mapped according to the modulation and coding scheme of each device type to generate a transmit symbol block and may be appended by a midamble to form a segment. The details of each modulation and coding scheme is described in 10.2.2.5, 10.2.3.5 and 10.3.2.5. The resulting segments shall be joined consecutively to form the PPDU payload. Figure 11 depicts this operation. Each segment of the PPDU payload may be sent using a different data rate mode. The least-significant bit (LSB) of an octet shall be the first bit transmitted. Type A devices shall be capable of transmission and reception of multi-segment frames. Type B devices may be capable of transmission and reception of multi-segment frames.

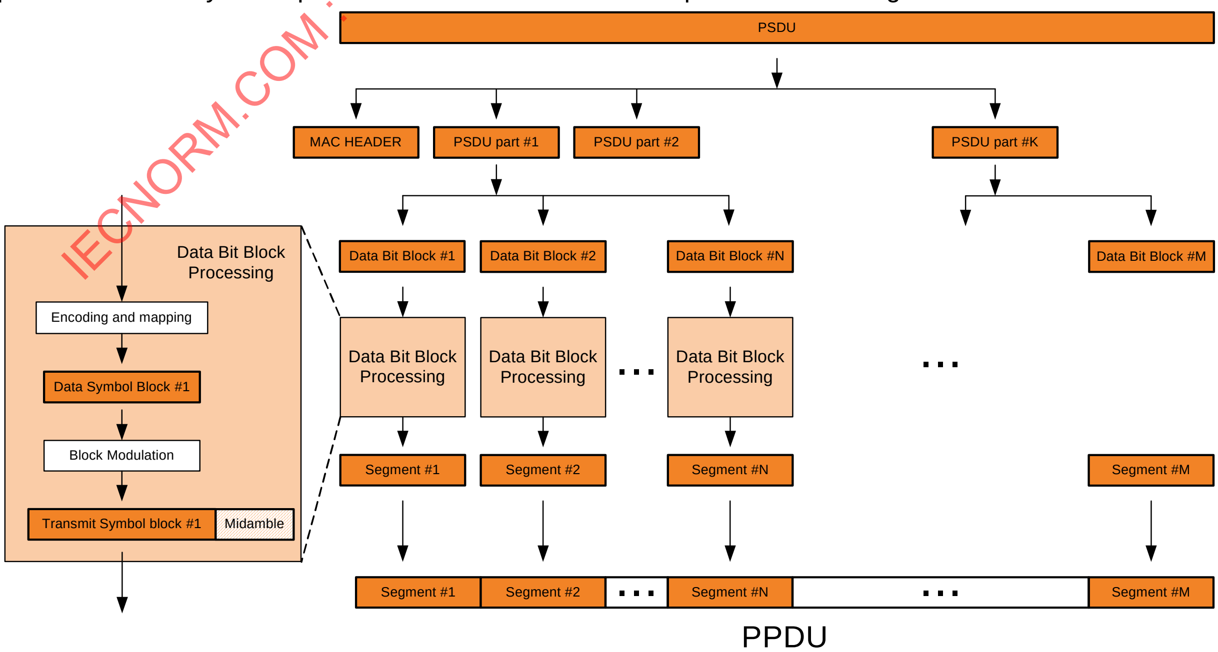

Figure 11 - Formation of the PPDU payload from the PSDU

10.1.3.1 Midamble

The midamble for Type A (SCBT and OFDM) and Type B are different. The details of the midamble for each case can be found in 10.2.2.5.4, 10.2.3.5.2, and 10.3.2.5.5.

10.1.4 Antenna training sequence

The Type A frames shall include the ATS, upon request. The Type B frames may include the ATS upon request. When included, the ATS shall be the concatenation of $N_{TXTS} + N_{RXTS}$ FZ sequences of length 256 (parameter $A_{FZ} = 16$). The FZ sequence is defined in 10.2.2.3, where N_{TXTS} and N_{RXTS} are the number of training sequences for transmitter and receiver antenna training, respectively (see Figure 12).

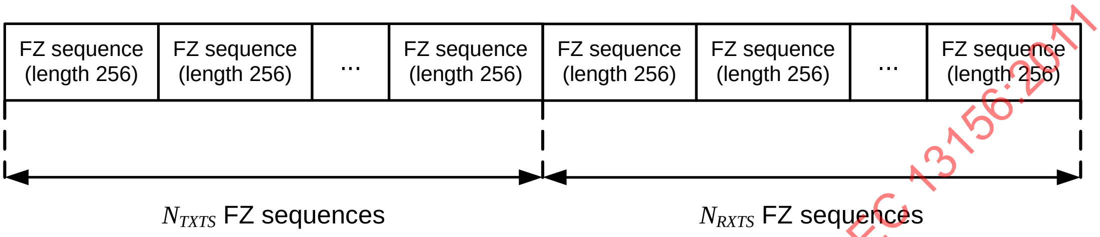

Figure 12 - Antenna training sequence

10.2 Type A PPDU

10.2.1 Mode dependent parameters

The PSDU mode dependent modulation parameters are listed in Table 1.

Table 1- PSDU mode dependent Parameters

Mode	Base Data Rate[1] (Gbps)				Modulation	Constellation	Encoding	CC Code Rate (R)	TDSF (N_{TDS})	Tail bits (N_{tail})
	No channel bonding	2 bonded channels	3 bonded channels	4 bonded channels						
A0	0.397	0.794	1.191	1.588	SCBT	BPSK	RS & CC	1/2	2	4
A1	0.794	1.588	2.381	3.175	SCBT	BPSK	RS & CC	1/2	1	4
A2	1.588	3.175	4.763	6.350	SCBT	BPSK	RS	1	1	0
A3	1.588	3.175	4.763	6.350	SCBT	QPSK	RS & CC	1/2	1	4
A4	2.722	5.443	8.165	10.886	SCBT	QPSK	RS & CC	6/7	1	6
A5	3.175	6.350	9.526	12.701	SCBT	QPSK	RS	1	1	0
A6	4.234	8.467	13.701	16.934	SCBT	NS8QAM	RS & TCM	5/6	1	8
A7	4.763	9.526	14.288	19.051	SCBT	NS8QAM	RS	1	1	0
A8	4.763	9.526	14.288	19.051	SCBT	TCM-16QAM	RS & TCM	2/3	1	6
A9	6.350	12.701	19.051	25.402	SCBT	16QAM	RS	1	1	0
A10[2]	1.588	3.175	4.763	6.350	SCBT	QPSK	RS & UEP-CC	R_{MSB}: 1/2	1	4
A11	4.234	8.467	12.701	16.934	SCBT	16QAM	RS & UEP-CC	R_{MSB}: 4/7, R_{LSB}: 4/5	1	4
A12	2.117	4.234	6.350	8.467	SCBT	UEP-QPSK	RS & CC	2/3	1	4
A13	4.234	8.467	12.701	16.934	SCBT	UEP-16QAM	RS & CC	2/3	1	4

Table 1- PSDU mode dependent Parameters (concluded)

Mode	Base Data Rate[1] (Gbps) No channel bonding	2 bonded channels	3 bonded channels	4 bonded channels	Modulation	Constellation	Encoding	CC Code Rate (R)	TDSF (N_{TDS})	Tail bits (N_{tail})
A14	1.008	N/A	N/A	N/A	OFDM	QPSK	RS & CC	1/3	1	6
A15	2.016	N/A	N/A	N/A	OFDM	QPSK	RS & CC	2/3	1	6
A16	4.032	N/A	N/A	N/A	OFDM	16QAM	RS & CC	2/3	1	6
A17	2.016	N/A	N/A	N/A	OFDM	QPSK	RS & UEP-CC	R_{MSB}: 4/7, R_{LSB}: 4/5	1	6
A18	4.032	N/A	N/A	N/A	OFDM	16QAM	RS & UEP-CC	R_{MSB}: 4/7, R_{LSB}: 4/5	1	6
A19	2.016	N/A	N/A	N/A	OFDM	UEP-QPSK	RS & CC	2/3	1	6
A20	4.032	N/A	N/A	N/A	OFDM	UEP-16QAM	RS & CC	2/3	1	6
A21[2]	2.016	N/A	N/A	N/A	OFDM	QPSK	RS & CC	R_{MSB}: 2/3	1	6

NOTE 1 Base data rates assume a cyclic prefix length of zero.

NOTE 2 In modes A10 and A21 only the four MSB bits of each octet shall be transmitted, while the four LSB bits of each octet are discarded by the transmitter.

All Type A devices shall support mode A0 (without channel bonding), and may support mode A0 with channel bonding or modes A1 through A21. In addition all Type A devices shall support mode B0 (without channel bonding). Type A devices may support modes B0 (with channel bonding), B1, B2, or B3 (with or without channel bonding). See 10.3 for Type B modes.

10.2.2 SCBT

10.2.2.1 Timing related parameters

The timing parameters associated with the SCBT PHY are listed in Table 2.

Table 2 - Timing related parameters

Parameter	Description	Value
f_{sym}	Symbol frequency	1.728 Gsps
T_{sym}	Symbol duration	0.5787 ns
N_{SCBTB}	Number of symbols per SCBT block	256
N_D	Number of data symbols per SCBT block	252
N_P	Number of pilot symbols per SCBT block	4
N_{CP}	Number of symbols in the CP	0 32 64 96
N_{SCBTS}	Number of symbols per SCBT block	256 288 320 352

Table 2 - Timing related parameters (concluded)

Parameter	Description	Value
T_{SCBTB}	SCBT block interval	148.148 ns
T_{CP}	CP interval	0 18.5185 ns 37.037 ns 55.5556 ns
T_{SCBTS}	SCBT symbol interval	148.148 ns 166.667 ns 185.185 ns 203.707 ns

10.2.2.2 Frame related parameters

The frame related parameters associated with the PHY are listed in Table 3.

Table 3 - Frame Related Parameters

Parameter	Description	Value
N_{sync}	Number of symbols in the frame synchronization sequence	2048
T_{sync}	Duration of the frame synchronization sequence	1185.19 ns
N_{CE}	Number of symbols in the channel estimation sequence	768
T_{CE}	Duration of the channel estimation sequence	444.444 ns
$N_{preamble}$	Number of symbols in the PLCP preamble	2816
$T_{preamble}$	Duration of the PLCP preamble	1629.63 ns
N_{ATS}	Number of symbols in the ATS	$256(N_{TXTS} + N_{RXTS})N_{DISCREP}$
T_{ATS}	Duration of the ATS	$N_{ATS}T_{sym}$
N_{frame}	Number of symbols in the frame	$N_{preamble} + N_{header} + N_{payload} + N_{ATS}$
T_{frame}	Duration of the frame	$(N_{preamble} + N_{header} + N_{payload} + N_{ATS})T_{sym}$

10.2.2.3 PLCP preamble

The preamble for Type A SCBT frames can be subdivided into two distinct portions: a frame synchronization sequence, $S_{sync,A,SCBT}[.]$, and a channel estimation sequence, $S_{CE,A,SCBT}[.]$. Figure 13 shows the structure of the PLCP preamble for Type A SCBT frames.

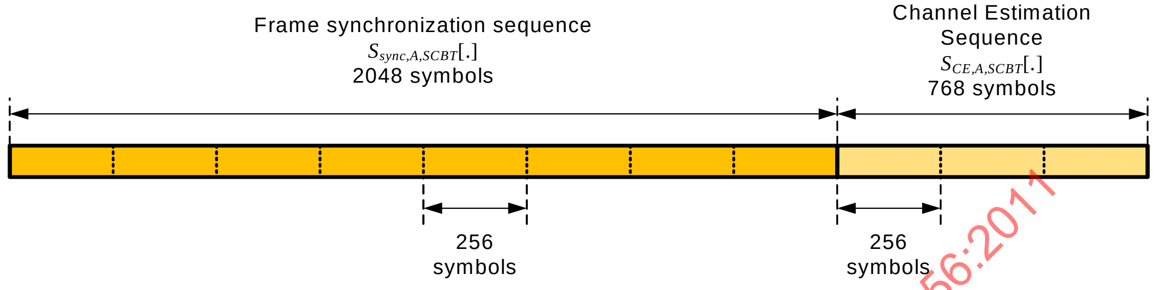

Figure 13 - Structure of the Type A PLCP preamble for Type A SCBT frames

Both the frame synchronization sequence and the channel estimation sequence are constructed based on the Frank-Zadoff (FZ) sequences. A Frank-Zadoff sequence with parameter A_{FZ}, which has length $N_{FZ}=A_{FZ}^2$, is defined as:

$$S_{FZ,A_{FZ}}[n] = \exp\left(j\frac{2\pi pq}{A_{FZ}}\right) \tag{4}$$

where $p=(n \bmod A_{FZ})+1$, and $q=\left\lfloor \frac{n}{A_{FZ}} \right\rfloor+1$, for $n=0,\ldots, A_{FZ}^2-1$, and $\lfloor . \rfloor$ denotes the floor function, which returns the largest integer value smaller than or equal to its argument.

10.2.2.3.1 Frame synchronization sequence

The frame synchronization sequence for Type A SCBT frames consists of eight repetitions of a hierarchical sequence $S_h[.]$, covered by sequence $S_{cover}[.]$, as described in equation (7). The hierarchical sequence, $S_h[.]$, is of length 256, and is defined as the Kronecker product of the Frank-Zadoff sequence of length 16, $S_{FZ,4}[.]$, with itself. That is,

$$S_h[n] = S_{FZ,4}[n \bmod 16]S_{FZ,4}\left[\left\lfloor \frac{n}{16} \right\rfloor\right] \tag{5}$$

for $n=0,\ldots,255$. The cover sequence is of length eight and is defined as

$$S_{cover}[n] = \begin{cases} 1 & n=0,\ldots,6 \\ -1 & n=7 \end{cases} \tag{6}$$

The frame synchronization sequence is defined as the Kronecker product of the hierarchical sequence $S_h[.]$ and the cover sequence $S_{cover}[.]$. That is,

$$S_{sync,A,SCBT}[n] = S_{cover}\left[\left\lfloor \frac{n}{256} \right\rfloor\right]S_h[n \bmod 256] \tag{7}$$

for $n=0,\ldots,2047$.

10.2.2.3.2 Channel estimation sequence

The channel estimation sequence consists of three repetitions of the Frank-Zadoff sequence of length 256, $S_{FZ,16}[.]$. That is, the channel estimation sequence is given by:

$$S_{CE,A,SCBT}[n] = S_{FZ,16}[n \bmod 256] \qquad (8)$$

for $n=0,\ldots,767$.

NOTE The frame synchronization sequence can be used for frame acquisition and detection, coarse carrier frequency estimation, coarse symbol timing, and for synchronization within the preamble. Whereas, the channel estimation sequence can be used for estimation of the channel frequency response, fine carrier frequency estimation, and fine symbol timing.

10.2.2.4 PLCP header

For Type A SCBT frames, the formed PLCP header (as described in 10.1.2) shall be first padded with N_{padhdr} bits where

$$N_{padhdr} = 192\left\lceil \frac{N_{hdrbits}}{192} \right\rceil - N_{hdrbits} \qquad (9)$$

where $N_{hdrbits}$ is the number of bits in the formed PLCP header (as described in 10.1.2). The padded header shall be encoded and modulated as described in 10.2.2.5, starting with the demultiplexing. Encoding and modulation parameters identical to mode A0 shall be used. The resulting data symbols shall be block modulated (see 10.2.2.5.3) in order to create the baseband signal.

10.2.2.4.1 PHY header

10.2.2.4.1.1 Fixed length PHY header

10.2.2.4.1.1.1 PLCP scrambler field (SCRAMBLER)

The bits A_0 and A_1 shall be set according to the scrambler seed identifier value. This two-bit value corresponds to the seed value chosen for the data scrambler (see 10.2.2.5.1.1).

10.2.2.4.1.1.2 Bit reversal field (BIT_REVERSAL)

The value of the BIT_REVERSAL bit shall be set to 0_B for all Type A SCBT frames.

10.2.2.4.1.1.3 ATIF existence field (ATIF_EX)

The ATIF_EX field determines whether or not an ATIF field exists at the end of the PHY header. The mapping between this field and existence of the ATIF is given in Table 4. The ATIF_EX value for Type A Beacon and Type B Beacon shall be 0

Table 4 – Mapping between the value of the ATIF_EX bit and existence of ATIF

ATIF_EX (C)	ATIF Existence
0	ATIF does not exist
1	ATIF does exist

10.2.2.4.1.1.4 CP length field (CP_LENGTH)

The mapping between the value of the CP_LENGTH field and N_{CP} for the current frame is given in Table 5. The default value of N_{CP} for the first frame, or for beacons, discovery mode frames and TRN frames shall be 96.

Table 5 - Mapping between the value of the CP_LENGTH field and N_{CP}

CP_LENGTH (D_1D_0)	N_{CP}
00	0
01	32
10	64
11	96

10.2.2.4.1.1.5 Requested CP length field (REQ_CP_LENGTH)

The mapping between the value of the REQ_CP_LENGTH field and N_{CP} for the following frames is given in Table 6.

Table 6 - Mapping between the value of the REQ_CP_LENGTH field and N_{CP}

REQ_CP_LENGTH (E_1E_0)	N_{CP}
00	0
01	32
10	64
11	96

10.2.2.4.1.1.6 Number of segments field (NUM_SEGMENTS)

Depending on the number of segments (NUM_SEGMENTS), bits F_0 to F_4 shall be set according to the values in Table 7. When a frame does not have a payload, this field shall be set to 0000_B.

Table 7 - Mapping between the value of the NUM_SEGMENTS field and number of segments

NUM_SEGMENTS ($F_4F_3F_2F_1F_0$)	Number of segments ($N_{segments}$)
00000	0
00001	1
00010	2
00011	3
00100	4
00101	5
00110	6
00111	7

Table 7 - Mapping between the value of the NUM_SEGMENTS field and number of segments (concluded)

NUM_SEGMENTS ($F_4F_3F_2F_1F_0$)	Number of segments ($N_{segments}$)
01000	8
01001	9
01010	10
01011	11
01100	12
01101	13
01110	14
01111	15
10000	16
10001	17
10010	18
10011	19
10100	20
10101	21
10110	22
10111	23
11000	24
11001	25
11010	26
11011	27
11100	28
11101	Reserved
11110	Reserved
11111	Reserved

NOTE Discovery mode frames and ACK frames may have zero segments (i.e. no payload).

10.2.2.4.1.1.7 Number of MSDUs field (NUM_MSDUS)

Depending on the number of MSDUs (NUM_MSDUS), bits P_0 to P_4 shall be set according to the values in Table 8. When a frame does not have a payload, this field shall be set to 00000_B.

Table 8 - Mapping between the value of the NUM_MSDUS field and number of MSDUs

NUM_MSDUS ($P_4P_3P_2P_1P_0$)	Number of MSDUs (N_{MSDUs})
00000	0
00001	1
00010	2
00011	3
00100	4
00101	5
00110	6
00111	7
01000	8
01001	9
01010	10
01011	11
01100	12
01101	13
01110	14
01111	15
10000	16
10001	17
10010	18
10011	19
10100	20
10101	21
10110	22
10111	23
11000	24
11001	25
11010	26
11011	27
11100	28

Table 8 - Mapping between the value of the NUM_MSDUS field and number of MSDUs (concluded)

NUM_MSDUS ($P_4P_3P_2P_1P_0$)	Number of MSDUs (N_{MSDUs})
11101	Reserved
11110	Reserved
11111	Reserved

10.2.2.4.1.2 Variable length PHY header

10.2.2.4.1.2.1 Mode field (MODE)

Depending on the mode (MODE), bits G_0 to G_5 shall be set according to the values in Table 9.

Table 9 - MODE field

Mode	MODE (G_5...G_0)	Mode	MODE (G_5...G_0)	Mode	MODE (G_5...G_0)	Mode	MODE (G_5...G_0)
A0	000000	A16	010000	Reserved	100000	Reserved	110000
A1	000001	A17	010001	Reserved	100001	Reserved	110001
A2	000010	A18	010010	Reserved	100010	Reserved	110010
A3	000011	A19	010011	Reserved	100011	Reserved	110011
A4	000100	A20	010100	Reserved	100100	Reserved	110100
A5	000101	A21	010101	Reserved	100101	Reserved	110101
A6	000110	Reserved	010110	Reserved	100110	Reserved	110110
A7	000111	Reserved	010111	Reserved	100111	Reserved	110111
A8	001000	Reserved	011000	Reserved	100111	Reserved	111000
A9	001001	Reserved	011001	B0	101001	Reserved	111001
A10	001010	Reserved	011010	B1	101010	Reserved	111010
A11	001011	Reserved	011011	B2	101011	Reserved	111011
A12	001100	Reserved	011100	B3	101100	Reserved	111100
A13	001101	Reserved	011101	Reserved	101101	Reserved	111101
A14	001110	Reserved	011110	Reserved	101110	Reserved	111110
A15	001111	Reserved	011111	Reserved	101111	Reserved	111111

10.2.2.4.1.2.2 Segment length field (LENGTH)

The segment length field shall be an unsigned 16-bit integer that indicates the number of octets in the uncoded segment payload (which does not include the tail bits, or the pad bits). The length of each segment shall not be larger than 65535 octets.

10.2.2.4.1.2.3 Midamble existence field (MID_EX)

The MID_EX field determines whether or not the segment is appended by a midamble. The mapping between this bit and existence of the midamble is given in Table 10.

Table 10 - Mapping between the value of the MID_EX bit and existence of midamble

MID_EX (J)	Midamble existence
0	Midamble does not exist
1	Midamble exists

10.2.2.4.1.2.4 Segment continued field (CONT)

The CONT field determines whether or not this segment is a fragment of the received PSDU, and should be reassembled at the receiver. The value of this bit shall be according to Table 11.

Table 11 - Value of the CONT bit

CONT (K)	Fragmentation
0	The segment is not a fragment or is the last fragment
1	The segment is a fragment

10.2.2.4.1.3 Antenna training indicator field

10.2.2.4.1.3.1 Discovery mode repetition field (DISC_REP)

The discovery mode repetition field represents the number of repetitions in the discovery mode. The mapping between the value of DISC_REP and $N_{DISCREP}$ is given in Table 12. For all non-discovery mode frames the value of DISC_REP shall be set to $L_3L_2L_1L_0 = 0000_B$.

Table 12 - Mapping between the value of the DISC_REP field and number of discovery mode repetitions

DISC_REP ($L_3L_2L_1L_0$)	Number of repetitions in discovery mode ($N_{DISCREP}$)
0000	1
0001	2
0010	4
0011	8
0100	16
0101	32
0110	64
0111	128
1000	Reserved
1001	Reserved
1010	Reserved

Table 12 - Mapping between the value of the DISC_REP field and number of discovery mode repetitions (concluded)

DISC_REP ($L_3L_2L_1L_0$)	Number of repetitions in discovery mode ($N_{DISCREP}$)
1011	Reserved
1100	Reserved
1101	Reserved
1110	Reserved
1111	Reserved

10.2.2.4.1.3.2 Number of training symbols for RX training field (NUM_RXTS)

The number of training symbols for RX training field shall be an unsigned 6-bit integer. Its value shall indicate the number of training symbols for RX training in the ATS. The NUM_RXTS value for Discovery Mode Beacon shall be 0.

10.2.2.4.1.3.3 Number of training symbols for TX training field (NUM_TXTS)

The number of training symbols for TX training field shall be an unsigned 6-bit integer. Its value shall indicate the number of training symbols for TX training in the ATS. The NUM_TXTS value for Discovery Mode Beacon is 0.

10.2.2.5 PPDU payload

The PPDU payload consists of one or more segments as described in 10.1. The segments are formed as described in 10.1.3. Each data bit block shall be processed as follows:

1.The resulting bits shall be encoded and mapped as described in 10.2.2.5.1 or 10.2.2.5.2.

2.The resulting symbols shall be modulated into SCBT symbols as described in 10.2.2.5.3.

10.2.2.5.1 FEC and mapping (Equal error protection)

Figure 14 depicts the general overview of the encoding and mapping scheme for modes A0 through A9. For the payload, first, the data bits shall be scrambled as specified in 10.2.2.5.1.1. Then, $N_{padbits}$ zero pad bits shall be appended to the end of the data bit block, where

$$N_{padbits} = 1792\left\lceil \frac{N_{bits}}{1792} \right\rceil - N_{bits} \quad (10)$$

and N_{bits} is the number of bits in the data bit block. The padded data bit block shall be encoded according to the Reed-Solomon code, as specified in 10.2.2.5.1.2.

The resulting RS coded payload bits or the formed header bits shall be demultiplexed into 4 groups, as described in 10.2.2.5.1.3. Depending on the data rate mode, each group of bits shall then be interleaved using the bit interleaver described in 10.2.2.5.1.4. The resulting bits shall be appended with N_{tail} tail bits (in each branch). They then shall be encoded using the convolutional code or trellis coded modulation described in 10.2.2.5.1.5 and 10.2.2.5.1.6. The resulting encoded bits from the 4 branches shall then be multiplexed together into one group of symbols as specified in 10.2.2.5.1.7. The multiplexed bits are then mapped to constellations as specified in 10.2.2.5.1.8. The data symbols shall be spread in time domain as described in 10.2.2.5.1.9. The resulting data symbols shall then be appended with N_{padsym} zero symbols, where

$$N_{padsyms} = 2N_D \left\lceil \frac{N_{sym}}{2N_D} \right\rceil - N_{sym} \qquad (11)$$

At the end of each segment, the resulting data symbols shall be interleaved with the Dual Helical Scan (DHS) symbol interleaver as described in 10.2.2.5.1.10.

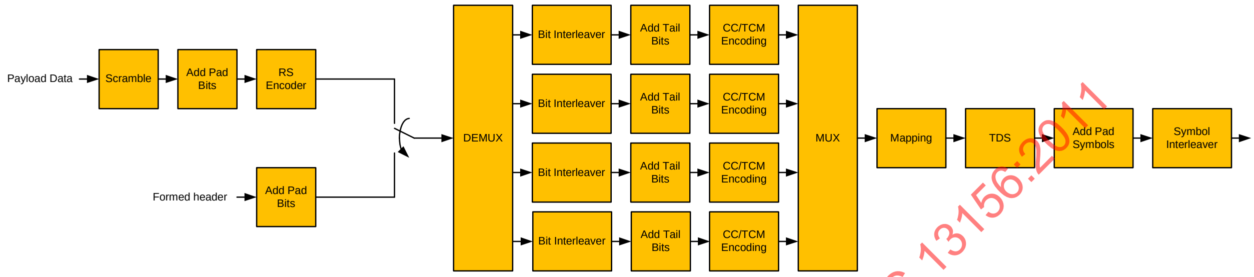

Figure 14 - General view of the encoding procedure

10.2.2.5.1.1 Data scrambler

A side-stream scrambler shall be used to whiten only portions of the PLCP header, i.e. the MAC header and HCS, and the entirety of each segment. In addition, the scrambler shall be initialized to a seed value specified by the MAC at the beginning of the MAC header and then re-initialized to the same seed value at the beginning of each segment.

The polynomial generator, $g(D)$, for the pseudo-random binary sequence (PRBS) generator shall be: $g(D) = 1 + D^{14} + D^{15}$, where D is a single bit delay element. Using this generator polynomial, the corresponding PRBS, $x[n]$, is generated as

$$x[n] = x[n-14] \oplus x[n-15], n = 0,1,2,... \qquad (12)$$

where $\oplus$ denotes modulo-2 addition. The following sequence defines the initialization vector, $\mathbf{x}_{init}$, which is specified by the parameter “seed value” in:

$$\mathbf{x}_{init} = [x_i[-1] \quad x_i[-2] \quad ... \quad x_i[-14] \quad x_i[-15]] \qquad (13)$$

where $x_i[-k]$ represents the binary initial value at the output of the k^{th} delay element. The scrambled data bits, v_m, are obtained as show in Figure 15.

$$v[n] = s[m] \oplus x[m], m = 0,1,2,... \qquad (14)$$

where $s[m]$ represents the non-scrambled data bits.

The 15-bit initialization vector or seed value shall correspond to the seed identifier as shown in Table 13. The seed identifier value shall be set to 00 when the PHY is initialized and this value shall be incremented in a 2-bit rollover counter for *each* frame sent by the PHY.

Table 13 - Scrambler seed selection

Seed Identifier (A_1A_0)	Seed Value $x_{init} = x_i[-1]\ x_i[-2] \ldots x_i[-15]$	PRBS Output First 16 bits $x[0]\ x[1] \ldots x[15]$
00	0011 1111 1111 111	0000 0000 0000 1000

Table 13 - Scrambler seed selection (concluded)

01	0111 1111 1111 111	0000 0000 0000 0100
10	1011 1111 1111 111	0000 0000 0000 1110
Seed Identifier (A_1A_0)	Seed Value $x_{init} = x_i[-1]\ x_i[-2] \ldots x_i[-15]$	PRBS Output First 16 bits $x[0]\ x[1] \ldots x[15]$
11	1111 1111 1111 111	0000 0000 0000 0010

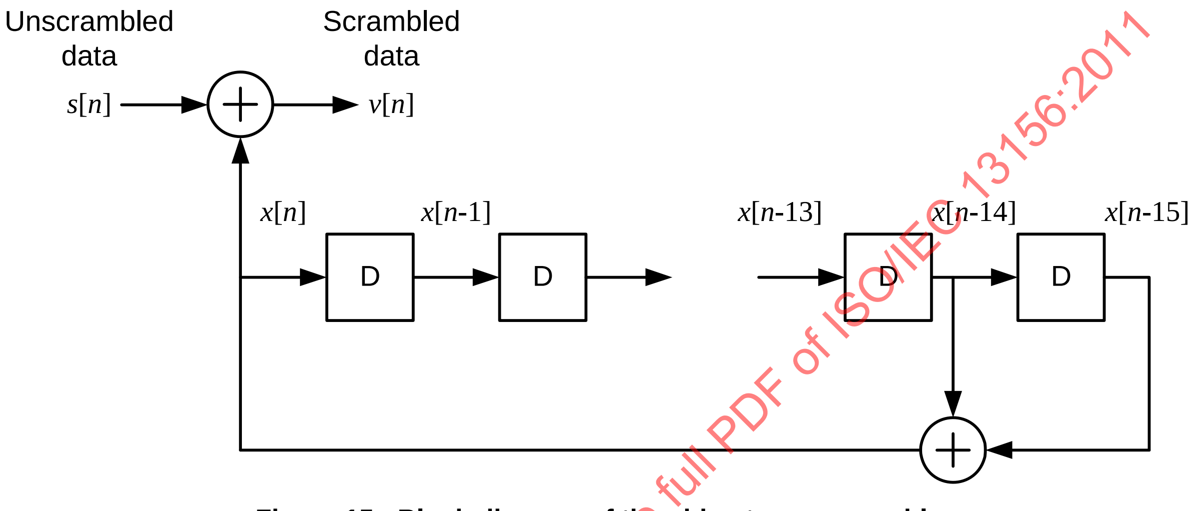

Figure 15 - Block diagram of the side-stream scramble

10.2.2.5.1.2 Reed-Solomon code

A systematic (N=255, K=239, T=8) Reed-Solomon code is shortened to different lengths to be used for the header and the payload. The RS(255,239) shall be shortened to RS($255\text{-}N_{RSpad}$,$239\text{-}N_{RSpad}$) by first padding a block of $239\text{-}N_{RSpad}$ data octets with N_{RSpad} zero octets. The resulting block shall be coded using the RS(255,239) Reed-Solomon code described below. The resulting 16 parity octets r_{15}, ..., r_0 shall be appended to the data block to form the coded data block of length $255\text{-}N_{RSpad}$.

For the header, the N_{RSpad} shall be equal to $237\text{-}4N_{segments}\text{-}3I_{ATIF}\text{-}N_{MACheader}$, where $N_{MACheader}$ is the length of the MAC header (refer to 15.2.1 and 15.2.2). For the payload, the N_{RSpad} shall be equal to 15.

The RS(255,239) code is defined over $GF(2^8)$ with a primitive polynomial $p(z) = z8+ z4+z3+z2+1$. As notation, the element $M = b_7z^7+b_6z^6+b_5z^5+b_4z^4+b_3z^3+b_2z^2+b_1z+b_0$, where M Î $GF(2^8)$, has the following binary representation $b_7b_6b_5b_4b_3b_2b_1b_0$, where b_7 is the MSB and b_0 is the LSB.

The code is specified by the generator polynomial $g(x)$ defined over $GF(2^8)$.

$$g(x) = \prod_{i=0}^{2T-1} (x - \alpha^i) \tag{15}$$

where a=02_H is the root of the polynomial $p(z)$.

The mapping of the information octets $m = (m_{238}, m_{237}, \ldots, m_0)$ to codeword octets $c = (m_{238}, m_{237}, \ldots, m_0, r_{15}, r_{14}, \ldots, r_0)$ is achieved by computing the remainder polynomial $r(x)$

$$r(x) = \prod_{i=0}^{15} r_i x^i = x^{16} m(x) \bmod g(x) \tag{16}$$

where $m(x)$ is the information polynomial

$$m(x) = \prod_{i=0}^{238} m_i x^i \tag{17}$$

and r_i, $i = 0, \ldots, 15$, and m_i, $i = 0, \ldots, 238$, are elements of $GF(2^8)$.

10.2.2.5.1.3 Demultiplexer

The Reed-Solomon encoded bits $b[.]$ shall be demultiplexed to create four bit streams $c_m[.]$, m=0,…,3. The $c_m[.]$ shall be obtained from $b[.]$ from $c_m[k] = b[4k+m]$.

10.2.2.5.1.4 Bit interleaver

The bit streams shall be interleaved by a bit interleaver of length 48. The interleaved bit $d_m[l]$ shall be equal to $c_m[k]$ by

$$l = \left[6 \left\lfloor \frac{k}{6} \right\rfloor + 7(k \bmod 6) \right] \bmod 48 \tag{18}$$

10.2.2.5.1.5 Convolutional code

The convolutional encoder shall use the rate R=1/2 code with generator polynomials, $g_0 = 23_O$, and $g_1 = 35_O$, as shown in Figure 16. The bit denoted as "A" shall be the first bit generated by the encoder, followed by the bit denoted as "B". Additional coding rates are derived from the "mother" rate R=1/2 convolutional code by employing "puncturing". Puncturing is a procedure for omitting some of the encoded bits at the transmitter (thus reducing the number of transmitted bits and increasing the coding rate) and inserting a dummy "zero" metric into the decoder at the receiver in place of the omitted bits. The puncturing patterns are illustrated in Figure 17, Figure 18, Figure 19, Figure 20, and Figure 21. In each of these cases, the tables shall be filled in with encoder output bits from left to right. For the last block of bits, the process shall be stopped at the point at which encoder output bits are exhausted, and the puncturing pattern applied to the partially filled block.

The encoder shall start from the all-zero state. After the encoding process for the PLCP header has been completed, the encoder shall be reset to the all-zero state before the encoding starts for each segment; in other words, the encoding of each segment shall also start from the all-zero state. The PSDU shall be encoded using the appropriate coding rate of R = 4/7, 2/3, 4/5, 5/6, or 6/7.

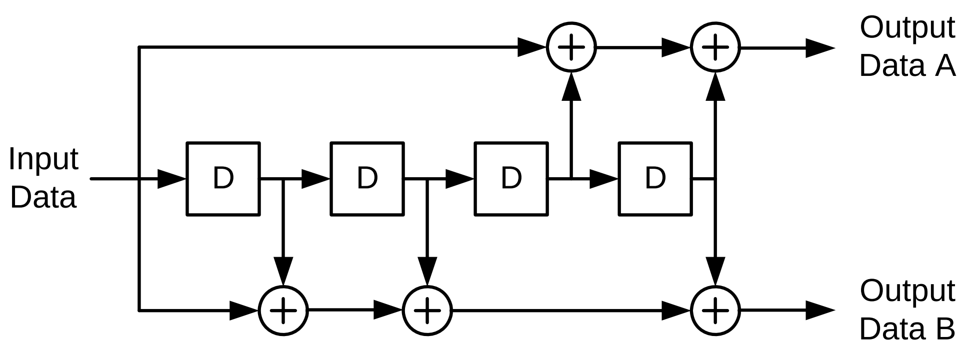

Figure 16 - Convolutional encoder: rate R=1/2, constraint length K=5

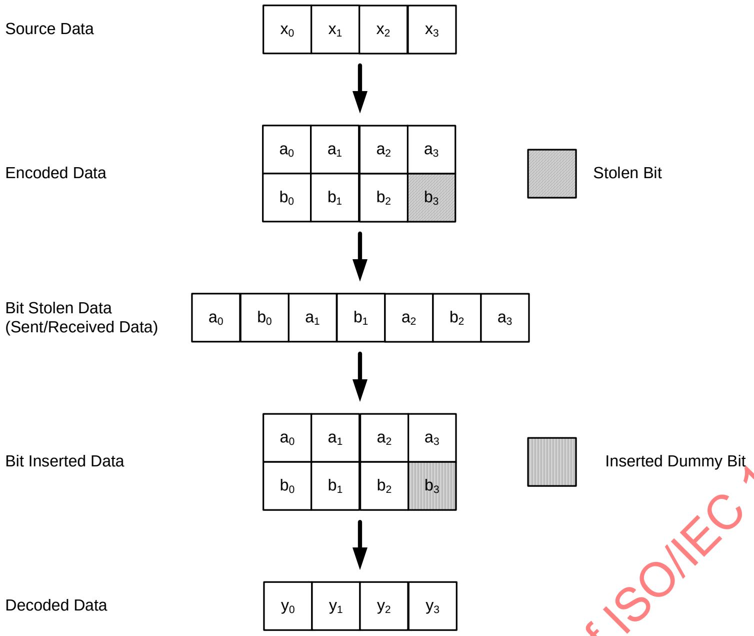

Figure 17 - An example of bit stealing and bit insertion for R=4/7 code

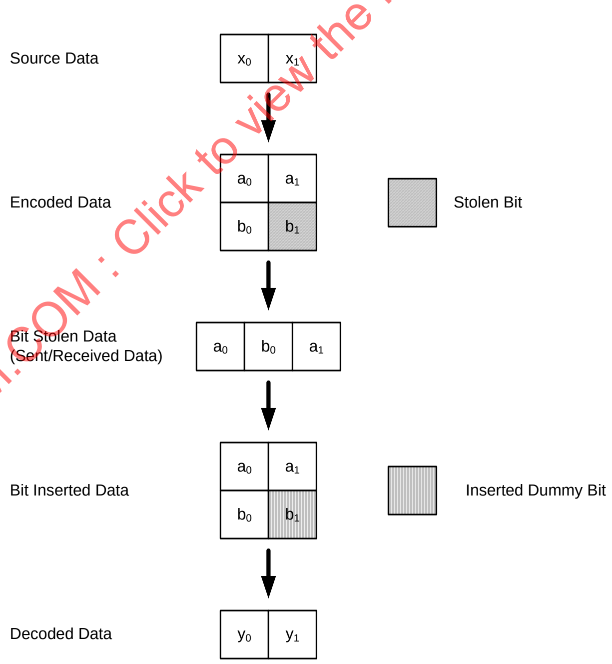

Figure 18- An example of bit-stealing and bit-insertion for R=2/3 code

Figure 19 - An example of bit-stealing and bit-insertion for R=4/5 code

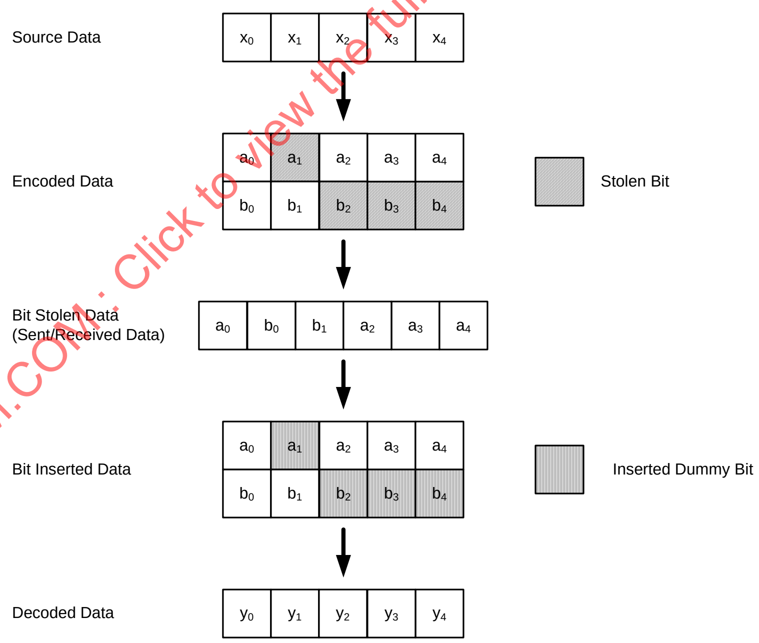

Figure 20 - An example of bit-stealing and bit-insertion for R=5/6 code

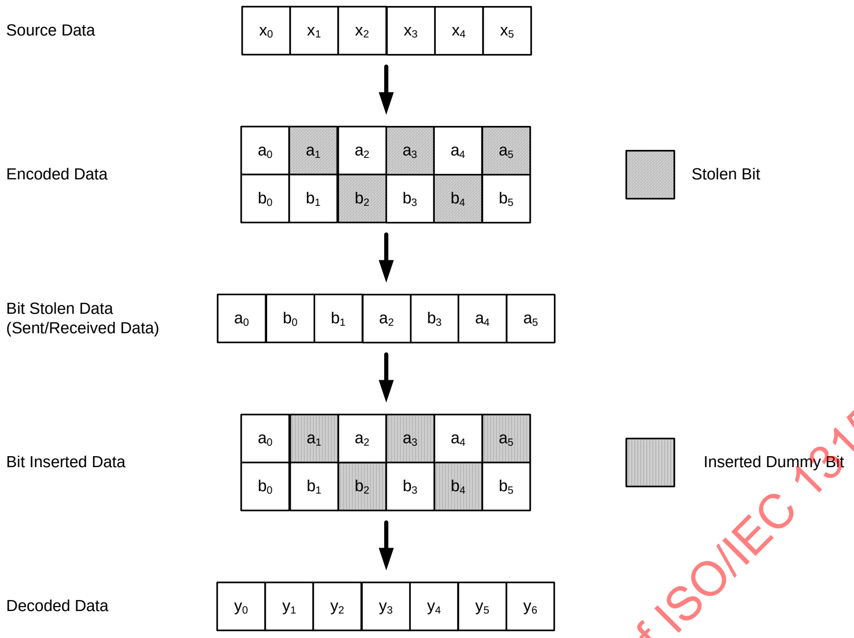

Figure 21 - An example of bit-stealing and bit-insertion for R=6/7 code

NOTE Decoding by the Viterbi algorithm is recommended

10.2.2.5.1.6 Trellis coded modulation

For NS8QAM and 16QAM constellations, a "pragmatic" trellis coded modulation (TCM) is employed by using the convolutional code and the puncturing patterns as described in 10.2.2.5.1.5.

Before the convolutional encoding, the bits $e_m[.]$ shall be demultiplexed into two groups $e_{c,m}[.]$ and $e_{u,m}[.]$, where

$$e_{u,m}[i] = e_m\left[\left\lfloor \frac{i}{U} \right\rfloor (U + C) + i \bmod U\right] \quad (19)$$

and

$$e_{c,m}[i] = e_m\left[\left\lfloor \frac{i}{C} \right\rfloor (U + C) + i \bmod C + U\right] \quad (20)$$

The bits $e_{c,m}[.]$ shall be encoded and punctured to produce the bits $p_m[.]$. The resulting bits, $p_m[.]$, shall then be multiplexed with bits $e_{u,m}[.]$ to generate $f_m[.]$, where

$$f_m[i] = \begin{cases} e_{u,m}[i/P] & i \bmod P = 0 \\ e_{c,m}[\lfloor i/P \rfloor (P-1) + i \bmod P - 1] & i \bmod P \neq 0 \end{cases} \quad (21)$$

For a given data rate mode, the convolutional code rate (including encoding and puncturing) and C, P and U shall be as described in Table 14.

Table 14 - Parameters for trellis coded modulation

Mode	R	U	C	P
A6	5/6	3	5	3
A8	2/3	1	2	4

An example of this operation for data rate mode A6 is depicted in Figure 22.

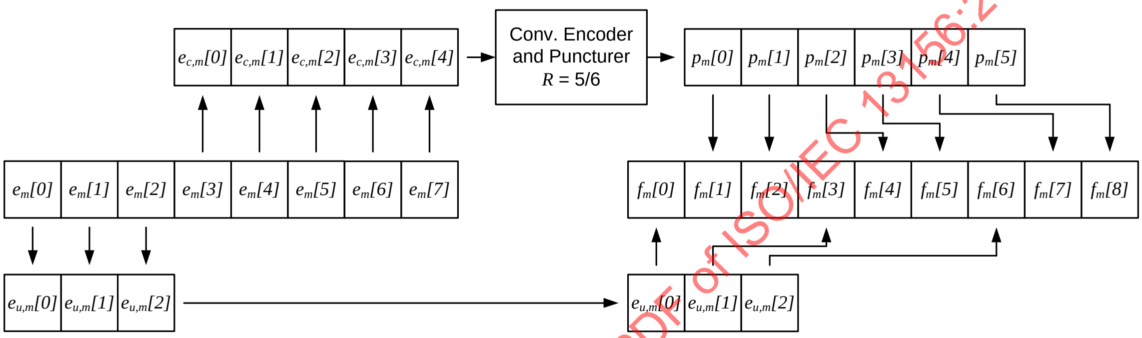

Figure 22 - An example of encoding for trellis coded modulation for mode A6

10.2.2.5.1.7 Multiplexer

The coded bits from the four branches $f_m[.]$, m=0,...,3 shall be multiplexed to create a single stream $g[.]$. The $g[.]$ shall be obtained from $f_m[.]$ from

$$g[k] = f_{\lfloor k/k_{MUX} \rfloor \bmod 4}\left[k \bmod k_{MUX} + k_{MUX}\left\lfloor \frac{k}{4k_{MUX}} \right\rfloor\right] \quad (22)$$

where k_{MUX} = 1 for modes A0 to A5, A7 and A9, k_{MUX} = 3 for mode A6, and k_{MUX} = 4 for mode A8.

10.2.2.5.1.8 Mapping

The data bits shall be mapped to the constellations according to the data rate mode. The constellation mappings are given in 10.2.4.

10.2.2.5.1.9 Time domain spreading (TDS)

The data symbols shall be consecutively repeated N_{TDS} times, that is the output of the spread symbols $s[.]$ shall be derived from the data symbols $v[.]$ by

$$s[n] = v\left[\left\lfloor \frac{n}{N_{TDS}} \right\rfloor\right] \quad (23)$$

where N_{TDS} is the time domain spreading factor (TDSF) defined for each mode in Table 48.

10.2.2.5.1.10 Symbol interleaver

The data symbols, $u[.]$, shall be interleaved using a 21 by 24 dual helical scan (DHS) interleaver. The DHS interleaver interleaves the data symbols by writing and reading the data symbols in a memory block with a helical scan pattern. Each block of 504 data symbols, $u[0],\ldots,u[503]$, shall be interleaved separately to form the interleaved symbols $w[0],\ldots,w[503]$.

The interleaving operation is described by defining the auxiliary symbols $r[m,n]$. The $r[m,n]$ are then given by $r[m,n] = u[k]$ where

$$m = \left[(k \bmod 21) + \left\lfloor \frac{k}{21} \right\rfloor \right] \bmod 24 \quad (24)$$

and

$$n = k \bmod 21 \quad (25)$$

and $w[k]$ are given by $w[k]=r[m,n]$, where

$$m = 23 - \left[(k \bmod 21) + \left\lfloor \frac{k}{21} \right\rfloor \right] \bmod 24 \quad (26)$$

and

$$n = k \bmod 21 \quad (27)$$

10.2.2.5.2 FEC and mapping (Unequal error protection)

Figure 23 depicts the general overview of the encoding and mapping scheme for modes A10 through A13. For the payload, first, the MSB and LSB bits of the each octet is separated as described in 10.2.2.5.2.1. Then, for each group, the data bits shall be scrambled as specified in 10.2.2.5.2.2. Then, $N_{padbits}$ shall be appended to the end of the data bit block, where

$$N_{padbits} = 1792 \left\lceil \frac{N_{bits}}{1792} \right\rceil - N_{bits} \quad (28)$$

and N_{bits} is the number of bits in the data bit block. The padded data bit block shall be encoded according to the Reed-Solomon code, as specified in 10.2.2.5.2.3.

The resulting RS coded payload bits or the formed header bits shall be demultiplexed into 4 groups, as described in 10.2.2.5.2.4. Depending on the data rate mode, each group of bits shall then be interleaved using the bit interleaver described in 10.2.2.5.2.5. The resulting bits shall be appended with N_{tail} tail bits (in each branch). Then they shall be encoded using the convolutional code described in 10.2.2.5.2.6. The resulting encoded bits from all 8 branches shall then be multiplexed together into one group of bits as specified in 10.2.2.5.2.7. The multiplexed bits are then mapped to constellations as specified in 10.2.2.5.2.8. The resulting data symbols shall then be appended with N_{padsym} zero symbols, where

$$N_{padsym} = 2N_D \left\lceil \frac{N_{sym}}{2N_D} \right\rceil - N_{sym} \quad (29)$$

At the end of each segment, the resulting data symbols shall be interleaved with the Dual Helical Scan (DHS) symbol interleaver as described in 10.2.2.5.2.9.

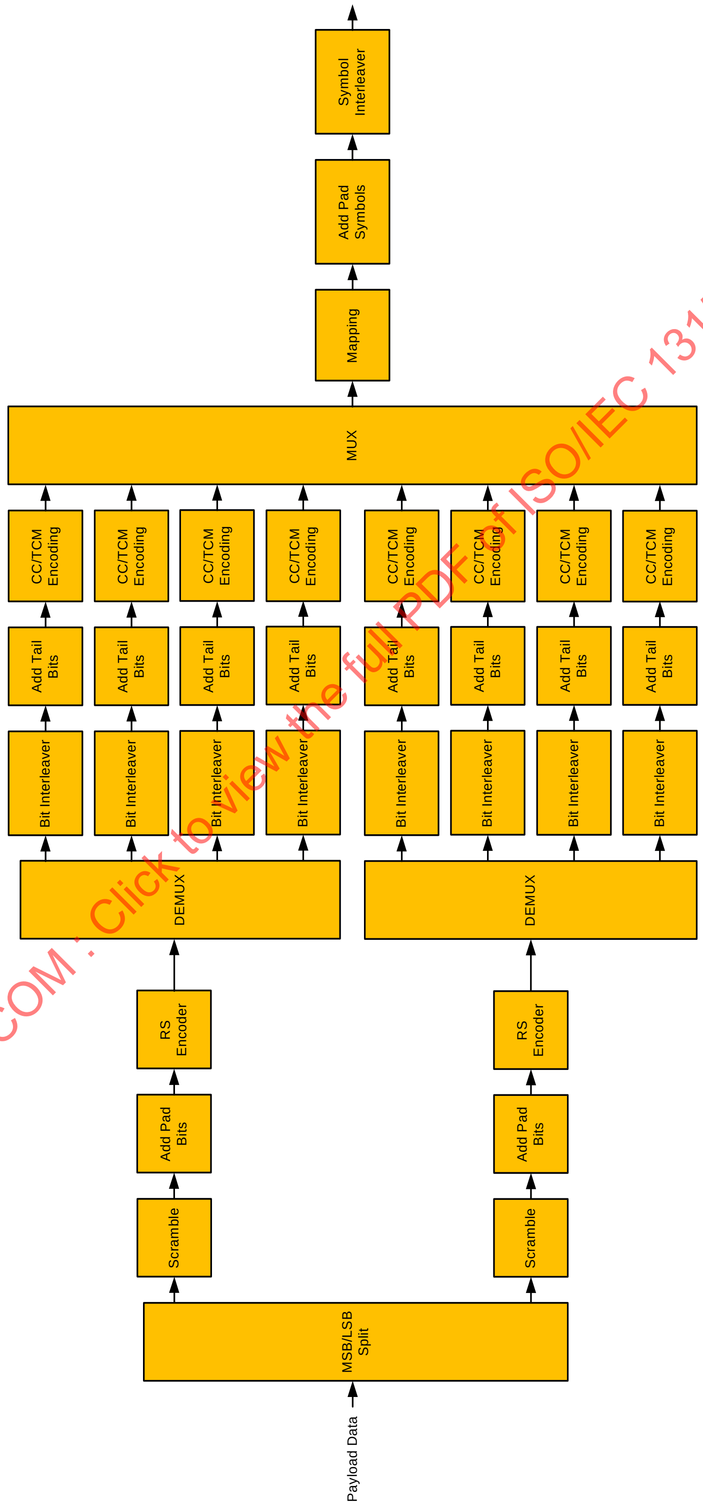

Figure 23 - General view of the encoding procedure

10.2.2.5.2.1 MSB/LSB separation

The incoming bit stream, $a[.]$, shall be demultiplexed into two streams, such that the MSBs and LSBs of each octet are separated. That is, the two bit streams $a_{LSB}[.]$ and $a_{MSB}[.]$ are created from the incoming bit stream $a[.]$ by

$$a_{LSB}[k] = a\left[n \bmod 4 + 8\left\lfloor \frac{k}{4} \right\rfloor\right] \tag{30}$$

and

$$a_{MSB}[k] = a\left[n \bmod 4 + 8\left\lfloor \frac{k}{4} \right\rfloor + 4\right] \tag{31}$$

10.2.2.5.2.2 Data scrambler

Refer to 10.2.2.5.1.1.

10.2.2.5.2.3 Reed-Solomon code

Refer to 10.2.2.5.1.2.

10.2.2.5.2.4 Demultiplexer

Refer to 10.2.2.5.1.3.

10.2.2.5.2.5 Bit interleaver

Refer to 10.2.2.5.1.4

10.2.2.5.2.6 Convolutional code

The convolutional encoding in each branch shall be as described in 10.2.2.5.1.5. For modes A10 and A11 the code rate for the branches 0 through 3 (MSB branches) shall be R_{MSB}, and the code rate for the branches 4 through 7 (LSB branches) shall be R_{LSB}, as given in Table 48. For mode A10 only MSBs shall be encoded, modulated and transmitted. The LSBs shall be discarded.

10.2.2.5.2.7 Multiplexer

For mode A10, the bits from the eight branches $f_m[.]$, m=0,...,7 shall be multiplexed to create a single stream as described in 10.2.2.5.1.7.

For mode A11, the bits from the eight branches $f_m[.]$, m=0,...,7 shall be multiplexed to create a single stream $g[.]$. The $g[.]$ shall be obtained from $f_m[.]$ from

$$g[k] = f_m\left[n\left\lfloor \frac{k}{48} \right\rfloor + p\right] \tag{32}$$

where m and n are given in Table 15 as a function of $k \bmod 48$.

Table 15 - Multiplexing parameters for mode A11

k mod 48	m	n	p	k mod 48	m	n	p	k mod 48	m	n	p
0	0	7	0	16	2	7	2	32	4	5	4
1	0	7	1	17	2	7	3	33	5	5	0
2	0	7	2	18	2	7	4	34	5	5	1
3	0	7	3	19	2	7	5	35	5	5	2
4	0	7	4	20	2	7	6	36	5	5	3
5	0	7	5	21	3	7	0	37	5	5	4

Table 15 - Multiplexing parameters for mode A11 (concluded)

k mod 48	m	n	p	k mod 48	m	n	p	k mod 48	m	n	p
6	0	7	6	22	3	7	1	38	6	5	0
7	1	7	0	23	3	7	2	39	6	5	1
8	1	7	1	24	3	7	3	40	6	5	2
9	1	7	2	25	3	7	4	41	6	5	3
10	1	7	3	26	3	7	5	42	6	5	4
11	1	7	4	27	3	7	6	43	7	5	0
12	1	7	5	28	4	5	0	44	7	5	1
13	1	7	6	29	4	5	1	45	7	5	2
14	2	7	0	30	4	5	2	46	7	5	3
15	2	7	1	31	4	5	3	47	7	5	4

For modes A12 and A13, the symbols from the eight branches $f_m[.]$, $m=0,...,7$ shall be multiplexed to create a single stream $g[.]$. The $g[.]$ shall be obtained from $f_m[.]$ from

$$g[k] = f_{\lfloor k/k_{MUX} \rfloor \bmod 4 + 4(\lfloor 2k/k_{MUX} \rfloor \bmod 2)}\left[k \bmod \frac{k_{MUX}}{2} + \frac{k_{MUX}}{2}\left\lfloor \frac{k}{4k_{MUX}} \right\rfloor \right] \quad (33)$$

where $k_{mux} = 2$ for mode A12, and $k_{mux} = 4$ for mode A13.

10.2.2.5.2.8 Mapping

Refer to 10.2.2.5.1.8.

10.2.2.5.2.9 Symbol interleaver

Refer to 10.2.2.5.1.10.

10.2.2.5.3 Block modulation

The multiplexed transmit symbols, $w[.]$ shall be formed into SCBT symbols. First the transmit data symbols shall be divided into blocks of length N_D=252. This transmit symbol block shall be appended with pilot symbol. The pilots symbols consist of a sequence of length N_P=4, $S_{pilot}[n]=(-1)^{n+1}$, n=0,1,...,3, to form an SCBT block. The resulting SCBT block shall be pre-appended with a cyclic prefix. This is done by prefixing the SCBT block with the last N_{CP} transmit symbols of the SCBT symbol. The cyclic prefix length, N_{CP}, shall be set based on the REQ_CP_LENGTH field of the last received frame from the current receiver. The N_{CP} may take the values 0, 32, 64 or 96. The N_{CP} for the first data frame or for data frames that have multiple destinations, e.g. beacon frames, is set to 96. Figure 24 depicts an example of the SCBT symbol formation.

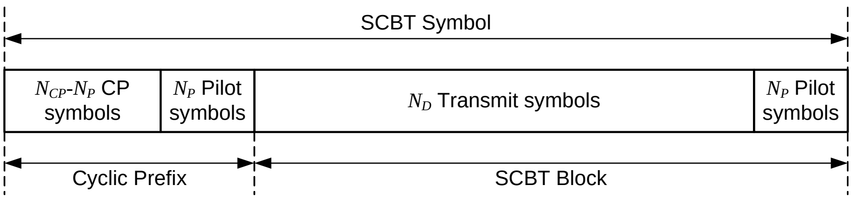

Figure 24 - Formation of the SCBT symbol

10.2.2.5.4 Midamble

The midamble sequence shall be identical to the channel estimation sequence, $S_{CE,A,SCBT}[.]$, as described in 10.2.2.3.2.

NOTE The midamble sequence can be used to update estimation of the channel frequency response, fine carrier frequency estimation, and fine symbol timing.

10.2.2.6 Channel bonding

Two, three or four adjacent channels may be bonded together. In this case, the formation of frames shall be as specified in 10.2.2.1 through 10.2.2.5. The symbol rate shall be increased according to Table 16.

Table 16 - Symbol rates for channel bonding

Number of bonded channels	f_{sym} (Gsps)
2	3.456
3	5.184
4	6.912

10.2.3 OFDM

10.2.3.1 Timing related parameters

The OFDM PHY timing related parameters are listed in Table 17.

Table 17 - OFDM timing related parameter

Parameters	Description	Value
f_{sym}	Symbol rate	2.592 Gsps
T_{sym}	Symbol time	0.386 ns
N_{FFT}	Number of subcarriers	512
T_{FFT}	FFT period	197.53 ns
N_D	Number of data carriers	360
N_{DC}	Number of DC carriers	3
N_P	Number of pilot carriers	16
N_N	Number of null carriers	133
N_{CP}	Cyclic prefix length	64
T_{CP}	Cyclic prefix duration	24.70 ns
$T_{sym,OFDM}$	OFDM symbol Duration	222.23 ns

10.2.3.2 Frame related parameters

The OFDM PHY frame related parameters are listed in Table 18.

Table 18 - OFDM frame related parameters

Parameters	Description	Value
N_{sync}	Number of symbols in frame synchronization sequence	1792
T_{sync}	Duration of the frame synchronization sequence	691.7 ns
N_{CE}	Number of symbols in the channel estimation sequence	1088

Table 18 - OFDM frame related parameters (concluded)

Parameters	Description	Value
T_{CE}	Duration of the channel estimation sequence	419.97 ns
$N_{preamble}$	Number of symbols in the PLCP preamble	2880
$T_{preamble}$	Duration of the frame preamble	1111.68 ns
N_{ATS}	Number of symbols in the ATS	$256(N_{\mathrm{TXTS}}+N_{\mathrm{RXTS}})N_{DISCREP}$
T_{ATS}	Duration of the ATS	$N_{ATS}T_{sym}$
N_{frame}	Number of Symbols in the frame	$N_{preamble}+N_{header}+N_{payload}+N_{ATS}$
T_{frame}	Duration of the frame	$N_{frame}T_{sym}$

10.2.3.3 PLCP preamble

The preamble includes 7 short training sequences and 2 long training sequences, which are used for frame synchronization and channel estimation, respectively. The short training sequence $S_{short,T}[.]$ is 256 symbols long, while the last sequence is rotated by 180 degree to indicate the end of short sequences and beginning of long sequences. The long training sequence $S_{long,T}[.]$ is 512 symbols long, with a cyclic prefix of length 64. The total length of the preamble is 5 OFDM symbols long. The preamble is normalized to unity.

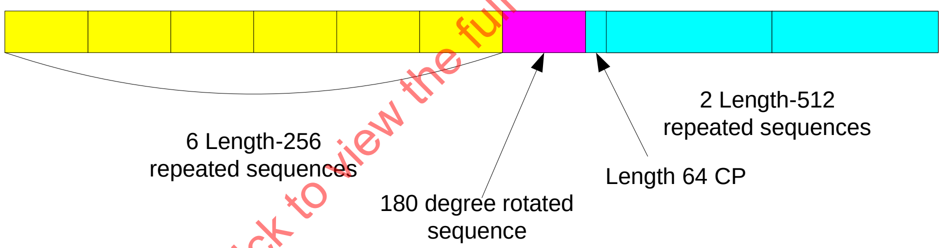

Figure 25 - PLCP preamble

The length-256 short training sequence shall be defined in the frequency domain using QPSK signal with inserted zeros. The frequency domain representation $S_{short,F}[.]$ is illustrated in Table 19.

Table 19 - Short training sequence in the frequency domain

n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$
0	0	0	128	1	-1	256	0	0	384	-1	-1
1	0	0	129	0	0	257	0	0	385	0	0
2	-1	1	130	-1	-1	258	0	0	386	1	-1
3	0	0	131	0	0	259	0	0	387	0	0
4	1	1	132	-1	1	260	0	0	388	-1	1
5	0	0	133	0	0	261	0	0	389	0	0
6	1	-1	134	-1	1	262	0	0	390	-1	-1
7	0	0	135	0	0	263	0	0	391	0	0

Table 19 - Short training sequence in the frequency domain (continued)

n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$
8	1	-1	136	1	-1	264	0	0	392	-1	1
9	0	0	137	0	0	265	0	0	393	0	0
10	1	-1	138	-1	1	266	0	0	394	1	-1
11	0	0	139	0	0	267	0	0	395	0	0
12	-1	1	140	1	1	268	0	0	396	-1	1
13	0	0	141	0	0	269	0	0	397	0	0
14	1	-1	142	-1	1	270	0	0	398	-1	-1
15	0	0	143	0	0	271	0	0	399	0	0
16	1	-1	144	1	1	272	0	0	400	-1	-1
17	0	0	145	0	0	273	0	0	401	0	0
18	-1	1	146	1	-1	274	0	0	402	1	1
19	0	0	147	0	0	275	0	0	403	0	0
20	-1	1	148	1	-1	276	0	0	404	-1	-1
21	0	0	149	0	0	277	0	0	405	0	0
22	-1	-1	150	1	1	278	0	0	406	1	1
23	0	0	151	0	0	279	0	0	407	0	0
24	1	1	152	1	-1	280	0	0	408	-1	1
25	0	0	153	0	0	281	0	0	409	0	0
26	-1	1	154	1	1	282	0	0	410	-1	1
27	0	0	155	0	0	283	0	0	411	0	0
28	1	-1	156	-1	1	284	0	0	412	1	1
29	0	0	157	0	0	285	0	0	413	0	0
30	-1	-1	158	-1	1	286	0	0	414	-1	1
31	0	0	159	0	0	287	0	0	415	0	0
32	-1	-1	160	1	-1	288	0	0	416	-1	-1
33	0	0	161	0	0	289	0	0	417	0	0
34	1	1	162	-1	-1	290	0	0	418	-1	-1
35	0	0	163	0	0	291	0	0	419	0	0
36	-1	-1	164	1	1	292	0	0	420	-1	-1
37	0	0	165	0	0	293	0	0	421	0	0
38	1	-1	166	1	-1	294	0	0	422	1	1
39	0	0	167	0	0	295	0	0	423	0	0
40	1	1	168	-1	1	296	0	0	424	1	-1
41	0	0	169	0	0	297	0	0	425	0	0
42	1	1	170	1	1	298	0	0	426	-1	1
43	0	0	171	0	0	299	0	0	427	0	0
44	-1	1	172	-1	1	300	0	0	428	1	-1

Table 19 - Short training sequence in the frequency domain (continued)

n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$
45	0	0	173	0	0	301	0	0	429	0	0
46	1	-1	174	-1	-1	302	0	0	430	1	1
47	0	0	175	0	0	303	0	0	431	0	0
48	-1	-1	176	1	-1	304	0	0	432	1	-1
49	0	0	177	0	0	305	0	0	433	0	0
50	1	-1	178	1	1	306	0	0	434	-1	1
51	0	0	179	0	0	307	0	0	435	0	0
52	1	1	180	0	0	308	0	0	436	-1	1
53	0	0	181	0	0	309	0	0	437	0	0
54	-1	-1	182	0	0	310	0	0	438	-1	-1
55	0	0	183	0	0	311	0	0	439	0	0
56	-1	-1	184	0	0	312	0	0	440	-1	-1
57	0	0	185	0	0	313	0	0	441	0	0
58	1	-1	186	0	0	314	0	0	442	-1	-1
59	0	0	187	0	0	315	0	0	443	0	0
60	1	-1	188	0	0	316	0	0	444	-1	-1
61	0	0	189	0	0	317	0	0	445	0	0
62	-1	1	190	0	0	318	0	0	446	1	-1
63	0	0	191	0	0	319	0	0	447	0	0
64	1	1	192	0	0	320	0	0	448	-1	1
65	0	0	193	0	0	321	0	0	449	0	0
66	-1	-1	194	0	0	322	0	0	450	1	1
67	0	0	195	0	0	323	0	0	451	0	0
68	1	-1	196	0	0	324	0	0	452	-1	1
69	0	0	197	0	0	325	0	0	453	0	0
70	1	-1	198	0	0	326	0	0	454	1	1
71	0	0	199	0	0	327	0	0	455	0	0
72	-1	-1	200	0	0	328	0	0	456	-1	1
73	0	0	201	0	0	329	0	0	457	0	0
74	-1	-1	202	0	0	330	0	0	458	1	-1
75	0	0	203	0	0	331	0	0	459	0	0
76	1	-1	204	0	0	332	0	0	460	-1	1
77	0	0	205	0	0	333	0	0	461	0	0
78	1	1	206	0	0	334	1	1	462	1	1
79	0	0	207	0	0	335	0	0	463	0	0
80	1	-1	208	0	0	336	1	1	464	1	-1
81	0	0	209	0	0	337	0	0	465	0	0

Table 19 - Short training sequence in the frequency domain (continued)

n	$\text{Re}\{S_{short,F}[n]\}$	$\text{Im}\{S_{short,F}[n]\}$	n	$\text{Re}\{S_{short,F}[n]\}$	$\text{Im}\{S_{short,F}[n]\}$	n	$\text{Re}\{S_{short,F}[n]\}$	$\text{Im}\{S_{short,F}[n]\}$	n	$\text{Re}\{S_{short,F}[n]\}$	$\text{Im}\{S_{short,F}[n]\}$
82	1	-1	210	0	0	338	-1	-1	466	1	1
83	0	0	211	0	0	339	0	0	467	0	0
84	1	-1	212	0	0	340	1	-1	468	-1	-1
85	0	0	213	0	0	341	0	0	469	0	0
86	1	-1	214	0	0	342	1	1	470	1	-1
87	0	0	215	0	0	343	0	0	471	0	0
88	-1	-1	216	0	0	344	-1	1	472	1	-1
89	0	0	217	0	0	345	0	0	473	0	0
90	1	1	218	0	0	346	-1	-1	474	1	-1
91	0	0	219	0	0	347	0	0	475	0	0
92	1	1	220	0	0	348	-1	1	476	-1	1
93	0	0	221	0	0	349	0	0	477	0	0
94	-1	-1	222	0	0	350	1	1	478	1	-1
95	0	0	223	0	0	351	0	0	479	0	0
96	-1	-1	224	0	0	352	1	-1	480	1	-1
97	0	0	225	0	0	353	0	0	481	0	0
98	1	-1	226	0	0	354	1	-1	482	1	-1
99	0	0	227	0	0	355	0	0	483	0	0
100	-1	-1	228	0	0	356	-1	-1	484	-1	-1
101	0	0	229	0	0	357	0	0	485	0	0
102	-1	-1	230	0	0	358	-1	-1	486	1	1
103	0	0	231	0	0	359	0	0	487	0	0
104	1	-1	232	0	0	360	1	1	488	1	-1
105	0	0	233	0	0	361	0	0	489	0	0
106	1	-1	234	0	0	362	-1	-1	490	1	1
107	0	0	235	0	0	363	0	0	491	0	0
108	-1	1	236	0	0	364	-1	-1	492	1	1
109	0	0	237	0	0	365	0	0	493	0	0
110	-1	1	238	0	0	366	1	-1	494	1	-1
111	0	0	239	0	0	367	0	0	495	0	0
112	-1	-1	240	0	0	368	-1	1	496	-1	1
113	0	0	241	0	0	369	0	0	497	0	0
114	1	1	242	0	0	370	1	1	498	-1	1
115	0	0	243	0	0	371	0	0	499	0	0
116	-1	-1	244	0	0	372	1	1	500	1	-1
117	0	0	245	0	0	373	0	0	501	0	0
118	1	1	246	0	0	374	-1	1	502	1	-1

Table 19 - Short training sequence in the frequency domain (concluded)

n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$	n	$\mathrm{Re}\{S_{short,F}[n]\}$	$\mathrm{Im}\{S_{short,F}[n]\}$
119	0	0	247	0	0	375	0	0	503	0	0
120	-1	1	248	0	0	376	-1	-1	504	-1	-1
121	0	0	249	0	0	377	0	0	505	0	0
122	1	-1	250	0	0	378	-1	1	506	1	-1
123	0	0	251	0	0	379	0	0	507	0	0
124	1	-1	252	0	0	380	1	-1	508	-1	-1
125	0	0	253	0	0	381	0	0	509	0	0
126	-1	1	254	0	0	382	1	1	510	-1	-1
127	0	0	255	0	0	383	0	0	511	0	0

The short training sequence, $S_{short,T}[.]$ is then determined by the inverse FFT of the $S_{short,F}[.]$. Specifically, let

$$S_{short,T}[m] = \sum_{k=0}^{511} \frac{S_{short,F}[k]}{\sqrt{356}} e^{\frac{j2\pi}{512}km} \tag{34}$$

where $S_{short,T}[m]$ represents the m^{th} element of the sequence $S_{short,T}[m]$ and $S_{short,F}[m]$ represents the k^{th} element of the sequence $S_{short,F}[m]$. Note that the $\sqrt{356}$ in (111) is the normalization factor due to the 356 non-zero QPSK symbols in the frequency representation.

The overall short training sequence shall be of length 1792 and be given by

$$S_{sync,A,OFDM}[n] = \begin{cases} T_{short,T}[(n \bmod 256)+1] & n = 0,\ldots,1535 \\ -T_{short,T}[(n \bmod 256)+1] & n = 1536,\ldots,1791 \end{cases} \tag{35}$$

The length 512 long training sequence is generated also in frequency domain using BPSK signal. The frequency domain representation $S_{long,F}[.]$ is illustrated in Table 20.

Table 20 - Frequency domain long training sequence

n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$
0	0	128	-1	256	0	384	-1
1	0	129	1	257	1	385	1
2	1	130	1	258	0	386	1
3	-1	131	-1	259	0	387	1
4	-1	132	1	260	0	388	-1
5	1	133	-1	261	0	389	1

Table 20 - Frequency domain long training sequence (continued)

n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$
6	1	134	-1	262	0	390	-1
7	1	135	-1	263	0	391	-1
8	1	136	1	264	0	392	-1
9	1	137	-1	265	0	393	1
10	1	138	-1	266	0	394	1
11	-1	139	-1	267	0	395	-1
12	-1	140	1	268	0	396	1
13	-1	141	1	269	0	397	1
14	-1	142	-1	270	0	398	1
15	1	143	-1	271	0	399	1
16	-1	144	-1	272	0	400	-1
17	1	145	1	273	0	401	-1
18	1	146	1	274	0	402	1
19	1	147	-1	275	0	403	-1
20	-1	148	1	276	0	404	-1
21	-1	149	-1	277	0	405	1
22	-1	150	1	278	0	406	-1
23	1	151	-1	279	0	407	1
24	1	152	-1	280	0	408	-1
25	1	153	-1	281	0	409	-1
26	1	154	1	282	0	410	-1
27	-1	155	-1	283	0	411	1
28	1	156	-1	284	0	412	-1
29	1	157	-1	285	0	413	-1
30	1	158	-1	286	0	414	1
31	1	159	1	287	0	415	-1
32	1	160	1	288	0	416	-1
33	-1	161	1	289	0	417	1
34	1	162	-1	290	0	418	-1
35	-1	163	-1	291	0	419	1
36	1	164	-1	292	0	420	-1
37	-1	165	1	293	0	421	1
38	1	166	-1	294	0	422	-1
39	1	167	1	295	0	423	1
40	1	168	1	296	0	424	-1
41	-1	169	-1	297	0	425	-1

Table 20 - Frequency domain long training sequence (continued)

n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$
42	-1	170	1	298	0	426	1
43	-1	171	1	299	0	427	-1
44	1	172	1	300	0	428	1
45	1	173	-1	301	0	429	1
46	1	174	1	302	0	430	1
47	-1	175	1	303	0	431	1
48	1	176	1	304	0	432	1
49	1	177	1	305	0	433	-1
50	1	178	-1	306	0	434	-1
51	-1	179	0	307	0	435	-1
52	-1	180	0	308	0	436	-1
53	1	181	0	309	0	437	1
54	1	182	0	310	0	438	-1
55	-1	183	0	311	0	439	-1
56	-1	184	0	312	0	440	-1
57	-1	185	0	313	0	441	1
58	1	186	0	314	0	442	-1
59	-1	187	0	315	0	443	1
60	1	188	0	316	0	444	1
61	-1	189	0	317	0	445	-1
62	-1	190	0	318	0	446	-1
63	-1	191	0	319	0	447	1
64	-1	192	0	320	0	448	1
65	1	193	0	321	0	449	-1
66	-1	194	0	322	0	450	-1
67	-1	195	0	323	0	451	-1
68	-1	196	0	324	0	452	1
69	-1	197	0	325	0	453	-1
70	1	198	0	326	0	454	-1
71	1	199	0	327	0	455	1
72	1	200	0	328	0	456	-1
73	-1	201	0	329	0	457	1
74	1	202	0	330	0	458	-1
75	-1	203	0	331	0	459	-1
76	1	204	0	332	0	460	-1
77	1	205	0	333	0	461	-1

Table 20 - Frequency domain long training sequence (continued)

n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$
78	1	206	0	334	1	462	1
79	-1	207	0	335	1	463	-1
80	-1	208	0	336	-1	464	-1
81	-1	209	0	337	-1	465	-
82	-1	210	0	338	-1	466	-1
83	1	211	0	339	1	467	1
84	-1	212	0	340	1	468	-1
85	1	213	0	341	1	469	1
86	-1	214	0	342	1	470	-1
87	1	215	0	343	1	471	-1
88	-1	216	0	344	-1	472	-1
89	-1	217	0	345	-1	473	1
90	1	218	0	346	1	474	-1
91	1	219	0	347	1	475	-1
92	-1	220	0	348	-1	476	1
93	-1	221	0	349	-1	477	1
94	1	222	0	350	-1	478	1
95	1	223	0	351	1	479	1
96	1	224	0	352	1	480	1
97	-1	225	0	353	1	481	-1
98	1	226	0	354	-1	482	-1
99	-1	227	0	355	-1	483	-1
100	1	228	0	356	-1	484	1
101	1	229	0	357	1	485	-1
102	-1	230	0	358	1	486	1
103	-1	231	0	359	-1	487	1
104	-1	232	0	360	1	488	1
105	-1	233	0	361	1	489	-1
106	1	234	0	362	1	490	1
107	-1	235	0	363	-1	491	-1
108	-1	236	0	364	-1	492	1
109	-1	237	0	365	1	493	-1
110	1	238	0	366	1	494	1
111	-1	239	0	367	1	495	-1
112	1	240	0	368	-1	496	1
113	1	241	0	369	-1	497	1

Table 20 - Frequency domain long training sequence (concluded)

n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$	n	$S_{long,F}[n]$
114	1	242	0	370	1	498	1
115	1	243	0	371	-1	499	1
116	-1	244	0	372	1	500	-1
117	-1	245	0	373	-1	501	-1
118	-1	246	0	374	1	502	-1
119	-1	247	0	375	1	503	1
120	-1	248	0	376	1	504	1
121	-1	249	0	377	-1	505	-1
122	1	250	0	378	-1	506	-1
123	-1	251	0	379	-1	507	-1
124	1	252	0	380	-1	508	1
125	1	253	0	381	-1	509	-1
126	1	254	0	382	-1	510	1
127	1	255	0	383	1	511	0

The long training sequence is determined by the inverse FFT of the $S_{long,F}$ given by

$$S_{long,T}[m] = \sum_{k=0}^{511} \frac{S_{long,F}[k]}{\sqrt{354}} e^{\frac{j2\pi}{512}km} \tag{36}$$

where $S_{long,T}[m]$ represents the m^{th} element of the sequence $S_{long,T}[.]$ and $S_{long,F}[.]$ represents the k^{th} element of the sequence $S_{long,F}[.]$. Note that the $\sqrt{354}$ in equation (36) is the normalization factor due to the 354 non-zero QPSK symbols in the frequency representation.

10.2.3.4 PLCP header

For Type A OFDM frames, the formed PLCP header has a length of N_{header}. The PLCP header shall be first padded with zeros where the number of zero bits is 24. The padded header shall be demultiplexed and then shall be encoded with four encoders of rate 2/3 convolutional coding (labelled A through D) in the MSB branch before QPSK modulation (see 10.2.3.5.1.10), OFDM symbol padding (see 10.2.3.5.1.11), tone interleaver (see 10.2.3.5.1.13) and OFDM modulation (see 10.2.3.5.1.12), as illustrated in Figure 26. Let A(0), A(1), …, A(n), … be the bit stream at the output of encoder A with A(0) being the first bit, B(0), B(1), …, B(n), … be the bit stream at the output of encoder B with B(0) being the first bit, let C(0), C(1), …, C(n), … be the bit streams at the output of encoder C with C(0) being the first bit, D(0), D(1), …, D(n), … be the bit stream at the output of encoder D with D(0) being the first bit, the multiplexer in Figure 26 shall be implemented such that the output bit stream of the multiplexer is A(0), B(0), C(0), D(0), A(1), B(1), C(1), D(1), ..., A(n), B(n), C(n), D(n),... with A(0) being the first bit.

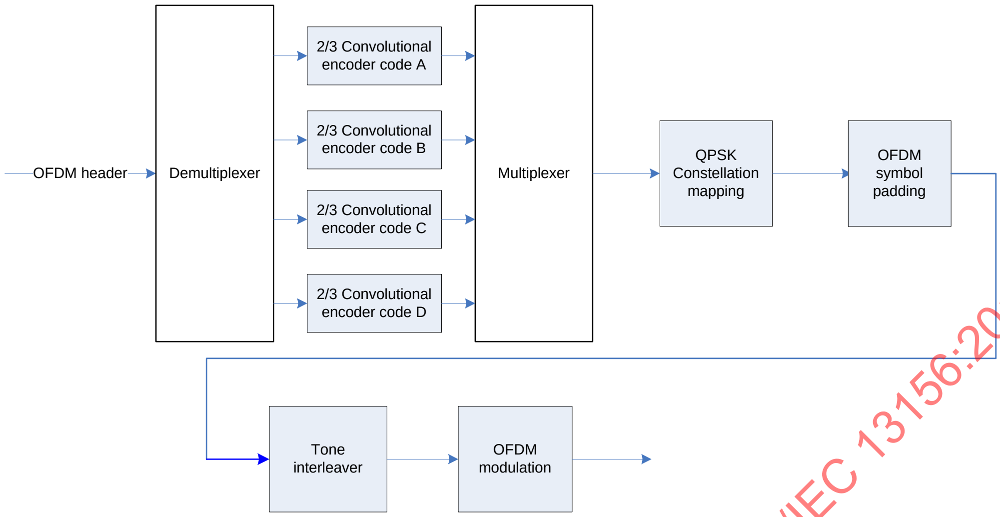

Figure 26 - OFDM block diagram for header operation

10.2.3.4.1 PHY header

10.2.3.4.1.1 Fixed length PHY header

10.2.3.4.1.1.1 PLCP SCRAMBLER FIELD (SCRAMBLER)

Refer to 10.2.2.4.1.1.1.

10.2.3.4.1.1.2 Bit reversal field (BIT_REVERSAL)

The BIT_REVERSAL bit shall be set to zero for all Type A OFDM frames.

10.2.3.4.1.1.3 ATIF existence field (ATIF_EX)

Refer to 10.2.2.4.1.1.3.

10.2.3.4.1.1.4 CP length field (CP_LENGTH)

The CP_LENGTH field D_1D_0 shall be set to 10_B, which corresponds to the cyclic prefix length of 64 for Type A OFDM.

10.2.3.4.1.1.5 Requested CP length field (REQ_CP_LENGTH)

The REQ_CP_LENGTH field D_1D_0 shall be set to 10_B.

10.2.3.4.1.1.6 Number of segments (NUM_SEGMENTS)

Refer to 10.2.2.4.1.1.6.

10.2.3.4.1.1.7 Number of MSDUs (NUM_MSDUS)

Refer to 10.2.2.4.1.1.7.

10.2.3.4.1.2 Variable length PHY header

10.2.3.4.1.2.1 Mode field (MODE)

Refer to 10.2.2.4.1.2.1.

10.2.3.4.1.2.2 Segment length field (LENGTH)

Refer to 10.2.2.4.1.2.2.

10.2.3.4.1.2.3 Midamble existence field (MID_EX)

Refer to 10.2.2.4.1.2.3.

10.2.3.4.1.2.4 Segment continued field (CONT)

Refer to 10.2.2.4.1.2.4.

10.2.3.4.1.3 Antenna training indicator field

10.2.3.4.1.3.1 Discovery mode repetition field (DISC_REP)

Refer to 10.2.2.4.1.3.1.

10.2.3.4.1.3.2 Number of training symbols for RX training field (NUM_RXTS)

Refer to 10.2.2.4.1.3.2.

10.2.3.4.1.3.3 Number of training symbols for TX training field (NUM_TXTS)

Refer to 10.2.2.4.1.3.3.

10.2.3.5 PPDU payload

The PPDU payload consists of one or more segments as described in 10.1. The segments are formed as described in 10.1.3. Each data bit block shall be processed as follows:

1. The bits shall be encoded and mapped as described in 10.2.3.5.1.

2. The resulting symbols shall be modulated into OFDM symbol as described in 10.2.3.5.1.12.

10.2.3.5.1 FEC and mapping

The block diagram of the OFDM PHY for modes A14 to A21 is shown in Figure 27. The information bits are first split into two streams as described in 10.2.3.5.1.1. Then the bits shall be scrambled as described in 10.2.3.5.1.2. Then, $N_{padbits,OFDM}$ zero pad bits shall be appended to the end of the data bit block as described in 10.2.3.5.1.3. The padded data bit block shall be encoded according to the Reed-Solomon code as described in 10.2.3.5.1.4.

The resulting RS coded payload bits shall be interleaved by the outer interleaver as described in 10.2.3.5.1.5. The resulting bits shall be encoded using eight parallel convolutional encoders as described in 10.2.3.5.1.7. Depending on the data rate mode, puncturing may be performed as described in 10.2.3.5.1.8.

For EEP modes (A14 to A16), the encoded bits from the 8 branches shall be multiplexed into a single data stream as described in 10.2.3.5.1.9.1 and interleaved as described in 10.2.3.5.1.9.2. For UEP coding (A17 and A18), the encoded bits from the 8 branches shall be multiplexed into a single data stream as described in 10.2.3.5.1.9.3 and interleaved as described in 10.2.3.5.1.9.4. For UEP mapping (A19 and A20), the encoded bits from the 8 branches shall be multiplexed into a single data stream as described in 10.2.3.5.1.9.5 and interleaved as described in 10.2.3.5.1.9.6. For MSB-only transmission (A21), the upper branch that consists of MSB RS encoder, outer interleaver, demultiplexer and convolutional encoders A through D as shown in Figure 27 shall be used. The data MUX shall only use the bits from the upper branch.

The resulting bits shall then be mapped to the appropriate constellation as described in 10.2.3.5.1.10 prior to the bit reversal tone interleaver as described in 10.2.3.5.1.13. The QPSK and 16 QAM symbols shall be modulated with OFDM modulation as described in 10.2.3.5.1.12.

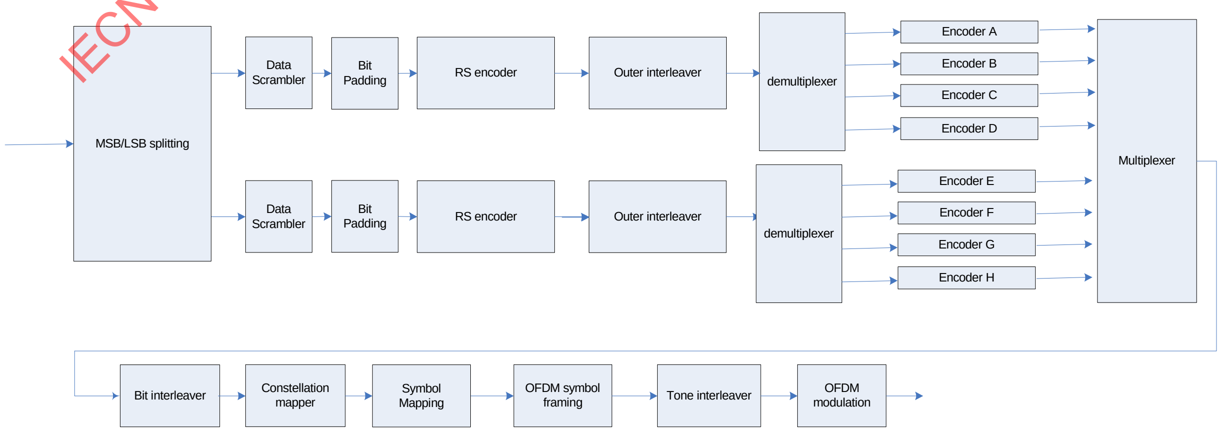

Figure 27 - Block diagram of the OFDM PHY baseband

10.2.3.5.1.1 MSB/LSB splitting

Refer to 10.2.2.5.2.1.

10.2.3.5.1.2 Data scrambler

Refer to 10.2.2.5.1.1.

10.2.3.5.1.3 Bit padding

$N_{padbits,OFDM}$ zero pad bits shall be appended to the end of the data bit block, where

$$N_{padbits,OFDM} = 1792\left\lceil \frac{N_{bits,OFDM}}{1792} \right\rceil - N_{bits,OFDM} \tag{37}$$

and $N_{bits,OFDM}$ is the number of bits in the data bit block. The padded data bit block shall be encoded according to the Reed-Solomon code, as specified in 10.2.3.5.1.4.

10.2.3.5.1.4 Reed-Solomon encoder

Refer to 10.2.2.5.1.2.

10.2.3.5.1.5 Outer interleaver

Let N_{OPi} be the interleaver size in bits. Interleaving encoders shall have an interleaving depth 4, leading to an interleaver size of N_{OPi} = $32N_{RS}$ = 8064 bits, where N_{RS} is number of bytes in a RS codeword.

Let $X = 0,\ldots,N_{OPi}\text{-}1$ be the ordering of the original bits entering the interleaver. Let $P = 2^{\lceil \log_2 N_{OPi} \rceil}$ be the smallest power of 2 that is greater than or equal to the interleaver size and let $S = 0,\ldots,P\text{-}1$. The ordering of bits leaving the interleaver shall be defined as the permutation

$$Y = prune\left(\left(\frac{S(S+1)}{2}\right) \bmod P\right) \tag{38}$$

where $prune(.)$ is the operation that removes all elements that are larger than $N_{OPi}\text{-}1$ from the input sequence without changing the order of the remaining elements.

10.2.3.5.1.6 Demultiplexing

For the first outer interleaver which accepts inputs from the first RS encoder, the output bit streams shall be demultiplexed to encoders A, B, C, D (see 10.2.3.5.1.7) in a round robin manner. Let bit 0 be the first incoming bit in the bit stream. With i=0, 1…,n, …, the bit streams at the input of encoder A shall be b[3], b[7], …, b[4n+3] where b[3] is the first incoming bit at the input of encoder A. The bit streams at the input of encoder B shall be b[2], b[6], …, b[4n+2] where b[2] is the first incoming bit at the input of encoder B. The bit streams at the input of encoder C shall be b[1], b[5], …, b[4n+1] where b[1] is the first incoming bit at the input of encoder C. The bit streams at the input of encoder D shall be b[0], b[4], …, b[4n] where b[0] is the first incoming bit at the input of encoder D.

For the second outer interleaver which accepts inputs from the second RS encoder, the output bit streams shall be demultiplexed to encoders E, F, G, H (see 10.2.3.5.1.7) in a round robin manner. Let bit 0 be the first incoming bit in the bit stream. With i=0, 1…,n, …, the bit streams at the input of encoder E shall be b[3], b[7], …, b[4n+3] where b[3] is the first incoming bit at the input of encoder E. The bit streams at the input of encoder F shall be b[2], b[6], …, b[4n+2] where b[2] is the first incoming bit at the input of encoder F. The bit streams at the input of encoder G shall be b[1], b[5], …, b[4n+1] where b[1] is the first incoming bit at the input of encoder G. The bit streams at the input of encoder H shall be b[0], b[4], …, b[4n] where b[0] is the first incoming bit at the input of encoder H.

10.2.3.5.1.7 Parallel convolutional codes

A parallel of 8 convolutional encoders shall be used, as depicted in Figure 28. The OFDM PHY transmitter uses eight parallel convolutional encoders, labelled A through H. The first four encoders, labelled A through D, is for the first outer Reed-Solomon coding branch and the last four encoders, labelled E through H is for the second outer Reed-Solomon coding branch.

For each of the 8 parallel convolutional codes, a mother code of rate 1/3 shall be used. The convolutional encoder shall use constraint length K = 7, delay memory 6, generator polynomial $g_0 = 133_O$, $g_1 = 171_O$, $g_2 = 165_O$, mother code rate 1/3. A detailed schematic diagram of the convolutional encoder is shown in Figure 29. The initial value of the delay register shall be zero at the beginning of every OFDM frame.

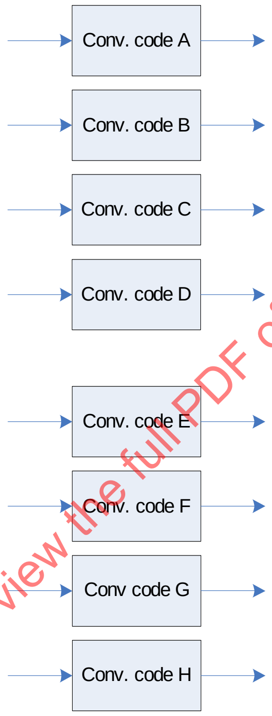

Figure 28 - Illustration of parallel convolutional codes

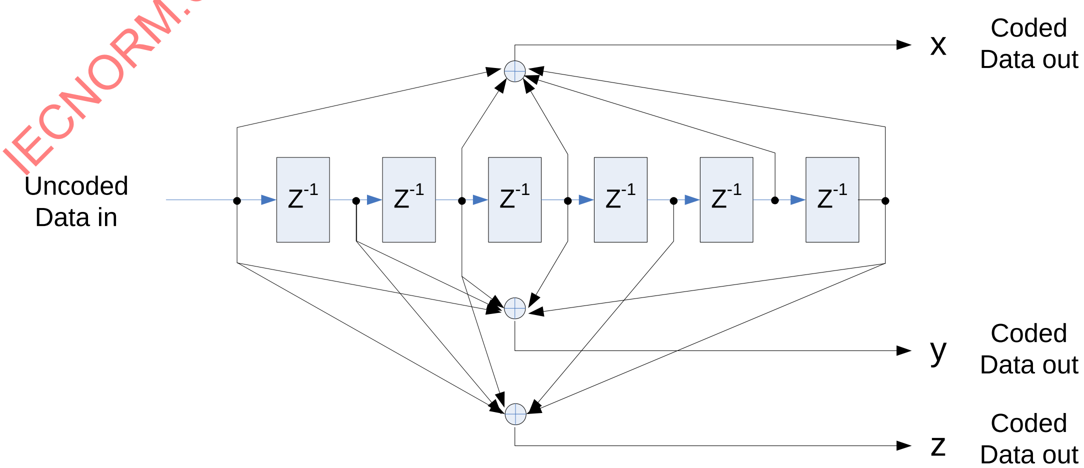

Figure 29 - Mother convolutional code generator

10.2.3.5.1.8 Code puncturing

Convolutional encoded data is punctured to generate code rates 4/7, 2/3, and 4/5. The puncturing patterns are illustrated in Figure 30, Figure 31, and Figure 32. In each of these cases, the tables shall be filled in with encoder output bits from left to right. For the last block of bits, the process shall be stopped at the point at which encoder output bits are exhausted and the puncturing pattern is applied to the partially filled block.

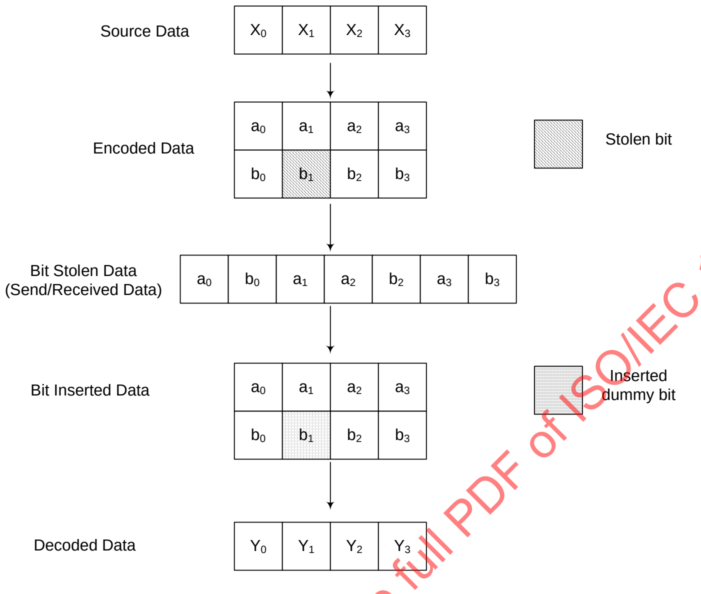

Figure 30 - An example of bit stealing and bit insertion for R=4/7 code

Figure 31 - An example of bit stealing and bit insertion for R=2/3 code

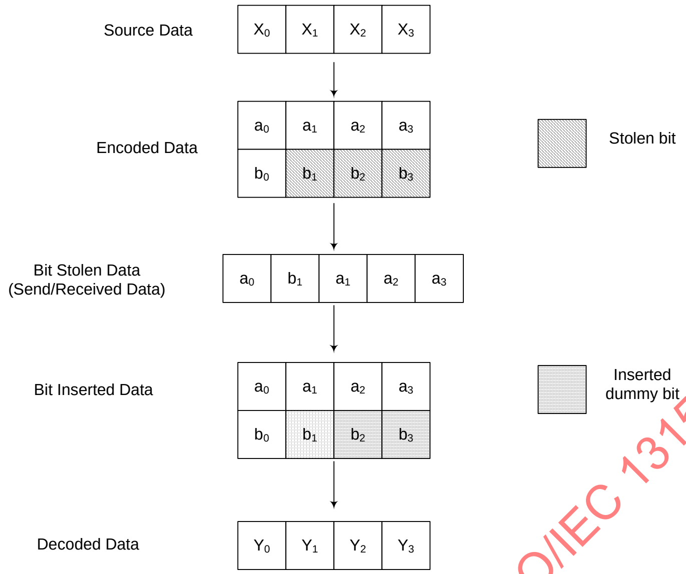

Figure 32 - An example of bit stealing and bit insertion for R=4/5 code

10.2.3.5.1.9 Data multiplexer and bit interleaver

The output of the 8 encoders, labelled A through H, shall be multiplexed to a single data stream prior to the bit interleaver, as illustrated in Figure 33. The method used to multiplex the encoded bits is dependent on the type of OFDM PHY mode.

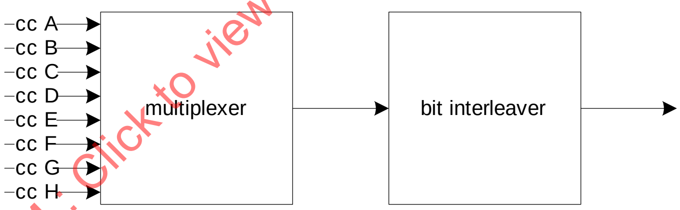

Figure 33 - Illustration of the multiplexer and bit interleaver

10.2.3.5.1.9.1 Multiplexer in EEP modes

In the EEP modes (A14 to A16), all the 8 encoders shall use the same code rate. The encoded bits shall be multiplexed and bit-interleaved every 48 bits. During the length 48 multiplexing/interleaving cycle, a group multiplexer shall be used first with fixed group size 6 for all eight encoders. Use A(1), A(2), A(3), A(4), A(5), A(6) to label the 6 encoded bits (in increasing order in time) from encoder A, and similarly B(1) through B(6), C(1) through C(6), D(1) through D(6), E(1) through E(6), F(1) through F(6), G(1) through G(6), and H(1) through H(6) from encoders B, C, D, E, F, G, and H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 21, where the numbering for each bit represents the position of the bit in the bit stream after the multiplexer and bit 0 is the first bit at the output of the multiplexer.

Table 21 - Input pattern for the EEP modes (A14 to A16)

Numbering	0	1	2	3	4	5	6	7	8	9	10	11
Labelling	A(1)	A(2)	A(3)	A(4)	A(5)	A(6)	B(1)	B(2)	B(3)	B(4)	B(5)	B(6)
Numbering	12	13	14	15	16	17	18	19	20	21	22	23

Table 21 - Input pattern for the EEP modes (A14 to A16) (concluded)

Labelling	C(1)	C(2)	C(3)	C(4)	C(5)	C(6)	D(1)	D(2)	D(3)	D(4)	D(5)	D(6)
Numbering	24	25	26	27	28	29	30	31	32	33	34	35
Labelling	E(1)	E(2)	E(3)	E(4)	E(5)	E(6)	F(1)	F(2)	F(3)	F(4)	F(5)	F(6)
Numbering	36	37	38	39	40	41	42	43	44	45	46	47
Labelling	G(1)	G(2)	G(3)	G(4)	G(5)	G(6)	H(1)	H(2)	H(3)	H(4)	H(5)	H(6)

10.2.3.5.1.9.2 Bit interleaver in EEP modes

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \ldots, 47$ and $y = 0, \ldots, 47$ be the index at the input and output of the bit interleaver, respectively. The bit interleaver in the EEP mode shall implement the following relation:

$$y = \left[6 \left\lfloor \frac{x}{6} \right\rfloor - 5(x \bmod 6) \right] \bmod 48 \qquad (39)$$

10.2.3.5.1.9.3 Multiplexer in UEP coding modes

In the UEP coding mode, top 4 encoders (or encoders A, B, C, D) shall use rate 4/7 convolutional codes, and bottom 4 encoders (or encoders E, F, G, H) shall use rate 4/5 convolutional codes. The encoded bits shall be multiplexed and bit interleaved every 48 bits. A group multiplexer with group size 7, 7, 7, 7, 5, 5, 5, 5 for all eight encoders, A through G, respectively, shall be used. Use A(1), A(2), A(3), A(4), A(5), A(6), A(7) to label the 7 encoded bits (in increasing order in time) from encoder A, and similarly B(1) through B(7), C(1) through C(7), D(1) through D(7), E(1) through E(5), F(1) through F(5), G(1) through G(5), and H(1) through H(5) from encoders B, C, D, E, F, G, and H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 22, where the numbering for each bit represents the position of the bit in the bit stream after the multiplexer and bit 0 is the first bit at the output of the multiplexer.

Table 22 - Input pattern for the UEP coding modes

Numbering	0	1	2	3	4	5	6	7	8	9	10	11
Labelling	A(1)	A(2)	A(3)	A(4)	A(5)	A(6)	A(7)	B(1)	B(2)	B(3)	B(4)	B(5)
Numbering	12	13	14	15	16	17	18	19	20	21	22	23
Labelling	B(6)	B(7)	C(1)	C(2)	C(3)	C(4)	C(5)	C(6)	C(7)	D(1)	D(2)	D(3)
Numbering	24	25	26	27	28	29	30	31	32	33	34	35
Labelling	D(4)	D(5)	D(6)	D(7)	E(1)	E(2)	E(3)	E(4)	E(5)	F(1)	F(2)	F(3)
Numbering	36	37	38	39	40	41	42	43	44	45	46	47
Labelling	F(4)	F(5)	G(1)	G(2)	G(3)	G(4)	G(5)	H(1)	H(2)	H(3)	H(4)	H(5)

10.2.3.5.1.9.4 Bit interleaver in UEP coding modes

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \ldots, 47$ be the index at the input of the bit interleaver, and $y = 0, \ldots, 47$ be the index at the output of the bit interleaver. The bit interleaver in the first half cycle of the UEP coding mode shall implement the following relation:

$$y = \left[6 \left\lfloor \frac{x}{6} \right\rfloor - 5(x \bmod 6) \right] \bmod 48 \qquad (40)$$

10.2.3.5.1.9.5 Multiplexer in UEP mapping modes

In the UEP mapping mode, all eight encoders shall use the same coding rate, but encoded bits from the top 4 encoders (or encoders A, B, C, D) shall be mapped to the I branch, and encoded bits from bottom 4 encoders (or encoders E, F, G, H) shall be mapped to the Q branch. In this case, the encoded bits shall be multiplexed and bit interleaved every 48 bits. During the length 48 multiplexing/ interleaving cycle, a group multiplexer shall be used with fixed group size 6 for all eight encoders. Use A(1), A(2), A(3), A(4), A(5), A(6) to label the 6 encoded bits (in increasing order in time) from encoder A, and similarly B(1) through B(6), C(1) through C(6), D(1) through D(6), E(1) through E(6), F(1) through F(6), G(1) through G(6), and H(1) through H(6) from encoders B, C, D, E, F, G, and H, respectively. At the output of the multiplexer, the 48 encoded bits shall be ordered and numbered as illustrated in Table 23, where the numbering for each bit represents the position of the bit in the bit stream after the multiplexer and bit 0 is the first bit at the output of the multiplexer.

Table 23 - Input pattern for the UEP mapping modes

Numbering	0	1	2	3	4	5	6	7	8	9	10	11
Labelling	A(1)	A(2)	A(3)	A(4)	A(5)	A(6)	B(1)	B(2)	B(3)	B(4)	B(5)	B(6)
Numbering	12	13	14	15	16	17	18	19	20	21	22	23
Labelling	C(1)	C(2)	C(3)	C(4)	C(5)	C(6)	D(1)	D(2)	D(3)	D(4)	D(5)	D(6)
Numbering	24	25	26	27	28	29	30	31	32	33	34	35
Labelling	E(1)	E(2)	E(3)	E(4)	E(5)	E(6)	F(1)	F(2)	F(3)	F(4)	F(5)	F(6)
Numbering	36	37	38	39	40	41	42	43	44	45	46	47
Labelling	G(1)	G(2)	G(3)	G(4)	G(5)	G(6)	H(1)	H(2)	H(3)	H(4)	H(5)	H(6)

10.2.3.5.1.9.6 Bit interleaver in UEP mapping modes

The multiplexed 48 bits are then sent to the bit interleaver. Let $x = 0, \ldots, 47$ and $y = 0, \ldots, 47$ be the index at the input and at the output of the bit interleaver. The bit interleaver in the UEP mapping mode shall implement the following relation

$$y = \begin{cases} 2\left(\left[6\left\lfloor \frac{x}{6} \right\rfloor - 5(x \bmod 6)\right] \bmod 24\right) & 0 \le x \le 23 \\ 2\left(\left[6\left\lfloor \frac{x}{6} \right\rfloor - 5(x \bmod 6)\right] \bmod 24\right) + 1 & 24 \le x \le 47 \end{cases} \quad (41)$$

10.2.3.5.1.10 Mapping

Refer to 10.2.4.

10.2.3.5.1.11 Symbol padding

The resulting data symbols shall then be appended with $N_{padsym,OFDM}$ zero symbols, where

$$N_{padsym} = 360\left\lceil \frac{N_{sym}}{360} \right\rceil - N_{sym} \quad (42)$$

10.2.3.5.1.12 OFDM PHY modulation

The subcarriers are numbered from -256 to 255. The subcarriers shall be arranged as indicated in Table 24.

Table 24 - OFDM subcarrier allocation

Subcarrier type	Subcarrier number, k
Null	k= -256,...,-190 and k=190,...,255
Pilots	k= -189, -164, -139, -114, -89, -64, -39, -14, 14, 39, 64, 89, 114, 139, 164, 189
DC	k= -1,0,1
Data	All other k

The stream of complex symbols from the modulation mapping is divided into groups of *360* complex numbers, numbered from $n = 0$ to $n = 359$ where $n = 0$ corresponds to the first complex number received in time. Each of the complex numbers are mapped sequentially to the subcarriers, skipping the pilots and DC subcarriers, beginning with $n = 0$ mapped to $k = -188$ and $n = 359$ mapped to $k = 188$.

10.2.3.5.1.13 Bit reversal tone interleaver

All modulated QPSK or QAM symbols shall be interleaved by a block interleaver with a block size corresponding to the size of FFT in a single OFDM symbol. The interleaver ensures that the adjacent data symbols are mapped onto separate subcarriers.

At the transmitter side, the interleaver permutation shall be defined as follows: Let k be the index of the tones (including data tones, pilot tones, DC tones and null tones) before permutation ranging between 0 and 511, and i be the index of the interleaved tones also ranging between 0 and 511 (including data tones, pilot tones, DC tones and null tones) after permutation. Let

$$k = \sum_{j=0}^{8} a_j 2^j \qquad (43)$$

with $a_8 a_7 \ldots a_0$ being the binary representation of integer k. Then the binary representation of integer i can be written as $a_0 a_1 \ldots a_8$, i.e.,

$$i = \sum_{j=0}^{8} a_j 2^{n-1-j} \qquad (44)$$

DC, null, and pilot tones shall be inserted in the bit-reversal position before the tone interleaver.

10.2.3.5.2 Midamble

The midamble sequence shall be identical to the channel estimation sequence, $S_{CE,A,OFDM}[.]$, which consists of the two long training sequence, with a cyclic prefix of length 64 to be added in front of the long training sequence, as described in 10.2.3.3.

10.2.4 Constellation mapping

This Clause describes the techniques for mapping the coded and interleaved binary data sequence onto a complex constellation. The constellation mapping shall be chosen according to Table 1.

10.2.4.1 BPSK

The coded and interleaved binary serial input data, $g[i]$ where i=0,1,2,..., shall be converted into a complex number representing one of the two BPSK constellation points. The output values, $v[k]$ where k=0,1,2,..., are formed by:

$$v[k] = K_{const}(I[k] + jQ[k]) \tag{45}$$

where $I[k]$ and $Q[k]$ are given by Table 25. The resulting constellation is illustrated in Figure 34. The normalization factor $K_{const}=1$ for a BPSK constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements.

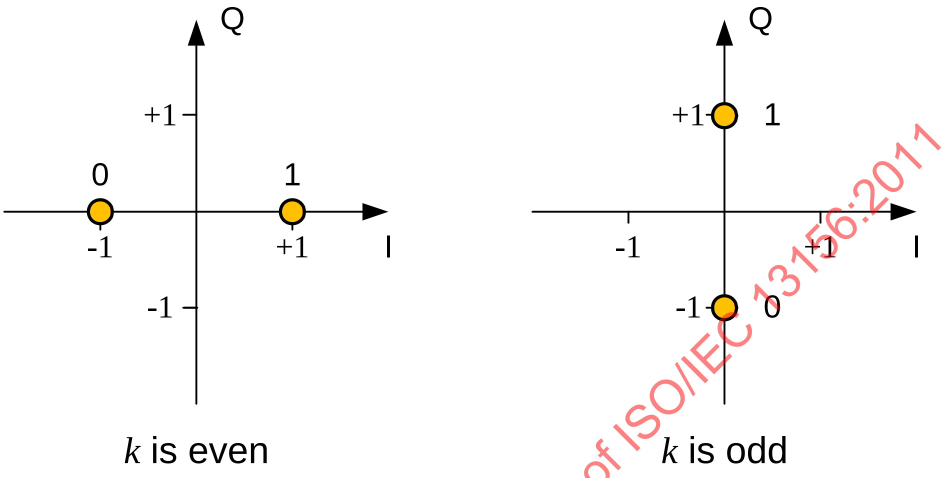

Figure 34 - BPSK constellation bit encoding

Table 25 - BPSK encoding table

Input bit $g[k]$	k is even		k is odd	
	$I[k]$	$Q[k]$	$I[k]$	$Q[k]$
0	-1	0	0	-1
1	1	0	0	1

10.2.4.2 QPSK

The coded and interleaved binary serial input data, $g[i]$ where $i=0,1,2,...$, shall be divided into groups of two bits and converted into a complex number representing one of the four QPSK constellation points. The conversion shall be performed according to the Gray-coded constellation mapping, illustrated in Figure 35 and described in Table 73, with the input bit, $g[2k]$ where $k=0,1,2,...$, being the earliest of the two in the stream. The output values, $v[k]$ where $k=0,1,2,...$, are formed by:

$$v[k] = K_{const}(I[k] + jQ[k]) \tag{46}$$

where $I[k]$ and $Q[k]$ are given by Table 73. The resulting constellation is illustrated in Figure 35. The normalization factor $K_{const} = 1/\sqrt{2}$ for a QPSK constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements. For QPSK, $g[2k]$ determines the I value, and $g[2k+1]$ determines the Q value, as illustrated in Table 26.

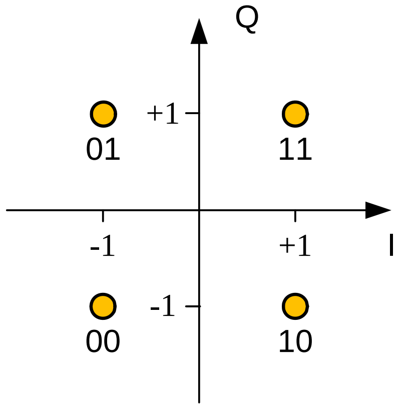

Figure 35 - QPSK constellation bit encoding

Table 26 - QSPK encoding table

Input bits ($g[2k]$,$g[2k+1]$)	$I[k]$	$Q[k]$
00	-1	-1
01	-1	1
10	1	-1
11	1	1

10.2.4.3 UEP-QPSK

The coded and interleaved binary serial input data, $g[i]$ where $i=0,1,2,...$, shall be divided into groups of two bits and converted into a complex number representing one of the four UEP-QPSK constellation points. The conversion shall be performed according to the Gray-coded constellation mapping, illustrated in Figure 36 and described in Table 27, with the input bit, $g[2k]$ where $k=0,1,2,...$, being the earliest of the two in the stream. The output values, $v[k]$ where $k=0,1,2,...$, are formed by:

$$v[k] = K_{const}\left(I[k] + jQ[k]\right) \qquad (47)$$

where $I[k]$ and $Q[k]$ are given by Table 27. The resulting constellation is illustrated in Figure 36. The value of a is 1.25. The normalization factor $K_{const} = 1/\sqrt{1+\alpha^2}$ for the UEP-QPSK constellation. An approximate value of the normalization factor may be use, as long as the device conforms to the modulation accuracy requirements. For UEP-QPSK, $g[2k]$ determines the I value, and $g[2k+1]$ determines the Q value, as illustrated in Table 27.

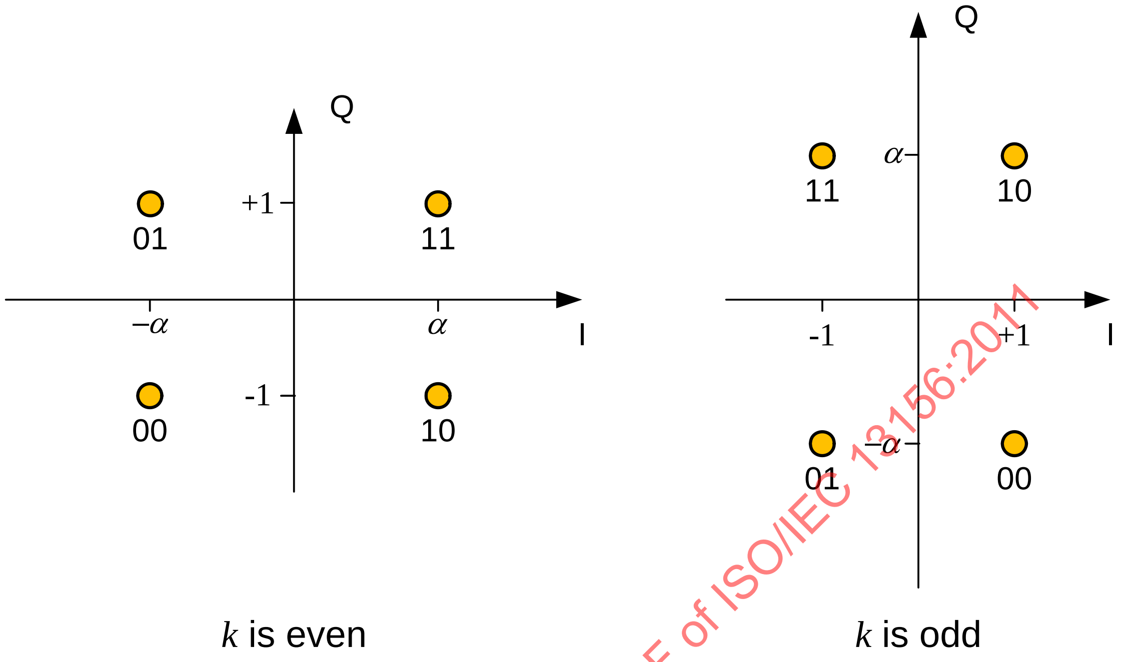

Figure 36 - UEP-QPSK constellation bit encoding

Table 27 - UEP-QPSK encoding table

Input bits ($g[2k]$,$g[2k+1]$)	*k* is even		*k* is odd	
	$I[k]$	$Q[k]$	$I[k]$	$Q[k]$
00	-a	-1	1	-a
01	-a	1	-1	-a
10	a	-1	1	a
11	a	1	-1	a

10.2.4.4 NS8QAM

The coded and interleaved binary serial input data, $g[i]$ where $i=0,1,2,...$, shall be divided into groups of three bits and converted into a complex number representing one of the four NS8QAM constellation points. The conversion shall be performed according to constellation mapping illustrated in Figure 37 and described in Table 28, with the input bit, $g[3k]$ where $k=0,1,2,...$, being the earliest of the three in the stream. The output values, $v[k]$ where $k=0,1,2,...$, are performed by:

$$v[k] = K_{const}\left(I[k] + jQ[k]\right) \quad (48)$$

where $I[k]$ and $Q[k]$ are given by Table 28. The resulting constellation is illustrated in Figure 37. The normalization factor $K_{const} = 1/\sqrt{10}$ for a NS8QAM constellation.

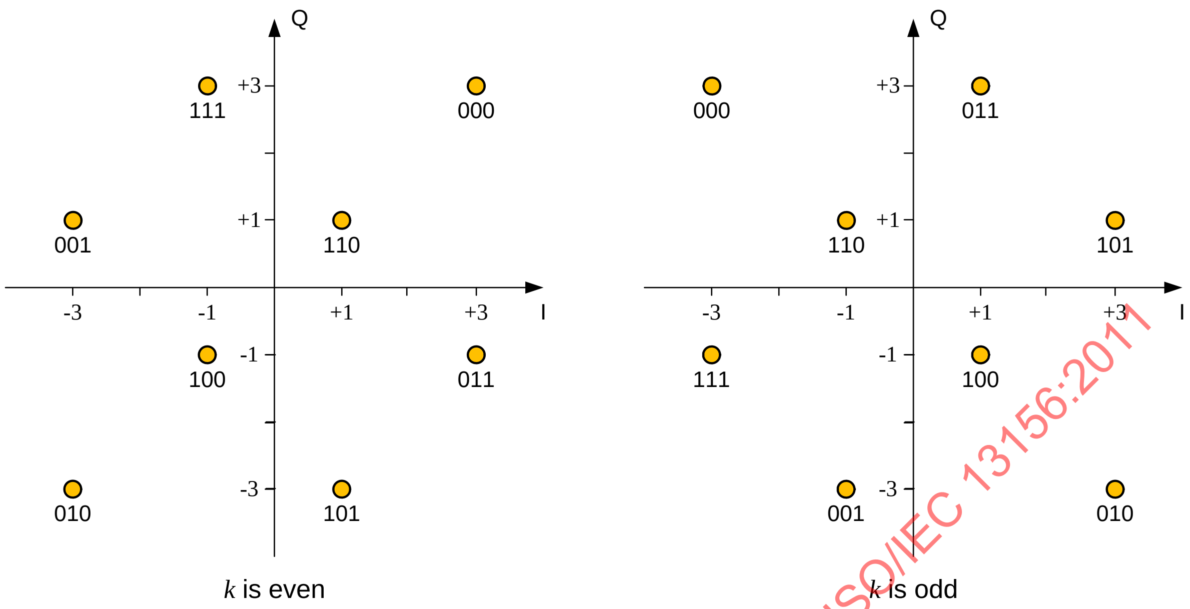

Figure 37 - NS8QAM constellation bit encoding

Table 28 - NSQAM encoding table

Input bits ($g[3k]$, $g[3k+1]$, $g[3k+2]$)	*k* is even $I[k]$	*k* is even $Q[k]$	*k* is odd $I[k]$	*k* is odd $Q[k]$
000	3	3	-3	3
001	-3	1	-1	-3
010	-3	-3	3	-3
011	3	-1	1	3
100	-1	-1	1	-1
101	1	-3	3	1
110	1	1	-1	1
111	-1	3	-3	-1

10.2.4.5 16QAM

The coded and interleaved binary serial input data, $g[i]$ where $i=0,1,2,...$, shall be divided into groups of four bits and converted into a complex number representing one of the four 16QAM constellation points. The conversion shall be performed according to the constellation mapping, illustrated in Figure 38 and described in Table 29, with the input bit, $g[4k]$ where $k=0,1,2,...$, being the earliest of the four in the stream. The output values, $v[k]$ where $k=0,1,2,...$, are performed by:

$$v[k] = K_{const}\left(I[k] + jQ[k]\right) \quad (49)$$

where $I[k]$ and $Q[k]$ are given by Table 29. The resulting constellation is illustrated in Figure 38. The normalization factor $K_{const} = 1/\sqrt{10}$ for a 16QAM constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements. For 16QAM, $g[4k]$ determines the I value, and $g[4k+1]$ determines the Q value, as illustrated in Table 29.

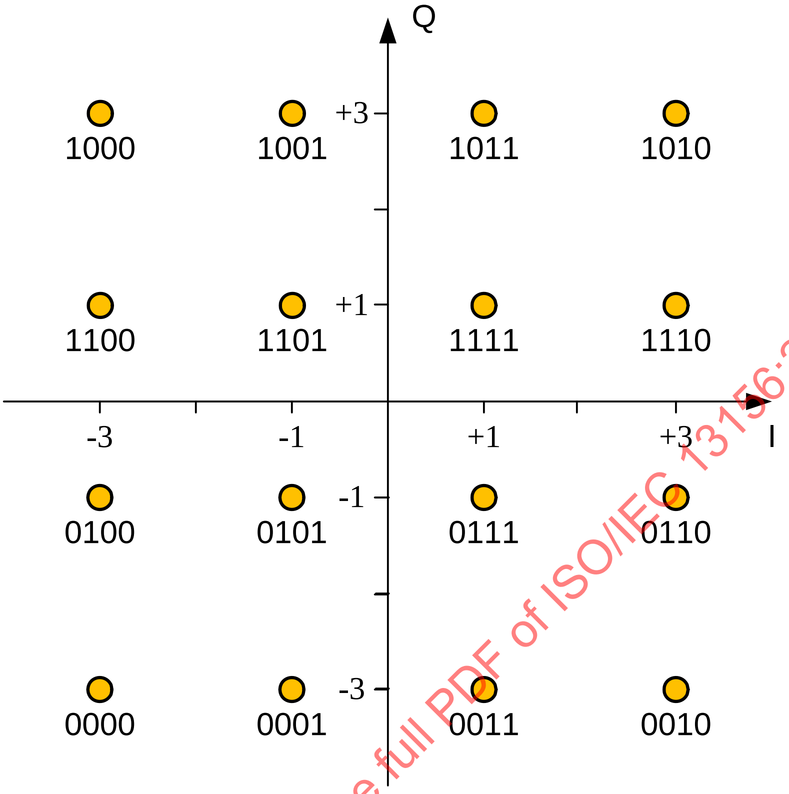

Figure 38 - 16QAM constellation bit encoding

Table 29 - 16QAM encoding table

Input bits ($g[4k]$,$g[4k+1]$,$g[4k+2]$,$g[4k+3]$)	$I[k]$	$Q[k]$
0000	-3	-3
0001	-1	-3
0010	3	-3
0011	1	-3
0100	-3	-1
0101	-1	-1
0110	3	-1
0111	1	-1
1000	-3	3
1001	-1	3
1010	3	3
1011	1	3
1100	-3	1
1101	-1	1
1110	3	1
1111	1	1

10.2.4.6 TCM-16QAM

The coded and interleaved binary serial input data, $g[i]$ where $i=0,1,2,...$, shall be divided into groups of four bits and converted into a complex number representing one of the four TCM-16QAM constellation points. The conversion shall be performed according to the constellation mapping, illustrated in Figure 39 and described in Table 30, with the input bit, $g[4k]$ where $k=0,1,2,...$, being the earliest of the four in the stream. The output values, $v[k]$ where $k=0,1,2,...$, are performed by:

$$v[k]=K_{const}(I[k]+jQ[k]) \quad (50)$$

where $I[k]$ and $Q[k]$ are given by Table 29. The resulting constellation is illustrated in Figure 39. The normalization factor $K_{const}=1/\sqrt{10}$ for a TCM-16QAM constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements. For TCM-16QAM, $g[4k]$ determines the I value, and $g[4k+1]$ determines the Q value, as illustrated in Table 30.

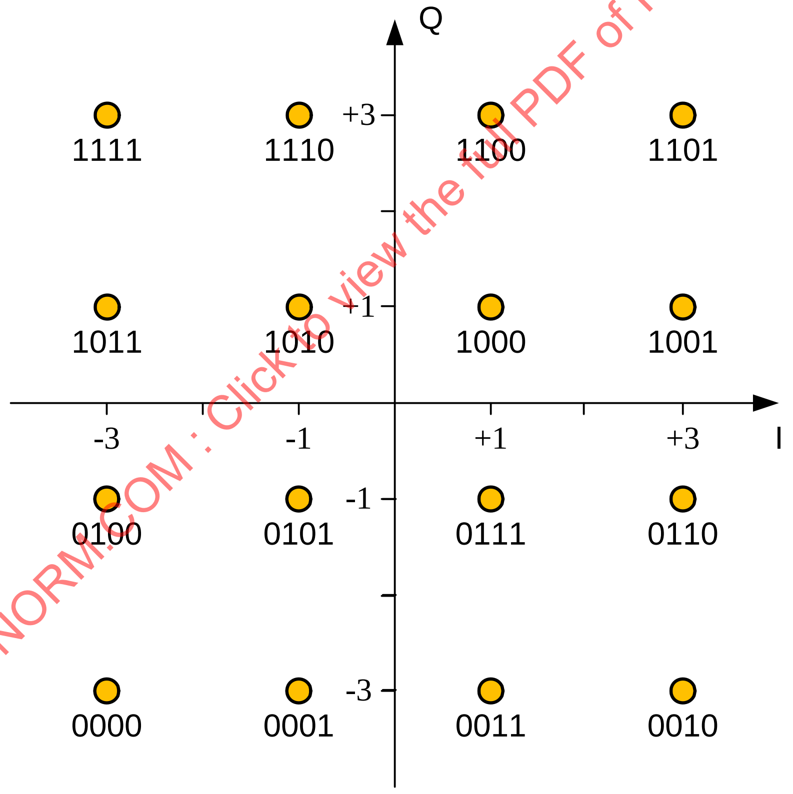

Figure 39 - TCM-16QAM constellation bit encoding

Table 30 - TCM-16QAM encoding table

Input bits ($g[4k]$,$g[4k+1]$,$g[4k+2]$,$g[4k+3]$)	$I[k]$	$Q[k]$
0000	-3	-3
0001	-1	-3
0010	3	-3
0011	1	-3
0100	-3	-1
0101	-1	-1
0110	3	-1
0111	1	-1
1000	1	1
1001	3	1
1010	-1	1
1011	-3	1
1100	1	3
1101	3	3
1110	-1	3
1111	-3	3

10.2.4.7 UEP-16QAM

The coded and interleaved binary serial input data, $g[i]$ where $i=0,1,2,...$, shall be divided into groups of four bits and converted into a complex number representing one of the four UEP-16QAM constellation points. The conversion shall be performed according to the Gray-coded constellation mapping, specified in Table 31 and illustrated in Figure 40, with the input bit, $g[4k]$ where $k=0,1,2,...$, being the earliest of the four in the stream. The output values, $v[k]$ where $k=0,1,2,...$, are formed by:

$$v[k] = K_{const}\left(I[k] + jQ[k]\right) \tag{51}$$

where $I[k]$ and $Q[k]$ are given by Table 31. The resulting constellation is illustrated in Figure 40. The value of α is 1.25. The normalization factor is $K_{const} = 1/\sqrt{5+5\alpha^2}$ for a UEP-16QAM constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements. For UEP-16QAM, $g[4k]$ determines the I value, and $g[4k+1]$ determines the Q value, as illustrated in Table 31.

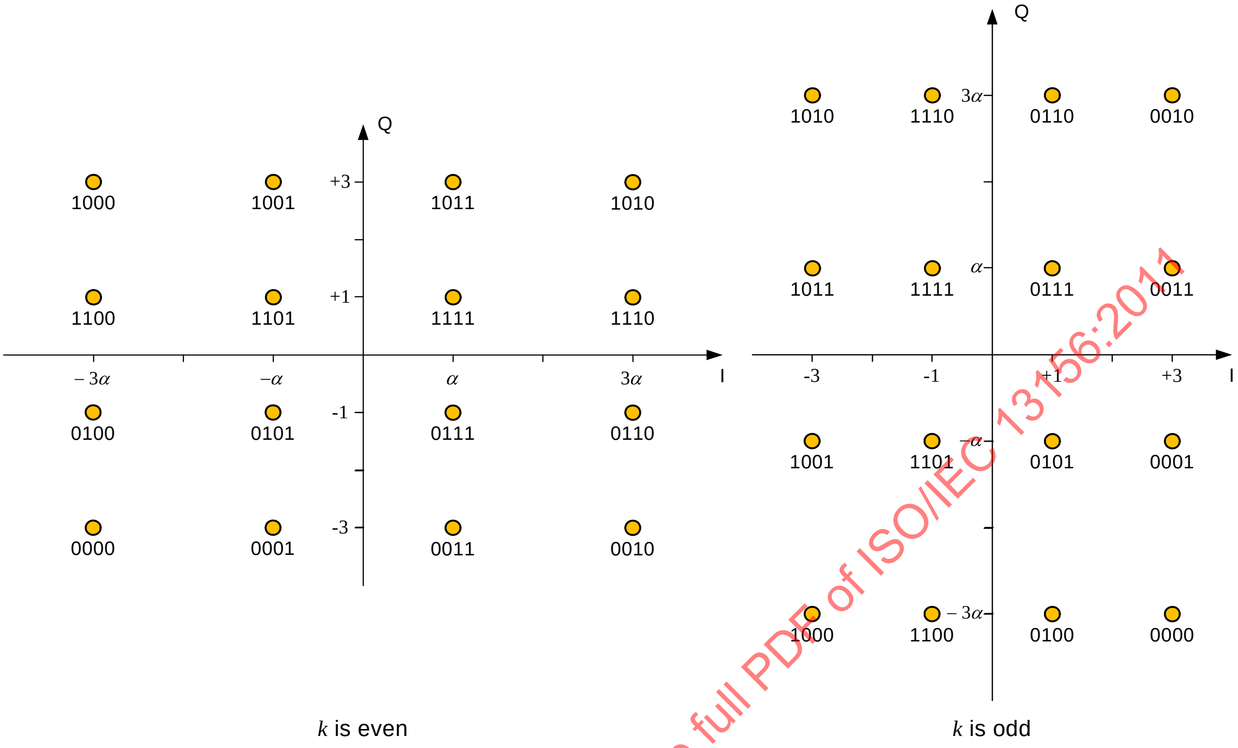

Figure 40 - UEP-16QAM constellation bit encoding

Table 31 - UEP-16QAM encoding table

Input bits ($g[4k]$,$g[4k+1]$,$g[4k+2]$,$g[4k+3]$)	k is even		k is odd	
	$I[k]$	$Q[k]$	$I[k]$	$Q[k]$
0000	-3a	-3	3	-3a
0001	-a	-3	3	-a
0010	3a	-3	3	3a
0011	a	-3	3	a
0100	-a	-1	1	-a
0101	-a	-1	1	-a
0110	3a	-1	1	3a
0111	a	-1	1	a
1000	-3a	3	-3	-3a
1001	-a	3	-3	-a
1010	3a	3	-3	3a
1011	a	3	-3	a
1100	-3a	1	-1	-3a
1101	-a	1	-1	-a

Table 31 - UEP-16QAM encoding table (concluded)

Input bits ($g[4k]$,$g[4k+1]$,$g[4k+2]$,$g[4k+3]$)	k is even		k is odd	
	$I[k]$	$Q[k]$	$I[k]$	$Q[k]$
1110	3a	1	-1	3a
1111	a	1	-1	a

10.2.5 Discovery mode

The discovery mode is used for communications prior to training of antenna arrays. To compensate for the lack of array gain prior to training, the discovery mode increases the time-bandwidth product by repetition. The frame format for discovery mode shall follow the general frame format described in 10.1. Eight modes, D0 through D7, are defined. These modes are identical except for the number of repetitions (see Table 32).

Table 32 - Discovery modes

Mode	$N_{DISCREP}$	Data Rate (Mbps)
D0	128	2.255
D1	64	4.510
D2	32	9.020
D3	16	18.041
D4	8	36.082
D5	4	72.164
D6	2	144.327
D7	1	288.655

The value of the $N_{DISCREP}$ shall be set based on the requested value by the receiver. The initial value of $N_{DISCREP}$ shall be equal to 128 (mode D0).

10.2.5.1 Discovery mode PLCP preamble

The discovery mode preamble shall consist of the concatenation of a narrowband preamble, $S_{NB}[.]$, and a wideband preamble, $S_{WB}[.]$, as shown in Figure 41.

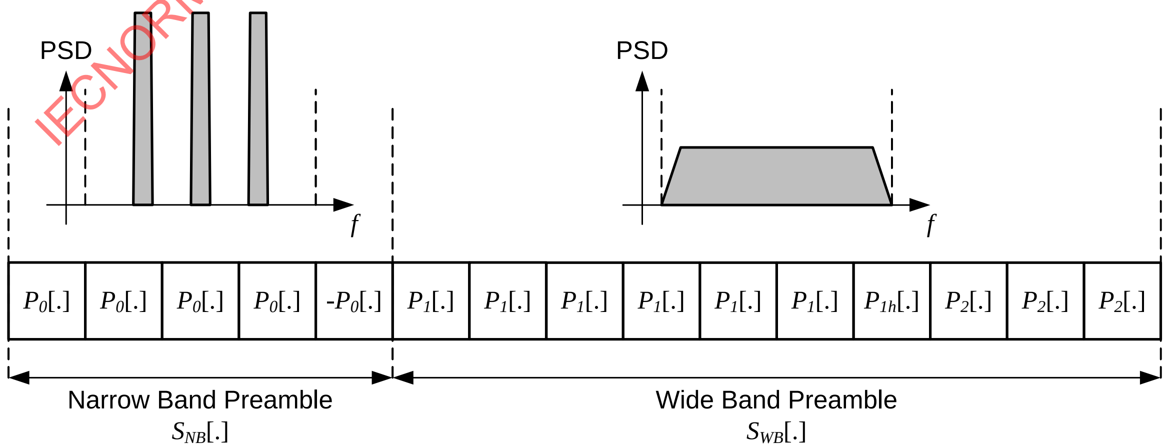

Figure 41 - Discovery mode preamble

The narrowband preamble, $S_{NB}[.]$, is defined as the Kronecker product of the sequence $S_0[.]$ and the a cover sequence $S_{cover,disc}[.]$. That is,

$$S_{NB}[n] = S_{cover,disc}\left[\left\lfloor \frac{n}{32768} \right\rfloor\right] S_0[n \bmod 32768] \qquad (52)$$

for n = 0,...,163839, where the cover sequence is of length five and is defined as

$$S_{cover,disc}[n] = \begin{cases} 1 & n = 0,\ldots,3 \\ -1 & n = 4 \end{cases} \qquad (53)$$

and $S_0[.]$ is a narrowband sequence of length 32768 and shall be obtained from $S_h[.]$ (defined in 10.2.2.3.1) by

$$S_0[n] = S_h\left[\left\lfloor \frac{n}{128} \right\rfloor\right] \qquad (54)$$

The wideband preamble, $S_{WB}[.]$, is the concatenation of six copies of $S_1[.]$, followed by one copy of sequence $S_{1h}[.]$, followed by three copies of sequence $S_2[.]$. That is,

$$S_{WB}[n] = \begin{cases} S_1[n \bmod 32768] & 0 \le n < 196608 \\ S_{1h}[n \bmod 32768] & 196608 \le n < 229376 \\ S_2[n \bmod 32768] & 229376 \le n < 327680 \end{cases} \qquad (55)$$

for n = 0,...,327679, where $S_1[.]$ is a wideband sequence of length 32768 and shall be obtained from $S_h[.]$ (defined in 10.2.2.3.1) by

$$S_1[n] = S_h[n \bmod 256] \qquad (56)$$

$S_{1h}[.]$ is a wideband sequence of length 65536 and shall be obtained from $S_h[.]$ and the cover sequence $S_{h2}[.]$ by

$$S_{1h}[n] = S_h[n \bmod 256] S_{hFZ}\left[\left\lfloor \frac{n}{256} \right\rfloor\right] \qquad (57)$$

where $S_{hFZ}[.]$ is a differentially encoded FZ sequence of length 256 (parameter A_{FZ}=16) defined by

$$S_{hFZ}[0] = S_{FZ,16}[0] \qquad (58)$$

and

$$S_{h2}[n] = S_{h2}[n-1] S_{FZ,16}[n] \qquad (59)$$

for n=1,...,255, and $S_2[.]$ is a wideband sequence of length 32768 and shall be obtained from $S_{FZ,16}[.]$ by

$$S_2[n] = S_{FZ,16}[n \bmod 256] \quad (60)$$

10.2.5.1.1 Three carrier modulation

The narrowband preamble, $S_{NB}[.]$, shall be modulated using three carriers, f_c, f_c+f_0, and f_c-f_0, as described in 9.2. The wideband preamble shall be modulated using one carrier, f_c, as described in 9.1.

10.2.5.2 Discovery mode PLCP header

The header for a discovery mode frame shall be formed based on the header of the Type A SCBT frames, as described in 10.2.2.4. After the formation of the header, each SCBT symbol of the header shall be repeated consecutively 128 times.

10.2.5.3 Discovery mode PPDU payload

The payload for a discovery mode frame is based on the payload of the Type A SCBT mode A0 frames, as described in 10.2.2.5. After the formation of the payload, each SCBT symbol of the payload shall be repeated consecutively $N_{DISCREP}$ times.

10.2.5.4 Discovery mode ATS

The ATS for a discovery mode frame is based on the ATS of the Type A frames, as described in 10.1.4. After the formation of the ATS, each block (of length 256) of the ATS shall be repeated consecutively $N_{DISCREP}$ times.

10.3 Type B PPDU

10.3.1 Mode dependent parameters

The Type B PSDU data rate-dependent modulation parameters are listed in Table 33.

Table 33 - PSDU mode dependent parameters

Mode	Base Data Rate (Gbps)				Constellation	Encoding	TDSF (N_{TDS})
	no channel bonding	2 bonded channels	3 bonded channels	4 bonded channels			
B0	0.794	1.588	2.381	3.175	DBPSK	RS & Diff	2
B1	1.588	3.175	4.763	6.350	DBPSK	RS & Diff	1
B2	3.175	6.350	9.526	12.701	DQPSK	RS & Diff	1
B3	3.175	6.350	9.526	12.701	UEP-QPSK	RS	1

All Type B devices shall support mode B0 (without channel bonding). Type B devices may support modes B0 (with channel bonding) or modes B1 through B3. In addition, all Type B devices shall support the transmission of mode A0 (without channel bonding). See 10.2 for Type A modes.

10.3.2 Single carrier (DBPSK, DQPSK, UEP-QPSK)

10.3.2.1 Timing related parameters

The timing parameters associated with the Type B PHY are listed in Table 34.

Table 34 - Timing related parameters

Parameter	Description	Value
f_{sym}	Symbol frequency	1.728 Gsps
T_{sym}	Symbol duration	0.5787 ns
N_B	Number of symbols per SC block	256
N_D	Number of data symbols per SC block	252
N_P	Number of pilot symbols per SC block	4
N_{CP}	Number of symbols in CP	0
T_{CP}	CP interval	0

10.3.2.2 Frame related parameters

The frame related parameters associated with the PHY are listed in Table 35.

Table 35 - Frame related parameters

Parameter	Description	Value
N_{sync}	Number of symbols in the frame synchronization sequence	2048
T_{sync}	Duration of the frame synchronization sequence	1185.19 ns
N_{CE}	Number of symbols in the channel estimation sequence	768
T_{CE}	Duration of the channel estimation sequence	444.444 ns
$N_{preamble}$	Number of symbols in the PLCP preamble	2816
$T_{preamble}$	Duration of the PLCP preamble	1629.63 ns
N_{ATS}	Number of symbols in the ATS	$256N_{RXTS}$
T_{ATS}	Duration of the ATS	$N_{ATS}T_{sym}$
N_{frame}	Number of symbols in the frame	$N_{preamble}+N_{header}+N_{payload}+N_{ATS}$
T_{frame}	Duration of the frame	$(N_{preamble}+N_{header}+N_{payload}+N_{ATS})T_{sym}$

10.3.2.3 PLCP preamble

The preamble for Type B frames can be subdivided into two distinct portions: a frame synchronization sequence, $S_{sync,B}[.]$, and a channel estimation sequence, $S_{CE,B}[.]$. Figure 42 shows the structure of the PLCP preamble for Type B frames.

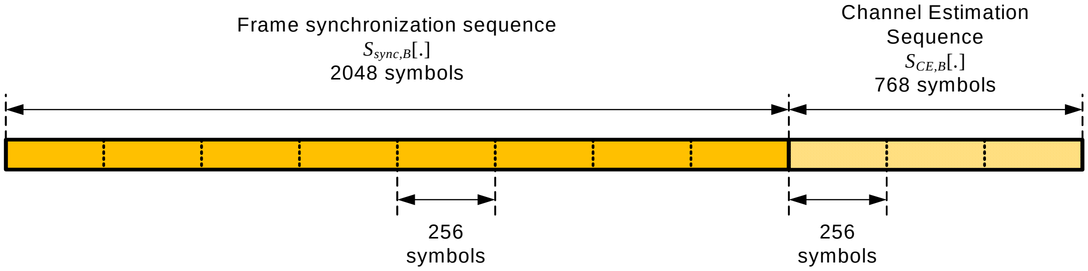

Figure 42 - Structure of the Type B PLCP preamble for Type B frames

Both the frame synchronization sequence and the channel estimation sequence are constructed based on the Frank-Zadoff sequences as defined in 10.2.2.3.

10.3.2.3.1 Frame synchronization sequence

The frame synchronization sequence for Type B is a modified hierarchical sequence $S_{h,B}[.]$, based on the hierarchical sequence $S_h[.]$ defined in 10.2.2.3.1.

The modified synchronization sequence $S_{h,B}[.]$ for Type B is obtained by adding the real and imaginary part of each term of sequence $S_h[.]$ as follows

$$S_{h,B}[n] = \mathrm{Re}\{S_h[n]\} + \mathrm{Im}\{S_h[n]\} \tag{61}$$

for $n=0,...,255$.

The frame synchronization sequence for Type B frames is defined as the Kronecker product of the hierarchical sequence $S_{h,B}[.]$ and the cover sequence $S_{cover}[.]$ defined in 10.2.2.3.1. That is,

$$S_{sync,B,SC}[n] = S_{cover}\left[\left\lfloor \frac{n}{256} \right\rfloor\right] S_{h,B}[n \bmod 256] \tag{62}$$

for $n=0,...,2047$.

10.3.2.3.2 Channel estimation sequence

The channel estimation sequence for Type B frames is a modified channel estimation sequence $S_{CE.B}[.]$, based on the channel estimation sequence $S_{SC,A,SCBT}[.]$ defined in 10.2.2.3.2.

The channel estimation sequence for Type B frames, $S_{CE,B}[n]$ is defined as:

$$S_{CE,B}[n] = \begin{cases} 1 & \mathrm{Re}\{S_{CE,A,SCBT}[n]\} > \mathrm{Im}\{S_{CE,A,SCBT}[n]\} \text{or} \, \mathrm{Re}\{S_{CE,A,SCBT}[n]\} = \mathrm{Im}\{S_{CE,A,SCBT}[n]\} > 0 \\ -1 & \mathrm{Re}\{S_{CE,A,SCBT}[n]\} < \mathrm{Im}\{S_{CE,A,SCBT}[n]\} \text{or} \, \mathrm{Re}\{S_{CE,A,SCBT}[n]\} = \mathrm{Im}\{S_{CE,A,SCBT}[n]\} < 0 \end{cases} \tag{63}$$

for $n=0,...,767$.

NOTE The frame synchronization sequence can be used for frame acquisition and detection, coarse carrier frequency estimation, coarse symbol timing, and for synchronization within the preamble. Whereas, the channel estimation sequence can be used for estimation of the channel frequency response, fine carrier frequency estimation, and fine symbol timing.

10.3.2.4 PLCP header

For Type B frames, the formed PLCP header (as described in 10.1.2) shall be

first padded with N_{padhdr} bits where

$$N_{padhdr} = 192\left\lceil \frac{N_{hdrbits}}{192} \right\rceil - N_{hdrbits} \tag{64}$$

where $N_{hdrbits}$ is the number of bits in the formed PLCP header (as described in 10.1.2). The padded header shall be encoded and modulated as described in 10.3.2.5, starting with the demultiplexing. Encoding and modulation parameters identical to mode B0 shall be used. The resulting data symbols shall be block modulated (as described in 10.3.2.5.4) in order to create the baseband signal.

10.3.2.4.1 PHY header

10.3.2.4.1.1 Fixed length PHY header

10.3.2.4.1.1.1 PLCP scrambler field (SCRAMBLER)

Refer to 10.2.2.4.1.1.1.

10.3.2.4.1.1.2 Bit reversal field (BIT_REVERSAL)

The value of the BIT_REVERSAL bit shall be set to 0_B for all Type B frames.

10.3.2.4.1.1.3 ATIF existence field (ATIF_EX)

Refer 10.2.2.4.1.1.3.

10.3.2.4.1.1.4 CP length field (CP_LENGTH)

The value of the CP_LENGTH field shall be set to 00_B for all Type B frames.

10.3.2.4.1.1.5 Requested CP length field (REQ_CP_LENGTH)

The value of the REQ_CP_LENGTH field shall be set to 00_B for all Type B frames.

10.3.2.4.1.1.6 Number of segments field (NUM_SEGMENTS)

Refer to 10.2.2.4.1.1.6.

10.3.2.4.1.1.7 Number of MSDUs field (NUM_MSDUS)

Refer to 10.2.2.4.1.1.7.

10.3.2.4.1.2 Variable length PHY header

10.3.2.4.1.2.1 Mode field (MODE)

Refer to 10.2.2.4.1.2.1.

10.3.2.4.1.2.2 Segment length field (LENGTH)

Refer to 10.2.2.4.1.2.2.

10.3.2.4.1.2.3 Midamble existence field (MID_EX)

Refer to 10.2.2.4.1.2.3.

10.3.2.4.1.2.4 Segment continued field (CONT)

Refer to 10.2.2.4.1.2.4.

10.3.2.4.1.3 Antenna training indicator field

10.3.2.4.1.3.1 Discovery mode repetition field (DISC_REP)

Refer to 10.2.2.4.1.3.1.

10.3.2.4.1.3.2 Number of training symbols for RX training field (NUM_RXTS)

The value of the NUM_RXTS field for all Type B frames transmitted by a Type A device shall be set to 000000_B. For the value of the NUM_RXTS field for all Type B frames transmitted by a Type B device refer to 10.2.2.4.1.3.2.

10.3.2.4.1.3.3 Number of training symbols for TX training field (NUM_TXTS)

The value of the NUM_TXTS field for all Type B frames transmitted by a Type B device shall be set to 000000_B. For the value of the NUM_TXTS field for all Type A frames transmitted by a Type A device refer to 10.2.2.4.1.3.3.

10.3.2.5 PPDU payload

The PPDU payload consists of one or more segments as described in 10.1.3. The each segment is formed as described in 10.1.3. Each data bit block shall be processed as follows:

1. The bits shall be mapped as described in 10.3.2.5.1 or 10.3.2.5.2.
2. The resulting symbols shall be modulated into SC blocks as described in 10.3.2.5.4.

10.3.2.5.1 FEC and mapping (Equal error protection)

Figure 43 depicts the general overview of the encoding and mapping scheme for modes B0, B1, and B2. For the payload, first, the data bits shall be scrambled as specified in 10.3.2.5.1.1. Then, $N_{padbits}$ zero pad bits shall be appended to the end of the data bit block, where

$$N_{padbits} = 1792\left\lceil \frac{N_{bits}}{1792} \right\rceil - N_{bits} \tag{65}$$

and N_{bits} is the number of bits in the data bit block. The padded data bit block shall be encoded according to the Reed-Solomon code, as specified in 10.3.2.5.1.2.

The resulting RS coded payload bits or the formed header bits shall be demultiplexed into 4 groups, as described in 10.3.2.5.1.3. Each group of bits shall then be interleaved using the bit interleaver described in 10.3.2.5.1.4. The resulting bits from the 4 branches shall then be multiplexed together into one group of symbols as specified in 10.3.2.5.1.5. The multiplexed bits are then mapped to constellations as specified in 10.3.2.5.1.6. The data symbols shall be spread in time domain as described in 10.3.2.5.1.7. The resulting data symbols shall then be appended with N_{padsym} "-1"valued symbols, where

$$N_{padsym} = N_D\left\lceil \frac{N_{sym}}{N_D} \right\rceil - N_{sym} \tag{66}$$

The differential encoding shall then be applied as described in 10.3.2.5.1.8.

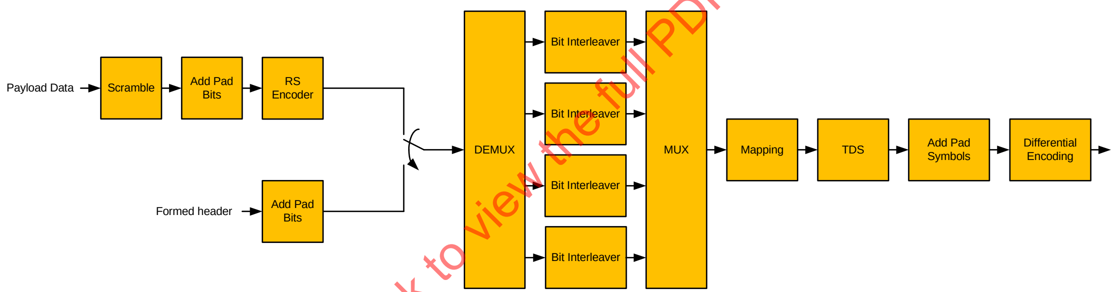

Figure 43 - Type B encoding procedure

10.3.2.5.1.1 Data scrambler

Refer to 10.2.2.5.1.1.

10.3.2.5.1.2 Reed-Solomon code

Refer to 10.2.2.5.1.2.

10.3.2.5.1.3 Demultiplexer

Refer to 10.2.2.5.1.3.

10.3.2.5.1.4 Bit interleaver

Refer to 10.2.2.5.1.4.

10.3.2.5.1.5 Multiplexer

Refer to 10.2.2.5.1.7.

10.3.2.5.1.6 Mapping

Refer to 10.3.2.5.3.

10.3.2.5.1.7 Time domain spreading (TDS)

Refer to 10.2.2.5.1.9.

10.3.2.5.1.8 Differential encoding

For modes B0, B1 and B2, the padded data symbols shall be differentially encoded. The differently encoded data symbols $t[n]$ are formed from data symbols $v[n]$ by

$$t[n] = \begin{cases} v[n] & n \bmod N_D = 0 \\ t[n-1]v[n]/|v[n-1]| & n \bmod N_D > 0 \end{cases} \quad (67)$$

NOTE As equation (67) specifies, the differential encoding shall be reset for every SC block. The last symbol of the channel estimation sequence or the last symbol of the pilot symbol sequence may be used to initialize the differential decoder.

10.3.2.5.2 FEC and mapping (Unequal error protection)

Figure 44 depicts the general overview of the encoding and mapping scheme for mode B3. For the payload, first, the MSB and LSB bits of the each octet is separated as described in 10.3.2.5.2.1. Then, for each group, the data bits shall be scrambled as specified in 10.3.2.5.2.2. Then, $N_{padbits}$ zero pad bits shall be appended to the end of the data bit block, where

$$N_{padbits} = 1792\left\lceil \frac{N_{bits}}{1792} \right\rceil - N_{bits} \quad (68)$$

and N_{bits} is the number of bits in the data bit block. The padded data bit block shall be encoded according to the Reed-Solomon code, as specified in 10.3.2.5.2.3.

The resulting RS coded payload bits or the formed header bits shall be demultiplexed in 4 groups, as described in 10.3.2.5.2.4. Depending on the data rate mode, each group of bits may then be interleaved using the bit interleaver described in 10.3.2.5.2.5. The resulting bits from all 8 branches shall then be multiplexed together into one group of symbols as specified in 10.3.2.5.2.6. The multiplexed bits are then mapped to constellations as specified in 10.3.2.5.3. The resulting data symbols shall then be appended with N_{padsym} zero symbols, where

$$N_{padsym} = 2N_D\left\lceil \frac{N_{sym}}{2N_D} \right\rceil - N_{sym} \quad (69)$$

At the end of each segment, the resulting data symbols shall be interleaved with the dual helical scan (DHS) symbol interleaver as described in 10.3.2.5.2.7.

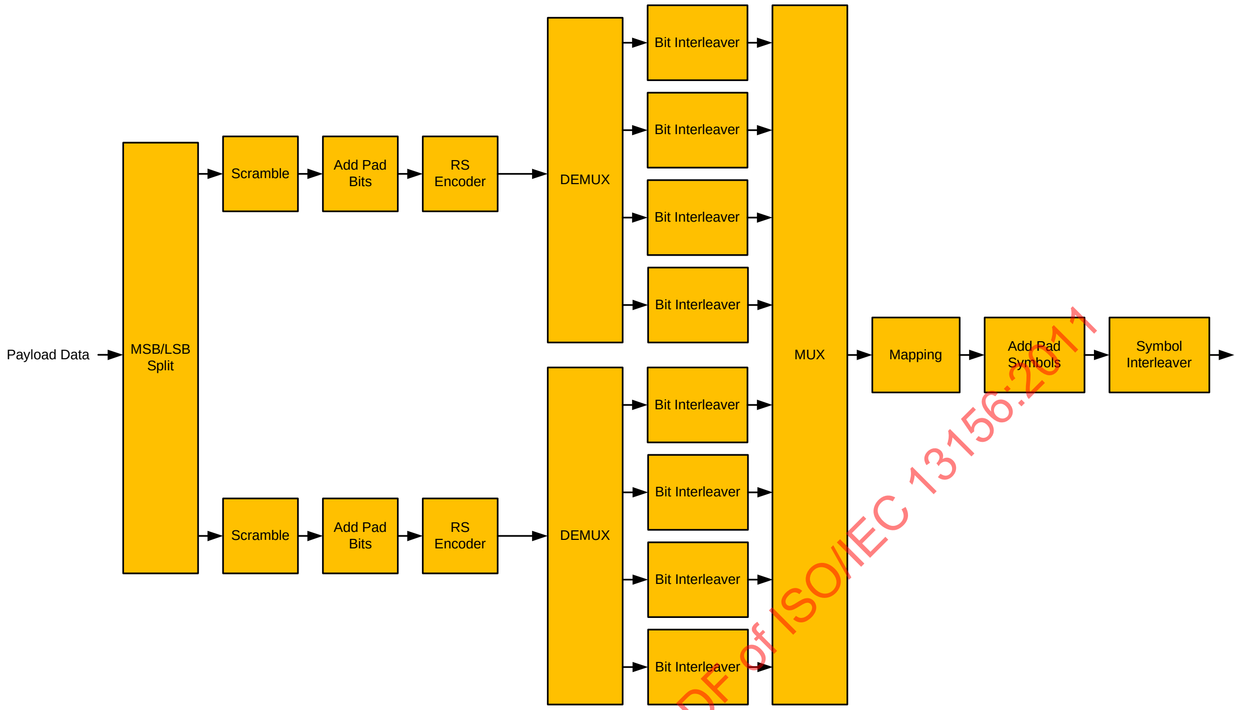

Figure 44 - General view of the encoding procedure

10.3.2.5.2.1 MSB/LSB separation

Refer to 10.2.2.5.2.1.

10.3.2.5.2.2 Data scrambler

Refer to 10.2.2.5.1.1.

10.3.2.5.2.3 Reed-Solomon code

Refer to 10.2.2.5.1.2.

10.3.2.5.2.4 Demultiplexer

Refer to 10.2.2.5.1.3.

10.3.2.5.2.5 Bit interleaver

Refer to 10.2.2.5.1.4.

10.3.2.5.2.6 Multiplexer

The multiplexer shall be identical to that described for mode A12 in 10.2.2.5.2.7.

10.3.2.5.2.7 Symbol interleaver

Refer to 10.2.2.5.1.8.

10.3.2.5.3 Constellation mapping

The constellation mapping for mode B0, B1, B2 and B3 shall be chosen according to Table 33.

10.3.2.5.3.1 DBPSK

The coded and interleaved binary serial input data, $g[i]$ where i=1,2,..., shall be converted into a complex number representing one of the two DBPSK constellation points. The output values, $v[k]$ where k=0,1,2,... are formed by:

$$v[k] = K_{const}\, I[k] \quad (70)$$

where $I[k]$ is given by Table 36. The resulting constellation is illustrated in Figure 45. The normalization factor $K_{const}=1$ for a DBPSK constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements. The differential encoding is described in 10.3.2.5.1.8 and is applied after the multiplexing and symbol padding as described in Table 36.

Table 36 - DBPSK encoding table

Input bits g[k]	I[k]
0	-1
1	1

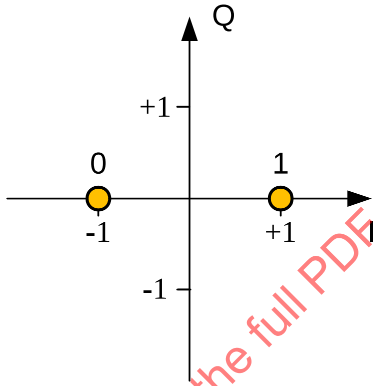

Figure 45 - DBPSK constellation bit encoding

10.3.2.5.3.2 DQPSK

The coded and interleaved binary serial input data, $g[i]$ where $i=1,2,\ldots$, shall be divided into groups of two bits and converted into a complex number representing one of the four DQPSK constellation points. The conversion shall be performed according to the Gray-coded constellation mapping, illustrated in Figure 46, with the input bit, $g[2k]$ where $k=1,2,\ldots$, being the earliest of the two in the stream. The output values, $v[k]$ where $k=1,2,\ldots$, are formed by:

$$v[k] = K_{const} (I[k] + jQ[k]) \tag{71}$$

where $I[k]$ and $Q[k]$ are given by Table 37. The resulting constellation is illustrated in Figure 46. The normalization factor $K_{const}=1$ for a DQPSK constellation. An approximate value of the normalization factor may be used, as long as the device conforms to the modulation accuracy requirements. For DQPSK, $g[2k]$ determines the I value, and $g[2k+1]$ determines the Q value, as illustrated in Table 37.

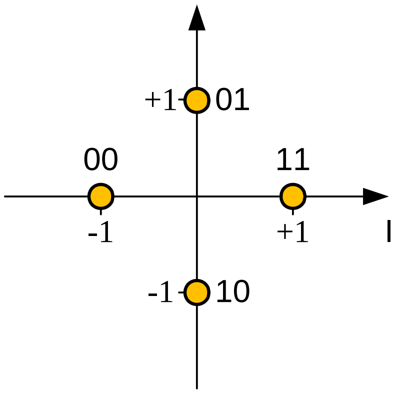

Figure 46 - DQPSK constellation bit encoding

Table 37 - DQPSK encoding table

Input bits (g[2k], g[2k+1])	I[k]	Q[k]
00	-1	0
01	0	1
10	0	-1
11	1	0

10.3.2.5.3.3 UEP-QPSK

Refer to 10.2.4.3.

10.3.2.5.4 Block modulation

The multiplexed transmit symbols, $t[.]$ shall be modulated formed into the SC block. First the transmit data symbols shall be divided into blocks of length N_D=252. This transmit symbol block shall be appended with pilot symbols. The pilot symbols consist of a sequence of length N_P=4, $S_{pilot}[n]=(-1)^{n+1}$, n=0,1,2,3. Figure 47 depicts an example of the SC block formation.

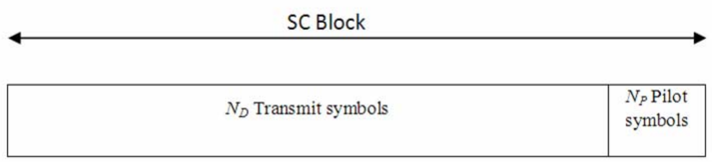

Figure 47 - Formation of the SC block

10.3.2.5.5 Midamble

The midamble sequence shall be identical to the channel estimation sequence, $S_{CE,B}[.]$, as described in 10.3.2.3.2.

NOTE The midamble sequence may be used to update estimation of the channel frequency response, fine carrier frequency estimation, and fine symbol timing.

10.3.3 Channel bonding

Type B devices may perform channel bonding. If a Type B device uses channel bonding, it shall be as described in 10.2.2.6.

10.3.4 Discovery mode

Type B devices shall not transmit frames in the discovery mode.

11 General requirements

11.1 Operating band frequencies

11.1.1 Operating frequency range

This PHY operates in the 57 - 66 GHz frequency band.

11.1.2 Channel numbering

The relationship between centre frequency, f_c, and BAND_ID number, n_b, is given in Table 38.

Table 38 - Band allocation

BAND_ID (n_b)	Channel Bonding	Lower Frequency (GHz)	Centre Frequency (GHz)	Upper Frequency (GHz)
1	No	57.240	58.320	59.400
2	No	59.400	60.480	61.560
3	No	61.560	62.640	63.720
4	No	63.720	64.800	65.880
5	Yes (1 & 2)	57.240	59.400	61.560
6	Yes (2 & 3)	59.400	61.560	63.720
7	Yes (3 & 4)	61.560	63.720	65.880
8	Yes (1, 2, & 3)	57.240	60.480	63.720
9	Yes (2, 3, & 4)	59.400	62.640	65.880
10	Yes (1, 2, 3, & 4)	57.240	61.560	65.880

This definition provides a unique numbering system for all channels within the band 57 - 66 GHz. The third channel (BAND_ID = 3) shall be used as the discovery channel. Figure 48 depicts the defined channels.

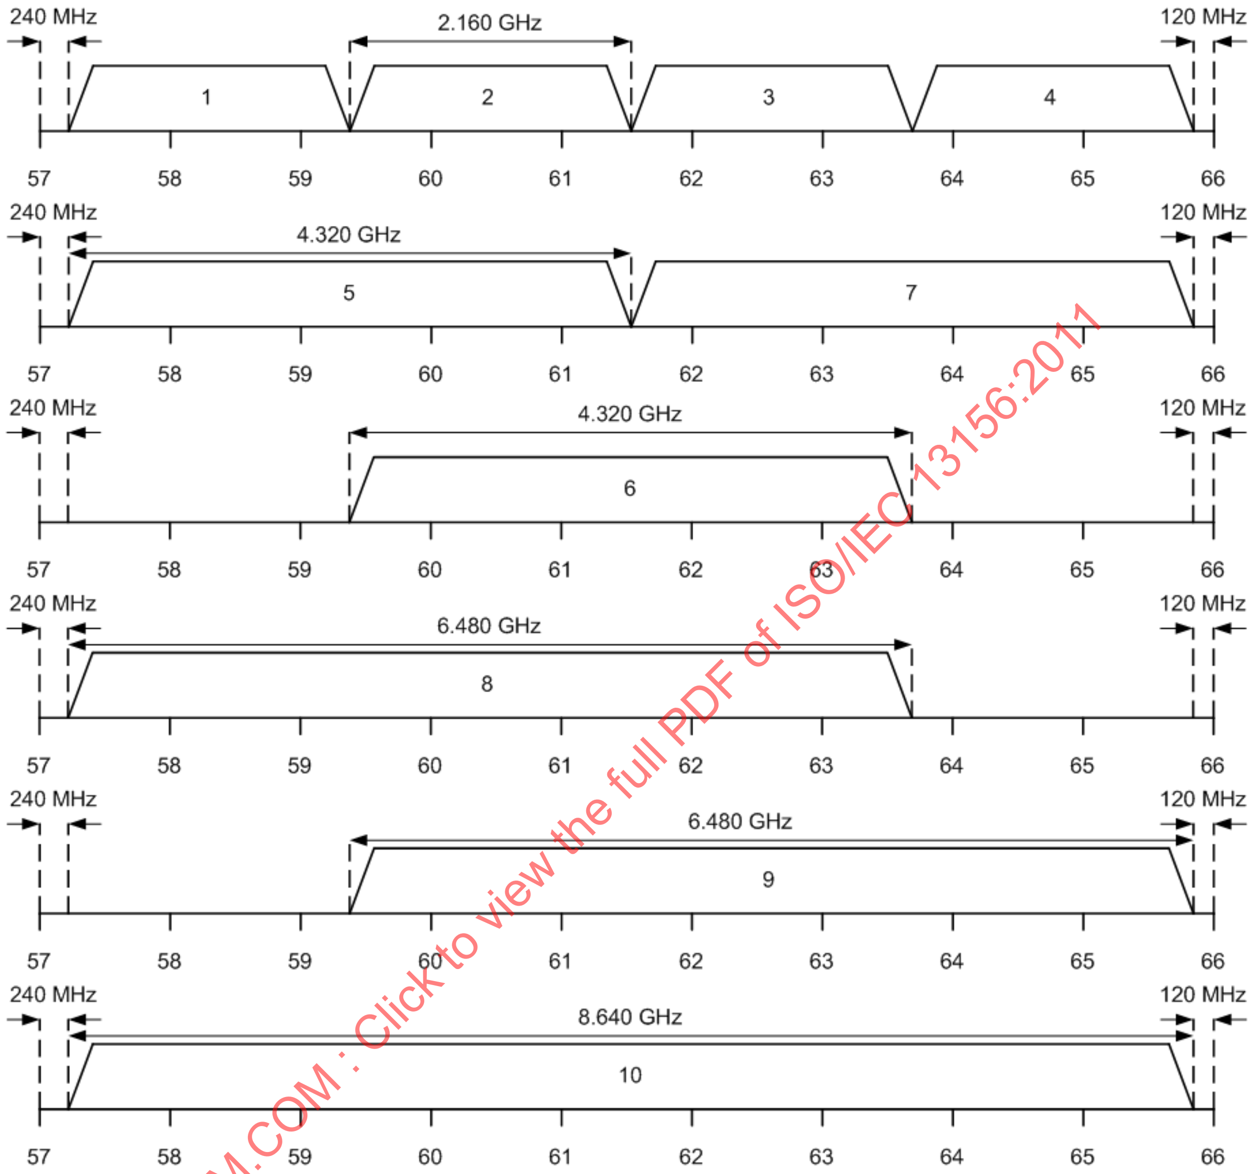

Figure 48 - Channel numbering

11.2 PHY layer timing

The values of PHY layer timing parameters are defined in Table 39.

Table 39 - PHY layer timing parameters

PHY Parameter	Value
pMIFS	888 ns
pSIFS	2666 ns
pCCADetectTime	5037 ns

11.2.1 Receive-to-transmit turnaround time

The RX-to-TX turnaround time shall not be greater than pSIFS. This turnaround time shall be measured at the air interface. The time elapsed from the trailing edge of the last received symbol to the leading edge of the first transmitted symbol of the PLCP preamble for the next frame shall not be greater than pSIFS.

11.2.2 Transmit-to-receive turnaround time

The TX-to-RX turnaround time shall not be greater than pSIFS. This turnaround time shall be measured at the air interface. The time elapsed from the trailing edge of the last transmitted symbol until the receiver is ready to begin the reception of the next PHY frame shall not be greater than pSIFS.

11.2.3 Time between successive transmissions

For uninterrupted successive transmissions by a device, the interframe spacing after the frame shall be pSIFS if Number of segments field is zero, and shall not be less than pMIFS if the Number of segments field is nonzero. The interframe spacing time shall be measured at the air interface. When the Number of segments field is zero, the time elapsed from the trailing edge of the last transmitted symbol to the leading edge of the first transmitted symbol of the PLCP preamble for the following packet shall be equal to pSIFS. When the PLCP length field is nonzero, the time elapsed from the trailing edge of the last transmitted symbol to the leading edge of the first transmitted symbol of the PLCP preamble for the following frame shall not be less than pMIFS.

12 Transmitter specifications

12.1 Transmit PSD mask

12.1.1 Transmit PSD

The transmit spectral mask shall conform to the values as indicated in the Figure 49 and Table 40.

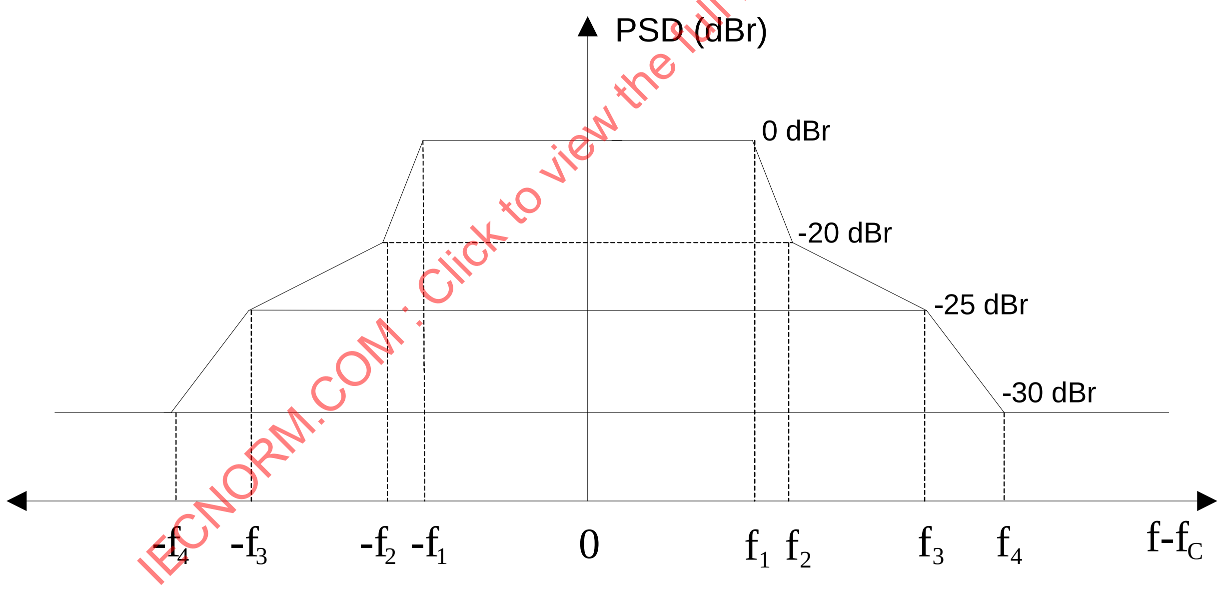

Figure 49 - Transmit spectral mask

Table 40 - Transmit spectral mask limit

Frequency	Relative Limit (dBr)
$\lvert f-f_c \rvert \le f_1$	0
$f_1 \le \lvert f-f_c \rvert \le f_2$	$-20(\lvert f-f_c \rvert-f_1)/(f_2-f_1)$
$f_2 \le \lvert f-f_c \rvert \le f_3$	$-20-5(\lvert f-f_c \rvert-f_2)/(f_3-f_2)$

Table 40 - Transmit spectral mask limit (concluded)

Frequency	Relative Limit (dBr)
$f_3 \le \lvert f-f_c \rvert \le f_4$	$-25-5(\lvert f-f_c \rvert -f_3)/(f_4-f_3)$
$\lvert f-f_c \rvert \ge f_4$	-30

where f_1, f_2, f_3 and f_4 for a single channel transmission, a two bonded channels transmission, a three bonded channels transmission and a four bonded channels transmission are given in Table 41.

Table 41 - Transmit spectral mask requirements

Channel Bonding	f_1 (MHz)	f_2 (MHz)	f_3 (MHz)	f_4 (MHz)
Single Channel Transmission	1050	1080	1500	2000
Two Bonded Channels Transmission	2100	2160	3000	4000
Three Bonded Channels Transmission	3150	3240	4500	6000
Four Bonded Channels Transmission	4200	4320	6000	8000

The transmit spectral mask shall be measured with 1 MHz resolution bandwidth. The transmit spectral mask requirement does not include any carrier leakage.

12.2 Transmit centre frequency tolerance

The transmitted centre frequency tolerance shall be ±20 ppm maximum.

12.3 Symbol clock frequency tolerance

The symbol clock frequency tolerance shall be ±20 ppm maximum.

12.4 Clock synchronization

The transmit centre frequencies and the symbol clock frequency shall be derived from the same reference oscillator.

12.5 Transmit power control

A device should provide support for transmit power control to optimize its power consumption and minimize interferences to other existing links, while still providing a reliable link for the transfer of information. The transmitter should change its transmission power at the receiver's request.

Transmit power shall be changed with a step size granularity of 2dB. The relative accuracy of change in transmit power shall be the maximum of ±1 dB or ±20% of the change (in the dB scale). As an example, for a change of 4 dB and a change of 8 dB, the allowed relative accuracy is ±1.0 dB and ±1.6 dB, respectively.

12.6 Transmitter EVM

12.6.1 Type A

12.6.1.1 SCBT

The Error Vector Magnitude (EVM) defines the average constellation error power with reference to the power of a highest constellation point of the modulation scheme.

Figure 50 illustrates EVM.

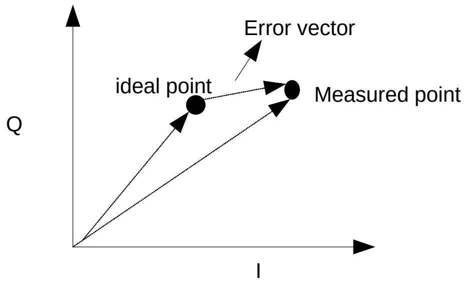

Figure 50 - Illustrative diagram for EVM for one constellation point

This EVM is defined in percentage as

$$EVM = \frac{\sqrt{\frac{1}{N} \cdot \sum_{1}^{N} (\Delta I_i^2 + \Delta Q_i^2)}}{R_{max}} \quad . \tag{72}$$

where N is the number of symbols used in measurement, $\Delta I_i^2 + \Delta Q_i^2$ is the error vector, and R_{max} is the magnitude of the maximum constellation point. The EVM shall be measured over the payload with identical modulation scheme.

Considering the error vector as noise added to thermal and channel noise, the required EVM can be estimated from the transmitter implementation margin. The transmitter implementation margin is the excess power needed at the transmitter to negate the effect of transmitter imperfections so that the received SNR of a real transmitter will be identical to that of an ideal transmitter.

The measurement of the EVM assumes the use of a raised cosine filter with 25% excess bandwidth at the transmitter and a near “ideal” receiver corresponding to an AWGN channel.

The EVM measurement includes imperfections of the transmitter due to transmitter filter inaccuracy, D/A converter, I/Q imbalances, phase noise, and non-linearity of amplifiers.

Based on the assumption on the required margin at transmitter shown in the following table, the EVM is computed using the required SNR and transmitter margin. The results are shown in Table 42. The requirement for EVM shall be as shown in this table based on a near ideal receiver and AWGN channel

.

Table 42 - Maximum allowable EVM values for Type A SCBT modes

Mode	Max. Allowed EVM (dB)	Max. Allowed EVM (%)
A0	-4.8	33.4
A1	-6.3	23.7
A2	-9.5	11.2
A3	-6.4	23.1
A4	-8.4	14.4
A5	-9.6	10.9

Table 42 - Maximum allowable EVM values for Type A SCBT modes (concluded)

Mode	Max. Allowed EVM (dB)	Max. Allowed EVM (%)
A6	-10.9	8.1
A7	-12.0	6.3
A8	-11.1	7.7
A9	-12.5	5.6
A10	-6.4	23.1
A11	-9.5	11.1
A12	-7.4	18.2
A13	-11.1	7.7

12.6.1.1.1 RMS error measurement and calculation

The relative constellation RMS error calculation shall be performed using a device capable of converting the transmitted signal into a stream of samples at 1.728 Gsps or more, with sufficient accuracy in the DC offset, phase noise, etc. The sampled signal shall then be processed in a manner similar to that of an ideal receiver. The necessary steps for receiver processing are listed below.

1. Detect the start of the packet and frame boundary.
2. Estimate the correct sampling time. Correct as needed.
3. Estimate the channel impulse response.
4. Equalize the sampled signal with the estimated channel impulse response.
5. For each of the sampled signal, find the closest constellation point and compute the Euclidean distance.
6. Compute the RMS error, averaged over all data symbols and over all frames, as follows:

$$EVM = \frac{1}{N_f}\sqrt{\frac{\sum_n |R[n]-C[n]|^2}{N_{datasyms}P_0}} \quad (73)$$

where N_f is the number of packets under test, $N_{datasyms}$ is the number of data symbols in the segment (excluding the pilot symbols), P_0 is the average power over all payload data symbols, $C[n]$ is the transmitted n^{th} data symbol, and $R[n]$ is the observed n^{th} data symbol, and the sum is calculated over all the data symbols in the segments.

The RMS error shall be computed over one segment of the frame only. P_0 is re-computed for each frame. The test shall be performed over a minimum of N_f = 100 frames, where the PSDU of each packet is at least 16384 symbols in length and is generated from random data.

12.6.1.2 OFDM

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 given in Table 43.

Table 43 - Permissible relative constellation error

Mode	Modulation	Multipath Margin (dB)	Accepted degradation due to inaccuracies in the constellation points (dB)	Relative constellation error (dB)
A14	QPSK (1/2)	1.5	2	-7.3
A15	QPSK (2/3)	1.5	3	-12.3
A16	16 QAM (2/3)	3	4	-19.6
A17	UEP-QPSK (Coding)	1.5	3	-13.1
A18	UEP-16 QAM (Coding)	3	4	-20.9
A19	UEP-QPSK (Mapping) 2/3	1.5	3	-13.1
A20	UEP-16 QAM (Mapping) 2/3	3	4	-20.9
A21	QPSK MSB-only 2/3	1.5	3	-12.3

The relative constellation error values for each of the data rates and modulations used are based on a multi-path margin listed in Table 43 with an implementation loss of 3 dB. In addition, it is assumed that the degradation due to the relative constellation error can be no more than 2.0 dB, 3.0 dB and 4.0 dB for data rates at 1.008 Gbps, 2.016 and 4.032 Gbps, respectively.

The relative constellation RMS error calculation shall be performed using a device capable of converting the transmitted signal into a stream of complex samples at 2.592 Gsps 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. The necessary steps necessary for receiver processing are listed below:

1. Detect the start of the packet and frame boundary.
2. Estimate the correct sampling time and frequency offset. Correct as needed.
3. Estimate the channel frequency response and equalize the channel.
4. For each of the data and pilot subcarriers, find the closest constellation point and compute the Euclidean distance.
5. Compute the RMS error, averaged over all the data and pilot subcarriers, and over all frames as given by

$$E_{RMS}=\frac{1}{N_f}\sum_{i=1}^{N_f}\sqrt{\sum_{n=N_{preamble}+N_{header}}^{N_{frame}}\left(\frac{\sum_{k=1}^{N_D}\left|R_{D,n}[k]-C_{D,n}[k]\right|^2+\sum_{k=1}^{N_P}\left|R_{P,n}[k]-C_{P,n}[k]\right|^2}{(N_D+N_P)N_{frame}P_0}\right)} \quad (74)$$

where N_f is the number of packets under test, N_{packet} is the number of symbols in the packet, $N_{preamble}$ is the number of symbols in the PLCP preamble, N_{header} is the number of symbols in the PLCP header, N_{frame} = N_{packet} - $N_{preamble}$ - N_{header} is the number of symbols in the PSDU, N_D is the number of data subcarriers, N_P is the number of pilot subcarriers, P_0 is the average power of the data and pilot constellations, $C_{D,n}[k]$ and $C_{P,n}[k]$ are the transmitted k^{th} data subcarrier and k^{th} pilot subcarrier for the n^{th} OFDM symbol, respectively, and $R_{D,n}[k]$ and $R_{P,n}[k]$ are the observed k^{th} data subcarrier and k^{th} pilot subcarrier for the n^{th} OFDM symbol, respectively. The values for N_D and N_P are defined in Table 16.

12.6.2 Type B

The relative constellation RMS error, averaged over all data and over all of the frames, shall not exceed the values given in Table 44 when using the method of 12.6.1.1.1.

The measurement of the EVM assumes the use of a raised cosine filter with 25% excess bandwidth at the transmitter and a near ideal receiver corresponding to an AWGN channel.

The EVM measurement includes imperfections of the transmitter due to transmitter filter inaccuracy, D/A converter, I/Q imbalances, phase noise, and non-linearity of amplifiers.

Table 44 - Permissible relative constellation error

Modes	Relative Constellation Error (dB)	Relative Constellation Error (%)
B0	-7	20
B1	-7	20
B2	-8.2	15
B3	-8.2	15

13 Receiver specification

13.1 Type A device

13.1.1 SCBT receiver sensitivity

For a packet error rate (PER) less than 8% with payload length of 2048 octets, the minimum receiver sensitivity numbers (expressed in EIR$_x$P) in AWGN for modes A0 through A13 are listed in Table 45.

Table 45 - Minimum receiver sensitivity for different SCBT modes

Mode	Minimum Receiver Sensitivity (dBm)
A0	-60.0
A1	-57.0

Table 45 - Minimum receiver sensitivity for different SCBT modes (concluded)

Mode	Minimum Receiver Sensitivity (dBm)
A2	-50.5
A3	-53.9
A4	-49.8
A5	-47.4
A6	-45.7
A7	-43.5
A8	-43.5
A9	-40.7
A10	-53.9
A11	-43.5
A12	-52.3
A13	-43.5

13.1.2 OFDM receiver sensitivity

For a packet error rate (PER) less than 8% with payload length of 2048 octets, the minimum receiver sensitivity numbers (expressed in EIR_xP) in AWGN for modes A14 up to and including A21 are listed in Table 46.

Table 46 - Minimum receiver sensitivities for the different modes

Mode	Minimum Receiver Sensitivity (dBm)
A14	-60.6
A15	-54.3
A16	-50.2
A17	-56.0
A18	-52.0
A19	-54.3
A20	-50.2
A21	-54.3

13.2 Type B device receiver sensitivity

For a packet error rate (PER) less than 8% with payload length of 2048 octets, the minimum receiver sensitivity numbers (expressed in EIR_xP) in AWGN for modes B0 up to and including B3 are listed in Table 47.

Table 47 - Minimum receiver sensitivity for different modes

Mode	Minimum Receiver Sensitivity (dBm)
B0	-60.7

Table 47 - Minimum receiver sensitivity for different modes (concluded)

Mode	Minimum Receiver Sensitivity (dBm)
B1	-57.7
B2	-54.6
B3	-54.6

13.3 Receiver CCA performance

Transmissions at a receiver at an EIR_xP a receiver level equal to or greater than -85.0 dBm shall cause CCA to indicate the channel is busy with a probability greater than 90% within pCCADetectTime.

14 Antenna training symbols and feedback methods

Devices transmit antenna training symbols to each other to perform antenna training or tracking. Training symbols are transmitted in the ATS field in variable length PHY header (10.1.4). Upon reception of training symbols, a training peer device determines 1) its receiving antenna weight settings that it will use subsequently, and, if requested, 2) the transmitting antenna weight settings of the device that sends the antenna training symbols.

For training receiving antenna, the transmitter shall include a number of training symbols in the ATS of a single training frame.

For training transmit antenna, there are two modes:

1.**closed-loop mode:** a device determines its transmitting antenna weight settings based on feedback from the other device;

2.**open-loop mode:** a device uses the same antenna weight settings for transmission as for reception.

14.1 Antenna training sequence transmission

The antenna training symbols are transmitted in the ATS field in the variable length PHY header of a training frame. The antenna training indicator field (ATIF) in the variable length PHY header 10.1.2.1 carries the information about ATS field such as transmission spreading factor (DISC_REP) and numbers of training symbols for receiving (NUM_RXTS) as well as transmitting (NUM_TXTS) antennas. Existence of the ATIF is indicated through the ATIF_EXISTENCE field in the fixed PHY header.

The transmission of training sequences from a source device to a destination device is impacted by the negotiated antenna training configurations of source and destination devices and the state of an iterative training process. Table 48 specifies the inclusion of ATS (for receiving and/or transmitting antenna) and training/tracking feedback in training frame exchange based on the antenna training configuations of the initiator(I) and responder(R) devices.

The antenna training configuraiton is negociated through the RTT/CTT handshake (16.18.2). The Request Training and Request Feedback fields in the ATIE carried in RTT/CTT frames determines inclusion of ATS in training frames and transmission of feedback frames.

Table 48 - Configurability of antenna training/tracking frame exchange

ATIE fields		*R configuration*		
		Request Training = ONE *Request Feedback = ONE*	*Request Training = ONE* *Request Feedback = ZERO*	*Request Training = ZERO*
I configuration	*Request Training = ONE* *Request Feedback = ONE*	I sends: ATS (receiver, transmitter), Feedback R sends: ATS (receiver, transmitter), Feedback	I sends: ATS(receiver, transmitter), R sends: ATS(receiver), Feedback	I sends: ATS(transmitter) R sends: ATS(receiver), Feedback
	Request Training = ONE *Request Feedback = ZERO*	I sends: ATS (receiver), Feedback R sends: ATS(transmitter, receiver)	I sends: ATS(receiver) R sends: ATS(receiver)	R sends: ATS(receiver)
	Request Training = ZERO	I sends: ATS (receiver), Feedback R sends: ATS (transmitter)	I sends: ATS(receiver)	no ATS and Feedback frames are sent

For transmitting antenna weight training in closed-loop mode, the transmitter shall include a number of training symbols in the ATS of a single frame. This number shall be equal to the number of columns of the training matrix (14.1.1), and shall be signaled via the NUM_TXTS field of the ATIF. Each of these symbols is sent with different antenna weights, according to the training matrix.

For receiving antenna weight training, the transmitter shall include a number of training symbols in the ATS of a single training frame. The number of training symbols shall be signalled via the NUM_RXTS field of the ATIF. These training symbols shall all be transmitted using the same transmit antenna weights.

14.1.1 Training matrix in closed-loop mode

A transmitter with N antenna elements shall train its transmit antenna weights using a training matrix T of size N by K as defined below. The antenna weight of the n^{th} antenna element in the transmission of the k^{th} training symbol for transmit antenna training ($n = 1,\ldots,N$; $k = 1,\ldots,K$) shall be equal to the matrix element T_{nk}. K is the number of training symbols and is determined as below.

The complex Hadamard matrices $H(K)$, listed below, satisfy $H(K)H(K)' = KI_K = H(K)'H(K)$ and have the property that all matrix elements take values in $\{1, j, -1, -j\}$. This eases the implementation of the corresponding antenna weights. These matrices exist when K is even.

For a transmitter with N antennas, K shall be chosen follows: if $N \leq 16$, K is chosen as the lowest even number that is not smaller than N, i.e., $K = N$ if N is even and $K = N+1$ if N is odd. If $N>16$, K is chosen as the lowest multiple of 4 that is not smaller than N. Finally, the training matrix T shall be chosen as the first N rows of $H(K)$.

The following lists the complex Hadamard matrix used to construct the training matrix.

$$H(2) = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \qquad (75)$$

$$H(6)=\begin{pmatrix}1&1&1&1&1&1\\1&-1&j&-j&-j&j\\1&j&-1&j&-j&-j\\1&-j&j&-1&j&-j\\1&-j&-j&j&-1&j\\1&j&-j&-j&j&-1\end{pmatrix} \quad (76)$$

$$H(10)=\begin{pmatrix}1&1&1&1&1&1&1&1&1&1\\1&-1&-j&-j&-j&-j&j&j&j&j\\1&-j&-1&j&j&-j&-j&-j&j&j\\1&-j&j&-1&-j&j&-j&j&-j&j\\1&-j&j&-j&-1&j&j&-j&j&-j\\1&-j&-j&j&j&-1&j&j&-j&-j\\1&j&-j&-j&j&j&-1&-j&-j&j\\1&j&-j&j&-j&j&-j&-1&j&-j\\1&j&j&-j&j&-j&-j&j&-1&-j\\1&j&j&j&-j&-j&j&-j&-j&-1\end{pmatrix} \quad (77)$$

$$H(14)=\begin{pmatrix}1&1&1&1&1&1&1&1&1&1&1&1&1&1\\1&-1&j&-j&j&j&-j&-j&-j&-j&j&j&-j&j\\1&j&-1&j&-j&j&j&-j&-j&-j&-j&j&j&-j\\1&-j&j&-1&j&-j&j&j&-j&-j&-j&-j&j&j\\1&j&-j&j&-1&j&-j&j&j&-j&-j&-j&-j&j\\1&j&j&-j&j&-1&j&-j&j&j&-j&-j&-j&-j\\1&-j&j&j&-j&j&-1&j&-j&j&j&-j&-j&-j\\1&-j&-j&j&j&-j&j&-1&j&-j&j&j&-j&-j\\1&-j&-j&-j&j&j&-j&j&-1&j&-j&j&j&-j\\1&-j&-j&-j&-j&j&j&-j&j&-1&j&-j&j&j\\1&j&-j&-j&-j&-j&j&j&-j&j&-1&j&-j&j\\1&j&j&-j&-j&-j&-j&j&j&-j&j&-1&j&-j\\1&-j&j&j&-j&-j&-j&-j&j&j&-j&j&-1&j\\1&j&-j&j&j&-j&-j&-j&-j&j&j&-j&j&-1\end{pmatrix} \quad (78)$$

The complex Hadamard matrix *H(K)* for K=4, 8, 12, 16, 20, 24, 28 and 32 is formed as the Kronecker product $H(K/2)\otimes H(2)$. Finally *H(36)* equals the Kronecker product $H(6)\otimes H(6)$. Here the

Kronecker product of a matrix *A* of size *p x q* and a matrix *B* of size *r x s* is the matrix $C = A \otimes B$ of size *pr x qs* with elements $C_{(k-1)r+m,(l-1)s+n} = A_{k,l}B_{m,n}$, where *k=1,...,p*, *l=1,...,q*, *m=1,...,r*, and *n=1,...,s*.

14.1.2 Tracking matrix

The training matrix for antenna tracking shall be of size *N* x 3 and shall be chosen as $T_{nk} = W_n^{current} \exp[-j2\pi(r_n - 1)(c_k - 1)/M]$, where *M* and the numbers r_n are as in the Fourier codebook, clause 14.2.1, and the numbers c_k are given by $c_1 = M, c_2 = 1, c_3 = 2$. Here $W_n^{current}$ stands for the transmit antenna weights that the transmitter is currently using. The three columns of the training matrix correspond to an angular sweep around the current antenna weights.

14.2 Antenna training feedback in closed-loop mode

In closed-loop mode specified in Table 48, after receiving the training frame containing the K training symbols, the receiver shall give feedback about the transmitting antenna weight settings the transmitter shall use in subsequent transmissions, Such feedback shall be sent in one of the following forms:

- an index of best received training symbols.
- an index of a codeword in a codebook
- quantised antenna weights

Table 49 summarizes the possible modes of feedback and indicates whether they are mandatory or optional for the transmitter, i.e., whether the transmitter shall be able to interpret that kind of feedback in closed loop mode, and the receiver, i.e., whether the receiver shall be able to give that kind of feedback when requested to do so.

Table 49 - Summary of feedback methods

Kind of feedback	Mandatory/optional for transmitter (interpreting feedback)	Mandatory/optional for receiver (giving feedback)
index of best received training symbol	mandatory for Type A	optional for Type B
transmission weights with Fourier codebook	mandatory for Type A	mandatory for Type A
transmission weights with Walsh codebook	optional for Type A	optional for Type A
transmission weights with quantised coefficients	optional for Type A	optional for Type A

14.2.1 Index feedback

In index feedback, the feedback is the index of best received training symbol. It is also the column of the training matrix which gave the best reception quality. This index i_{max} is encoded using 8 bits as specified in Table 50.

14.2.2 Codebook based feedback

There are two types of codebook, the Fourier codebook and the Walsh codebook. Their constructions are described below.

The transmit beamforming vector ν is estimated at the receiver side. In order to feed the knowledge of ν back to the transmitter, a vector quantization approach is used. Let ***C*** be the Fourier or Walsh codebook matrix of size *NxM* as defined below and let c_i denote the i^{th} column of ***C***. The receiver determines the index i_{max} as $i_{\max} = \arg\max_{1 \le i \le M} \left| c_i^{\ H} v \right|$, i.e., i_{max} is the index of the codeword that is most correlated to the estimated beamforming vector v. Then the binary representation of the index i_{max} shall be fed back to the transmitter, using 8 bits as specified in Table 50.

Table 50 - Representation of the beamforming index

i_{max}	$b_7,...,b_0$	i_{max}	$b_7,...,b_0$	i_{max}	$b_7,...,b_0$	i_{max}	$b_7,...,b_0$
1	00000000	33	00100000	65	01000000	97	01100000
2	00000001	34	00100001	66	01000001	98	01100001
3	00000010	35	00100010	67	01000010	99	01100010
4	00000011	36	00100011	68	01000011	100	01100011
5	00000100	37	00100100	69	01000100	101	01100100
6	00000101	38	00100101	70	01000101	102	01100101
7	00000110	39	00100110	71	01000110	103	01100110
8	00000111	40	00100111	72	01000111	104	01100111
9	00001000	41	00101000	73	01001000	105	01101000
10	00001001	42	00101001	74	01001001	106	01101001
11	00001010	43	00101010	75	01001010	107	01101010
12	00001011	44	00101011	76	01001011	108	01101011
13	00001100	45	00101100	77	01001100	109	01101100
14	00001101	46	00101101	78	01001101	110	01101101
15	00001110	47	00101110	79	01001110	111	01101110
16	00001111	48	00101111	80	01001111	112	01101111
17	00010000	49	00110000	81	01100000	113	01110000
18	00010001	50	00110001	82	01010001	114	01110001
19	00010010	51	00110010	83	01010010	115	01110010
20	00010011	52	00110011	84	01010011	116	01110011
21	00010100	53	00110100	85	01010100	117	01110100
22	00010101	54	00110101	86	01010101	118	01110101
23	00010110	55	00110110	87	01010110	119	01110110
23	00001110	47	00101110	79	01001110	111	01101110
24	00010111	56	00110111	88	01010111	120	01110111
25	00011000	57	00111000	89	01011000	121	01111000
26	00011001	58	00111001	90	01011001	122	01111001
27	00011010	59	00111010	91	01011010	123	01111010
28	00011011	60	00111011	92	01011011	124	01111011
29	00011100	61	00111100	93	01011100	125	01111100

Table 50 - Representation of the beamforming index (concluded)

i_{max}	$b_7,...,b_0$	i_{max}	$b_7,...,b_0$	i_{max}	$b_7,...,b_0$	i_{max}	$b_7,...,b_0$
30	00011101	62	00111101	94	01011101	126	01111101
31	00011110	63	00111110	95	01011110	127	01111110
32	00011111	64	00111111	96	01011111	128	11111111

Upon receiving the index i_{max}, the transmitter shall use the corresponding codeword $c_{i_{\max}}$ to approximate the actual beamforming vector ν.

14.2.2.1 Construction of Fourier codebook

The Fourier codebook is constructed based on a Fourier transformation matrix and a difference set. Let N be the number of transmit antennas. Compute $M = 2^{2+\lfloor \log_2(N) \rfloor} - 1$ and $K = 2^{1+\lfloor \log_2(N) \rfloor} - 1$.

Define the Fourier matrix $\mathbf{F}_M$ of size MxM as $[\mathbf{F}_M]_{n,m} = \frac{1}{\sqrt{M}} \exp[-j2\pi(n-1)(m-1)/M]$. The Fourier codebook matrix $\mathbf{C}_F$ shall be chosen as the normalized NxM submatrix of $\mathbf{F}_M$ such that the chosen column indices are {1,2,...,M} and the chosen row indices are {$r1,r2,\ldots,r_N$}, where $R_K = \{r_1, r_2, \ldots, r_K\}$ is a modulo M difference set as given in Table 51, i.e.

$$[\mathbf{C}_F]_{n,m} = \frac{1}{\sqrt{N}} \exp[-j2\pi(r_n-1)(m-1)/M].$$

Table 51- Parameters of Fourier codebook matrix

N	M	K	$r_1,...,r_K$
2,3	7	3	2, 3, 5
4,...,7	15	7	1, 6, 8, 11, 12, 14, 15
8,...,15	31	15	2, 3, 4, 5, 7, 9, 13, 16, 17, 18, 24, 25, 28, 30, 31
16,...31	63	31	1, 8, 10, 12, 15, 16, 19, 23, 26, 28, 29, 31, 32, 36, 37, 38, 40, 45, 46, 48, 50, 51, 52, 55, 56, 57, 59, 60, 61, 62, 63
32,...,36	127	63	2, 3, 4, 5, 6, 7, 9, 10, 11, 12, 13, 16, 17, 18, 19, 21, 23, 24, 25, 30, 31, 33, 34, 35, 37, 40, 41, 45, 47, 49, 50, 56, 58, 59, 60, 61, 65, 66, 67, 69, 70, 72, 73, 76, 79, 81, 84, 89, 92, 93, 94, 97, 99, 100, 102, 106, 110, 111, 114, 115, 117, 119, 121

14.2.2.2 Construction of Hadamard codebook

The Hadamard codebook is constructed based on the Hadamard matrix $H(2^k)$ of size $2^k x 2^k$, as given in 14.1.2, with the addition that $H(64) = H(32) \otimes H(2)$. Let N be the number of transmit antennas. Compute $M = 2^{\lceil \log_2 N \rceil}$. The Hadamard codebook $\mathbf{C}_H$ shall be chosen as the NxM submatrix of the Hadamard matrix $H(M)$ such that the chosen column indices are {1,2,3,...,M} and the chosen row indices are {1,2,3,...,N}. The (n,m) element of C_π is specified as $[\mathbf{C}_H]_{n,m} = \sqrt{\frac{1}{N}} \cdot [H(M)]_{n,m}$.

14.2.3 Quantised weights

The phases d_i and amplitudes a_i of the N components of the estimated transmit beamforming vector ν are uniquely determined by $v_i = a_i \exp(j\pi d_i/180)$, where $a_i \geq 0$ and $-11.25^\circ \leq d_i < 348.75^\circ$. The receiver calculates $n_i = \left[\frac{d_i}{22.5^\circ}\right]$, where [.] denotes rounding to the nearest integer. Each n_i is an integer in {0,1,,…,15}.

The feedback shall consist of N octets, the first four bits, $b_0,\ldots,b_3$ of the i[th] octet shall comprise the 4-bit representation of n_i, as specified in Table 52, the last four bits $b_4,\ldots,b_7$ are reserved for future use and shall be set to zero.

Table 52 - Quantised weights feedback

n_i	$b_3,\ldots,b_0$	n_i	$b_3,\ldots,b_0$
0	0000	8	1000
1	0001	9	1001
2	0010	10	1010
3	0011	11	1011
4	0100	12	1100
5	0101	13	1101
6	0110	14	1110
7	0111	15	1111

15 MAC frame formats

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. The final Clause contains a list of information elements that may appear in beacon frames and some command frames. Octets of the EUI [3] are passed to the PHY SAP in ascending index-value order.

15.1 Frame format conventions

The following conventions and definitions apply throughout this Clause.

15.1.1 Figures

MAC frames are described as a sequence of fields in a specific order. Figures in this clause depict fields in the order they are delivered to the PHY SAP, from left to right, where the left-most field is transmitted first in time. In field figures, bits within the field are numbered from the least-significant bit on the right to the most-significant bit on the left.

An example sequence of fields is specified in Figure 51.

octets: 2	1	…	4
First field transmitted (2 octets)	Second field transmitted (1 octet)	…	Last field transmitted (4 octets)

Figure 51 - Example sequence of fields

15.1.2 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.

An example of a bitmap specification for a two-octet field is specified in Figure 52.

bits: b15-b13	b12-b8	b7-b0
Most-significant bits of second octet transmitted	Least-significant bits of second octet transmitted	First octet transmitted

Figure 52 - Example bitmap specification for a field

15.1.3 Encoding

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

A bitmap is a sequence of bits, labelled 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.

15.2 General MAC frame format

Frame transactions that do not employ aggregation shall use unaggregated MAC frames, while frame transactions that employ aggregation shall use aggregated MAC frame format.

15.2.1 Unaggregated MAC frame

An unaggregated MAC frame consists of a fixed-length MAC Header and an optional variable-length MAC Frame Body.

The Fixed-Length MAC header is specified in Figure 53.

octets: 2	2	2	2	2
Frame Control	DestAddr	SrcAddr	Access Information	Sequence Control

Figure 53 - Fixed-length MAC Header format

The overall length of the MAC Header of an unaggregated MAC frame is 10 octets.

The MAC frame body, when present, contains a Frame Payload and Frame Check Sequence (FCS) as shown in Figure 54.

octets: L	4
Frame Payload	FCS

Figure 54 - MAC frame body format

In secure frames the frame payload includes security fields as shown in Figure 55. The left-most four fields in Figure 55 are collectively referred to as the security header.

octets: 3	1	2	6	P	16
Temporal Key Identifier (TKID)	Security Reserved	Encryption Offset (EO)	Secure Frame Number (SFN)	Secure Payload	Message Integrity Code (MIC)

Figure 55 - Frame payload field format for secure frames

The length of the frame body shall be limited to pMaxFrameBodySize. 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 this Clause, a reference to the payload of a frame indicates the Frame Payload field of a non-secure frame, or the Secure Payload field of a secure frame. The payload is a sequence of octets labelled as payload[0] through payload[P-1]. Octets are passed to the PHY SAP in ascending index-value order.

15.2.2 Aggregated MAC frame

An aggregated MAC frame consists of a variable-length MAC header and a variable-length MAC Frame Body. The MAC header of an aggregated MAC frame is specified in Figure 56.

octets: 2	2	2	2	2+4N
Frame Control	DestAddr	SrcAddr	Access Information	Aggregation Header

Figure 56- Variable-length MAC header of an aggregated MAC frame

The ACK Policy of the Frame Control field in the MAC header of any aggregated MAC frames shall not be set to B-ACK (See 15.2.3.3). The aggregation header field is specified in Figure 57.

octets: 2	2	2	...	2	2
MSDU Count	Sequence Control of MSDU 1	Length of MSDU 1	...	Sequence Control of MSDU N	Length of MSDU N

Figure 57 - Aggregation header

The MSDU Count field contains the number of MSDUs included in the aggregated frame.

The Length fields in the Aggregation Header field indicate the length in octets of the corresponding MSDUs. The lengths do not include the Pad octets.

The overall length of the MAC Header of an aggregated MAC frame is 10+N*4 octets.

The MAC Frame Body of an aggregated MAC frame contains multiple MSDUs, each aligned to a 4-octet boundary, and their corresponding FCS. The aggregated MAC frame body is specified in Figure 58.

octets: M_1	4	0-3	M_2	4	...	0-3	Mn	4
MSDU 1	FCS of MSDU 1	Pad to 4-octet boundary	MSDU 2	FCS of MSDU 2	...	Pad to 4-octet boundary	MSDU N	FCS of MSDU N

Figure 58 - Frame body for aggregated MAC frames

The length of the frame body shall be limited to pMaxFrameBodySize.

15.2.3 Frame control

The Frame Control field is specified in Figure 59.

bits: b15	b14	b13	b12-b9	b8-b6	b5-b4	b3	b2-b0
Tracking Indication	Aggregation Request	Retry	Frame Subtype / Delivery ID	Frame Type	ACK Policy	Secure	Protocol Version

Figure 59 - Frame control field format

15.2.3.1 Protocol version

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

15.2.3.2 Secure

The Secure bit is set to ONE in a secure frame, which is protected using the temporal key specified by the Temporal Key Identifier (TKID). The Secure bit is set to ZERO otherwise. Frames with the Secure bit set to ONE use the Frame Payload format for secure frames as shown in Figure 55. Valid settings for the Secure bit in each frame type are listed in Table 103 in 17.2.

15.2.3.3 ACK policy

The ACK Policy field is set to the type of acknowledgement requested by the transmitter. Acknowledgement procedures are described in 16.12. The allowed values for the ACK Policy field are defined in Table 53.

Table 53 - 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 regard for the actual result. The use of this policy is defined in 16.12.1.
1	Imm-ACK	The addressed recipient returns an Imm-ACK frame after correct reception, according to the procedures defined in 16.12.2.
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 defined in 16.12.3.
3	B-ACK Request	The addressed recipient returns a B-ACK frame after reception, according to the procedures defined in 16.12.3.

15.2.3.4 Frame type

The Frame Type field is set to the type of frame that is being sent. Table 54 lists the valid frame type values, descriptions, and the Clauses that describe the format and use of each of the individual frame types.

Table 54 - Frame type field encoding

Value	Frame type	Clause
0	Beacon frame	15.3

Table 54 - Frame type field encoding (concluded)

Value	Frame type	Clause
1	Discovery frame	15.4
2	Control frame	15.5
3	Command frame	15.6
4	Data frame	15.7
5	Aggregated MAC frame	15.8
6 - 7	Reserved	

If the Frame type field is set to Aggregated data frame, a variable-length MAC header that is described in 15.2.2 shall be used.

15.2.3.5 Frame subtype / delivery ID

The Frame Subtype / Delivery ID field is used to assist a receiver in the proper processing of received frames. In control or command frames, this field is used as Frame Subtype, as defined in Table 60 in 15.5 and Table 62 in 15.6. In data frames and aggregated data frames, this field is used as Delivery ID and set to Stream Index. This field is reserved in all other frame types.

15.2.3.6 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 in all other frame types.

15.2.3.7 Aggregation request

The Aggregation Request bit is set to ONE in a data frame to initiate an instance of aggregation mechanism as specified in 16.10. The Aggregation Request is set to ZERO in data frames transmitted at all other times. It is reserved in all other frame types.

15.2.3.8 Tracking indication

The Tracking Indication bit is set to ONE in a control frame, or an aggregated tracking frame to indicate the inclusion of ATTIE and/or AFIE in the last MSDU for antenna tracking purpose. The bit is set to ZERO otherwise, and is reserved in all other types of frames.

15.2.4 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. DevAddr values are described in 16.1.1.

15.2.5 SrcAddr

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

15.2.6 Sequence control

The Sequence Control field identifies the order of MSDUs/MCDUs and their fragments. The Sequence Control field is specified in Figure 60. The Sequence Control field is reserved in control frames.

bits: b15	b14	b13-b3	b2-b0
Reserved	More Fragments	Sequence Number	Fragment Number

Figure 60 - Sequence control field format

15.2.6.1 Fragment number

The Fragment Number field is set to the number of the fragment within the MSDU or MCDU. The fragment number is zero in the first or only fragment of an MSDU or MCDU and is incremented by one for each successive fragment of that MSDU or MCDU.

15.2.6.2 Sequence number

The Sequence Number field is set to the sequence number of the MSDU or MCDU, as defined in 16.1.9.3.

The Sequence Number field is used for duplicate frame detection, as described in 16.1.7, and to preserve frame order when using the B-ACK mechanism, as described in 16.12.3.

The Sequence Number field is reserved in control frames.

15.2.6.3 More fragments

The More Fragments field is set to ZERO to indicate that the current fragment is the sole or final fragment of the current MSDU or MCDU; otherwise, the field is set to ONE.

15.2.7 Access information

The Access information field is specified in Figure 61.

bits: b15	b14	b13-b0
Access Method	More Frames	Duration

Figure 61 - Access information field format

15.2.7.1 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 in units of microseconds. The duration value is set as defined in 16.1.9.1.

15.2.7.2 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, discovery and control frames. Additional rules regarding the More Frames field are specified in 16.1.9.2.

15.2.7.3 Access method

The Access Method bit is set to ONE in all frames transmitted within a hard or private DRP reservation block by the reservation owner or target prior to the release of the reservation block.

The Access Method bit in an Imm-ACK, B-ACK or CTT 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 discovery frames.

The Access Method bit is reserved in beacon frames.

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

15.2.8 Frame payload

The Frame Payload field is a variable length field that carries the information that is to be transferred to a device or group of devices. In a secure frame, it includes the required security fields as shown in Figure 55 and defined below.

15.2.8.1 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.

15.2.8.2 Security reserved

The Security Reserved field is reserved, but included in authentication of the frame.

15.2.8.3 Encryption offset (EO)

The Encryption Offset field indicates where encryption starts, in octets, relative to the beginning of the Secure Payload, as shown in Figure 55. A value of zero indicates that the entire Secure Payload is encrypted. A non-zero value in this field indicates that the first EO octets of the Security Payload are not encrypted. Regardless of the value of this field, the entire Secure Payload, along with other appropriate fields, is authenticated by the MIC.

15.2.8.4 Secure frame number (SFN)

The SFN field provides message freshness as a defence against replay attacks. The SFN field in a secure frame is set to the next value of the sender's secure frame counter (SFC) for the temporal key used by this frame. SFC setting and replay protection are described in 17.4.2.

15.2.8.5 Secure payload

The Secure Payload field in 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.

15.2.8.6 Message integrity code (MIC)

The MIC field contains an 16-octet cryptographic checksum used to protect the integrity of the MAC Header and Frame Payload.

15.2.9 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:

$$G(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^{8} + x^{7} + x^{5} + x^{4} + x^{2} + x + 1 \quad (79)$$

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

- The remainder resulting from $x^k \times (x^{31} + x^{30} + \ldots + x + 1)$ 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. Figure 62 illustrates the encoding of the FCS field for the CRC polynomial:

$$a_{31}x^{31} + a_{30}x^{30} + a_{29}x29 + \ldots + a_2x2 + a_1x + a_0 \quad (80)$$

bits: b31	b30	b29	…	b2	b1	b0
a_0	a_1	a_2	…	a_{29}	a_{30}	a_{31}

Figure 62 - FCS field encoding

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:

$$x^{31} + x^{30} + x^{26} + x^{25} + x^{24} + x^{18} + x^{15} + x^{14} + x^{12} + x^{11} + x^{10} + x^8 + x^6 + x^5 + x^4 + x^3 + x + 1 \quad (81)$$

15.3 Beacon frames

MAC header field settings for beacon frames are described in Table 55. Beacon frames are also referred to as beacons throughout this specification.

Table 55 - MAC header field values for beacon frames

Header field	Value
Protocol Version	0
Secure	0
ACK Policy	0 (No-ACK)
Frame Type	0 (beacon frame)
Frame Subtype / Delivery ID	Reserved
Retry	Reserved
Aggregation Request	Reserved
Tracking Indication	Reserved
DestAddr	BcstAddr
SrcAddr	DevAddr of the transmitter
Sequence Control	As defined in 15.2.6 and 16.1.9.3
Duration	As defined in 15.2.7.1 and 16.1.9.1
More Frames	Reserved
Access Method	Reserved

The beacon frame payload is specified in Figure 63.

octets: 9	L_1	...	L_N
Beacon Parameters	Information Element 1	...	Information Element N

Figure 63 - Payload format for beacon frames

The information elements (IEs) that may be included in a beacon frame are listed in Table 74 in 15.9. IEs are included in order of increasing Element ID, except for ASIEs. ASIEs do not appear prior to any IE with Element ID zero through seven, but may appear anywhere after those IEs. DRP IEs that have the same Target DevAddr and Stream Index are adjacent to each other in the beacon.

The beacon parameters field is specified in Figure 64.

octets: 6	1	1	1
Device Identifier	Beam Identifier	Beacon Slot Number	Device Control

Figure 64 - Beacon parameters field format

The device identifier field is set to the EUI-48 [3] of the device sending the beacon. A device may use a NULL EUI-48 value (all bits set to ONE) to indicate it does not have a unique EUI-48 value.The EUI is a sequence of 6 octets, labelled as eui[0] through eui[5]. The first three octets (eui[0] through eui[2]) are the manufacturer's OUI, and the last three octets (eui[3] through eui[5]) are the manufacturer-selected extension identifier.

The Beam Identifier field is set to a value that uniquely identifies the beam that the device uses to transmit this beacon frame.

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], except in beacons sent in signalling slots. In signalling slots it is set to the number of the device's non-signalling beacon slot.

The Device Control field is specified in Figure 65.

bits: b7-b6	b5	b4	b3-b2	b1	b0
Security Mode	Reserved	Forced Sync	Status	Signaling Slot	Movable

Figure 65 - Device control field format

The Movable bit is set to ONE if the beacon is movable according to 16.5.3.10, and is set to ZERO otherwise.

The Signalling Slot bit is set to ONE if the beacon frame is sent in a signalling beacon slot according to 16.5.3.6, and is set to ZERO otherwise.

The Status field is set to the following values to indicate the status or function of the beacon frame.

Table 56 - Status field meaning

Value	Meaning
0	Ready

Table 56 - Status field meaning (concluded)

Value	Meaning
1	Preemptive
2	Dual

The Forced Sync bit is set ONE to indicate the device is a forced synchronization device, according to 16.8, and is set to ZERO otherwise.

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

15.4 Discovery frames

MAC header field settings for discovery frames are described in Table 57.

Table 57 - MAC header field values for discovery frames

Header field	Value
Protocol Version	0
Secure	0
ACK Policy	0 (No-ACK)
Frame Type	1 (discovery frame)
Frame Subtype / Delivery ID	Reserved
Retry	Reserved
Aggregation Request	Reserved
Tracking Indication	Reserved
DestAddr	BcstAddr
SrcAddr	DevAddr of the transmitter
Sequence Control	As defined in 15.2.6 and 16.1.9.3
Duration	As defined in 15.2.7.1 and 16.1.9.1
More Frames	Reserved
Access Method	Reserved

The discovery frame payload is specified in Figure 66.

octets: 7	L_1	...	L_N
Discovery Parameters	Information Element 1	...	Information Element N

Figure 66 - Payload format for discovery frames

The information elements (IEs) that may be included in a discovery frame are listed in Table 74 in 15.9. IEs are included in order of increasing Element ID, except for ASIEs. ASIEs do not appear prior to any IE with Element ID zero through seven, but may appear anywhere after those IEs. DRP IEs that have the same Target DevAddr and Stream Index are adjacent to each other in the discovery frame.

The Discovery Parameters field is specified in Figure 67.

octets: 6	1
Device Identifier	Device Control

Figure 67 - Discovery parameters field format

The device identifier field is set to the EUI-48 [3] of the device sending the beacon. A device may use a NULL EUI-48 value (all bits set to ONE) to indicate it does not have a unique EUI-48 value.The EUI is a sequence of 6 octets, labelled as eui[0] through eui[5]. The first three octets (eui[0] through eui[2]) are the manufacturer's OUI, and the last three octets (eui[3] through eui[5]) are the manufacturer-selected extension identifier.

The Device Control field is specified in Figure 68.

bits: b7-b6	b5	b4-b0
Security Mode	Status	Reserved

Figure 68 - Device control field format

The security mode field is set to the security mode at which the device is currently, or will be, operating in a data channel.

The status field is set to the following values to indicate the function of the discovery frame.

Table 58 - Status field meaning

Value	Meaning
0	Discovery
1	Response

15.5 Control frames

Default MAC Header field settings for control frames are listed in Table 59. Specific MAC Header field settings and payload descriptions for each of the control frames are shown in the following Clauses.

Table 59 - MAC header field values for control frames

Header field	Value
Protocol Version	0
Secure	As defined in 15.2.3.2
ACK Policy	0 (No-ACK)
Frame Type	2 (control frame)
Frame Subtype	Value from Table 60
Retry	Reserved
Aggregation Request	Reserved

Table 59 - MAC header field values for control frames (concluded)

Header field	Value
Tracking Indication	As described in 15.2.3.8
DestAddr	DevAddr of the recipient
SrcAddr	DevAddr of the transmitter
Sequence Control	Reserved
Duration	As described in 15.2.7.1 and 16.1.9.1
More Frames	Reserved
Access Method	As described in 15.2.7.3

Table 60 lists valid values for the frame subtype field for control frames.

Table 60 - 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	RTT	Announces to a recipient device that the sender is ready to initiate antenna training and requests confirmation of the ability to perform antenna training
3	CTT	Responds to a RTT control frame that the recipient is able to perform antenna training
4	TRN	Provides antenna training status and/or feedback and requests to change training parameters
5	B-Poll	Provides Master services to Type B slave device
6	B-Poll Response	Acknowledges B-Poll sent by a Type A Master device
7-13	Reserved	Reserved
14	Application-specific	At discretion of application owner
15	Reserved	Reserved

15.5.1 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 with Tracking Indication set to ZERO have no frame payload.

When used in implicity antenna tracking, the Tracking Indication bit in the MAC header is set to ONE to indicate inclusion of ATTCIE and/or AFIE. Its payload format is specified in Figure 69.

Octets: 0, or 4	0, or L
ATTCIE	AFIE

Figure 69 - Payload format of ACK frames

15.5.2 Block acknowledgement (B-ACK)

In B-ACK frames, the DestAddr field is set to the SrcAddr of the frame that requested the 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 16.12.3. The B-ACK frame payload is specified in Figure 70.

octets: 2	1	1	2	o-n	0, or 4	0, or L
Buffer Size	Frame Count	Reserved	Sequence Control	Frame Bitmap	ATTCIE	AFIE

Figure 70 - Payload format for B-ACK frames

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.

If the Fixed Size B-ACK Buffer Element bit of the MAC capability IE of the device is set to ONE, the fixed size of a buffer element in the device is the Buffer Size value divided by the Frame Count value specified in the B-ACK frame.

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.

bits: b15-14	b13-b3	b2-b0
Reserved	Sequence Number	Fragment Number

Figure 71 - Sequence control field format

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 may 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.

The Tracking Indication bit in the MAC header of B-ACK frames indicates the existance of the ATTCIE and/or AFIE at the end of the frame. It is set ONE, if any of ATTCIE or AFIE is included for the purpose of antenna tracking; otherwise, it is set to ZERO.

15.5.3 Application-specific

The payload format for Application-specific control frames is specified in Figure 72

Octets: 2	N
Specifier ID	Data

Figure 72- Payload format for application-specific control frames

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

15.5.4 B-Poll

A B-Poll Frame shall be used by a Type A device to start Master-Slave operation with a Type B slave device and to direct the Type B Slave when to send its dual-beacon.

The payload of a B-Poll frame is specified in Figure 73.

octets: 4	L_1	...	L_N
Directed Beacon Parameters	Information Element 1	...	Information Element N

Figure 73 - Payload format for B-poll frames

The information elements (IEs) that may be included in a B-Poll frame are listed in Table 74 in 15.9. IEs are included in order of increasing Element ID, except for ASIEs. ASIEs do not appear prior to any IE with Element ID zero through seven, but may appear anywhere after those IEs. DRP IEs that have the same Target DevAddr and Stream Index are adjacent to each other in the beacon. In a B-Poll frame, the Type A master shall only include the IEs that need to be included in the directed dual-beacon of the Type B slave. ATTCIE may be included in a B-Poll frame to initiate implicit antenna tracking.

The Directed Beacon Parameters field is specified in Figure 74.

octets: 2	1	1
Directed Beacon Offset	Beacon Slot Number	Device Control

Figure 74 - Directed Beacon Parameters field format

The Directed Beacon Offset field is set to the time when Type B is directed to send its dual-beacon, measured from the end of the PLCP header of the B-Poll frame in units of microseconds.

The Beacon Slot Number field is set to the number of the beacon slot where the Type B slave is directed to send its dual-beacon within the beacon period (BP), in the range of [mSignalSlotCount, mMaxBPLength-1].

The Device Control field is specified in Figure 75.

b7-b2	b1	b0
Reserved	Signalling Slot	Movable

Figure 75 - Device Control field format

The Movable bit is set to ONE if the beacon is movable according to 16.5.3.10, and is set to ZERO otherwise.

The Signalling Slot bit is set to ONE if the Type B master is directed to send an additional beacon in a signalling beacon slot according to 16.5.3.6, and is set to ZERO otherwise.

15.5.5 B-Poll response frame

In B-Poll response frame, the DestAddr field is set to the SrcAddr of the received B-Poll frame. The payload of B-Poll response frames may carry a number of IEs, which is specified in Figure 76.

Octets: L_1	...	L_N
Information Element 1	...	Information Element N

Figure 76 - payload format for B-Poll response frame

AFIE may be included in a B-Poll response frame to send back antenna tracking feedback.

15.5.6 Antenna training/tracking control frames

RTT and CTT control frames shall be used to initiate and negotiate antenna training between source and destination devices. RTT is transmitted by the source device whenever it wants to initiate an antenna training procedure with a destination device. Upon receipt of a RTT frame, the addressed device shall respond with a CTT frame provided it is available to engage in the antenna training procedure. The antenna training procedure is performed by exchanging TRN frames.

15.5.6.1 Request to train (RTT)

In RTT frames, the DestAddr field is set to the DevAddr of the device to which antenna training is desired. RTT frames shall carry in their frame payload the ATIE with the transmitter device settings.

15.5.6.2 Clear to train (CTT)

In CTT frames, the DestAddr field is set to the SrcAddr of the received RTT frame. CTT frames shall carry in their frame payload the ATIE with the receiver device settings.

15.5.6.3 Training control (TRN)

A TRN frame contains an ATTCIE and/or an AFIE in its payload. A TRN frame may also contain ATIE to confirm training configuration in the reversed direction.

The payload of a TRN frame is specified in Figure 77.

Octets: 4	0, or 5	0, or L
ATTCIE	ATIE	AFIE

Figure 77 - Payload format of a TRN frame

15.6 Command frames

Default MAC header settings for command frames are shown in Table 61.

Table 61 - Default MAC header field values for command frames

Header field	Value
Protocol Version	0
Secure	As defined in 15.2.3.2
ACK Policy	0 (No-ACK) or 1 (Imm-ACK)
Frame Type	3 (command frame)
Frame Subtype	Value from Table 62
Retry	As defined in 15.2.3.6
Aggregation Request	Reserved
Tracking Indication	Reserved
DestAddr	DevAddr of the recipient
SrcAddr	DevAddr of the transmitter
Sequence Control	As defined in 15.2.6
Duration	As defined in 15.2.7.1 and 16.1.9.1
More Frames	As defined in 15.2.7.2
Access Method	As defined in 15.2.7.3

Table 62 specifies the values for the frame subtype field for command frames.

Table 62 - Frame subtype field encoding for command frames

Value	Command frame subtype	Description
0	DRP Reservation Request	Used to request creation or modification of a DRP reservation
1	DRP Reservation Response	Used to respond to a DRP 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	Link Feedback	Used to exchange link feedback information and/or to recommend transmit power/rate change
6	Relay	Used to request for or respond with information elements
7	Transmit Switched Diversity (TSD) Set Request	Used to initiate or to terminate the TSD operation. In addition, used to acquire channel state information

Table 62 - Frame subtype field encoding for command frames (concluded)

Value	Command frame subtype	Description
8	Transmit Switched Diversity (TSD) Set Response	Used to respond to the TSD Set Request command
9	FUCA	Informs a recipient device of FUCA map information on the next frame
10	TSD Switch	Requests to switch the transmit antenna to an arbitrary antenna during the TSD operation
11	Channel Selection	Used to select initial channel to send beacons
12 - 13	Reserved	Reserved
14	Application-specific	At discretion of application owner
15	Reserved	Reserved

15.6.1 DRP reservation request

The DRP Reservation Request command frame is used to create or modify a DRP reservation. The DRP Reservation Request command frame payload is specified in Figure 78.

octets: M_1	M_2	...	M_N
DRP IE-1	DRP IE-2	...	DRP IE-N

Figure 78 - Payload format for DRP reservation request command frames

Each DRP IE field included in the command frame corresponds to a reservation request identified by the Target/Owner DevAddr, Stream Index, and Reservation Type in the IE. The DRP IE is defined in 15.9.12.

15.6.2 DRP reservation response

The DRP Reservation Response command frame is used to respond to a DRP Reservation Request command frame. The DRP Reservation Response command frame payload is specified in Figure 79.

octets: M_1	M_2	...	M_N	0, or 2 to 34
DRP IE-1	DRP IE-2	...	DRP IE-N	DRP Availability IE

Figure 79 - Payload format for DRP Reservation Response command frames

The DRP Reservation Response command frame includes all the DRP IEs from the reservation request. The DRP Availability IE is included according to the rules defined in 16.6.

15.6.3 Channel selection

The set of channel selection command frames are used to request channel scanning, to respond with channel scanning results and to select channel to exchange MPDUs.

The ACK policy field in the MAC header of channel selection command frame shall always be set to Imm-ACK.

The payload of channel selection command frame is specified in Figure 80.

octets: 2	L_1	...	L_N
Channel Selection Control	Information Element 1	...	Information Element N

Figure 80 - Payload format for Channel selection command frames

The Channel Selection Control field is specified in Figure 81.

bits: b15-b14	b13-b12	b11-b10	b9-b0
Command ID	Reserved	Reason Code	Channel Bitmap

Figure 81 - Channel Selection Control field format

The Command ID field is set to the value as listed in Table 63 that identifies the Channel selection command.

Table 63 - Channel selection command ID

Command ID	Channel selection command
0	Channel Scanning Request
1	Channel Scanning Response
2	Channel Change Request
3	Channel Change Response

The encoding of Channel Bitmap field is specified in in Figure 82.

Bit	Band ID	Description
0	1	Channel 1
1	2	Channel 2
2	3	Channel 3
3	4	Channel 4
4	5	Channel 5
5	6	Channel 6
6	7	Channel 7
7	8	Channel 8
8	9	Channel 9
9	10	Channel 10

Figure 82 - Channel Bitmap Field Encoding

15.6.3.1 Channel Scanning Request

The Channel Bitmap field is set such that the bits corresponding to the channels requested to be scanned are set to ONE.

The Reason Code field in a Channel Scanning Request is reserved.

The first Information Elements field shall include a Scan Timing IE to indicate the time when the device returns back to Discovery Channel to listen for scanning response to the scanning request.

15.6.3.2 Channel Scanning Response

The Channel Bitmap field is set to the same as in the Channel Scanning Request.

The Reason Code field in a Channel Scanning Response is reserved.

The Information Elements fields shall include a number of DRP availability IEs. The number of DRP availability IEs shall be the same as the number of channels requested to be scanned.

Following a number of DRP availability IEs, the Information Elements fields may also include a number of channel measurement IEs. The number of channel measurement IEs shall be either the same as the number of channels requested to be scanned or the number of channels with no superframe structure.

If a scanned channel has no superframe structure, the corresponding DRP availability IE shall have zero length DRP availability bitmap and a channel measurement IE (for channels without superframe structure, see 15.9.11) shall be included in the Channel Scanning Response.

15.6.3.3 Channel Change Request

The first Information Elements field shall include a Channel Change IE.

The Reason Code and Channel Bitmap fields in a Channel Change Request are reserved.

15.6.3.4 Channel Change Response

The first Information Element field shall include a Channel Change IE that is the same as the one in the Channel Change Request to which it is responding. The Reason Code field shall be set appropriately as listed in Table 64.

Table 64 - Reason code field encoding

Value	Code	Meaning
0	Accepted	The channel change request is agreed upon
1	Unavailable	The channel change request is rejected because the channel in concern does not have enough MASs
2	Conflict	The channel change request is rejected because the channel change request conflicts with the scanning response sent.
3	Invalid	The device does not support channel

The Channel Bitmap field is reserved in the Channel Change Response.

15.6.4 Link feedback

The set of Link Feedback command frames are used to exchange feedback information pertaining to power control, rate control and certain link quality metrics between two devices. The payload of the Link Feedback command is specified in Figure 83.

octets: 1	0 or N
Feedback Control	Link Feedback

Figure 83 - Payload format for link feedback command frames

The Feedback Control field is specified in Figure 84.

bits: b7-b6	b5-b3	b2-b0
Command ID	Reserved	Request Bitmap

Figure 84 - Feedback Control field format

The command ID field is set to the value as listed in Table 65 that identifies the link feedback command

Table 65 - Feedback Control field format

Command ID	Link feedback command	Description
0	Transmit power and rate control	Recommends transmit power and rate change to a source device.
1	Link Feedback Response	Provides, in response to a link feedback request, link quality metrics and/or recommendation of rate and transmission power change.
2	Link Feedback Request	Requests a sink device to provide link feedback.

The Request Bitmap field is set according to Table 66.

Table 66- Request bitmap

Bit	Meaning	Description
0	TPRC	Bit set to ONE to indicate requesting for, or the existence of sink's recommendation on transmit rate and power change
1	LQI	Bit set to ONE to indicate requesting for, or the existence of feedback on LQI of the link
2	FER	Bit set to ONE to indicate requesting for, or the existence of feedback on frame loss/error rate of the link

A link feed back command may contain 0-3 link feedback fields depending on setting of the Feedback Control field.

15.6.4.1 Link feedback request

There shall be none Link Feedback field in the Link feedback request frame.

15.6.4.2 Transmit power and rate control (TPRC)

The Feedback Control field of the TPRC is specified in Figure 85.

Bits: b7-b6	b5-b3	b2-b0
=0	Reserved	=0

Figure 85 - Setting of Feedback Control field of TPRCs

The Link Feedback field of the TPRC is specified in Figure 86.

bits: b7-b4	B3-b0
Data rate	Transmit Power Level Change

Figure 86 - Link Feedback field format of TPRCs

The Transmit Power Level Change sub field is set to the change in transmit power level that the sink device recommends according to Table 87.

The Data Rate sub field is set to the data rate that the sink device recommends according to Table 88.

15.6.4.3 Link feedback response

The Request Bitmap field of a link feedback response is set to the same as in the link feedback request to which the device is responding.

The Link Feedback field of a link feedback response is specified in Figure 87.

octets: 1	1	7
TPRC	LQI	FER

Figure 87 - Link feedback field format

The existence of TPRC, LQI and FER sub fields in the Link Feedback field is determined by the Request Bitmap field. A bit set to ONE in the Request Bitmap field indicates that the corresponding sub field exists in the Link Feedback field. The order the sub fields appear is specified in Figure 87.

The TPRC sub field is set according to Figure 86.

The LQI sub field is set to the Link Quality Indicator (LQI) of the link between the source and sink devices. Its encoding is specified in Table 67.

Table 67- LQI field encoding

LQI value	SNR/SINR value
0	0.0 db
1	0.5 db
2	1.0 db
...	...

Table 67- LQI field encoding (concluded)

LQI value	SNR/SINR value
62	31.0 db
63	31.5 db
64-255	Reserved

The format of the FER sub field is specified in Figure 88.

octets: 1	1	2	2	1
Recipient Frame Loss Count	Recipient Frame Error Count	Recipient Frame Count	Source Frame Count	Measurement Window Size

Figure 88 - FER field format

The Measurement Window Size field is set to the amount of time in milliseconds during which the measurements as included in FER field were taken.

The Source Frame Count field is set to the total number of frames, including retransmissions that were transmitted by the sink device to the source device.

The Recipient Frame Count field is set to the total number of frames that were correctly received by the sink device from the source device.

The Recipient Frame Error Count field contains the total number of frames that were received with FCS errors but not with HCS errors by the sink device from the source device.

The Recipient Frame Loss Count field is set to the number of frames as observed by the sink device to have been lost.

15.6.5 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 Figure 89.

Octets: M_1	M_2	...	M_N
Information Element 1	Information Element 2	...	Information Element N

Figure 89 - Payload format for probe command frames

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

15.6.6 Pairwise temporal key (PTK)

The PTK command frame is used in a 4-way handshake by a pair of devices, as described in 17.3.1, 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 Figure 90.

octets: 1	1	3	11	16	16	8
Message Number	Status Code	PTKID	Reserved	MKID	I-Nonce / R-Nonce	PTK MIC

Figure 90 - Payload format for PTK command frames

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 68.

Table 68 - 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 in 17.3.1.

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 payload of this command using the KCK generated from the first two messages of the 4-way handshake as specified in 17.3.1.

The MAC Header for the PTK command frame is set as indicated in Table 61, with the ACK Policy set to Imm-ACK.

15.6.7 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 Figure 91.

octets: 1	1	3	3	2	6	16
Message Number	Status Code	GTKID	Reserved	GroupAddr	GTK SFC	GTK

Figure 91 - Payload format for GTK command frames

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 69.

Table 69 - 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 MAC Header for the GTK command frame is set as indicated in 15.6.7, with the ACK Policy set to Imm-ACK.

15.6.8 Application-specific

The payload format for Application-specific command frames is specified in Figure 92.

octets: 2	N
Specifier ID	Data

Figure 92 - Payload format for application-specific command frame

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

15.6.9 Relay

The Relay command frame is used to request information from a device or respond to a Relay request. The payload is the Relay IE as specified in 15.9.24.

15.6.10 Transmit switched diversity (TSD) request

The TSD Set Request command frame is used to initiate the TSD operation, to scan channels or to terminate the TSD operation by acquiring the information on the destination's capability of the TSD operation. This command frame also indicates whether the TSD operation is in scan mode or not. The payload format for TSD Set Request command frame payload is specified in Figure 93.

Octet: 1	2
TSDStatus	Reserved

Figure 93 - Payload format for TSD set request command frames

The TSDStatus field indicates whether the TSD set operation begins or ends. This field also indicates whether the TSD set operation is in scan mode or not. The TSDStatus field is specified in Table 94.

bits: b7-b1	b0
Reserved	TSD Status indicator

Figure 94 - TSD status field format

The TSD status indicator field is set to 0b00 to indicate that this TSD set procedure is to begin the TSD operation, to 0b01 to indicate that this TSD set procedure is to end the TSD operation or to 0b10 to indicate that this TSD set procedure is to request channel status measured from this command frame. The value of 0b11 is reserved. The TSD status indicator field is encoded as Table 70.

Table 70 - TSD status indicator field encoding

TSD Status indicator	Message
0b00	BEGIN
0b01	END
0b10	SCAN
0b11	Reserved

15.6.11 Transmit switched diversity (TSD) set response

The TSD Set Response command frame is used to respond TSD Set Request command frame. The TSD Set Request command frame payload is specified in Figure 95 .

Octet: 1	1	1
TSDStatus	TSDPermit	ChannelStatus

Figure 95 - Payload format for TSD set response command frames

The TSDPermit field indicates whether the device supports the TSD operation or not. The TSDPermit field is specified in Figure 96.

bits: b7-b1	b0
Reserved	TSD support indicator

Figure 96 - TSD permit field format

The TSDPermit field is set to one to indicate that the device supports the TSD operation or to zero to indicate that the device does not support the TSD operation. The TSD support indicator field is encoded as Table 71.

Table 71 - TSD support indicator field encoding

TSD support indicator	Message
0	NO_PERMIT
1	PERMIT

The ChannelStatus field indicates the link quality of the channel, where the link quality shall be defined as an estimate of the received SNR. When reporting the estimates of the link quality, the device shall quantize these values to the nearest values in Table 72 and report them as the link quality estimates. The standard encoding, summarized in Table 72, is used to report the estimates in the range from 0 dB (0b0000 001) to +25 dB (0b0001 1010). The all-ZERO bit implies that reporting of the estimates is not supported by the device, or that the estimate is too small to be measured accurately. Additionally, the range from 0b0001 to 0b1111 1111 are reserved for further use.

Table 72 - Channel status field encoding

Bits: b7-b0	Description
0b0000 0000	Link quality estimate is too low to be measured
0b0000 0001 - 0b0001 1010	Estimate of the link quality = ChannelStatus - 1 dB
0b0001 1011 - 0b1111 1111	Reserved

15.6.12 Transmit switched diversity (TSD) switch

This frame is used to feedback the transmit antenna switching information for the TSD operation. The TSD Switch control frame commands the recipient of the frame to switch the transmit antenna. This frame is sent from the destination to the source when the source is required to switch the TX antenna. A payload of the TSD Switch control frame is specified in Figure 97.

Octet: 1
TSD Switch

Figure 97 - TSD switch control frame payload

The TSD Switch field specifies the antenna switching information for the recipient of the control frame. The TSD Switch field is specified in Figure 98.

bits: b7-b1	b0
Reserved	Antenna Switch

Figure 98 - TSD switch field format

The Antenna Switch field indicates that the source switches the transmit antenna. The Antenna Switch field is encoded as shown in Table 73.

Table 73 - Antenna switch field encoding

Antenna Switch	Message
0	No Switch
1	Switch

15.6.13 Fast uplink channel allocation (FUCA)

The FUCA command frame is used to inform the reservation target the duration of the next data frame, after which the target may send a short frame back to the reservation owner in the current private DRP reservation.

The Next Frame Duration field is set to the duration of the next data frame, in unit of microsecond, which is sent by reservation owner following the FUCA command frame.

The Sequence Number field is set to the sequence number of the next data frame to be transmitted by reservation owner. FUCA map field contains the information generated by HDMI PAL in accordance with the next data frame.

The FUCA frame payload is shown in Figure 99.

Octets: 2	2	5
Next Frame Duration	Sequence Number	FUCA Map

Figure 99 - Payload format for FUCA frames

15.7 Data frames

MAC Header and Frame Payload fields in data frames are set as specified in 15.2.1 and 15.2.3-15.2.9.

15.8 Aggregated MAC frames

Aggregated MAC frames are aggregated data frames and tracking frames. The Frame Payload field in aggregated MAC frames are set as specified in 15.2.2.

15.8.1 Aggregated data frames

In additonal to settting MAC Header fields as specified in 15.2.2 -15.2.9, the ACK Policy of the Frame Control Field in any aggregated data frames shall not be set to Imm-ACK.

15.8.2 Aggregated tracking frames

Aggregated tracking frames are aggregated data frames that carry AFIE or ATTCIE in its last MSDU and may carry ATS in its PLCP header according to antenna tracking configuration (16.18.2).

In addition to setting the MAC header fields specified in 15.2.2 -15.2.9, the Tracking Indication bit of the Frame Control field of any aggregated tracking frames shall be set to ONE. And the ACK Policy of the Frame Control field in any tracking frame shall not be set to Imm-ACK unless the tracking frame has only two MSDUs including the one that contains AFIE or ATTCIE.

In Frame Payload, the AFIE or ATTCIE shall be the last MSDU in the aggregated MAC frame. The MAC Frame Body of such a tracking frame is specified in Figure 100.

Figure 100 - Payload format of an aggregated tracking frame

octets: M_1	4	0-3	...	0-3	M_n	4
MSDU 1	FCS of MSDU 1	Pad to 4-octet boundary	...	Pad to 4-octet boundary	AFIE or ATTCIE	FCS of MSDU N

15.9 Information elements

This Clause defines the information elements (IEs) that can appear in beacons and certain control or command frames.

The general format of all IEs is specified in Figure 101.

octets: 1	**1**	**N**
Element ID	Length (=N)	IE-specific fields

Figure 101 - General IE format

The Element ID field is set to the value as listed in Table 74 that identifies the information element.

The Length field is set to the length, in octets, of the IE-specific fields that follow.

The IE-specific fields contain information specific to the IE.

Table 74 contains a list of IEs defined in this International Standard.

Table 74 - Information elements

Element ID	Information element	Description
0	Beacon Period Occupancy IE (BPOIE)	Provides information on neighbours' BP occupancy in the previous superframe
1-7	Reserved	Reserved
8	ACIE	Provides information on antenna types and training/tracking feedback capabilities.
9	AFIE	Provides antenna training/tracking feedback

Table 74 - Information elements (concluded)

Element ID	Information element	Description
10	Application-specific Probe IE	Indicates a device is requesting an Application-specific IE from another device
11	ATIE	Configures antenna training and tracking processes.
12	ATTCIE	Sets antenna training parameters or initiates antenna tracking
13	BP Switch IE	Indicates the device will change its BPST at a specified future time
14	Channel Bonding IE (CBOIE)	Negotiates and indicate the use of channel bonding
15	Channel Change IE	Indicates a device will change to another channel
16	Channel Measurement IE	Provides directional channel measurement report
17	Distributed Reservation Protocol (DRP) IE	Indicates a reservation with another device
18	DRP Availability IE	Indicates a device's availability for new DRP reservations
19	Hibernation Anchor IE	Provides information on devices in hibernation mode
20	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
21	Identification IE	Provides identifying information about the device, including a name string
22	Link Feedback IE	Provides data rate and power control feedback
23	MAC Capabilities IE	Indicates which MAC capabilities a device supports
24	Master Key Identifier (MKID) IE	Identifies some or all of the master keys held by the transmitting device
25	Multicast Address Binding (MAB) IE	Indicates an address binding between a multicast EUI-48 and a McstAddr
26	PHY Capabilities IE	Indicates which PHY capabilities a device supports
27	Probe IE	Indicates a device is requesting one or more IEs from another device or/and responding with requested IEs
28	Relay IE	Indicates a device requesting relay information element from another device or/and responding with requested IE
29	Relinquish Request IE	Indicates that a neighbour requests that a device release one or more MASs from its reservations
30	Scan Timing IE	Announces the start and duration of a scanning period
31	UEP IE	Provide information on supported UEP types and UEP MCS modes
32 - 249	Reserved	Reserved
250 - 254	Reserved	Reserved for allocation in assigned numbers as specified in [D4]
255	Application-Specific IE (ASIE)	Use varies depending on the application

15.9.1 Application-specific IE (ASIE)

The ASIE is specified in Figure 102.

octets: 1	1	2	N
Element ID	Length (=2+N)	ASIE Specifier ID	Application-specific Data

Figure 102 - ASIE format

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

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

15.9.2 Application-specific probe IE

The Application-specific Probe IE is used to request an application-specific IE from a device. It is specified in Figure 103.

octets: 1	1	2	2	N
Element ID	Length (=4+N)	Target DevAddr	ASIE Specifier ID	Application-specific Request Information

Figure 103 - Application-specific probe IE format

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 [C.4].

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

15.9.3 Antenna Capabilities IE (ACIE)

The ACIE provides information about the device type and the training and feedback capabilities of the RX and TX antenna of the transmitting device. The ACIE is specified in Figure 104.

Octets: 1	1	2
Element ID	Length (= 2)	Antenna Capabilities

Figure 104 - ACIE format

The Antenna Capabilites field is specified in Figure 105.

bits: b15-8	b7-b6	b5-b4	b3-b0
Reserved	Device Type and ATS Capability	Antenna Type	Feedback Capability

Figure 105 - Antenna Capabilities field format

The encoding of Antenna Capabilitities field is specified in Table 75.

Table 75- Antenna Capabilities field bitmap

Capabilities	Bit	Attribute	Description
Reserverd	7		
ATS Capability	6	ATS	Bit set to ONE to indicate the transmitting device is capable of transmitting attenna training sequence.
Antenna Types	5	Trainable TX PAA	Bit set to ONE to indicate the transmitting device has trainable transmit Phased Array Antenna (PAA)
	4	Trainable RX PAA	Bit set ONE to indicate the transmitting device has trainable receiving Phased Array Antenna (PAA)
Feedback Capabilities	3	Index	Bit set to ONE to indicate the device is capable to provide index feedback
	2	Fourier codebook	Bit set to ONE to indicate the device is capable to provide Fourier code book feedback
	1	Walsh codebook	Bit set to ONE to indicate the device is capable to provide Walsh code book feedback
	0	Quantised weights	Bit set to ONE to indicate the device is capable to provide Quantised weights feedback

A Type A device shall set the ATS attribute to ONE.

15.9.4 ATIE

The ATIE is used to negotiate and configure the antenna training and tracking. It is specified in Figure 106.

Octets: 1	1	4	1
Element ID	Length (=5)	Antenna Properties	Training Configuration

Figure 106 - ATIE frame format

The Antenna Properties field is specified in the Figure 107.

Bits: 16	6	6	4
Antenna Capabilities	Number of TX Elements	Number of RX Elements	Initial Discovery Mode

Figure 107 - Antenna Properties field format

The Antenna Capabilities field is set to the same value as in the ACIE (see 15.9.3).

The Number of TX Elements field is set to the number of antenna elements the device's the transmission antenna array has, in the range of [1, 36].

The Number of of RX Elements field is set to the number of antenna elements the device's receiving antenna has, in the range of [1, 36].

Initial Discovery Mode field is set to the discovery mode to be used for the first iteration of antenna training for the device's receiving antenna.

The Training Configuration field is specified in Figure 108.

Bits: 1	3	1	1	2
Request Feedback	Feedback Type	Request Training	Continuition	Status

Figure 108 - Training Configuration field format

The Request Feedback bit is set to ONE to indicate closed-loop mode training and subsequently to request antenna training feedback, or to ZERO to indicate open-loop mode training and feedback is not needed. This field has an implication in the subsequent training sequence exchange as specified in Table 48.

The Feedback Type field is set to what type of feedback is expected in closed-loop mode training when the Request Feedback bit is set to ONE. Otherwise if the Request Feedback bit is set to ZERO, the Feedback Type field is reserved. The encoding of the Feedback Type field is specified in Table 76.

Table 76 - Encoding of Feedback Type field

Value	Code	Meaning
0	Index	Requesting Index feedback
1	Fourier	Requesting feedback using Fourier codebook
2	Walsh	Requesting feedback using Walsh codebook
3	Quantised	Requesting feedback of quantised weights
4-7	Reserved	Reserved

The Request Training bit is set to ONE to request the training of the device's transmitting antenna and the target's receiving antenna. Otherwise, it is set to ZERO to indicate no training is needed for that direction. This field has an implication in the subsequent training sequence exchange as specified in Table 48

The Continuation bit is set to ONE if the device is requesting continuation of the previous training session, or set to ZERO if the device is requesting a new training session.

The Status field is set by the antenna training initiator or responder to indicate the status of the negotiation. Its encoding is specified in Table 77.

Table 77 - Status field Encoding

Value	Code	Meaning
0	Request	The device is requesting antenna training
1	Accepted	The device has accepted the training request with the specified training configuration
2	Unsupported	The requested training configuration is not supported
3	Rejected	The device rejects the training request

15.9.5 ATTCIE

Antenna Training/Tracking Control IE (ATTCIE) is used to set antenna trainig/tracking parameters. The format of ATTCIE is specified in Figure 109.

octets: 1	1	2
Element ID	Length (=2)	ATT Control

Figure 109 - ATTCIE format

The ATT Control field is specified in Figure 110.

bits: 1	3	4	1	1	6
Tracking	Iteration	Discovery Mode	TX Training Status	RX Training Status	RX Training Size

Figure 110 - ATT Control field format

The Tracking field is set to ONE, if the ATTCIE is in used in antenna tracking; or ZERO, if the IE is used in antenna training.

The Iteration field counts the iteration number of the iterative training process, in the range of [0, 4]. It is set to 0 if used in antenna tracking.

The TX Training Status field is set to ZERO, if the transmitter is performing transmitting antenna tracking or has completed training of its transmitting antenna. It is set to ONE if the transmitter is still continuing training its transmitting antenna.

The RX Training Status field is set to ZERO, if the transmitter is performing receiving antenna tracking or has completed training of its receiving antenna. It is set to ONE if the transmitter is still continuing training its receiving antenna.

The Discovery Mode field is set to the discovery mode to be used for the next TRN/tracking frame in the reverse direction. Its encoding is specified in Table 78.

Table 78 - Discovery Mode field encoding

Value	Code	Meaning
0	mode D7	request to use mode D7 for the next TRN/tracking frame in the reverse direction

Table 78 - Discovery Mode field encoding (concluded)

Value	Code	Meaning
1	mode D6	request to use mode D6 for the next TRN/tracking frame in the reverse direction
2	mode D5	request to use mode D5 for the next TRN/tracking frame in the reverse direction
3	mode D4	request to use mode D4 for the next TRN/tracking frame in the reverse direction
4	mode D3	request to use mode D3 for the next TRN/tracking frame in the reverse direction
5	mode D2	request to use mode D2 for the next TRN/tracking frame in the reverse direction
6	mode D1	request to use mode D1 for the next TRN/tracking frame in the reverse direction
7	mode D0	request to use mode D0 for the next TRN/tracking frame in the reverse direction
8-15	Reserved	

The RX Training Size field is set to number of receiving antenna training symbols to be transmitted in the next TRN/tracking frame, in the range of [1, 36].

15.9.6 AFIE

Antenna Feedback IE (AFIE) is used to indicate the status of the antenna training procedure and to provide the feedback to the transmitte on the transmitter weights. AFIE is specified in Figure 111.

octets: 1	1	1	L
Element ID	Length (=1+L)	Feedback Control	Antenna Feedback

Figure 111 - AFIE format

The Feedback Control field is specified in the Table 112.

bits: 5	3
Reserved	Feedback Type

Figure 112 - Feedback control field format

The Feedback Type field is set to indicate the type of antenna feedback that is included in the AFIE. Its encoding is specified in Table 76.

The Antenna Feedback field and its length depend on the Feedback Type in the Feedback Control field and are specfied in Table 79.

Table 79- Length and content of the Antenna Feedback field

Feedback Type	L (in Octets)	content
Index	1	index of best received training symbol

Table 79- Length and content of the Antenna Feedback field (concluded)

Feedback Type	L (in Octets)	content
Fourier	1	index of selected codeword
Walsh	1	index of selected codeword
Quantised	L (= number of TX Elements, as in 15.9.4)	Quantised transmitter weights

When Feedback Type is set to Index, Fourier or Walsh, the encoding of the Antenna Feedback field is specified in Table 50. Otherwise, when the Feedback Type is set to Quantised, the Antenna Feedback format is specified in Table 113 .

Octets: 1	1	...	1
Weights 0	Weight 1	...	Weight L

Figure 113 - Antenna Feedback field foramt (Feedback Type set as Quantised)

The field Weight i is set to the quantised weight of transmitter antenna element i. And its encoding is specified in Table 52.

15.9.7 Beacon period occupancy IE (BPOIE)

The BPOIE provides information on the BP observed by the device sending the IE and the corresponding receiving beams. The BPOIE is specified in Figure 114.

octets: 1	1	1	1	L
Element ID	Length (=2+L)	BP Length	Number of Beams	Beacon Slot Info

Figure 114 - BPOIE format

The BP Length field is set to the length of the BP, measured in beacon slots, as defined in 16.5.3.3.

The Number of Beams field is set the total number of beams that the device uses to transmit beacons.

The format of Beacon Slot Info field is specified in the figure below.

octets: K	2	...	2	M_1	...	M_N
Beacon Slot Info Bitmap	DevAddr 1	...	DevAddr N	Receiving Beams 1	...	Receiving Beams N

Figure 115 - Beacon Slot Info field format

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 80. Element zero is the least-significant two bits of the field. Unused elements, if any, are set to zero.

Table 80 - 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 HCS was not a beacon frame.
1	Occupied & non-movable A beacon frame was received with a valid HCS 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 16.16.
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 HCS and FCS.
3	Occupied & movable A beacon frame was received with a valid HCS 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 HCS 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 HCS 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.

The Receiving Beams fields correspond to beacon slots encoded as occupied in the Beacon Slot Info Bitmap field. The Receiving Beams field lists the beam identifiers of the antenna beams, from which the device can receive the corresponding beacon.

The format of Receiving Beams field is specified in the figure below,

octets: 1	M
Number of Receiving Beams (=M)	Beam Identifiers

Figure 116 - Receiving beams field format

The Number of Receiving Beams field is set to the number of beam identifiers that follow.

The Beam Identifiers field contains a list of the beam identifiers, each occupying 1 octet, of the antenna beams from which the device can receive the beacon sent at the specific beacon slot.

15.9.8 BP Switch IE

The BP Switch IE indicates a device will change its BPST to align with an alien BP. It is specified in Figure 117.

octets: 1	1	1	1	2
Element ID	Length (=4)	BP Move Countdown	Beacon Slot Offset	BPST Offset

Figure 117 - BP Switch IE format

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.

15.9.9 Channel bonding IE (CBOIE)

A CBOIE is used to negotiate bonding of two or more channels and to announce the channels that are involved in the channel bonding. The CBOIE is specified in Figure 118.

octets: 1	1	2	1
ElementID	Length (=3)	Target DevAddr	Bonding Control

Figure 118 - Channel Bonding IE format

The Target DevAddr field is set to the DevAddr of the recipient device to which the device is sending the channel bonding request.

The Bonding Control field is specified in Figure 119.

bits: b7	b6	b5-b4	b3-b0
Bonding Status	Reserved	Reason Code	Bonded Channel Bit Map

Figure 119 - Bonding control field format

The Bonding Status bit indicates the status of the channel bonding negotiation process. The Bonding Status bit is set to ZERO in a CBOIE for a channel bonding request. It is set to ONE by a device granting or accepting the bonding request, which is then referred to as an established channel bonding.

The Reason code field is used by the targeted device to indicate whether it accepts a channel bonding request and is encoded as shown in Table 81.

Table 81 - Reason code field encoding

Value	Code	Meaning
0	Accepted	The channel bonding request is granted
1	Unavailable	The channel bonding request is rejected because one of the channels is not idle
2	Conflict	The channel bonding request is rejected because one of the channels is already bonded.
3	Invalid	The band ID is not bonded channel band ID

The Bonded Channel Bitmap field is specified in Figure 120.

Bit	Band ID	Description
0	1	Bonding channel 1
1	2	Bonding channel 2
2	3	Bonding channel 3
3	4	Bonding channel 4

Figure 120 - Bonded channel bitmap

15.9.10 Channel change IE

A Channel Change IE announces that a device is preparing to change to another channel.

The Channel Change IE is specified in Figure 121.

octets: 1	1	1	1
Element ID	Length (=2)	Channel Change Countdown	New Channel Number

Figure 121 - Channel change IE format

The Channel Change Countdown field is set to the number of superframes remaining until the device changes to the new channel. If this field is zero, the device will change to the new channel at the end of the current superframe.

The New Channel Number field is set to the channel number of the new channel to which the device will change.

The Channel Change Countdown field in the Channel Change IE included in discovery frames or channel selection command frames shall be set to Zero; and the New Channel Number field in the Channel Change IE included in discovery frames shall be set to the channel in which the device sets up its superframe structure.

15.9.11 Channel measurement IE

The Channel Measurement IE is used to report the result of directional channel measurement. The format of the Channel Measurement IE is specified in Figure 122.

octets: 1	1	1	1	L_1	...	L_m
Element ID	Lenght(=2+L_1+...+L_m)	Beam Identifier	Number of Measurement Reports (m)	Measurement Report 1	...	Measurement Report m

Figure 122 - Channel measurement IE format

The Beam Identifier field is set to a value that uniquely identifies the beam that the device uses to measure the channel condition.

A Channel Measurement IE contains one or more Measurement Report fields. The Number of Measurement Reports field is set to the number of Measurement Report fields included in the IE.

Each Measurement Report 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 Measurement Report field is specified in Figure 123.

bits: 16	16	4	...	4	0 or 4
Zone Bitmap	MAS Bitmap	$NILI_0$	...	$NILI_{n-1}$	Pad to octet boundary

Figure 123 - Measurement Report field format

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

The MAS Bitmap specifies which MASs in the zones identified by the Zone Bitmap field are reported. 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 report, where bit zero corresponds to MAS zero within the zone.

The Measurement Report field consists of one or more NILI fields. Each NILI field corresponds to a MAS that is indicated by Zone Bitmap and MAS Bitmap fields. NILI fields are placed in the ascending order of MAS numbers. The NILI for each MAS is set to the Noise plus Interference Level Indicator (NILI) that indicates the average noise plus interference power measured on the MAS with the receiving beam corresponding to the value of the Beam Identifier field. The encoding of the NILI fields is specified in Table 82.

Table 82 - NILI field encoding

Value	NILI level (dBm)
0	<=-110
1	(-110, -102]
2	(-102, -94]
3	(-94, -86]

Table 82 - NILI field encoding (concluded)

Value	NILI level (dBm)
4	(-86, -78]
5	(-78, -70]
6	(-70, -62]
7	(-62, -54]
8	(-54, -46]
9	(-46, -38]
10	(-38, -30]
11	(-30, -22]
12	(-22, -14]
13	(-14, -6]
14	(-6, 2]
15	>2

When the Channel Measurement IE is used to report the condition of a channel that does not have a superframe structure, the Number of Measurement Reports field shall be set to “1” and only one Measurement Report field shall be included. In the Measurement Report field, both Zone Bitmap and MAS Bitmap fields shall be set to 0x0 and the (4 bit) NILI field shall be set to the average NILI level during the entire scanning period. The format of Channel Measure IE for a channel without superframe structure is specified in Figure 124.

octets: 1	1	1	1	2	2	1
Element ID	Lenght (=7)	Beam Identifier	(=1)	Zone Bitmap (=0x0)	MAS Bitmap (=0x0)	NILI + padding (4bits)

Figure 124 - Channel measurement IE format for channels without superframe structure

15.9.12 Distributed reservation protocol (DRP) IE

A DRP IE is used to negotiate a reservation or part of a reservation for certain MASs and to announce the reserved MASs. The DRP IE is specified in Figure 125.

octets: 1	1	2	1	2	4		4
Element ID	Length (=5+4xN)	DRP Control	Beam Identifier	Target/Owner DevAddr	DRP Allocation 1	...	DRP Allocation N

Figure 125 - DRP IE format

The DRP Control field is specified in Figure 126.

bits: b15	bit14	b13	b12	b11	b10	b9	b8-b6	b5-b3	b2-b0
Reserved	Waveform	Unidirectional	Unsafe	Conflict Tie-breaker	Owner	Reservation Status	Reason Code	Stream Index	Reservation Type

Figure 126 - DRP control field format

The Reservation Type of the reservation and is encoded as shown in Table 83.

Table 83 - Reservation type field encoding

Value	Reservation Type
0	Alien BP
1	Hard
2	DCA
3	Private
4	Absence
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 Alien BP.

The Reason Code is used by a reservation target to indicate whether a DRP reservation request was successful and is encoded as shown in Table 84. The Reason Code is set to zero in a DRP 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 DRP IEs have been combined, split or both. The Reason Code is set to 5 by a reservation target if some of MAS(s) in a DRP IE sent during negotiation by a reservation owner have significant interference. The field is reserved if the Reservation Type is Alien BP.

Table 84 - Reason code field encoding

Value	Code	Meaning
0	Accepted	The DRP reservation request is granted
1	Conflict	The DRP reservation request or existing reservation is in conflict with one or more existing DRP reservations
2	Pending	The DRP reservation request is being processed
3	Denied	The DRP reservation request is rejected or existing DRP reservation can no longer be accepted
4	Modified	The DRP reservation is still maintained but has been reduced in size or multiple DRP IEs for the same reservation have been combined
5	Interfered	Some of MAS(s) included in the DRP reservation request have significant interference
6 - 7	Reserved	Reserved

The Reservation Status bit indicates the status of the DRP negotiation process. The Reservation Status bit is set to ZERO in a DRP 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.

The Owner bit is set to ONE if the device transmitting the DRP IE is the reservation owner, or to ZERO if the device transmitting the DRP IE is a reservation target. The bit is reserved if the Reservation Type is Alien BP.

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 DRP IEs that represent the same reservation, the Conflict Tie-breaker bit is set to the same value.

The Target/Owner DevAddr field is set to the DevAddr of the reservation target if the device transmitting this DRP IE is the reservation owner. The reservation target may be a unicast or multicast DevAddr. The field is set to the DevAddr of the reservation owner if the device transmitting the DRP IE is a reservation target. The field is reserved if the Reservation Type is Alien BP.

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

The Unidirectional bit is set to ONE if the owner only transmits MSDUs with ACK Policy set to No-ACK in the reservation.

The Waveform bit encoding for Hard and DCA reservations is specified in Table 85. This bit is reserved for reservations of other types

Table 85 - Waveform field encoding

	Device Type	
Value	Type A	Type B
0	SCBT	SC
1	OFDM	Reserved

The Beam Identifier field is set to the same as the Beam Identifier field in the beacon transmitted to the reservation target.

A DRP IE contains one or more DRP Allocation fields. Each DRP 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 DRP Allocation field is specified in Figure 127.

octets: 2	2
Zone Bitmap	MAS Bitmap

Figure 127 - DRP allocation field format

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.

15.9.13 DRP availability IE

The DRP Availability IE is used by a device to indicate its view of the current utilization of MASs in the current superframe. The DRP Availability IE is specified in Figure 128.

octets: 1	1	N (0 to 32)
Element ID	Length (=N)	DRP Availability Bitmap

Figure 128 - DRP availability IE format

The DRP 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 DRP reservation in the corresponding MAS, or is set to ZERO otherwise. If the DRP Availability Bitmap field is smaller than 32 octets, the bits in octets not included at the end of the bitmap are treated as ZERO.

15.9.14 Hibernation anchor IE

The Hibernation Anchor IE is specified in Figure 129.

octets: 1	1	3	...	3
Element ID	Length (=3×N)	Hibernation Mode Device Information 1	...	Hibernation Mode Device Information N

Figure 129 - Hibernation anchor IE format

The Hibernation Mode Device Information field is specified in Figure 130.

octets: 2	1
Hibernation Mode Neighbour DevAddr	Wakeup Countdown

Figure 130 - Hibernation mode device information field format

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.

15.9.15 Hibernation mode IE

The Hibernation Mode IE is specified in Figure 131.

octets: 1	1	1	1
Element ID	Length (=2)	Hibernation Countdown	Hibernation Duration

Figure 131 - Hibernation mode IE format

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.

15.9.16 Identification IE

The Identification IE provides identifying information about the device, including a name string. The Identification IE is specified in Figure 132.

octets:1	1	M_1	…	M_N
Element ID	Length (=M_1+…+ M_N)	Device Information 1	…	Device Information N

Figure 132 - Identification IE format

The general format of the Device Information field is specified in Figure 133.

octets: 1	1	N
Device Information Type	Device Information Length (=N)	Device Information Data

Figure 133 - Device information field format

The encoding for the Device Information Type field is shown in Table 86.

Table 86 - Device information type field encoding

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

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

The Device Information Data field, if Device Information Type is Vendor ID, is specified in Figure 134.

octets: 3
Vendor ID

Figure 134 - Device information data field format for vendor ID

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

The Device Information Data field, if Device Information Type is Vendor Type, is specified in Figure 135.

octets: 3	3
Vendor ID	Device Type ID

Figure 135 - Device information data field format for vendor type

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 Device Information Data field, if Device Information Type is Name String, contains the name of the device encoded in Unicode UTF-16LE format, and is specified in Figure 136.

octets: 2	...	2
Name String Unicode Char 1	...	Name String Unicode Char N

Figure 136- Device information data field format for name string

15.9.17 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 Figure 137.

octets: 1	1	3		3
Element ID	Length (=3xN)	Link 1	...	Link N

Figure 137 - Link feedback IE format

The Link field is specified in Figure 138.

bits: b23-b20	b19-b16	b15-b0
Data Rate	Transmit Power Level Change	DevAddr

Figure 138 - Link field format

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 87.

Table 87 - Transmit power level change field encoding

Value	Power level change	Power level change (dB)
1000 - 1101	Reserved	Reserved

Table 87 - Transmit power level change field encoding (concluded)

Value	Power level change	Power level change (dB)
1110	-2	-4
1111	-1	-2
0000	no change	no change
0001	+1	+2 dB
0010	+2	+4 dB
0011 - 0111	Reserved	Reserved

The Data Rate field is set to the PHY mode of the data rate that the recipient device recommends that the source device use. The Data Rate field is encoded as shown in Table 88.

Table 88 - Data rate field encoding

Value	Type A SCBT modes	Type A OFDM modes	Type B modes
0	A0	A14	B0
1	A1	A15	B1
2	A2	A16	B2
3	A3	A17	B3
4	A4	A18	Reserved
5	A5	A19	Reserved
6	A6	A20	Reserved
7	A7	A21	Reserved
8	A8	Reserved	Reserved
9	A9	Reserved	Reserved
10	A10	Reserved	Reserved
11	A11	Reserved	Reserved
12	A12	Reserved	Reserved
13	A13	Reserved	Reserved
14-15	Reserved	Reserved	Reserved

15.9.18 MAC capabilities IE

The MAC Capabilities IE is specified in Figure 139.

octets: 1	1	2	2
Element ID	Length (=4)	MAC Capability Bitmap	Reserved

Figure 139 - MAC capabilities IE format

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 89. Subsequent octets are reserved and may or may not be present.

Table 89 - MAC capability bitmap

octet	Bit	Attribute	Description
0	0	Hard DRP	Capable of being the owner and target of Hard DRP reservations
	1	Aggregation	Capable of transmitting and receiving aggregated data frames, =1 for Type A; 0/1 for Type B
	2	Block ACK	Capable of transmitting and acknowledging frames using the B-ACK mechanism
	3	Fixed size B-ACK buffer element	To inidcate that B-ACK frames report available buffer in terms of number of fixed size buffer elements in the Frame Count field (Annex F).
	4	Explicit DRP negotiation	Capable of negotiating a DRP reservation using command frames
	5	Hibernation anchor	Capable of acting as a hibernation anchor
	6	Probe	Capable of responding to Probe IEs received in command frames
	7	Reserved	Reserved
1	0	TPRC	Capable of receiver-driven transmit power and rate control
	1	Link Feedback	Capable of transmitter-driven link feedback
	2	Relay Capability	Capable of being a relay device
	3	Relay Support Capability	Capable of being a device that send or receive data via relay device
	4	TSD Capability	Capable of TSD (Transmit Switched Diversity) operation
	5	FUCA Capability	Capable of sending or receiving FUCA command frame
	6	OOB Capability	Capable of OOB operation
	7	Reserved	Reserved

15.9.19 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 Figure 140.

octets: 1	1	16	...	16
Element ID	Length (=16×N)	MKID 1	...	MKID N

Figure 140 - MKID IE format

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

15.9.20 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.

The format of the MAB IE is shown in Figure 141.

octets: 1	1	8	...	8
Element ID	Length (=8×N)	Multicast Address Binding Block 1	...	Multicast Address Binding Block N

Figure 141 - MAB IE format

The format of the Multicast Address Binding Block field is shown in Figure 142.

octets : 6	2
MEUI	MDevAddr

Figure 142 - Multicast address binding block format

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.

15.9.21 PHY capabilities IE

The PHY Capabilities IE pertaining to the PHY is specified in Figure 143.

octets: 1	1	9	2
Element ID	Length (=11)	PHY Capability Bitmap	Reserved

Figure 143 - PHY capabilities IE format

The PHY Capability Bitmap field indicates capabilities supported by the PHY, as defined in the Physical Layer Clauses of this specification (Clauses 7 - 15). 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 90. Subsequent octets are reserved and may or may not be present.

Table 90 - PHY capability bitmap

Octet	Bit	Attribute	Description	Type A	Type B
0	0	Type A	Bit set to one indicates that the device type is A	1	0
	1	Type B	Bit set to one indicates that the device type is B	0	1
	2	Channel 1	Bit set to one indicates support for channel 1 (BAND ID=1)	0/1	0/1
	3	Channel 2	Bit set to one incates support for channel 2 (BAND ID=2)	0/1	0/1
	4	Channel 4	Bit set to one incates support for channel 4 (BAND ID=4)	0/1	0/1
	5	Channel 5	Bit set to one incates support for channel 5 (BAND ID=5)	0/1	0/1
	6	Channel 6	Bit set to one incates support for channel 6 (BAND ID=6)	0/1	0/1
	7	Channel 7	Bit set to one incates support for channel 7 (BAND ID=7)	0/1	0/1
1	0	Channel 8	Bit set to one incates support for channel 8 (BAND ID=8)	0/1	0/1
	1	Channel 9	Bit set to one incates support for channel 9 (BAND ID=9)	0/1	0/1
	2	Channel 10	Bit set to one incates support for channel 10 (BAND ID=10)	0/1	0/1
	3	A1 (TX)	Bit set to one indicates TX support for mode A1	0/1	0
	4	A1 (RX)	Bit set to one indicates RX support for mode A1	0/1	0
	5	A2 (TX)	Bit set to one indicates TX support for mode A2	0/1	0/1
	6	A2 (RX)	Bit set to one indicates RX support for mode A2	0/1	0
	7	A3 (TX)	Bit set to one indicates TX support for mode A3	0/1	0
2	0	A3 (RX)	Bit set to one indicates RX support for mode A3	0/1	0
	1	A4 (TX)	Bit set to one indicates TX support for mode A4	0/1	0
	2	A4 (RX)	Bit set to one indicates RX support for mode A4	0/1	0
	3	A5 (TX)	Bit set to one indicates TX support for mode A5	0/1	0/1
	4	A5 (RX)	Bit set to one indicates RX support for mode A5	0/1	0/1
	5	A6 (TX)	Bit set to one indicates TX support for mode A6	0/1	0
	6	A6 (RX)	Bit set to one indicates RX support for mode A6	0/1	0
	7	A7 (TX)	Bit set to one indicates TX support for mode A13	0/1	0

Table 90 - PHY capability bitmap (continued)

Octet	Bit	Attribute	Description	Type A	Type B
3	0	A7 (RX)	Bit set to one indicates RX support for mode A14	0/1	0
	1	A8	Bit set to one indicates TX support for mode A8	0/1	0
	2	A8	Bit set to one indicates RX support for mode A8	0/1	0
	3	A9 (TX)	Bit set to one indicates TX support for mode A9	0/1	0
	4	A9 (RX)	Bit set to one indicates RX support for mode A9	0/1	0
	5	A10 (TX)	Bit set to one indicates TX support for mode A10	0/1	0
	6	A10 (RX)	Bit set to one indicates RX support for mode A10	0/1	0
	7	A11 (TX)	Bit set to one indicates TX support for mode A11	0/1	0
4	0	A11 (RX)	Bit set to one indicates RX support for mode A11	0/1	0
	1	A12 (TX)	Bit set to one indicates TX support for mode A12	0/1	0
	2	A12 (RX)	Bit set to one indicates RX support for mode A12	0/1	0
	3	A13 (TX)	Bit set to one indicates TX support for mode A13	0/1	0
	4	A13 (RX)	Bit set to one indicates RX support for mode A13	0/1	0
	5	A14 (TX)	Bit set to one indicates TX support for mode A14	0/1	0
	6	A14 (RX)	Bit set to one indicates RX support for mode A14	0/1	0
	7	A15 (TX)	Bit set to one indicates TX support for mode A15	0/1	0
5	0	A15 (RX)	Bit set to one indicates RX support for mode A15	0/1	0
	1	A16 (TX)	Bit set to one indicates TX support for mode A16	0/1	0
	2	A16 (RX)	Bit set to one indicates RX support for mode A16	0/1	0
	3	A17 (TX)	Bit set to one indicates TX support for mode A17	0/1	0
	4	A17 (RX)	Bit set to one indicates RX support for mode A17	0/1	0
	5	A18 (TX)	Bit set to one indicates TX support for mode A18	0/1	0
	6	A18 (RX)	Bit set to one indicates RX support for mode A18	0/1	0
	7	A19 (TX)	Bit set to one indicates TX support for mode A19	0/1	0
6	0	A19 (RX)	Bit set to one indicates RX support for mode A19	0/1	0
	1	A20 (TX)	Bit set to one indicates TX support for mode A20	0/1	0
	2	A20 (RX)	Bit set to one indicates RX support for mode A20	0/1	0
	3	A21 (TX)	Bit set to one indicates TX support for mode A21	0/1	0
	4	A21 (RX)	Bit set to one indicates RX support for mode A21	0/1	0
	5	B1 (TX)	Bit set to one indicates TX support for mode B1	0/1	0/1
	6	B1 (RX)	Bit set to one indicates RX support for mode B1	0/1	0/1
	7	B2 (TX)	Bit set to one indicates TX support for mode B2	0/1	0/1

Table 90 - PHY capability bitmap (concluded)

Octet	Bit	Attribute	Description	Type A	Type B
7	0	B2 (RX)	Bit set to one indicates RX support for mode B2	0/1	0/1
	1	B3 (TX)	Bit set to one indicates TX support for mode B3	0/1	0/1
	2	B3 (RX)	Bit set to one indicates RX support for mode B3	0/1	0/1
	3	B4 (TX)	Bit set to one indicates TX support for mode B4	0/1	0/1
	4	B4 (RX)	Bit set to one indicates RX support for mode B4	0/1	0/1
	5	C1 (TX)	Bit set to one indicates TX support for mode C1	0/1	0/1
	6	C1 (RX)	Bit set to one indicates RX support for mode C1	0/1	0/1
	7	C2 (TX)	Bit set to one indicates TX support for mode C2	0/1	0/1
8	0	C2 (RX)	Bit set to one indicates RX support for mode C2	0/1	0/1
	1	Multi-Segment	Bit set to one indicates support for multi-segment PPDU	1	0/1
	2	Reserved			
	3	Reserved			
	4	Reserved			
	5	Reserved			
	6	Reserved			
	7	Reserved			

15.9.22 Probe IE

The Probe IE is used to request information from a device. It is specified in Figure 144.

octets: 1	1	2	1	...	1
Element ID	Length (=2+N)	Target DevAddr	Requested Element ID 1	...	Requested Element ID N

Figure 144 - Probe IE format for standard IEs

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.

15.9.23 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 Figure 145.

octets: 1	1	2	2	4		4
Element ID	Length (=4+4×N)	Relinquish Request Control	Target DevAddr	Allocation 1	...	Allocation N

Figure 145 - Relinquish request IE format

The Relinquish Request Control field is specified in Figure 146.

bits: b15-b4	b3-b0
Reserved	Reason Code

Figure 146 - Relinquish request control field format

The Reason Code field indicates the reason for the request, and is encoded as shown in Table 91.

Table 91 - 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 Figure 147.

octets: 2	2
Zone Bitmap	MAS Bitmap

Figure 147 - Allocation field format

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.

15.9.24 Relay IE

Relay IE is used to request or respond for relay operation such as Relay Reservation, Relay Set, Relay Complete and Relay Switch or is also used for a device to inform a corresponding device or relay device of setting parameters for relay operation. Relay IE is shown in Figure 148.

octets: 1	1	1	1	1	1	1	2	2	2	2
Element ID	Length	Relay Command Type	Relay Control	R-D Link LQI	Relay Mode	Path Change Interval	Detour Start Duration	Relay DevAddr	Source DevAddr	Destination DevAddr

Figure 148 - Relay IE

The Relay Command Type field is specified in Table 92.

Table 92 - Relay command type field

Value	Relay Command Type
0	Relay Reservation Request

Table 92 - Relay command type field (concluded)

Value	Relay Command Type
1	Relay Reservation Response
2	Relay Set Request
3	Relay Set Response
4	Relay Complete Request
5	Relay Complete Response
6	Relay Switch Request
7	Relay Switch Response

The Relay Control field is specified in Figure 149.

bits: b7-b6	b5	b4-b3	b2	b1-b0
Path ID	Result Code	Relay Reservation	Start Relay	Relay Antenna Complete

Figure 149 - Relay control field

The Path ID field indicates the ID of the path requested by the source device in Relay Switch Request command.

The Result Code field is set to ONE if the request command is successful and is set to ZERO if the request command is failed.

The Relay Reservation field indicates the reason codes of relay reservation request. The possible values and its interpretations are:

0: Relay reservation is accepted by a relay device.

1: Relay reservation is rejected because the relay has been already reserved by another device.

2: Relay reservation request is invalid because the requested device does not support relaying

3: Reserved.

The Start Relay field is set to ONE if a device received Relay Set Request command is able to do antenna training with source and destination device. Otherwise it is set to ZERO.

The Relay Antenna Complete field indicates the reason codes of Relay Complete request. The possible values and its interpretations are:

0: Relay Complete request is accepted because the antenna training between relay device and destination device is successful. In this case, a number of DRP availability IEs which informs the source of the result of Channel Scanning Response shall be included in the Relay Complete Response frame.

1: Relay Complete request is rejected because the antenna training between relay device and destination device is failed.

2 - 3: Reserved.

The R-D Link LQI field indicates LQI of the link between a relay device and a destination device. It is set only for Relay Complete Response command frame in order for the destination device to inform the source of the LQI of the relay-destination link if antenna training of the link is successfully done.

The Relay Mode field is specified in Figure 150.

bits: b7-b5	b4-b1	b0
Reserved	Path Order	Transmission Mode

Figure 150 - Relay mode field

The Transmission Mode field indicates whether transmission mode in DRP is Normal or Alternation. The Transmission Mode bit is set to ZERO for the Normal mode and is set to ONE for the Alternation mode, as described in 20.2.2, respectively.

The Path Order field is set to the order of path usage when Alternation mode is set. The Path Order field encoding is shown in Table 93.

Table 93 - Path order field

Value	Path Order
0000	0 - 1
0001	0 - 2
0010	0 - 1 - 2
0011	0 - 2 - 1
0100 - 1111	Reserved

Each number in the path order specifies the Path's ID. 0 means direct path between source device and destination device. Successive numbers such 1, 2 are named as the relay paths in order of antenna training between relay and two devices. Direct path is used first and relay paths are used consecutively. After one cycle of path order is used, the path order is repeated.

The Path Change Interval field indicates the instant when the path of data transmission between source device and destination device is changed. From the start position of one reserved consecutive MASs, every instant of Path Change Interval can have an opportunity to change the path, Within one Path Change Interval, only one path is used for data transmission. The unit of this field is MAS.

The Detour Start Duration field indicates when data transmission via detour path starts in the current DRP. It is set only for Relay Switch Request/Response Command.

The Relay DevAddr field is set to the DevAddr of the device which is requested from a source device for relaying. This is used only for Relay Complete Request command frame in order to let destination device inform the address of a relay device.

Source DevAddr field is set to the DevAddr of the device which requests relay operation and Destination DevAddr is set to the DevAddr of the recipient device to which the device is sending the relay operation request. Two fields are used for informing to the relay device that these two devices intend to use it.

15.9.25 Scan Timing IE

The Scan Timing IE announces the start and duration of a period of time when a device listens for specific frame transmissions. The format of Scan Timing IE is specified in Figure 151.

octets: 1	1	5
ElementID	Length (=5)	Scan Timing

Figure 151 - Scantiming IE format

The Scan Timing field is set to the start time and duration of the scanning time period and is specified in Figure 152.

bits: b39-b20	b19-b0
Scan Duration	Scan Start-time

Figure 152 - Scan timing field

Scan Start-time field is set to the starting time of scanning using the same antenna beam measured from the end of the PLCP header of the beacon frame in units of microseconds.

Scan Duration is set to the duration of the scanning in units of microseconds, using the same antenna beam.

15.9.26 UEP information IE

UEP information IE shall be used to indicate the supported UEP types and UEP MCS modes as specified in UEP information response command.

octets: 1	1	1	1	...	1
Element ID	Length	UEP Information	UEP MCS block-1	...	UEP MCS block-n

UEP Information:

bits: 2	1	1	4
UEP mode	UEP coding	UEP mapping	Reserved

UEP MCS block:

bits: 6	2
MCS	Reserved

Figure 153 - UEP information response command

UEP mode definition is specified in Table 94.

Table 94 - UEP mode

UEP Mode	Description
0b00	EEP
0b01	S-UEP
0b10	P-UEP
0b11	Reserved

UEP coding bit shall set to "1" if UEP coding is supported.

UEP mapping bit shall set to "1" if UEP mapping is supported.

Each UEP MCS block is used to indicate a supported UEP MCS mode.

16 MAC sublayer functional description

This Clause specifies MAC sublayer functionality. The rules for transmission and reception of MAC frames, including setting and processing MAC header fields and information elements, are specified in 16.1.

Devices discover each other through the transmission of discovery frames in the Discovery Channel, as specified in 16.3; Devices perform antenna training frames in the Discovery channel using the antenna training and tracking protocol (ATTP) specified in 16.18. Devices coordinate the selection of a Data Channel using the channel selection protocol specified in 16.4. CSMA/CA with random backoff is used to access the Discovery Channel for the purposes of device discovery, antenna training and channel selection, and is specified in 16.2.

For scheduled channel access, channel time is divided into superframes, with each superframe composed of two major parts, the beacon period (BP) and the data period. Beacon transmission and reception in the BP and merging of BPs are specified in 16.5.

During the data period devices send and receive data using reservations established using the distributed reservation protocol (DRP). The DRP enables a device to gain scheduled access to the medium within a negotiated reservation, and is specified in 16.6.

Clause 16.7 specifies the rules and mechanisms to support coexistence among mixed types of devices and data exchange between heterogeneous types of devices.

Device synchronization is specified in 16.8. The fragmentation and reassembly of MSDUs is specified in 16.9. Aggregation of multiple MSDUs in a single frame is specified in 16.10. Channel Bonding is specified in 16.11. Acknowledgement mechanisms are specified in 16.12. Clauses 16.13 through 16.17 specify probe commands, channel selection, multi-rate support, transmit power control, power management mechanisms and use of ASIEs. Clause 16.19 specifies dynamic relay transmission for blocked or shadowed links. Clause 16.20 specifies values for all MAC sublayer parameters.

16.1 Frame processing

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

16.1.1 Frame addresses

Frames are addressed using DevAddrs. There are four types of DevAddrs: Private, Generated, Multicast, and Broadcast. Table 95 shows the range for each type of DevAddr.

Table 95 - DevAddr types and ranges

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

A device shall associate a DevAddr of either type Private or type Generated with its local MAC sublayer. A device may use a NULL EUI-48 value (all bits set to ONE) to indicate it does not have a unique EUI_48 value. 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 should ensure that the generated value is unique among all devices in its extended beacon group. Private DevAddres are reserved for devices outside the scope of this specification. Selection and conflict resolution for Private DevAddrs is out of scope of this International 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 16.1.10.20. 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.

16.1.1.1 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.

16.1.2 Frame reception

Unless otherwise indicated, a frame is considered to be received by the device if it has a valid header check sequence (HCS) and segment check sequence (FCS) as defined in 15.2.9 and indicates a protocol version that is supported by the device. The HCS is validated by the PHY, which indicates whether or not a header error occurred.

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

16.1.3 Antenna training frame transaction

An antenna training frame transaction in the discovery channel consists of an RTT/CTT frame exchange and a sequence of training frames.

Figure 154 shows a frame transaction example where the initiator requests to train its receiving antenna and respondor requests to train both its transmitting and receiving antenna with 4 training iterations. Figure 155 shows a frame transaction example where both the initiator and the respondor request training of both transmitting and receiving antennas. The initiator and respondor devices shall set the Duration field to protect the training iterations as specified in 16.1.9.1. Other devices receiving these frames shall update their NAV as described in 16.2.2.

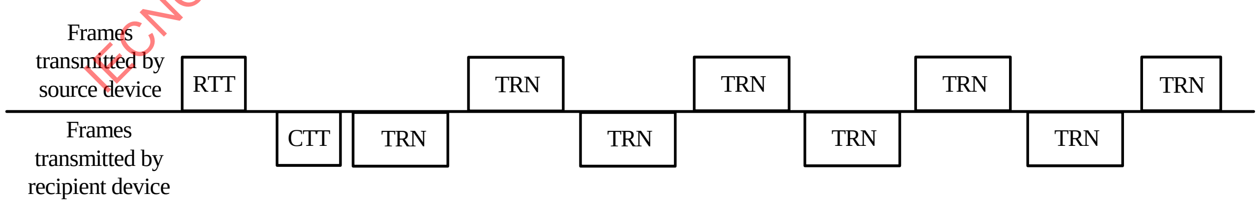

Figure 154 - Frame transaction of an antenna training example with 4 iterations

16.1.4 Frame transfer

A source device shall transmit MSDUs associated with the same Delivery ID (15.2.3.5) 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, labelled 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 were 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.

16.1.5 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 (15.2.3.6) 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.

16.1.6 Inter-frame space (IFS)

Three types of IFS are used in this International Standard: the long inter-frame space (LIFS), the minimum inter-frame space (MIFS), the short inter-frame space (SIFS). The actual values of the MIFS, SIFS 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.

16.1.6.1 MIFS

The length of MIFS is given by the pMIFS parameter defined in Table 39.

16.1.6.2 SIFS

Within a frame transaction, all frames shall be separated by a SIFS interval. The length of SIFS is given by the pSIFS parameter defined in Table 39.

16.1.6.3 LIFS

When transmitting beacon frames in discovery channel, LIFS is the minimum time that a device defer access to the medium after it determines the medium to have become idle.

16.1.7 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.

16.1.8 RTT/CTT use

An RTT/CTT exchange is used to initiate antenna training between a source and a destination device, and shall be used only on the discovery channel. With an appropriately set Duration field as specified in 16.1.9.1, the RTT and CTT frames prevent the neighbours of the source and recipient devices from accessing the medium while the source and recipient are exchanging the following frames.

If a device receives an RTT frame addressed to it in the discovery channel, it shall transmit a CTT frame SIFS after the end of the received frame if and only if its NAV is zero.

On receiving an expected CTT response, the source device shall transmit the antenna training frames for which it transmitted the preceding RTT frame SIFS after the end of the received CTT frame. If the source device does not receive the expected CTT frame SIFS plus the CTT frame transmission time after the end of the RTT frame transmission, and it transmitted the RTT in the discovery channel, it shall invoke a backoff as specified in 16.2.6.

16.1.9 MAC header fields

16.1.9.1 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;

— Zero.

A device shall set the Duration field in discovery, 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 (Imm-ACK, or B-ACK frame), if any;

— The transmission time of subsequent frames, if any, to be sent to the same recipient up to and including (a) the next frame with ACK Policy set to Imm-ACK or B-ACK Request or (b) the last frame in the DCA TXOP or reservation block, whichever is earlier; or, alternatively, the transmission time of the next frame in the DCA TXOP or reservation block to be sent to the same recipient, if any; and

— All the IFSs separating the frames included in the Duration calculation.

A device shall set the Duration field in RTT and TRN frames to the sum of:

— The transmission time of the frame body of the current frame;

— The transmission time of the expected CTT;

— The transmission time of subsequent TRN frames of one training iteration if any; and

— All the IFSs separating the frames included in the Duration calculation.

For the purpose of calculation of duration field for RTT and TRN frames, a device shall assume the same spreading factor as current iteration for the next training iteration. In addition, in calculation of the duration field for a RTT frame, the device shall assume the training peer device also requires training feedback and training of receiving antenna, if the peer device is capable of beamforming.

A device shall set the Duration field in CTT, responding TRN, 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 SIFS, minus the transmission time of the frame body of the received frame to which the CTT, Imm-ACK or B-ACK is responding, minus the transmission time up to the end of the PLCP header of this CTT, Imm-ACK or B-ACK frame.

Figure 155 illustrates setting of Duration field for RTT, CTT and TRN frames.

A device shall round a fractional calculated value for Duration in microseconds up to the next integer.

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.

The following exceptions to previous rules are allowed:

— For frames with ACK Policy set to B-ACK Request, a device may 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 to zero.

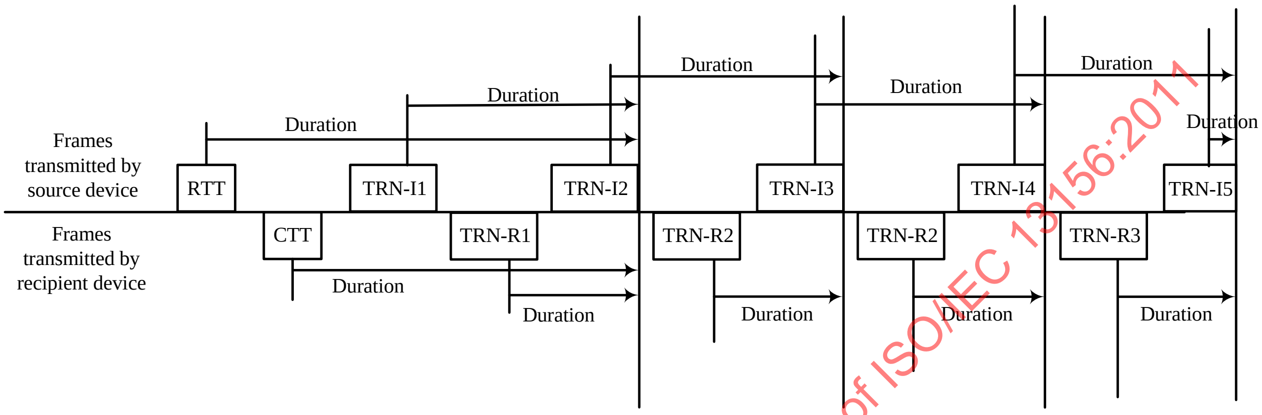

Figure 155 - Duration field values of an example of antenna training with 4 iterations

16.1.9.2 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.

16.1.9.3 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 device shall use a dedicated counter for MCDUs.

A device shall use a dedicated counter for discovery frames.

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.

16.1.10 Information elements

IEs are contained in beacon, discovery, command and control 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 Clause describes when each IE is generated.

16.1.10.1 Application-specific IE (ASIE)

A device may include an ASIE in its beacon for each of its applications which have made the request, as described in 16.17. The scope of the ASIE is dependent on the application that requested the inclusion of the ASIE.

16.1.10.2 Application-specific probe IE

A device may send an Application-specific Probe IE in order to request a specific ASIE. The scope and required response is dependent on the application that defines the ASIE.

16.1.10.3 Antenna Capability IE (ACIE)

A device shall include a ACIE in all the discovery frames when participating in the discovery process as speficified in 16.18.1. A device may inlcude ACIE in its beacon.

16.1.10.4 Antenna Training IE (ATIE)

As specified in 16.18.2. the ATIE shall be included in the RTT and CTT control frames used in the antenna training/tracking configuration, and may be included in a TRN frame sent by a training initiator to respond to the configuration request sent by a training responder.

16.1.10.5 Antenna Training/Tracking Control IE (ATTCIE)

ATTCIE shall be included in TRN frames in iterative antenna training (16.18.3) or explicit tracking (16.18.4.1) and aggredated training frames, Imm-ACK or B-ACK frame in implict tracking (16.18.4.2).

16.1.10.6 Antenna Feedback IE (AFIE)

At request, AFIE shall be included in TRN frames in antenna training (16.18.3) and explict tracking (16.18.4.1) and B-ACK, Imm-ACK and aggregated training frames in implicit antenna tracking (16.18.4.2).

16.1.10.7 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 previous superframe, as well as information retained based on hibernation mode rules.

16.1.10.8 BP Switch IE

A device should include a BP Switch IE in its beacon prior to changing its BPST, as specified in 16.5.3.11.3.

16.1.10.9 Channel bonding IE (CBOIE)

A CBOIE is used to negotiate bonding of two or more channels and to announce the channels that are involved in the channel bonding.

When using channel bonding, a device shall include a CBOIE in its beacons transmitted in a designated beaconing channel to indicate the (proposed) bonded channels as defined in 16.11.

16.1.10.10 Channel change IE

A device should include a Channel Change IE in its beacon prior to changing to a different channel. A device that includes a Channel Change IE should change channels as indicated in the IE. A Type A device shall include a Channel Change IE in its discovery frame to announce its home channel to potential newly discovered device when the device returns to Discovery Channel to discover additional device.

16.1.10.11 Channel measurement IE

A device shall include a channel measurement IE in a Channel Scanning Response command frame for each scanned channel without superframe structure in explict channel scan procedure (16.4.1.1). A device may include channel measurement IEs in a Channel Scanning Response command frames for all the scanned channels with or without superframe structure (15.6.3.2).

16.1.10.12 Distributed reservation protocol (DRP) IE

A device shall include DRP IEs in its beacon for all reservations in which it participates as a reservation owner or target, as described in 16.6.

16.1.10.13 DRP availability IE

A device shall include a DRP Availability IE in its beacon as required to support DRP reservation negotiation, as described in 16.6.

16.1.10.14 Hibernation anchor IE

A device that indicates it is capable of acting as a hibernation anchor should include a Hibernation Anchor IE in its beacon to provide information on neighbours that are currently in hibernation mode as described in 16.16.4.

16.1.10.15 Hibernation mode IE

A device shall include a Hibernation Mode IE in its beacon before entering hibernation mode, as specified in 16.16.3. 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.

16.1.10.16 Identification IE

A device may include an Identification IE in its beacon to provide its own identifying information to neighbours.

16.1.10.17 Link feedback IE

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

16.1.10.18 MAC capabilities IE

A device may include a MAC Capabilities IE in its beacon.

16.1.10.19 Master key identifier (MKID) IE

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

16.1.10.20 Multicast address bind (MAB) IE

A device may include a MAB IE for any active multicast bindings between multicast EUI-48s and McstAddrs. A device should include a MAB IE in its beacon 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 sublayer 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.

16.1.10.21 PHY capabilities IE

A device may include a PHY Capabilities IE in its beacon.

16.1.10.22 Probe IE

A device may include a Probe IE in its beacon to request certain IEs from another device.

16.1.10.23 Relinquish Request IE

A device may include a Relinquish Request IE in its beacon 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 DRP Availability IE and a Relinquish Request IE identifying those MASs with the Target DevAddr field set to the DevAddr of the reservation owner.

16.1.10.24 Relay IE

If a device supports the relay procedure as specified in 20.1.2, 20.1.4, and 20.1.5; it shall include a Relay IE in its beacon or command frame to support the relay procedures.

If a device supports the relay procedure as specified in 20.2.7.1, it shall include a Relay IE in its beacon or command frame to detour to another reliable path when a current path is considered unavailable as described in 20.2.6.

16.1.10.25 Scan Timing IE

A device shall always include a Scan Timing IE in a discovery frame. In the Scan Timing IE, the device shall indicate the time and duration the device will listen for response to the beacon using the same antenna beam. A device shall include Scan Timing IE in a Channel Scanning Request command frame to indicate the time when the device returns to Discovery Channel to listen for Channel Scanning Response command frame.

16.1.10.26 UEP IE

A device may include an UEP IE in its beacon to indicate the UEP types and UEP MCS modes that it can support.

16.2 Distributed contention access (DCA)

The DCA mechanism provides distributed contention access to the medium for the following types of frame exchanges in the discovery channel:

- Discovery frames
- RTT/CTT control frames
- Antenna training frames
- Channel selection command frames

In addition, DCA is also the medium access method in DCA reservation blocks.

A frame transferred over the wireless medium using DCA is referred to as a DCA frame.

All devices shall use both physical and virtual carrier sensing before accessing the medium. The RTT and CTT frames shall be transmitted using PHY mode D0 (see 10.2.5). The training frames shall be transmitted using one of the discovery PHY modes with adaptive spreading factor.

16.2.1 DCA medium availability

A device shall consider the medium available for DCA at the following times:

- All the times in the discovery channel.
- DCA reservation blocks except for a zero-length interval at the start of DCA reservation blocks with Reservation Status set to one if a neighbor is the reservation owner, for purposes of determining TXOP limits.

At all other times, a device shall consider the medium unavailable for DCA.

16.2.2 NAV

A device that transmits and receives frames using DCA shall maintain a network allocation vector (NAV) that contains the remaining time that a neighbour device has indicated it will access the medium. The 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 HCS outside shall update its NAV as if the frame were correctly received with Duration equal to ZERO.

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

16.2.3 Medium status

For DCA 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 is transmitting or receiving a frame on the medium;
- When the Duration announced in a previously transmitted frame has not yet expired;
- When the device's NAV is greater than zero, if the device maintains a NAV;
- When the medium is unavailable for DCA.

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

16.2.4 Obtaining a TXOP

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

- The device has one or more newly generated DCA frames;
- The device had a backoff counter of zero value and had no DCA frames prior to the generation of the new DCA frames;
- The device determines that the medium has been idle for LIFS or longer.

The device shall start transmitting a frame, which may be an RTT frame, as soon as the above conditions are satisfied. 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 if it meets the following conditions:

- The device has one or more DCA frames buffered for transmission, including retry;
- The device set the backoff counter to zero in the last backoff and determines that the medium has been idle for LIFS since that backoff at the end of the current DCA slot, or the device decrements its backoff counter from one to zero in the current DCA slot.

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

A device shall ensure that the TXOP it has obtained ends SIFS plus mGuardTime before the medium becomes unavailable for DCA.

Figure 156 illustrates the timing relationship in obtaining a TXOP.

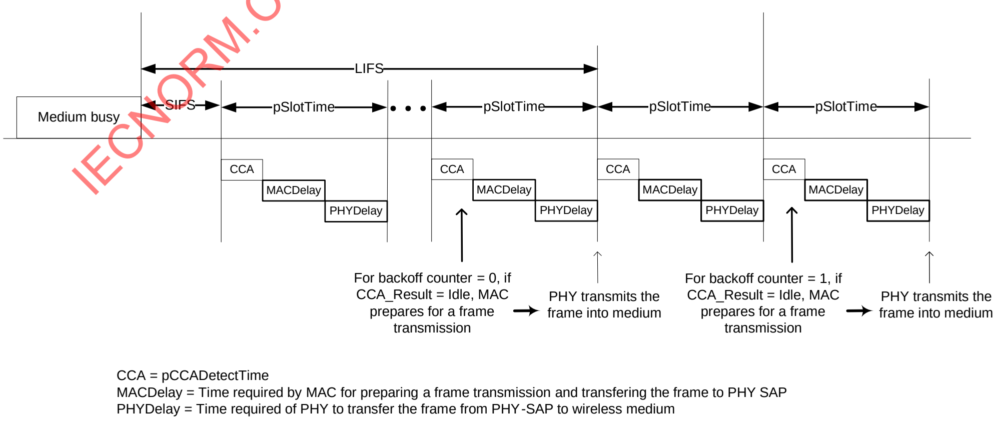

Figure 156 - DCA timing relationships

16.2.5 Using a TXOP

A device that has obtained a TXOP is referred to as a TXOP owner. A TXOP owner shall initiate one or more frame transactions without backoff, in the TXOP it has obtained, 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 recipient device shall not transmit a CTT frame in response to a received RTT frame if its NAV is greater than zero. A recipient device shall not transmit a CTT, Imm-ACK or B-ACK response to a received frame requiring such a response if the response will not be completed SIFS plus mGuardTime before the medium becomes unavailable for its DCA.

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 discovery frame, the TXOP owner shall transmit a subsequent discovery frame MIFS after the end of that transmitted frame.
- 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 RTT frame or a non-RTT frame with the ACK Policy set to Imm-ACK or B-ACK Request, the recipient device shall transmit a CTT frame or an Imm-ACK or B-ACK frame SIFS after the end of the received frame.
- After receiving an expected CTT, 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, SIFS after the end of the received frame.
- After receiving a requested B-ACK frame with a valid HCS but an invalid FCS, the TXOP owner shall retransmit the last frame it transmitted, or transmit the next frame, SIFS after the end of the B-ACK frame.

If a device cannot transmit its next frame according to these timing requirements, it shall consider the TXOP ended.

16.2.6 Invoking a backoff procedure

A device shall maintain a backoff counter to transmit frames using DCA.

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

The device shall set CW back to mCWmin after receiving a CTT or Imm-ACK frame or the MAC header of a B-ACK frame expected in response to the last transmitted frame, or upon transmitting a frame with ACK Policy set to No-ACK. A device shall also set CW back to mCWMin, but shall not select a new backoff counter value, after discarding a buffered DCA frame.

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

1. A device shall invoke a backoff, with CW set to mCWmin, when it has a DCA frame arriving at its MAC SAP, or a DCA frame generated at the MAC sublayer entity under the following conditions:

   - The device had a backoff counter of zero value but is not in the middle of a DCA frame transaction; and the device determines that the medium is busy.

2. A device shall invoke a backoff, with CW set to mCWmin, at the end of transmitting a DCA frame with the ACK policy set to No-ACK or B-ACK, or at the end of receiving an expected Imm-ACK or B-ACK response to its last transmitted DCA frame, under the following condition:

- The device has no other DCA frames for transmission in the current TXOP obtained.

3. A device shall invoke a backoff, with CW set to mCWmin, at the end of transmitting a DCA frame with the ACK policy set to No-ACK or B-ACK, or at the end of currently receiving the DCA frame header of an expected Imm-ACK or B-ACK response frame to its last transmitted DCA frame, under the following conditions:

   - The device has one or more DCA frames 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 to complete the next DCA frame transaction.

4. A device shall invoke a backoff, with CW (but not the backoff counter in general) kept to the same value, at the start of a TXOP obtained under the following condition:

   - The device finds that there is not enough time to complete a pending DCA frame transaction in the obtained TXOP.

5. A device shall invoke a backoff, with CW set to the smaller of mCWmax or 2×CW+1 (the latter CW being the last CW value), at SIFS plus the Imm-ACK frame transmission time after the end of the last DCA frame it transmitted, under the following condition:

   - The device does not receive an expected CTT or Imm-ACK frame, or does not correctly receive the MAC header of a requested B-ACK frame by this time.

16.2.7 Decrementing a backoff counter

Upon invoking a backoff, a device shall ensure that the medium is idle for LIFS before starting to decrement the backoff counter. To this end, a device shall define the first DCA slot to start at the time SIFS after the medium has been idle, as defined in Figure 157, with subsequent DCA slots following successively until the medium becomes busy. All DCA slots have a length of pSlotTime.

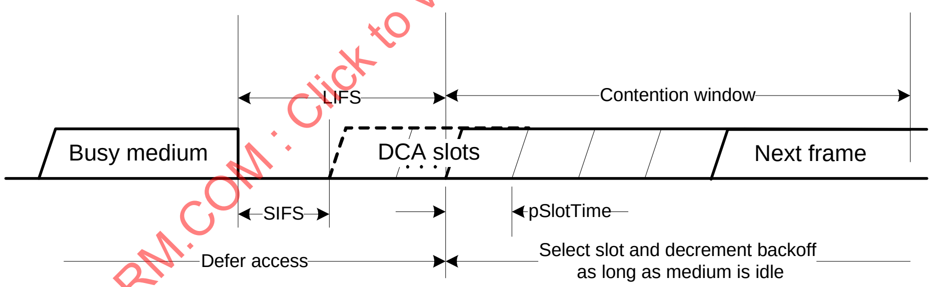

Figure 157 - Transmission of beacons in the discovery channel

A device shall treat the CCA result at pCCADetectTime after the start of a DCA slot to be the CCA result for that DCA slot. If the medium is idle in a DCA slot, and the medium has been idle for at least LIFS, the device shall decrement the backoff counter by one at that time. This procedure is also defined in Figure 157.

The device shall freeze the backoff counter when 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, following the above procedure to resume decrementing the backoff counter.

16.3 Device discovery

To discover another device with which it intends to exchange MPDUs, a device shall transmit discovery frames in the Discovery Channel follwowing an initial scan upon power-up.

Discovery frames shall be transmitted using a specified PHY mode according to its device type. To discover devices of different types it supports, a device shall transmit discovery frames using different PHY modes corresponding to the types of the devices it intends to discover, as specified in the following sub clauses. To discover Type A devices, discovery frames shall be transmitted using mode-D0, and are referred to as A-Discovery frames. To discover Type B devices, discovery frames shall be transmitted using mode-B0, and are refered to as B-Discovery frames.

16.3.1 Power-up scan

After powering up, a device shall perform an initial scan according to its device type, as specfiied below.

16.3.1.1 Type A

A Type A device shall scan for Type A beacons in the Discovery Channel for at least one superframe after powering up. If such beacons are received, the Type A device shall send a beacon with Status set to Preemptive in a signalling slot that is randomly chosen in the BP indicated by the received beacon. If the Status field of the received beacon is Ready, the preempting beacon shall be of Type A; Otherwise if the Status field is Dual, the preempting beacon shall be of Type B. In addition, the Type A device may also send the same preempting beacons in any slots in the BP indicated by the received beacon. If no Type A beacons are received during the scan, the Type A device shall start device discovery process in the discovery channel as described in 16.3.2. In order to discover other devices after starting the transmission in a Data Channel, a Type A device shall switch back to the Discovery Channel and send a discovery block set as described in 16.3.2.2.

If a discovey frame is received correctly in the Discovery Channel, the Type A device may start the antenna training with the device from which the discovery frame is received as described in 16.18, or may start the explicit channel selection with the device from which the discovery frame was received, as specified in 16.4.1, or may transmit a discovery frame with Status field set to Response.

16.3.1.2 Type B

A Type B device shall scan each channel for at least one superframe after powering up. After the scanning, the Type B device shall not transmit any frames unless the scanning indicates that one of the following conditions is true:

- If a channel is detected as busy and at least one Type B beacon is received correctly, the Type B device may join the device from which the Type B beacon is received by sending a dual-beacon as described in 16.5.2.
- If a channel is detected as busy and at least one B-Discovery frame is received correctly, the Type B device may join the device from which the B-Discovery frame is received by transmitting a B-Discovery frame with Status set to Response during the time period as indicated in the Scan Timing IE in the received B-Discovery frame.
- If it has detected an idle channel during the scanning and has not joined another device by responding to a Type B beacon or discovery frame, the Type B device shall continue the discovery process specified below:

  A Type B device shall scan for Type B beacons in the Discovery Channel for at least one superframe. If such beacons are received, the Type B device shall send a Type B beacon with Status set to Preemptive in a signalling slot that is randomly chosen in the BP indicated by the received beacon. In addition, the Type B device may also send the same preempting beacons in any slots in the BP indicated by the received beacon. If no Type B frames are received during the scan, the device shall start the device discovery process in the discovery channel as described in 16.3.2.1.

If a B-Discovery frame is received correctly in the discovery channel, the Type B device may start an explicit channel selection process with the device from which the B-Discovery frame is received, as described in 16.4.1, or may transmit a discovery frame with the Status field set to Response, or if the device from which the discovery frame was received is of Type A, it may start the antenna training procedure with the originating device, as specified in 16.18.

16.3.2 Transmission and reception of discovery frames

For the purposes of device discovery, a device shall transmit discovery frames in Discovery Channel as specified in this subclause. DCA shall be used to transmit discovery frames in Discovery channel.

16.3.2.1 Type B discovery block set

A Type B device shall transmit B-Discovery frames only in the Discovery Channel. After completion of a B-Discovery frame transmission, the Type B device shall scan for responses for a duration of B-SCAN. The scanning shall start a SIFS after the end of the B-Discovery frame transmission. The transmission of a B-Discovery frame and the scanning afterwords form a B-Discovery block. If it has multiple antenna that cover multiple sectors, after the completion of the scanning in the first sector, the Type B device shall switch to the next sector and repeat the same B-Discovery block for every sector that the device can cover using a different antenna block, or beam. A Type B discovery block set is a series of B-Discovery blocks that are performed in turn sweeping though all the sectors that the device has. The timing of a Type B discovery block set is depicted in Figure 158.

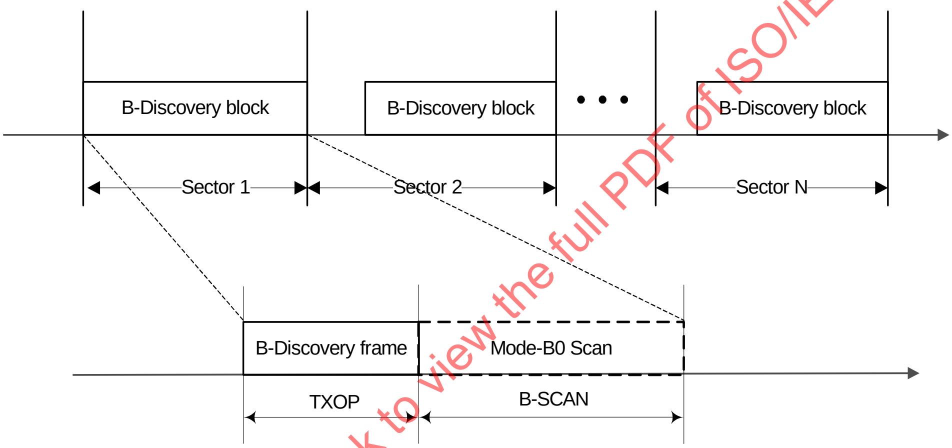

Figure 158 - Type B discovery block set

Norminal duration of a Type B DBS (with no medium access collision) is specified in Table 96.

Table 96 - Norminal duration of a Type B DBS

PHY-dependent Parameter	Value
pDBSDuration	573.3 μs

16.3.2.2 Type A discovery block set

A Type A device shall transmit A-Discovery frames only in the Discovery Channel.

After the completion of an A-Discovery frame transmission, the Type A device shall perform one or more B-Discovery blocks. The number of B-Discovery blocks shall be equal to the number of sectors that the device is capable of covering using a number of antenna blocks or beams. The transmission of the A-Discovery and B-discovery frames in series in the same sector shall be separated by MIFS and completed in one TXOP obtained. The scanning shall start SIFS time after the end of the B-Discovery frame transmission. The Type A device shall first scan for responses to the transmitted B-Discovery frame for a duration of B-SCAN, after which the Type A device shall scan for responses to the previously transmitted A-Discovery frame for a duration of D-SCAN. If the Type A device has multiple antennas that cover multiple sectors, after the completion of the scanning in the first sector, the Type A device shall switch to the next sector to perform the next B-Discovery block in which the transmission of the B-Discovery frame shall be completed in one TXOP obtained. A Type A discovery block set consists of an A-Discovery frame and a Type B discovery block set with a scanning for Type A discovery responses in between. The timing of a Type A discovery block set is specified in Figure 159.

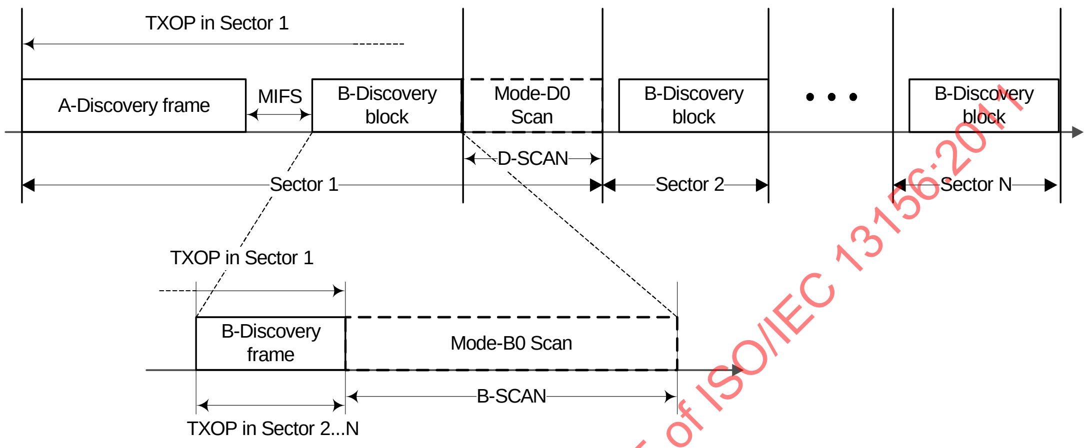

Figure 159 - Type A discovery block set

Norminal duration of a Type A DBS (with no medium access collision) is specified in Table 97.

Table 97 - Norminal duration of a Type A DBS

PHY-dependent Parameter	Value
pDBSDuration	10003.2 μs

16.3.2.3 Discovery scanning while performing DBS

When accessing Discovery Channel for the purpose of device discovery, the channel time shall be divided into Discovery Intervals with equal durations of mDILength, where the Discovery Interval is repeated in time.

In each Discovery Interval, a device shall perform one and only one discovery block set. When transmitting the first discovery frame of a discovery block set as described in 16.3.2.2 and 16.3.2.1, a device shall schedule its transmission at a time randomly drawn from a uniform distribution over the set {0, mDBSMax} measuring from the begining of the Discovery Interval.

When performing the next discovery block set, a Type A device shall switch to the next sector to transmit the A-Discovery frame using a different antenna block or beam, if it covers multiple sectors using a number of antenna blocks, or beams. The order of antenna blocks, or beams used when switching shall remain the same in repetition.

During each Discovery Interval except for the time to perform a DBS, a Type A device shall scan for discovery responses transmitted in mode-D0, and a Type B device shall scan for discovery responses transmitted in mode-B0. A device shall scan all the sectors in turn and repeat in the same order if no response to a transmitted discovery frame is received. For each sector it covers, a Type A shall scan for A-DIS-SCAN Discovery Intervals (embedded with all DBSs) before switching to scan the next sector, and a Type B device shall scan for B-DIS-SCAN Discovery Intervals (embedded with all DBSs) before switching to scan the next sector. Such processes of discovery scanning while performing DBS are illustrated in Figure 160 and Figure 161 for Type A and Type B devices respectively.

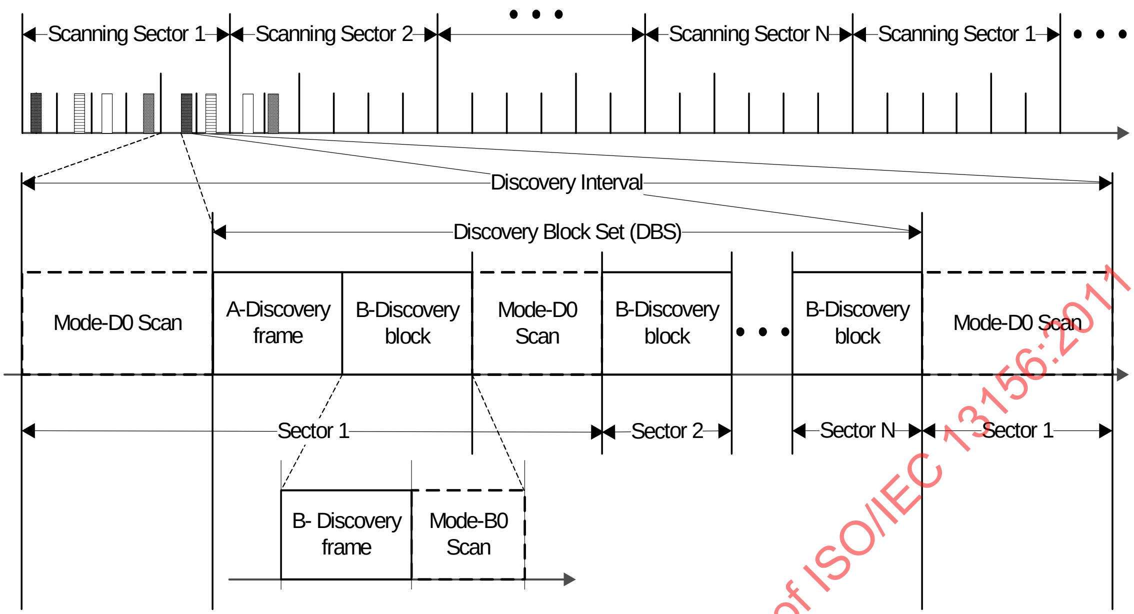

Figure 160 - Discovery intervals of a Type A device

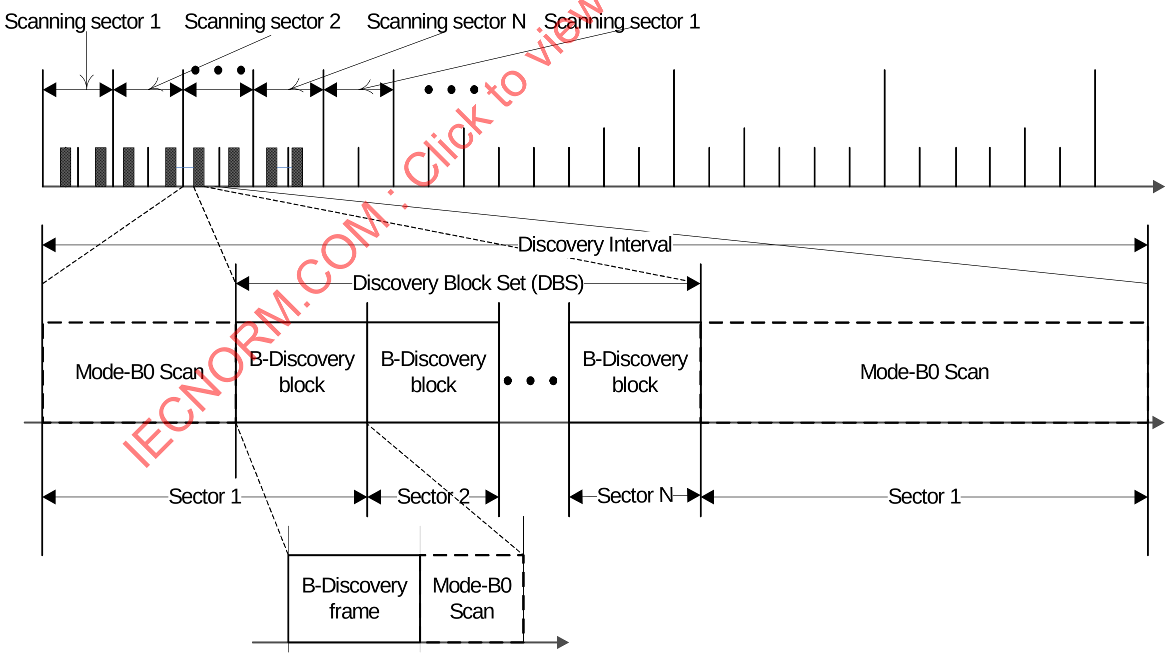

Figure 161 - Discovery intervals of a Type B device

16.4 Channel selection

A device shall use the explicit channel selection process specified in 16.4.1 to select a channel to send its beacons or B-Poll, before it sets up or joins a beacon group in that channel.

A device shall send channel selection command frames with a mode that its intended recipient supports. A device shall not send channel selection command frames with any discovery mode.

After devices started or joined a beacon group by sending beacons in a selected data channel, an implicit channel selection procedure described in 16.4.2 allows a group of devices in the beacon group to change channels in a coordinated manner.

16.4.1 Explicit channel selection

A device shall not start explicit channel selection unless the device has received a discovery frame in Discovery Channel from another device with which it intends to exchange MPDUs. The device shall transmit channel selection command frames only in Discovery Channel using DCA. In the explicit channel selection process, the device shall first coordinate explicit channel scanning with the newly discovered device as described in 16.4.1.1 followed by explicit channel switch as described in 16.4.1.2.

16.4.1.1 Explicit channel scan

A device shall request newly discovered device to perform channel scanning by sending a Channel Scanning Request. Once the device receives an ACK frame to its Channel Scanning Request, the device shall leave the Discovery Channel to scan the same channel(s) as specified in its Channel Scanning Request. At the time it specified in the Scan Timing IE in its Channel Scanning Request, the device shall return to Discovery Channel to listen for Channel Scanning Response.

Upon reception of such a Channel Scanning Request, a device shall respond with an ACK frame, after which it shall remain in Discovery Channel for a duration of mChannelScanRequestRetry. If it does not receive additional Channel Scanning Request from the same device during this period, the device shall perform the requested scanning in the channels indicated in the Channel Scanning Request. Otherwise, the device shall respond with another ACK frame. The device shall return back to Discovery Channel at the time indicated in the Channel Scanning Request and send a Channel Scanning Response which includes DRP availability IE(s) to indicate the scanning result. In addition, for each channel without superframe structure, the device shall include a channel measurement IE in the Channel Scanning Response. And the device may include channel measurement IEs for all the channels with or without superframe structure in the Channel Scanning Response as specified.

The formats of Channel Scanning Request and Response are specified in 15.6.3.1 and 15.6.3.2 respectively.

16.4.1.2 Explicit channel switch

Upon reception of a Channel Scanning Response, a device shall send Channel Change Request to the device from which it received the Channel Scanning Response. After sending a Channel Change Request, the device shall listen for a Channel Change Response. Once it receives an Channel Change Response with Reason Code field set to Accepted, It shall wait for a period time that is randomly chosen over the range [0, mSuperframeLength] before switching to the accepted channel to set up a superframe structure as specified in 16.5.1 and 16.5.2 with the exception that the pair of devices will engage in a Master-Slave operation, as specified in 16.7.2, in the agreed channel. If the device receives a frame with Reason Code field set to a value other than Accepted, the device shall transmit a revised channel change request. The device shall not switch to a channel that is not accepted by the recipient.

Upon reception of a Channel Change Request, a device shall respond with a Channel Change Response frame with Reason Code field appropriately set as in Table . A device shall not switch to a channel until an Ack frame to its accepted Channel Change Response is received. Before it switches to an agreed channel to set up a superframe structure, the device shall wait for a period of time that is randomly chosen over the range [0, mSuperframeLength] with the exception that the pair of devices will engage in a Master-Slave operation, as specified in 16.7.2, in the agreed channel.

The formats of Channel Change Request and Response are specified in 15.6.3.3 and 15.6.3.4 respectively.

16.4.2 Implicit channel selection

A device may initiate implicit channel selection after it has performed a channel scan. If a device initiates implicit channel selection, it shall include a Channel Change IE in its beacon sent in the current channel, as described in 15.9.10.

In a Channel Change IE, the device shall set the New Channel Number field to the number of the new channel. It shall set the Change Channel Count field to the remaining number of superframes before the device moves to another channel. In successive superframes, the Change Channel Count field should be decremented.

If the value set in the Change Channel Count field is zero, the device shall move to the new channel at the end of the current superframe.

On reception of the Channel Change IE, a device that also intends to change channels in a coordinated manner should include a Channel Change IE with the same field values in its beacon.

16.5 Transmission and reception of beacons

For data exchange, the device shall set up a superframe structure by transmitting beacon(s) as specified in 16.5.3.

16.5.1 Transmission and reception of Type A beacons

A Type A device that has unreserved MAS, shall not transmit any beacons with Status set to Ready in any channel unless it has (1) discovered another device with which the Type A device intends to exchange MPDUs, as described in 16.3 and (2) completed the antenna training if the newly discovered device is capable of antennna training.

Before a Type A device transmits any beacon using the antenna beam determined in 16.3, it shall scan for Type A beacons for at least one superframe, or at least two superframes if no Type A beacon frame is received, using the same antenna beam. The device shall follow the following rules based on the scanning result:

A. If the device receives one or more Type A beacon headers, but no beacon frames with a valid FCS during the scan, the device should scan for an additional superframe.

B. If the device has not selected a beacon slot to transmit Type A beacons to exchange a MPDU with another device using any of its antenna beams, it shall follow the following rules:

- If the device receives no Type A beacon frame headers during the scan, it shall create a new BP and send a Type A beacon in the first beacon slot after the signalling slots;
- otherwise, if the device receives one or more Type A beacons during the scan, it shall not create a new BP. Instead, prior to communicating with the newly discovered device, the device shall transmit a Type 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.

C. If the device already selected a beacon slot to transmit Type A beacons to exchange a MPDU with another device using a different antenna beam, it shall follow the following rules:

- If the device receives one or more Type A beacons that indicate a BPST that is not aligned with its own, it shall start the beacon relocation process as specified in 16.5.3.11;
- otherwise, if either of the following conditions (in the following sub-bullets) is true, it shall transmit an additional Type A beacon, using the same antenna beam as the one used in the scan, 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 the device receives no Type A beacon frame headers during the scan; or
    - The Type A beacon frame received indicates an aligned BPST to its own.

Figure 162 illustrates the Type A beacon transmission procedure after device discovery and beamforming using a flow chart.

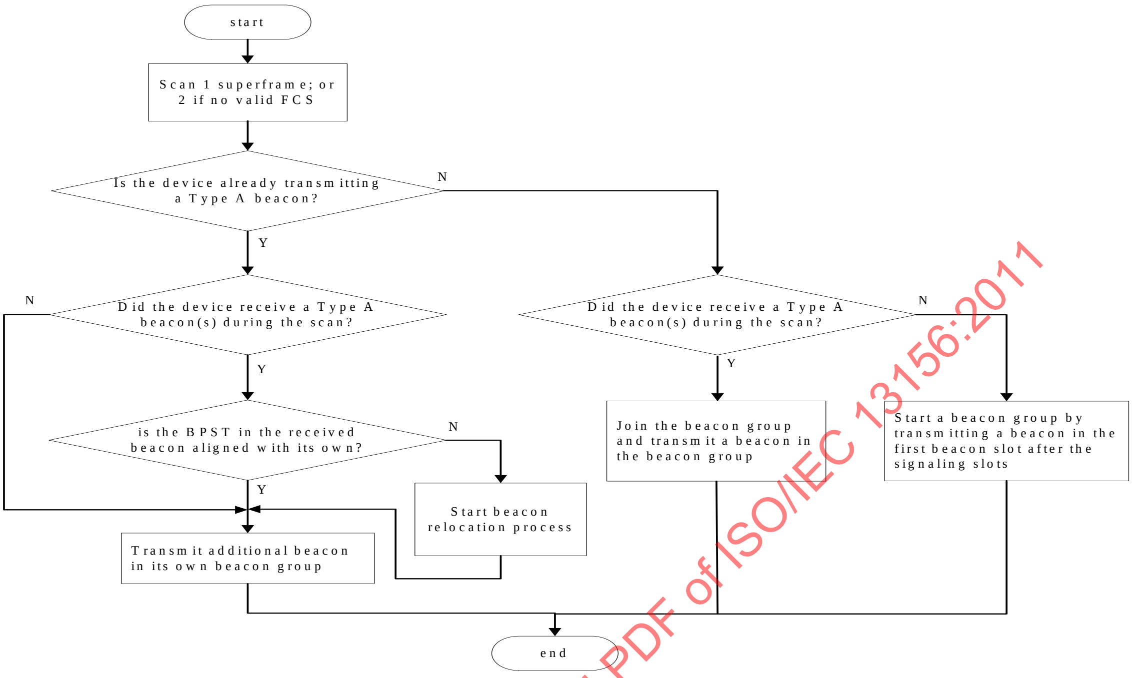

Figure 162 - Transmission of Type A beacons

When transmitting any Type A beacons, a Type A device shall follow the rules specified in 16.5.3.

With the exception of transmitting its own beacon as described in 16.5.1 and 16.5.3, a device shall not transmit frames in the current superframe during the BP length indicated in the most recent beacon received from any neighbour in the previous mMaxLostBeacons+1 superframes. A device shall not change beacon slots to a slot earlier than the highest-numbered unavailable beacon slot in the last superframe.

16.5.1.1 Discovery of additional devices

After starting the transmission of Type A beacons in a channel, a Type A shall make a DRP reservation with Reservation Type set to Absence no later than mMaxDiscoveryLatency superframes. The length of this DRP reservation shall be greater than the minimal time needed by the Type A device to transmit a Type A discovery block set. Within this reservation, the Type A device shall change to Discovery Channel and perform a Type A Discovery Block Set as described in 16.3.2.2. The discovery frame sent in the Discovery Channel after starting the transmission of Type A beacons in a channel shall include a Channel Change IE to indicate which channel the device sends its Type A beacons.

16.5.2 Transmission and reception of Type B beacon frames

A Type B device shall not transmit any beacons with Status set to Ready in any channel unless it has located another device with which the Type B device will exchange MPDUs. Before a Type B device transmits any beacon, it shall scan for Type B beacons for at least one superframe or at least two superframes if no Type B beacon frame is received. If the channel is detected as idle, it shall create a new BP and send a dual-beacon in the first beacon slot after the signalling slots. If the device receives one or more Type B beacon headers only, but no beacon frames with a valid FCS during the scan, the device should scan for an additional superframe. If the device receives only Type B beacons during the scan, it shall not create a new BP. Instead, prior to communicating with another device, the device shall transmit a dual-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. Otherwise, the Type B device shall not transmit a Type B beacon with Status set to Ready in that channel.

When transmitting its dual-beacon, a Type B device shall follow the rules specified in 16.7.1.

With the exception of transmitting its own beacon as described in this sub clause and 16.5.3., a Type B device shall not transmit any frames during the announced BP length of any of its neighbours.

A Type B device shall listen for duration of one superframe when it skips its beacon transmission as specified in 16.5.3.8. If transmission other than dual beaconing is detected using Physical sensing, the Type B device shall not move its (dual-) beacon slots.

16.5.3 Superframe

Refer to 7.2.4.4 for the timing structure of the superframe. A device shall use superframe timing structure for any frame exchange in a Data Channel. Superframe timing structure may be used in the discovery channel for MPDU exchange subject to preemption by device discovery or antenna training frame exchanges. A device shall not use superframe timing structure for MPDU exchange in Discovery Channel, if the initial scanning indicates that Discovery Channel is busy. A device shall suspend all transmission using superframe timing structure in Discovery Channel for at leas mMaxNeighbourDetectionInterval, If the device receives a beacon with Status set to preemptive. The device shall access the channel after the suspension as it is powered up for the first time as described in 16.3.

16.5.3.1 Beacon period

Each superframe starts with a BP, which has a maximum length of mMaxBPLength beacon slots. The length of each beacon slot is mBeaconSlotLength. Beacon slots in the BP are numbered in sequence, starting at zero. The first mSignalSlotCount beacon slots of a BP are referred to as signalling slots and are used to extend the BP length of neighbours. Beaconing devices shall transmit beacons in the BP and listen for neighbour's beacons in all beacon slots specified by its BP length in each superframe, except as described in 16.5.1. 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 using the PHY modes based on the type of beacons to be sent as described in 16.5. 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. Figure 163 illustrates an example of a BP observed by a device in a given superframe.

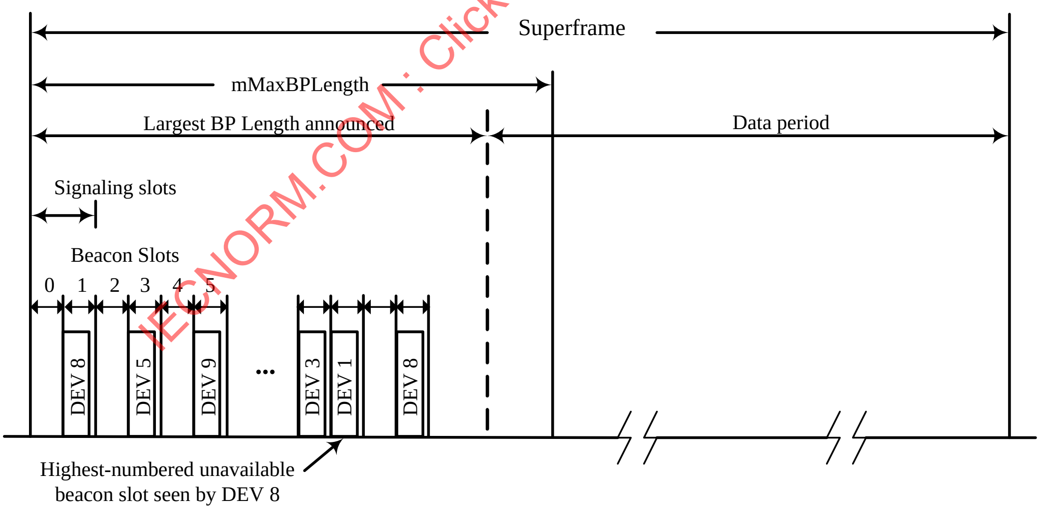

Figure 163 - Example BP structure

When a device selects an initial beacon slot after scanning for beacons as described in 16.5.1 and 16.5.2, the device shall transmit a beacon only if it selected a slot within mMaxBPLength after the BPST.

16.5.3.2 Beacon slot state

A device shall consider a beacon slot unavailable if in any of the latest mMaxLostBeacons+1 superframes:

- The beacon slot was considered to be occupied (according to Table 80); or
- The beacon slot was encoded as occupied (according to Table 80) in the BPOIE of any beacon received by the device.

A device shall consider a beacon slot available in all other cases.

16.5.3.3 BP length

A device shall consider a beacon slot to be monitored if in any of the latest mMaxLostBeacons+1 superframes:

- The device received a beacon frame in that beacon slot that is aligned to its BPST;
- The device received a beacon frame with an invalid FCS within 2×mGuardTime of that beacon slot boundary; or
- The beacon slot was encoded as occupied (according to Table 80) with a DevAddr not equal to BcstAddr in the BPOIE of any beacon received by the device.

A device shall announce its BP length in its beacon as a count of beacon slots starting from the BPST. The announced BP length shall include a) the device's own beacon slot in the current superframe, b) all monitored beacon slots in the BP of the prior superframe, and c) the beacon slot indicated in any beacon received in a signalling slot in the prior superframe.

The announced BP length shall not include more than mBPExtension beacon slots after the latest of a, b, and c, above. The announced BP length shall not exceed mMaxBPLength. Power-sensitive devices generally should not include any beacon slots after the last monitored 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 slot collision resolution or BP contraction.

16.5.3.4 Neighbours

A device shall consider another device to be a neighbour if it has received a beacon from that device within the last mMaxLostBeacons+1 superframes, and the latest beacon from the device indicated a BPST aligned with its own. If a device has not received a beacon from another device for the last mMaxLostBeacons+1 superframes, it shall not consider the device a neighbour. A device shall not consider a received beacon with the Signalling Slot bit set to one as received from a neighbour.

16.5.3.5 Beacon slot collision

If a device detects a beacon collision as described in 16.5.3.9 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.

16.5.3.6 Use of signalling slots

If the beacon slot in which a device will transmit its beacon in the current superframe is located beyond the BP length indicated in any beacon the device received from a neighbour in the previous superframe, the device shall also transmit the same beacon, except with the Signalling Slot bit set to one, in a randomly selected signalling slot, except as follows:

- A device should follow recommendations in 16.5.3.11.4, if applicable.
- If a device transmits a beacon in a signalling slot for mMaxLostBeacons+1 consecutive superframes, it shall not transmit a beacon in a signalling slot in the next mMaxLostBeacons+1 superframes, and it should not transmit a signalling slot beacon for an additional aperiodic interval that does not exceed mMaxSignalingSlotBackoff superframes

Subject to the preceding exceptions, a device also may send a beacon in a signalling slot in response to abnormal conditions, such as failure to receive a beacon from a neighbour that previously did not include the device's beacon slot in its BP Length, or failure of a neighbour to report reception of the device's beacon in its BPOIE.

A device may consider a beacon received in a signalling slot as if it were not a received beacon, except to report reception as required in 16.1.10.7 and to process the Beacon Slot Number field as required in 16.5.3.3 and 16.5.3.9.

16.5.3.7 Required reception interval

An active mode device shall listen for neighbours' beacons in the first N beacon slots in each superframe, where N is the greater of its BP Length values for the current and previous superframes, as defined in 16.5.3.3. At a minimum, the device shall listen for intervals such that it would receive a frame with a reception time within mGuardTime of the start of any of the N beacon slots.

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 HCS in a signalling slot in the previous superframe, no BP Length adjustment is required, but it shall listen for beacons for an additional mBPExtension beacon slots after its BP length indicated in the current superframe, but not more than mMaxBPLength beacon slots.

16.5.3.8 Skipping beacon transmission

An active mode device shall transmit a beacon in each superframe, except as follows: In order to detect beacon slot 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, but not any associated signalling slot beacon, at least every mMaxNeighbourDetectionInterval. When a device skips beacon transmission, it shall act as if the skipped beacon were transmitted.

16.5.3.9 Beacon slot collision detection

A device shall consider itself involved in a beacon slot 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 BcstAddr nor its own DevAddr used in the previous superframe.
- After skipping beacon transmission in the previous superframe, its beacon slot is reported as occupied in the BPOIE in any beacon it receives in the current superframe, and the corresponding DevAddr is not BcstAddr.
- When skipping beacon transmission in the current superframe, it receives a MAC header of type beacon frame in its beacon slot.
- It receives a signalling slot beacon aligned with one of its own signalling slots, with the Beacon Slot Number field set to its own beacon slot.

Certain events indicate a potential beacon slot collision. A device should consider the possibility of a beacon slot collision and take appropriate action if one or more of the following anomalous events occurs, or occurs consistently over multiple superframes:

- The device's beacon slot was reported as occupied and the corresponding DevAddr was BcstAddr in the BPOIE of a beacon it received in the current superframe, and it sent a beacon in its beacon slot in the previous superframe.
- After skipping beacon transmission in the previous superframe, its beacon slot is reported as occupied in the BPOIE in any beacon it receives in the current superframe and the corresponding DevAddr is BcstAddr.
- When skipping beacon transmission in the current superframe, it receives a PHY indication of medium activity in its beacon slot that does not result in correct reception of a frame header.

In reaction to events that indicate a potential beacon slot collision, a device should:

- consider itself involved in a beacon slot collision and change slots as required in 16.5.3.5;

- skip beacon transmission; or
- send a beacon in a signalling slot, subject to requirements in 16.5.3.6.

At a minimum, a device shall execute at least one of these recommended reactions in the next superframe if in mMaxBeaconSlotCollisionDetectionLatency consecutive superframes one or more of the anomalous events described above occurs, and the device has not executed a recommended reaction in those mMaxBeaconSlotCollisionDetectionLatency superframes.

Other events can also indicate a potential beacon slot collision. For example, if a device's beacon slot is frequently reported as unoccupied in the BPOIE of a beacon it receives, it could indicate a collision, and the device may take action as described above.

16.5.3.10 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 signalling slots 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:

- The device's beacon was movable; and

- 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

- The device did not receive a beacon from a neighbour that contained a BPOIE that encoded a beacon slot after its own as Movable (per Table 80).

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 164 shows some examples of BP contraction.

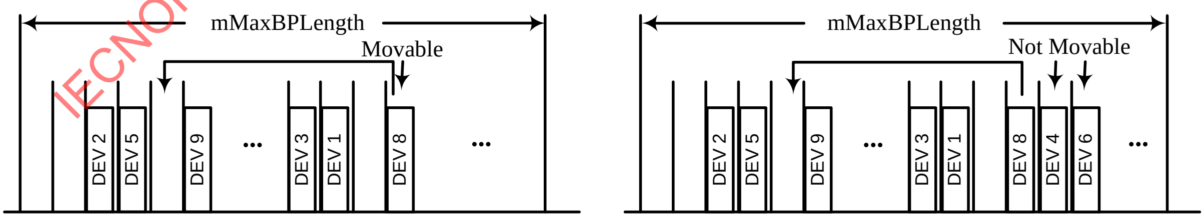

Figure 164 - Illustration for BP contraction by example device

16.5.3.11 Merger of multiple BPs

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 2xmGuardTime. 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 beacon that has the Signalling Slot bit set to one 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.

16.5.3.11.1 Overlapping BPs

If the BPST of a Type 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 follow normal BP join rules as specified in 16.5.1 and 16.5.2 to relocate its beacon to the alien BP.
3. The device shall not send further beacons in its previous BP.

16.5.3.11.2 Non-overlapping BPs

If a Type 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 DRP IE with Reservation Type set to Alien BP for the alien BP. Since the MAS boundaries may 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 DRP IE accordingly.
2. The device shall follow start the relocation process to the alien BP, according to 16.5.3.11.3, within mBPMergeWaitTime if the alien BPST falls within the first half of the superframe, or within 1.5×mBPMergeWaitTime if the alien BPST falls within the second half of the superframe, but shall not start the relocation process if a beacon received in that alien BP includes a BP Switch IE.

A Type A device that transmits or receives a beacon in its own BP that contains a DRP IE with Reservation Type set to Alien BP shall listen for beacons during the MASs indicated in the reservation.

A Type B shall not relocate to an alien BP unless the received alien beacon is of Type B with Status set to Master. When relocating its beacon to an alien, the Type B shall use the same procedure as defined for Type A device above.

16.5.3.11.3 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, minus mSignalSlotCount; or

b. Zero to indicate the device will join the alien BP using normal BP join rules as specified in 16.5.1 and 16.5.2.

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×mGuardTime, the device shall set the BP Move Countdown field, the BPST Offset field, and the Bea con 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×mGuardTime and smaller than the device's BPST Offset field + 2×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 would 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 neighbor'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 neighbor's BP Switch IE.

C2. The device shall set the BPST Offset field to the value of the same field contained in the neighbor'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 neighbor's BP Switch IE, based on the Beacon Slot Number field and BPOIE, minus mSignalSlotCount; or the Beacon Slot Offset field contained in the neighbor's beacon; or

b. Zero, to indicate the device will join the alien BP using normal BP join rules as specified in 16.5.1 and 16.5.2.

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 neighbor's beacon is larger than the device's BPST Offset field + 2×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 neighbor's beacon and the device's BPST Offset field is smaller than 2×mGuardTime, the device shall modify its BP Switch IE as follows:

a. If the Beacon Slot Offset field contained in the neighbor'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 neighbor'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 BPMoveCountdown field contained in the neighbor's beacon is less than the device's BPMoveCountdown 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 beacon of the previous superframe 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 beacon of the previous superframe.

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.

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. In the superframe after sending a BP Switch IE with BPST Offset set to 65535 and BP Move Countdown set to zero, the device shall remove the BP Switch IE from its beacon, but shall not change its beacon slot and shall continue to synchronize to current neighbors.

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 DRP 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 to relocate its beacon to the alien BP.

16.5.3.11.4 Use of signalling slots after BP merge

After changing its BPST, regardless of whether due to overlapping or non-overlapping BPs, if a device is required to send a beacon in a signalling slot according to 16.5.3.6, it should wait for a random number of superframes before sending a beacon in a signalling slot. The device should 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.

16.5.3.11.5 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.

16.6 Distributed reservation protocol (DRP)

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

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.

16.6.1 Reservation type

Each DRP IE, whether included in a beacon or separately transmitted during explicit DRP negotiation, specifies a reservation type. A device shall decode all DRP 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 98.

Table 98 - Reservation types

Reservation Type	Description	Reference
Alien BP	Prevents transmission during MASs occupied by an alien BP.	16.6.1.1
Hard	Provides exclusive access to the medium for the reservation owner and target.	16.6.1.2
DCA	Reserves time for DCA. No device has preferential access	16.6.1.3
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.	16.6.1.4
Absence	Indicates the medium time where the reservation owner returns to Discovery Channel to find additional device	16.5.1.1

16.6.1.1 Alien BP reservations

A device shall announce an alien BP reservation to protect alien BPs as described in 16.5.3.11.2. A device shall not transmit frames during an alien BP reservation except possibly to send a beacon in the alien BP.

16.6.1.2 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.

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. The reservation owner may transmit any command or control frame in a hard reservation.

16.6.1.3 DCA reservations

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

16.6.1.4 Private reservations

The channel access method and frame exchange sequences used during a private reservation are out of the scope of this International 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.

16.6.2 Reservation waveform

Each DRP IE, whether included in a beacon or separately transmitted during explicit DRP negotiation, specifies the waveform that shall be used for all the frame transactions in the reservation.

16.6.3 Medium access

A device shall not transmit a unicast frame within a reserved MAS in a hard, or private reservation in the current superframe unless:

— it included a DRP IE with the Reservation Status bit set to ONE that included that MAS in its beacon in the previous superframe;

— the destination device is a neighbor; and

— the most-recently received beacon from the destination device included a DRP IE with the Reservation Status bit set to ONE that included that MAS.

A device shall not transmit a multicast frame within a reserved MAS in a hard, or private reservation in the current superframe unless:

— it included a DRP IE with the Reservation Status bit set to ONE that included that MAS in its beacon in the previous superframe.

16.6.4 DRP availability IE

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

The combination of information from DRP Availability IEs and DRP IEs allows an owner to determine an appropriate time for a new DRP reservation.

A device shall mark a MAS unavailable if the device includes it in a DRP IE with the Reservation Status bit set to 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.

If a neighbour includes a MAS in a DRP IE with a target other than the device itself, whether the Reservation Status bit is zero or one, the device shall mark the MAS unavailable if one of the following cases is true:

a. The device's transmission in the MAS would interfere with the neighbour's reception, which entails: (1) the antenna beam to be used by the device for the new DRP reservation is the same beam that is used to transmit a beacon perceived by the neighbour, (2) the Beam Identifier in the neighbour's DRP IE is also in the neighbour's BPOIE's Receiving Beams field corresponding to the same (device's) beacon, (3) the neighbour is the target of its reservation or the neighbour's reservation is NOT unidirectional, and (4) the new DRP request is not unidirectional either

b. The neighbour's transmission in the MAS would interfere with the device's reception, which entails (1) the Beam Identifier in the neighbour's DRP IE is the same as in the neighbour's beacon perceived by the device, (2) the Beam Identifier to be used for the new DRP reservation by the device also appears in the device's BPOIE's Receiving Beams field corresponding to the same (neighbour's) beacon, and (3) the neighbour is the owner of its reservation or the neighbour's reservation is NOT unidirectional.

16.6.5 DRP reservation negotiation

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

A device shall not negotiate for MASs that are marked as unavailable, unless the MASs are referenced only in a DRP IE with Reason Code set to Denied.

A Type B device shall not negotiate for a new reservation or more MASs for an existing reservations if non-Type B transmission is detected during the device's beacon period.

A device shall announce in the MAC Capabilities IE in its beacon whether it is capable of explicit DRP negotiation. A device shall not initiate an explicit DRP 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 DRP IE with Reservation Status set to ONE on detection of an alien BP, as specified in 16.6.1.1.

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

16.6.5.1 Negotiation

When negotiating a reservation, the reservation owner shall set the Target/Owner DevAddr field of the DRP 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 DRP 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. The device shall set the BeamIdentifier as the BeamIdentifier in the beacon transmitted to the reservation target.

A reservation owner shall not transmit unicast frames within reserved MASs in a hard, DCA or private reservation unless it and the recipient included DRP 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 Target/Owner DevAddr field of the DRP IE to the DevAddr of the reservation owner. The device shall set the BeamIdentifier as the BeamIdentifier in the beacon transmitted to 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 DRP 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 cannot be granted due to some interference, the reservation target shall set the Reason Code to Interfered. 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 DRP Allocation fields to match those in the request. For a multicast reservation, it shall set the DRP Allocation fields to match the request, or to include a subset of the MASs included in the request.

16.6.5.2 Explicit negotiation

To start explicit DRP negotiation, the reservation owner shall send a DRP Reservation Request command frame to the target device, as defined in 15.6.1.

On reception of a DRP Reservation Request command the reservation target shall send a DRP Reservation Response command, as defined in 15.6.2, to the reservation owner. The fields in the DRP IE shall be set according to 16.6.5.1. If the reservation cannot be granted due to a conflict with its own or its neighbours' reservations or interference, the reservation target shall include a DRP Availability IE in the DRP Reservation Response command frame.

In a DRP Reservation Response command frame for a multicast reservation, the reservation target shall include a DRP Availability IE for a Reason Code other than Denied. Final multicast reservations are established implicitly, as described in 16.6.5.3.

16.6.5.3 Implicit negotiation

Implicit negotiation is carried out by transmitting DRP IE(s) in beacon frames. A device that supports the DRP shall parse all beacons received from neighbours for DRP IE(s) whose Target/Owner DevAddr field matches 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 DRP IE(s) that are new with respect to DRP IE(s) included in the most recently received beacon from the same device as a DRP reservation request or a DRP reservation response.

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

On reception of a unicast DRP reservation request in a beacon, the reservation target shall include a DRP reservation response in its beacon no later than the next superframe, with fields set as described in 16.6.5.1. If the Reason Code indicates Conflict or Interfered, the reservation target shall include a DRP Availability IE in its beacon.

As long as the reservation owner includes a unicast DRP reservation request in its beacon, the reservation target shall continue to include the DRP 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 DRP reservation request, a reservation target shall include a reservation response DRP IE in its beacon no later than the next superframe if it is a member of the targeted multicast group. The fields in the DRP IE shall be set according to . If the Reservation Status bit in the response is ZERO, the reservation target shall include a DRP 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 DRP 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 DRP 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 DRP Availability IE in the beacon.

16.6.5.4 Negotiation conclusion

To conclude negotiation for a unicast reservation, the reservation owner shall set Reservation Status to ONE in the DRP IE in its beacon after receiving a beacon from the reservation target that contains a corresponding DRP IE with Reservation Status set to ONE. To conclude negotiation for a multicast reservation, the reservation owner may set Reservation Status to ONE in a DRP IE in its beacon in the next superframe after transmitting the same DRP 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 16.6.7.

16.6.6 DRP reservation announcements

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

16.6.7 Resolution of DRP reservation conflicts

Devices engaged in independent DRP negotiation could attempt to reserve the same MAS, or due to mobility, devices could have reserved the same MAS. A device is considered to conflict with a neighbour's DRP reservation if a MAS included in the device reservation is also included in the neighbour's reservation and one of the following cases is true:

a. The neighbour's transmission in the MAS interferes with the device's reception, which entails: (1)the Beam Identifier in the neighbour's DRP IE is the same in the neighbour's beacon perceived by the device, (2) the Beam Identifier in the device's DRP IE that includes the MAS also appears in the device's BPOIE's Receiving Beams field corresponding to the same (neighbour's) beacon, (3) the device is the target of its reservation or the device's reservation is not unidirectional, and (4) the neighbour is the owner of its reservation or the neighbour's DRP reservation is NOT unidirectional.

b. The device's transmission in the MAS interferes with the neighbour's reception, which entails: (1) the Beam Identifier in the device's DRP IE is the same in thedevice's beacon perceived by the neighbour, (2) the Beam Identifier in the neighbour's DRP IE also appears in the neighbour's BPOIE's Receiving Beams field corresponding to the same (device's) beacon, (3) the neighbour is the target of its reservation or the neighbour's reservation is not unidirectional, and (4) the device is the owner of its reservation or the device's reservation is NOT unidirectional.

A device might detect a conflict during a DRP negotiation or after a reservation has been established. Reservations of type Alien BP never conflict with other reservations of type Alien BP.

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

1. If the device's reservation is of type Alien BP, the device shall maintain the reservation.
2. 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 DRP IE to Conflict.
3. If the device is a Type B device, the device shall maintain the reservation
4. If the neighbour is a Type B device, 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 DRP IE to Conflict.
5. If the device's DRP IE has the Reservation Status bit set to ZERO and the neighbour's DRP 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 DRP IE to Conflict.
6. If the device's DRP IE has the Reservation Status bit set to ONE and the neighbour's DRP IE has the Reservation Status bit set to ZERO, the device may maintain the reservation.
7. If the device's DRP IE and neighbour's DRP 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.
   a. The device's DRP IE and neighbour's DRP 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

   b. The device's DRP IE and neighbour's DRP 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.
8. If the device's DRP IE and neighbour's DRP 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 DRP IE to Conflict.
   a. The device's DRP IE and neighbour's DRP 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

   b. The device's DRP IE and neighbour's DRP 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.
9. If the device's DRP IE and neighbour's DRP 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 DRP IE to Conflict.
   a. The device's DRP IE and neighbour's DRP 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

   b. The device's DRP IE and neighbour's DRP 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 mDRPBackoffWinMin. When the backoff algorithm is invoked, the device shall select a random number N uniformly from [0, 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 mDRPBackoffWinMax. After a negotiation completes, the device shall generate a new backoff N. If a device does not request any MASs for 4×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 DRP IE in its beacon, it shall include a DRP Availability IE in the same beacon.

16.6.8 BPST realignment and existing DRP reservations

A device that realigns its BPST as described in 16.5.3.11 may assert new DRP reservations with Reservation Status bits set to ONE in the new beacon so long as they are equivalent to its old DRP reservations with the Reservation Status bit set to ONE in the prior BP. For this purpose, two DRP reservations are equivalent if their corresponding Target/Owner 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 DRP reservations with MASs that conflict with any BP it announced or detected. The device shall not assert DRP 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 16.6.7.

16.6.9 Modification and termination of existing DRP reservations

A reservation owner may reserve additional MASs for a stream by negotiating an addition to the reservation using a DRP IE with the same Target/Owner DevAddr, Stream Index, and Reservation Type. Once negotiation has completed successfully, the reservation owner should combine the DRP IEs. When combining DRP IEs, the reservation owner shall set the Reason Code to Modified until a DRP 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 DRP IE. If a reservation owner removes some MASs from an established reservation, it shall set the Reason Code in its DRP IE to Modified until the reservation target has changed its DRP IE to match.

A reservation target may remove MASs from an established reservation without changing the Reservation Status bit in the DRP IE due to a conflict, as described in 16.6.7 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 in the current or following superframe.

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

If a reservation owner changes or removes a DRP IE, the reservation targets shall update or remove the corresponding DRP 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. The reservation owner shall terminate the corresponding reservation or set the corresponding Reservation Status bit to ZERO in the current or following superframe.

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 DRP IE(s) from its beacon.

16.6.10 Retransmit procedures in DRP reservations

In a hard DRP 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.

A device shall not retransmit a frame earlier than pSIFS after the end of an expected acknowledgement, 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.

16.7 Coexistence and interoperability

This clause specifies mechanisms to prevent potential interference among devices of different types and to facilitate transmission between devices of different types.

16.7.1 Coexistence

A Type B device shall send dual-beacon to announce its DRP reservation.

When it is required to transmit a dual-beacon in this specification, a Type B device shall select a beacon slot according to the rules specified in 16.5.3 and 16.7.2. The Type B device shall transmit a Type B beacon with Status set to Ready and a Type A beacon of the same payload with Status set to Dual separated by MIFS in the selected beacon slot. The timing of the dual beaconing is specified in Figure 165.

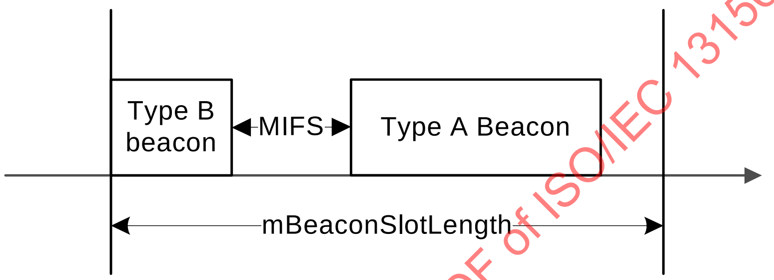

Figure 165 - Timing of a dual beacon

16.7.2 Interoperability

Interoperability between different types of devices is accomplished using Master-Slave operation. The MSPr shall select a channel to exchange MSDUs using the explicit channel selection process before the master device starts the transmission of Type A beacons with the Status set to Ready in a channel. If the master device has already started or joined a beacon group by transmitting Type A beacons in a selected channel, the slave device shall switch to the selected channel as indicated in the Channel Change IE in the Type A discovery frame sent by the master device. A Type A device shall discover additional slave devices as specified in 16.7.2.1, if there is sufficient medium time available for frame exchange of discovery block set (as specified in 16.3.2.2).

16.7.2.1 Discovery of slave devices in a data channel

After starting the transmission of Type A beacons with Status set to Ready in a channel, a Type A device shall make a DRP reservation of type Private to transmit Type A Discovery Blocks if there are MASs available. The Type A device shall scan for responses to the transmitted Type A Discovery blocks as specified in 16.3.2.2. If a Type B discovey frame with Status set to Response is received correctly in a channel, the Type A device shall start a Type A-B MSPr following the procedure specified in 16.7.2.2.

A Type A device may adjust the reservation used for transmission of a DBS to release the reserved MASs for data transmission. The Type A device may terminate an existing reservation used for transmission of a DBS, if all MASs are needed for data transmission.

16.7.2.2 Type A-B MSPr

After starting the transmission of Type A beacons in a channel, a Type A master device shall make a DRP reservation of type Private to transmit B-Poll frames in every superframe. The Type A device shall follow the same rules specified in 16.5.3 to select a beacon slot for its Type B slave device to send its dual-beacon and indicate the timing of the selected beacon slot in the B-Poll frame it sends. The Type A master shall indicate the selected slot as Occupied in the BPOIE in the Type A beacon it sends before sending the B-Poll frame in the current superframe.

After a Type B slave device switches to a selected channel to exchange MPDUs with its master, the Type B device shall not transmit any frames until it receives a B-Poll frame from the Type A master device. The Type B device shall send a B-Poll Response frame a SIFS after the reception of a B-Poll frame. In addition, the Type B device shall transmit a dual-beacon at the time indicated in the received B-Poll frame.

16.8 Synchronization of devices

Each beaconing 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 neighbours, through the reception of neighbours' beacons. Since a Type B device cannot decode beacons transmitted by Type A devices, the Type B device cannot determine timing information needed for synchronization with Type A devices. Therefore, a Type A device shall always synchronize with its slowest Type B neighbour, if it has any. The Type A device determines the difference between the actual reception time and the expected reception time of the mode-A0 beacon transmitted by a Type B device. 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. Otherwise, the neighbour is faster. A Type A device that synchronizes with a Type B device is a referred to as a forced synchronization device. A forced synchronization device shall set the Forced Sync field of each transmitted beacon to ONE.

If a Type A device is not a forced synchronization device and receives Type A beacons only from Type A neighbours, the device shall synchronize with a neighbour that is a forced synchronization. If none of the Type A device's neighbours is forced synchronization device, the Type A device shall synchronize with its slowest neighbour. A Type B device shall synchronize with its slowest Type B neighbour.

To maintain superframe synchronization with a slower neighbour, the device shall delay its BPST by the difference. To maintain superframe synchronization with a faster neighbour, the device shall advance its BPST by the difference. Any adjustment of the BPST shall be limited to a maximum 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.

A device shall not use a beacon with the Signalling Slot bit set 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 may be applied for up to mMaxLostBeacons consecutive superframes. Beacon transmit time and measured beacon receive time shall be accurate to at least mClockResolution.

A Type B device that is operating in the slave mode shall derive all times for communication with its master from the B-Poll frames transmitted by the master device.

16.8.1 Clock accuracy

MAC sublayers 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.

16.8.2 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.

16.8.3 Guard times

Due to inaccuracy in the superframe synchronization and drift between synchronization events, the MAS start times of 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 DRP reservation blocks and BPs.

Figure 166 is an illustration of how a device uses the guard interval to maintain a SIFS interval between transmissions in adjacent reservation blocks. 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:

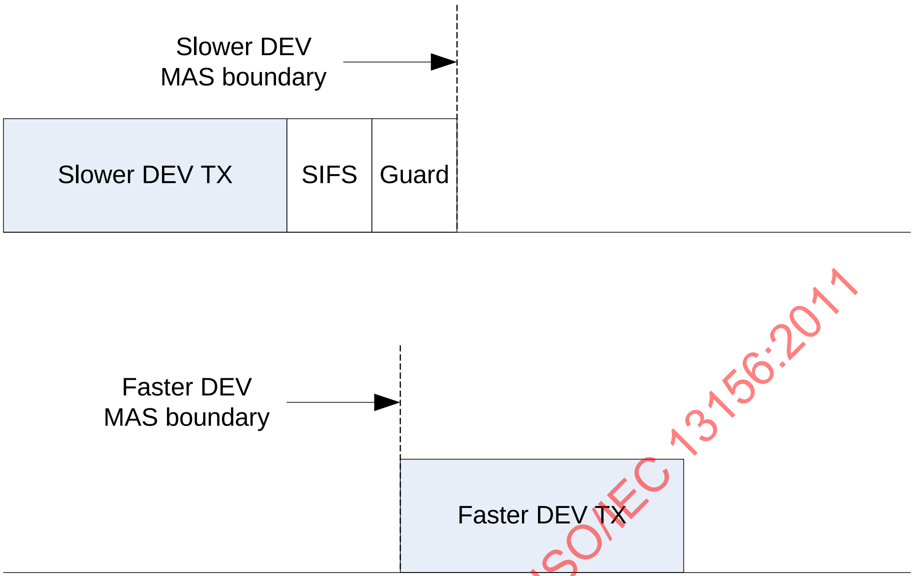

Figure 166 - Guard time

$$\text{mGuardTime} = \text{MaxSynchronizationError} + \text{MaxDrift} \quad (82)$$

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 16.6. 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:

$$\text{MaxDrift} = 2 \times \text{mClockAccuracy (ppm)} \times 1\text{E-6} \times \text{SynchronizationInterval} \quad (83)$$

where SynchronizationInterval = (mMaxLostBeacons+1) × 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.

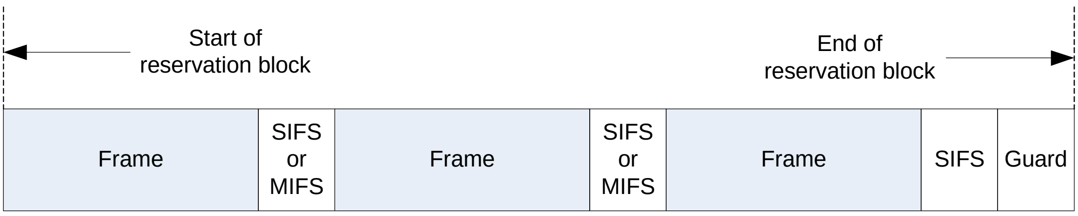

Figure 167 - SIFS and guard time in a DRP reservation block - No-ACK

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 168. 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 DRP reservation block, the start of a BP, or the start of a MAS in which the device announced it would be available.

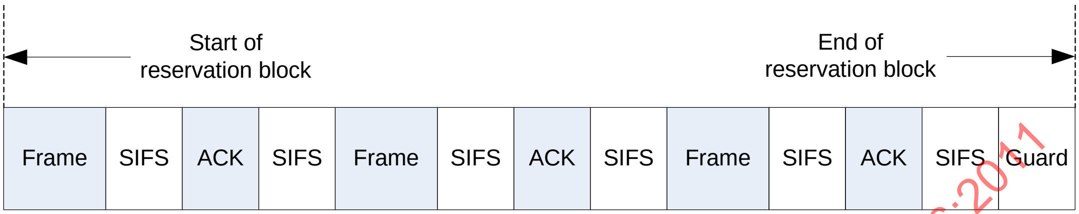

Figure 168 - SIFS and guard time in a DRP reservation block - Imm-ACK

16.9 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.

16.10 Aggregation

A transmitter may aggregate multiple MSDUs with identical Delivery ID into a single data frame. A device shall aggregate no more than mAggregationLimit MSDUs into an aggregated data frame. The aggregated data frame format (15.2.2) shall be used in frame transactions that employ aggregation.

A source device initiates the use of an aggregation 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 aggregation mechanism, it indicates the maximum number and size of the frames it can buffer. The source device includes a number of MSDUs in the aggregated frame, limited by the announced buffer size and maximum number of frames.

On receipt of such an aggregated frame that requires an acknowledgement, the recipient device returns a B-ACK frame giving feedback on the MSDUs received and indicating the buffer space available for the next aggregated frame.

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

A source device may transmit an aggregated frame to any potential recipient device advertising aggregation capability in its MAC Capabilities IE. A source device shall initiate use of the aggregation mechanism by transmitting an aggregation request frame from the same stream or of the same user priority. An aggregation request frame is a data frame with Aggregation Request set to ONE and Ack Policy set to B-ACK Request. After transmitting the aggregation request frame, the source device shall follow the rules of operation as described below.

When receiving an aggregation request frame from a source device for a specific stream or user priority, the recipient device shall respond as follows:

- To acknowledge receipt of the frame but reject the request for starting transmission of a new aggregation frame, the recipient device shall respond with a B-ACK frame with no frame payload.
- To accept the request for starting transmission of a new aggregation frame, 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 aggregation frame. The recipient shall acknowledge the received frame by indicating its reception in the acknowledgement window.

After transmitting the aggregation request frame , 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 operation, the source device shall retransmit the request frame using medium access rules as described in 16.6.
- 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 aggregation mechanism to be rejected.
- 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 requesting aggregation operation, the source device shall do one of the following: (1) if the frame was indicated as received correctly, the source shall transmit a new frame or a frame with zero payload with Aggregation Request bit set to ONE and ACK Policy set to B-ACK Request. (2) Otherwise, the source shall retransmit the same aggregation request frame.
- 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.

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

- MSDUs being held for retransmission with a sequence number earlier than the one indicated by the Sequence Control field were not received correctly from the reception of the last aggregated sequence, but shall not be retransmitted.
- MSDUs 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 MSDUs being held for retransmission should be retransmitted in the next aggregated frame, 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 new aggregated frame. The total number of MSDUs included in the new aggregated frame shall 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. Within an aggregated frame, the MSDUs shall be ordered by increasing sequence and fragment numbers. Due to retransmissions, this ordering might not hold from one aggregated frame to the next and MSDUs transmitted in an aggregated frame might not have consecutive sequence and fragment numbers.

When the recipient device receives a frame with Frame Type set to Aggregation data frame and 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

HCS 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. If any subframe in the received aggregated frame has an invalid FCS and has not timed out, the recipient device 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 the buffer resources to be freed.

16.11 Channel bonding

A device may use channel bonding as described in 10.2.2.6 to exchange MPDU or MSDUs with devices that also support channel bonding. The device shall access the bounded channel as accessing an unbounded channel defined in this clause, except for transmission of beacon frames.

When using channel bonding, a device shall transmit beacons in a designated beaconing channel. For each bonded channel, the designated beaconing channel shall be selected as given in Table 99 only as defined in 16.5. The device shall include a CBOIE in its beacons that describes the proposed bonded channels as defined in 15.9.9.

Table 99 - Designated beaconing channels for bonded channels

Bonded Channel BAND_ID	Designated beaconing channel BAND_ID
5	2
6	3
7	3
8	3
9	3
10	3

Before accessing any of the channels, 1, 2 or 4, a device shall scan for beacons in all designated beaconing channels for which their corresponding bonded channels overlap that channel along the direction determined by the antenna training in 16.18. If a CBOIE with Status set to One is included in the received beacons during the scan, the device shall consider the channels indicated in the Bonded Channel Bitmap field of the CBOIE busy and shall not access those channels, except the designated beaconing channel.

A device shall not initiate or accept channel bonding unless the channels to be bonded, excluding the designated beaconing channel, are free of any transmission along the direction determined by the antenna training.

A device shall not initiate or accept a new channel bonding unless the new bonded channel is identical to the existing bonded channel initiated or accepted by the device. Two bonded channel are considered identical if they are composed of the same channels.

To initiate channel bonding, a source device shall include a CBOIE that describes the proposed bonded channel with Status set to Zero in its beacon. The device should continue to include the CBOIE for at least mMaxLostBeacons+1 consecutive superframes or until a response is received.

A recipient device shall include a CBOIE in its beacon no later than the next superframe after receiving a CBOIE with Target DevAddr field matches the device's DevAddr. If the channels indicated by the Bonded Channel Bitmap field of the received CBOIE are available for channel bonding as describe above, the recipient device may set the Status to One in the CBOIE. Otherwise, the device shall set the

Reason Code to an appropriate value as described in Table 91. A source device shall not change the Status in the CBOIE from Zero to One until the recipient device has set the Status to One in the CBOIE in its beacons. The recipient device shall include the CBOIE in its beacons until the source device removes CBOIE from its beacons regardless of the value of Bonding status field in its CBOIE.

16.12 Acknowledgement policies

This Clause defines three acknowledgement policies: 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.

A device shall not set ACK Policy of any aggregated MAC frames to Imm-ACK, or B-ACK.

16.12.1 No-ACK

A frame with ACK policy set to No-ACK, as defined in 15.2.3.3, shall not be acknowledged by the recipient. The transmitting device MAC sublayer 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.

16.12.2 Immediate ACK

On reception of a frame with ACK Policy set to Imm-ACK, a device shall respond with an Imm-ACK frame, as defined in 15.5.1**,** transmitted pSIFS after the end of the received frame.

16.12.3 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.

16.12.3.1 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 16.12.3.2.

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.

16.12.3.2 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 16.6.
- 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 shall 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 DRP 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 HCS 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.

16.12.3.3 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.

16.13 Probe

The Probe IE and Application-specific Probe IE may be 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 should 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 requester, that includes no Requested Element IEs to reject the probe; or
- One or more requested IEs.

16.14 Multi-rate support

In device discovery (16.3) or antenna training, device shall transmit beacons or control frames using one of discovery modes as specified in 10.2.5.

In frame exchange other than device discovery or antenna training, a device shall transmit beacons using one of common PHY modes according to its device type, hence at the rate of the corresponding mode. More specifically, in a BP, a Type A device shall transmit beacon using mode-A0; while a Type B device shall transmit its beacons using both mode-A0 and mode-B0 (dual beaconing).

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 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 using explicit or implicit transmit rate control (TRC) mechanisms. For explicit TRC, the recipient sends the TPRC command frame to the source device. In addition, a source device may send a Link Feedback Request command frame to request a recipient device provide feedback on the quality of the link. The recipient sends a Link Feedback

Response command frame to the source device SIFS after the reception of a Link Feedback Request command frame from the source device. For implicit TRC, the recipient includes a Link Feedback IE in its beacon. The data rate in the Link Feedback IE or the received link feedback command frame should be interpreted as the maximum data rate that the source device should use for this particular link, for an acceptable frame error rate. The source device should either follow the recommendation, or determine a data rate based on the received feedback on the quality of the link, which should not exceed the data rate recommended in the received Link Feedback command frame. The method to determine the optimal data rate in the recipient is beyond the scope of this International Standard.

16.15 Transmit power control (TPC)

A device shall not transmit frames at a higher transmit power level than that used for its most-recently transmitted beacon.

A recipient device may recommend a transmit power level change to be used by a source device using explicit or implicit TPC mechanisms. For explicit TPC, the recipient sends a TPRC command frame to the source device to recommend transmit power change. In addition, a source device may send a Link Feedback Request command frame to request a recipient device provide feedback on the quality of the link. The recipient sends a Link Feedback Response command frame to the source device SIFS after the reception of a Link Feedback Request command frame from the source device. For implicit TPC, the recipient includes a Link Feedback IE in its beacon to recommend change in transmit power. A device that receives a Link Feedback IE or a Link Feedback command frame should either follow the recommendation of change, or should determine its transmit power based on the feedback on the quality of the link in the received Link Feedback Response. The method to determine transmit power is out of the scope of this International 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.

16.16 Power management mechanisms

This Clause specifies the power management mechanism of a device during its various operational modes. Part of this clause should be considered as informative.

16.16.1 Power management modes

A device may be in one of two power management modes during a superframe.

16.16.1.1 Active Mode

In the active mode the device sends and receives beacon frames in the current superframe. A device in active mode may switch between two power states during a superframe:

(a) Awake: device is able to transmit or receive frames.

(b) Sleep: device does not transmit or receive frames. Most of its Tx and Rx units are turned-off to save power except the timing and control units that should recover the system in relatively short time. The peer device(s) are not informed of the device transition to Sleep state.

16.16.1.2 Hibernation Mode

In the hibernation mode the device does not send or receive beacons or other frames in the current superframe; however, a device shall announce in previous superframe(s) that it plans to enter the hibernation mode.

16.16.2 Power state transitions at active mode

16.16.2.1 Power state transition from awake to sleep state

During the discovery process (device discovery, Antenna training or association), a device shall not switch to Sleep state.

A device may switch to Sleep state in the following cases:

1. After the end of BP till beginning of its first DRP reservation block in the current superframe.

2. Between two DRP reservation blocks in the current superframe

3. The device completes its DRP session in the current superframes.

A device may transmit a frame with Access field set to ONE and More Frame field set to ZERO to indicate the last frame in current DRP. And a device may transmit a frame with Access field set to ZERO and More Frame field set to ZERO to inicate that the device has no more frames to transmit in current superframe.

16.16.2.2 Power state transition from sleep to awake state

A device shall transient from Sleep to Awake state according to the following rules:

(1) A device shall be in the Awake state mGuardTime prior to its BPST in every superframe to participate in the transmission and reception of beacons.

(2) If a device has any pending frame to be transmitted in DRP reservations in the current superframe, it shall be in Awake state mGuardTime prior the start of each relevant DRP reservation block to start its transmission.

(3) If a device expects to receive transmissions from other devices in a DRP 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.

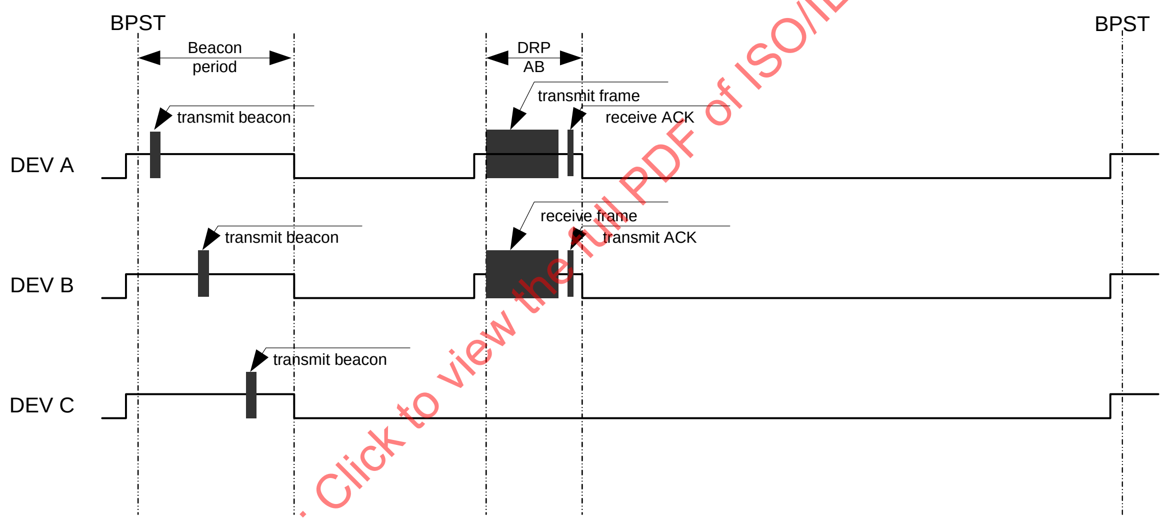

Figure 169 - Power state transition for devices in active mode

Figure 169 illustrates the power state transition for devices in active mode.

(a) DEV A is a device that has pending frames to be transmitted in a DRP reservation block in the current superframe.

(b) DEV B is a device that is expecting to receive a planned transmission from DEV A in a DRP reservation block in the current superframe.

(e) DEV C is a device that does not have any traffic pending in its transmission queues, and is not expecting any planned transmission from other devices.

16.16.3 Hibernation mode operation

A device using hibernation mode shall transient to and from hibernation mode according to the following rules:

(1) The device shall signal its intent to go into hibernation mode by including a Hibernation Mode IE in its beacon, as defined in 15.9.15. The Hibernation Duration field in the Hibernation Mode IE shall contain a non-zero value that specifies the duration of the hibernation period.

(2) The 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.

(3) During hibernation mode, the device shall not send a beacon or other frame. The device should terminate all established DRP reservations before entering hibernation.

(4) A device may leave hibernation mode prior to the end of its announced hibernation period by sending its beacon.

(5) 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.

(6) 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:

(a) If an active mode device receives a neighbour's beacon that includes a Hibernation Mode IE, the device shall consider all DRP 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.

(b) 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.

(a) 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.

(b) 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.

16.16.4 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. A device that indicates such capability should include a Hibernation Anchor IE 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 should remove indication of the capability from its MAC Capabilities IE.

Devices, such as those that were recently off or in hibernation mode, may 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 should 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 should include that device in the Hibernation Anchor IE.

16.17 ASIE operation

Zero or more ASIEs may be included in each beacon. ASIEs may appear within the IE area in a beacon as defined in 16.1.10. Unrecognized ASIEs shall be ignored. The format of the ASIE payload is defined by the owner of the value in the ASIE Specifier ID field and is outside the scope of this document.

16.18 Antenna training and tracking

After performing device discovery as specified in 16.3, a device shall perform antenna training (16.18.3) with the device which it instends to exchange MPDUs with to obtain antenna weight settings that increase its transmitting and receiving antenna gains if the device is capable of beamforming. After moving to a data channel through channel selection procedure specified in 16.4, the devices may update their antenna weight settings. This is called antenna tracking (16.18.4).

A Type A device shall support the training/tracking of its own antenna and of the antenna of another Type A device. A Type B device may support the training/tracking of the receiving and transmitting antenna of a Type A device.

A device may have multiple phased arrays (PAA) anntennas to cover wider angle with each array covering one sector. The number of sectors each device may have shall not exceed 4.

16.18.1 Announcement of antenna capabilities

A device shall announce its antenna capabilities in an ACIE that shall be included in all discovery frames when performing device discovery (16.3)

16.18.2 Antenna training/tracking configuration

A device shall use a RTT/CTT handshake to initiate antenna training with another device. The ATIE included as the payload of RTT/CTT and subsequent TRN frames configures the antenna training process that follows the RTT/CTT handshake. The initiating device is referred to as the I device. The helping device is refered to as the respondor device or the R device. Single RTT/CTT exchange configures antenna training for the I device. The R device may train its own antenna while helping training I's antenna. In the case when the R device also desires to train its antenna, R shall include an additional ATIE in the CTT frame or subsequent TRN frame to configure training parameters for the R device.

The I device shall initiate the training negociation by sending an RTT frame to the R device. The I device shall set the Status field of the ATIE included in the RTT frame to Request and shall set the Continuition bit to ZERO when initiating the RTT/CTT handshake for the first time.

SIFS after the reception of the RTT frame, the R device shall respond with a CTT frame with an ATIE. If the R device accepts the configuration specified in the received ATIE, it shall include the same ATIE except for the Status set to Accepted. Moreover, if it desires to train its own antenna, the R device shall include an additional ATIE with training configuratin for its own antenna and set its Status to Request.

SIFS after the reception of an CTT frame, the I device shall respond with a TRN frame. If the CTT frame contains an addtional ATIE that requests to train the R device's antenna, the TRN frame shall also include the same ATIE that configures the R device's anntenna training and indicate the acceptance or rejection of the R device's training configuration in the Status field.

The agreed configurations for the I and R devices shall be used in antenna tracking as well after the devices move to a data channel.

The Request Training fields in the ATIEs of the I and R devices carried in RTT/CTT frames determine the inclusion of ATS field in the subsequent frame exchange in antenna training and tracking processes. Devices shall transmit training symbols in subsequent frame exchange in antenna training as well as tracking according to Table 48.

16.18.3 Iterative antenna training

In general, the optimal receiving antenna weight settings for a device depend on the transmitting antenna weight settings of the peer device. Hence, devices transmits TRN frames back and forth to train their transmitting and/or receiving antenna weights. And the antenna training process consequently consists of one or more consecutive iterations. Each iteration of training involves transmission of a TRN frame, a responding TRN frame from training peer device SIFS after, and optionally another TRN frame that contains feedback if the training peer device requests feedback.

A device keeps track of iterations using the Iteration field of the ATTIE contained in TRN frames it transmits and receives. Upon reception of a TRN frame from the training peer device with a matching iteration value as that in the previous TRN frame it sends, the device shall increment the Iteration field of the ATTCIE contained in the next TRN frame it transmits to the training peer device. The maximum number of iterations shall not exceed 4. As antenna training progresses, in each iteration the training device should use a discovery mode with a smaller number of repetitions than that in previous iteration. The device shall indicate the desired number of repetitions by specifying the corresponding discovery mode in the Discovery Mode field of the ATTCIE contained in the TRN frame being sent to its training peer device. In addition, in the RX Training Size field of the ATTCIE, the device shall also specify the number of training symbols to be used in the next iteration for its receiving antenna training. In the next iteration, the training peer device shall use the specified discovery mode to transmit the training symbols for both the (training peer device's) transmitting antenna and the (device's) receiving antenna. And the number of training symbols transmitted for training the device's receiving antenna shall be equal to value in the RX Training Size field of the ATTCIE it received in the previous iteration.

Once the a device has completed 4 iterations and it still desires to continue training its antenna, it shall start another RTT/CTT exchange and set the Continuition bit of the ATIE included in the RTT/CTT frames to ONE. The transmission mode used in the following training iterations shall be the one specified in the previous iteration. And the device shall set the Initial Discovery Mode field of the ATIE contained in this RTT/CTT exchange to the same value. In an antenna training procedure between two devices, the I device that transmits the first RTT shall also be the one that transmits any other RTTs for that training pair.

16.18.3.1 Transmission of TRN frames

SIFS after the RTT/CTT exchange, the I device shall first send a TRN frame in all configurations except for the one specified in Table 100.

Table 100 - Training configuation that requires MIFS

Training device	ATIE fields	
	Request Training	Request Feedback
I	ONE	ZERO
R	ZERO	ZERO

In the configuration specified in Table 100, the R device shall first send a TRN frame MIFS after CTT frame.

Except for the first TRN frame following the CTT frame, all other TRN frames shall be transmitted SIFS after the reception of the TRN frame sent by the training peer device.

If a device requests feedback as indicated by the Request Feedback field in the pre-agreed ATIE, the training peer device shall include AFIE in all TRN frames it sends in response to the TRN frames sent by the device.

16.18.4 Antenna tracking

After it starts or joins a beacon group in a data channel, a device may update its antenna weight settings using either explict or implict tracking. Antenna tracking does not involve the RTT/CTT handshake. However, the devices shall use the same configuration as agreed in the last RTT/CTT exchange in antenna training. Either device of a training pair may initiate tracking. In this subclause, the device initiating tracking is referred to as the I device despite of the role it plays in antenna training.

When closed-loop mode is used in tracking and the device providing feedback is Type B, then index feedback method shall be used; if the device giving feedback is Type A, the Fourier codebook based feedback method shall be used.

16.18.4.1 Explicit tracking

In explict antenna tracking, the I and R device shall use TRN frame to perform tracking. To start explict tracking, the I device shall send a TRN frame to R device. SIFS after the reception of the TRN frame, the R device shall respond with another TRN frame. And the responding TRN frame shall contain an AFIE if the configuration of the I device requests feedback. If the configuration of R device does not require feedback, the second TRN frame concludes the explict tracking and the I device may resume regular MSDU exchange afterwords. Otherwise if the configuration of R device requires feedback, the I device shall respond with another TRN frame that includes an AFIE in its payload to conclude the tracking procedure.

16.18.4.2 Implicit tracking

Either a reservation owner or target may initiate implicit tracking. In addition, implicit tracking may also be carried out employing B-Poll and response frame exchange.

16.18.4.2.1 Initiation by a reservation owner

To initiate implicit antenna tracking, a reservation owner shall aggregate an ATTCIE with a (aggregated) data frame into an aggregated tracking frame (15.8.2) before sending to the target. The ACK Policy of the aggregated training frame shall be set to B-ACK Request. If the configuration of the owner requires feedback, the target shall include AFIE in the payload of the responding B-ACK frame and set the Tracking Indication bit to ONE. If the configuration of target does not require feedback, the B-ACK frame concludes the implicit tracking. Otherwise if the configuration of the target requires feedback, to conclude this implicit tracking, the owner shall aggregate an AFIE with the next MSDU(s) into an aggregated training frame before sending to the target.

16.18.4.3 Initiation by a reservataion target

To initiate implicit antenna tracking, a reservation target device shall include ATTCIE in an Imm- or B-ACK frame sent to the reservation owner. The Tracking Indication bit of the Imm- or B-ACK frame shall be set to ONE. And if the configuration of the target requires feedback, the owner shall aggregate an AFIE with MSDU(s) into an aggregated tracking frame and set the Tracking Indication bit to ONE before transmitting it to the target. If the configuration of owner does not require feedback, the transmission of the aggregated training frame concludes the implicit tracking. Otherwise if the configuration of owner requires feedback, the ACK Policy field of the aggregated training frame shall be set to B-ACK Request. To conclude this implicit tracking in this case, the target shall include an AFIE in the next B-ACK frame that it sends to the owner SIFS after the reception of the aggregated training frame.

16.18.4.4 Implicit tracking using B-Poll frames

If a Type B slave is capable of assisting its Type A master in antenna training, the Type A master device may include ATTCIE in a B-Poll frame to initiate implicit tracking. If the training configuration of the Type A master device requires feedback, the Type B slave shall include an AFIE in its B-Poll response frame sent to the Type A master.

16.19 Transmit switched diversity (TSD) operation

The transmit switched diversity (TSD) is used to achieve diversity gain from shadowing or blockage. To perform the TSD operation, the source must have multiple antennas either sharing one common RF chain or using independent RF chains, respectively, as shown in Figure 170.

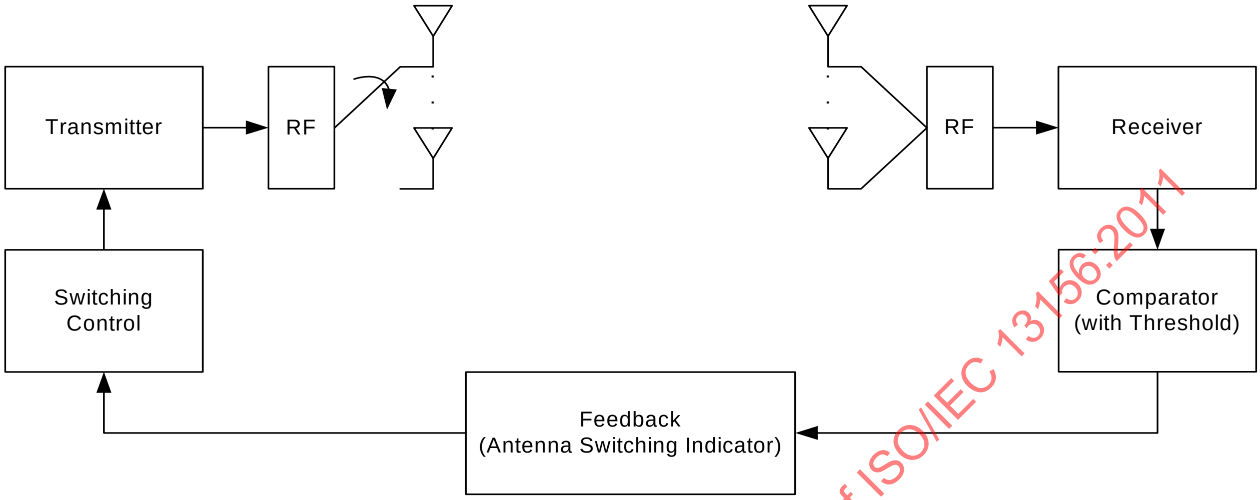

Figure 170 - Structure of transmit switched diversity system

16.19.1 TSD initiating procedure

The source shall send a TSD SET Request command to the destination using an arbitrarily selected antenna among multiple TX antennas, with the TSD Status set to BEGIN, as defined in 15.6.10. The destination, upon receiving the TSD SET Request command, shall send a TSD SET Response command, as defined in 15.6.11 to the source, proving the capability of the TSD operation of the destination.

16.19.2 Antenna switching

There are two mechanisms for switching TX antenna: reactive and proactive switching. In case of reactive switching, the destination sends the TSD Switch command frame with the TSD Switch field set to “Switch” to the source and then the source unconditionally switches to any other available TX antenna. In case of proactive switching, the source individually collects the channel status from every logical channel (link) which corresponds to each of the TX antennas, and switches the TX antenna to an optimal TX antenna which provides the maximum data rate among all TX antennas. The reactive switching is a simple way of the TX antenna switching which can overcome the link interruption, while the proactive switching can be used along with adaptive modulation and coding schemes to provide the maximum data rate, but it requires relatively complicated operating procedure.

16.19.2.1 Reactive switching

Before commencing data transmission, the source may arbitrarily select one antenna from multiple transmit antennas. At each frame transmission, the destination shall compare a certain metric based on the received SNR with a predetermined threshold after receiving data transmitted by the source. If the received SNR is smaller than the threshold, the destination shall feedback a TSD Switch command frame with the TSD Switch field set to “Switch” to the source. Note that the threshold may be selected by the receiver to satisfy the required criterion for a given data rate. Determining threshold may be decided by the implementers to get the required performance.

16.19.2.2 Proactive switching

In order to support proactive switching, the source should have the status information of each channel which corresponds to each of the TX antennas. To collect channel status for every TX antenna, the source shall be periodically operated in channel scan mode to acquire the status information of the channel which corresponds to each of the inactive antennas. When the data transmission is idle, the source and the destination may be operated in channel scan mode. During channel scan mode, the source shall stop using the current active TX antenna and switch to any TX antenna to acquire the channel state information by sending TSD SET Request command frame with the TSD Status field set to "SCAN" and by receiving TSD SET Response command frame with channel state information. Based upon the channel information acquired from all TX antennas, the source shall choose the optimal TX antenna which provides the best performance. The proactive switching may occur at every event of data rate change due to bad channel condition. However, if the current data rate is below the predetermined threshold, the proactive switching shall be periodically done to prevent staying with the current TX antenna despite bad channel conditions.

16.20 MAC sublayer parameters

Table 101 contains the values for the MAC sublayer parameters.

Table 101 - MAC sublayer parameters

Parameter	Value
mAggregationLimit	28
mBeaconSlotLength	21.3 *μs*
mBPExtension	6 beacon slots
mBPMergeWaitTime	128 superframes
mClockAccuracy	20 *ppm*
mClockResolution	1 *μs*
mChannelScanRequestRetry	pSIFS+10.2 *μs*
mDILength	4 × mSuperframeLength
mDRPBackoffWinMax	16 superframes
mDRPBackoffWinMin	2 superframes
mGuardTime	4.7 *μs*
mInitialMoveCountdown	3 × mMaxLostBeacons
mMasLength	64 *μs*
mMaxBeaconLength	mBeaconSlotLength - pSIFS - mGuardTime
mMaxBeaconSlotCollisionDetectionLatency	16
mMaxBPLength	72 beacon slots
mMaxDiscoveryLatency	128
mMaxFragmentCount	8
mMaxHibernationProtection	128 superframes
mMaxLostBeacons	3
mMaxMovableLatency	32
mMaxNeighbourDetectionInterval	128 superframes
mMaxSignalingSlotBackoff	128
mMaxSynchronizationAdjustment	4 *μs*

Table 101 - MAC sublayer parameters (concluded)

Parameter	Value
mMinFragmentSize	1
mSignalSlotCount	2 beacon slots
mSuperframeLength	256 x mMASLength
mCWMin	3
mCWMax	255
mDBSMax	mDILength-mDBSDuration
A-DIS-SCAN	6
B-DIS-SCAN	2
B-SCAN	128 *µs*
D-SCAN	8192 *µs*
LIFS	10666 *ns*
MIFS	pMIFS
SIFS	pSIFS
mDBSDuration	pDBSDuration

Table 102 contains the values of the PHY dependent parameters used by the MAC sublayer for the PHY.

Table 102 - PHY-dependent MAC sublayer parameters for the PHY

Parameter	Value
pCCADetectTime	Defined in Tables 39
pClockAccuracy	20 *ppm*
pMaxFrameBodySize	65535 octets
pSlotTime	8 *µs*

17 Security

This Clause specifies the security mechanisms needed to provide the security service introduced in 7.2.4.9. 17.1 reviews these security mechanisms. 17.2 defines security modes that govern the security operation of devices. 17.3 specifies the 4-way handshake procedure for two devices to establish pair-wise temporal keys (PTKs) and a secure relationship. This Clause also describes how a device may solicit or distribute group temporal keys (GTKs) within a secure relationship. 17.4 describes the procedures for frame reception and replay prevention. 17.5 provides the parameters needed in applying the AES-128 GCM cryptography to compute the message integrity code (MIC) and encrypt the secure payload for secure frames.

17.1 Security mechanisms

The security mechanisms specified in this International 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 GCM 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.

17.1.1 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 certain frames.

A device announces its selected security mode in the Beacon Parameters field in its beacons.

17.1.2 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).

17.1.3 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 the GTK commands and is sent in encrypted form.

17.1.4 Freshness protection

Freshness protection insures that no parties can successfully replay previously captured messages as an attack. This International Standard defines secure frame counters and replay counters on a per-temporal key basis to provide freshness protection.

17.1.5 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.

AES-128 counter mode is used for data encryption in this International Standard.

17.1.6 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.

AES-128 cipher block chaining - message authentication code (GCM) is used for MIC calculation in this International Standard.

17.2 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 17.3.

Once two devices establish a secure relationship, they shall use secure frames for frame transfers between them as specified in Table 103 and Table 104. 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 associated 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 103, "N" indicates a non-secure frame, and "S" indicates a secure frame.

Table 103 - 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.
RTT control frame	N	RTT frames shall be sent as non-secure frames.
CTT control frame	N	CTT 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.
DRP Reservation Request command frame	N, S	DRP Reservation Request frames may be sent as secure or non-secure frames.
DRP Reservation Response command frame	N,S	DRP 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.
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.
TRN control frame	N, S	The TRN control frame may be sent as secure or non-secure.

Table 104 specifies the values of the Encryption Offset (EO) field in secure frames

Table 104 - EO values in secure frames

Frame type or subtype	EO value
Application-specific control frame	Application defined
DRP Reservation Request command frame	Length of Secure Payload
DRP 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 + Aggregation Header Pad octets)

17.2.1 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 sublayer shall discard the frame.

17.2.2 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 103 and Table 104 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 or may not 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 sublayer 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 103, the MAC sublayer shall discard the frame.

A device that chooses to enable security mode 1 must understand and accept the responsibility that comes with receiving non-secure frames. The device shall instruct the higher layers to handle the received non-secure frames in a safe and secure manner.

A compliant MAC sublayer shall never use security mode 1 by default. Security mode 1 shall be entered from either mode 0 or mode 2. Requiring that a device explicitly select this mode serves as an indication that the device is aware of the security responsibilities it accepts when enabling security mode 1.

17.2.3 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 103 and Table 104 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 sublayer shall discard the frame.

If a device operating in mode 2 receives a non-secure frame that is required to have frame protection per Table 103, regardless of whether the device has a secure relationship with the device transmitting the frame, the MAC sublayer shall discard the frame.

17.3 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 Clause**.** 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 Clause.

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.

17.3.1 Mutual authentication and PTK derivation

This International 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 consists of four messages, called message 1, message 2, message 3, and message 4, that are sent back and forth between the two devices. The device sending message 1 becomes the initiator. The other device becomes the responder.

17.3.1.1 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 sublayer 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 17.3.4.
3. Construct and send message 2 in a PTK command.

17.3.1.2 4-way handshake message 2

The responder shall send message 2 to the initiator as specified in 17.3.1.1. 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 17.3.5. 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 17.3.4.

2. Recalculate the PTK MIC for the received message using the KCK according to 17.3.5. 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.

17.3.1.3 4-way handshake message 3

The initiator shall send message 3 to the responder as specified in 17.3.1.2. 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 17.3.5.

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 17.3.5. 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 using the appropriate MLME primitives.

17.3.1.4 4-way handshake message 4

The responder shall send message 4 to the initiator as specified in 17.3.1.3. 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 17.3.5.

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 17.3.5. 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; otherwise, install the PTK.

17.3.2 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 or may not 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 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. 17.3.6 provides an example means of generating a fresh GTK.

17.3.3 Pseudo-random function (PRF) definition

A PRF is used in several places in the security specification. Depending on the use, the PRF may need to output values of 64 bits, 128 bits, and 256 bits. This Clause 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 || denotes concatenation. Blocks are each 16 octets long, and are defined as inputs to the AES-128 GCM for the MIC generation as specified in 17.5.

GCM-MAC-FUNCTION(*K, N, A, B, Blen*)

begin

Form authentication block B_0 from flags = 0x59, *N*, and l(m) = 0

Form authentication block B_1 from l(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 ← MIC (*K, B_0, B_1, ..., A_0*)

return *R*

PRF(*K, N, A, B, Blen, Len*)

for *i* ← 1 **to** (*Len* + 63)/64 **do**

R ← *R* || GCM-MAC-FUNCTION(*K, N, A, B, Blen*)

N ← *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*)

17.3.4 PTK and KCK derivation

PRF-256 shall be employed to generate the PTK and KCK associated with a 4-way handshake as used in 17.3.1 based on the following parameters as defined in Table 105.

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

Blen - 32

Table 105 - 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:

KeyStream ← PRF-256(*K, N, A, B, Blen*)

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 106.

Table 106 - KCK and PTK source

Key	Source
KCK	KeyStream octets 0 through 15
PTK	KeyStream octets 16 through 31

17.3.5 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 105:

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

Blen - Length in octets of *B* = 48

PTK MIC ← PRF-64(*K, N, A, B, Blen*)

17.3.6 Random number generation

To implement the cryptographic mechanisms outlined in this International Standard, devices need to generate cryptographic grade random numbers. ISO/IEC 8802-11 Amendment 6 gives a detailed explanation 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 can provide a means for an implementation to create an unpredictable seed for a pseudo-random generation function. The example below shows how to distil 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 can 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 can 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

17.4 Frame reception steps and replay prevention measures

A recipient device shall carry out the reception steps and replay prevention measures as specified in this Clause.

17.4.1 Frame reception

The MAC sublayer 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 17.2. 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 17.5. If this validation fails, discard the frame; otherwise, proceed to the next step.
5. Detect frame replay as specified in 17.4.2. 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 17.4.2, and proceed to the next step.
6. Process the frame as specified in Clause 16, 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.

17.4.2 Replay prevention

Each transmitting MAC sublayer 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 sublayer 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 sublayer shall set up a 48-bit replay counter when a temporal key, PTK or GTK, is installed to it. The MAC sublayer 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 sublayer 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.

17.4.3 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.

17.5 AES-128 GCM inputs

AES-128 GCM provides confidentiality, authentication, and integrity for secure frames defined in this International Standard. This Clause specifies the various fields required for AES-128 GCM operation.

17.5.1 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. GCM is specified in ‘NIST Special Publication 800-38D’. GCM employs counter mode for encryption and authentication. AES-128 GCM combines AES-128 with GCM 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.

GCM has two input parameters - M (number of octets in authentication field) and L (number of octets in length field). For this International Standard, M = 8 and L = 2.

GCM 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 can 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 Clause showing the format of Nonce and GCM blocks, the most-significant octet is represented to the left of the other octets.

17.5.2 Nonce

The GCM 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 GCM. It is also used as a component of input block A_i for GCM encryption. It provides the uniqueness that GCM requires for each instance of authentication/ encryption. The GCM Nonce shall be formatted as shown in Figure 171. 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.

octets: 2	2	3	6
SrcAddr	DestAddr	TKID	SFN

Figure 171 - Nonce input to the GCM algorithm

17.5.3 GCM blocks

The GCM authentication blocks shall be formatted as shown in Figure 172 and further described below.

octets: 1	13	2	2	10	2	1	1	EO	0-15	P – EO	0-15
Flags (= 0x59)	Nonce	Encrypted data length l(m) = P – EO	Additional authenticated data length l(a) = 14 + EO	MAC Header	Encryption Offset (EO)	Security Reserved	0	Secure Payload portion not to be encrypted	Zero padding	Secure Payload portion to be encrypted	Zero padding
B_0			B_1					B_2, …, B_(M-1)		B_M, …, B_N	

Figure 172 - Input to GCM authentication blocks

17.5.3.1 Authentication block B_0

Authentication block B_0 is the first input block to the GCM algorithm. It shall be formatted as shown in Figure 173. 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.

octets: 1	13	2
Flags = 0x59	Nonce	l(m)

Figure 173 - Format of authentication block B_0

17.5.3.2 Authentication block B_1

Authentication block B_1 is the second input block to the GCM algorithm. It shall be formatted as shown in Figure 174. In this block, the l(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.

octets: 2	10	2	1	1
l(a)	MAC Header	EO	Security Reserved	0

Figure 174 - Format of authentication block B_1

17.5.3.3 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 GCM algorithm. They shall be formatted as shown in Figure 175. 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.

octets: EO	0-15	P – EO	0-15
Secure Payload portion not to be encrypted	Zero padding	Secure Payload portion to be encrypted	Zero padding
B_2, …, B_(M-1)		B_M, …, B_N	

Figure 175 - Format of authentication blocks beginning from B_2

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.

17.5.3.4 Encryption blocks A_0, A_1, …, A_m

GCM uses encryption blocks A_0, A_1, …, A_m to generate key stream blocks that are used to encrypt the GCM and the Secure Payload portion to be encrypted. These blocks shall be formed as shown in Figure 176. 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.

octets: 1	13	2
Flags = 0x01	Nonce	Counter *i*

Figure 176 - Format of A_i blocks

17.6 Token authentication

Tokens are used by devices to authenticate with each other indicating they belong to the same security domain. The user designates one device as the domain authorizer, and uses the domain authorizer to issue tokens to devices in which the user trusts.

17.6.1 Token issuance

A device transmits a token request to the domain authorizer when it needs to join the domain. The request contains the device's public key. The user verifies the request. The method of how to interact with the user is implementation related. A simple example is a flashing light flash button which the user can push to confirm the authorization. After the verification, the domain authorizer generates a token. The token is the signature of a hash value signed using the domain authorizer's private key. The hash value is the hash of the requesting device's public key. The domain authorizer then transmits the token, together with the domain authorizer's public key, back to the device. The device then verifies the token and stores it.

Two devices shall authenticate with each other that they belong to the same security domain before transmission starts. The tokens are exchanged and verified between two devices and the Revocation List is checked before key exchange. The Revocation List is described in 17.6.2.

17.6.2 Token revoke

The User uses the domain authorizer to revoke tokens by generating a Revocation List (RL). The Revocation List is the list of the hash values of each revoked device's public key and is signed by the domain authorizer's private key. After one device is added or removed from the list, a new RL is generated with a higher version number. After a new RL is generated, it is broadcast to all the devices in the domain. Devices receiving a higher version numbered RL will update its RL.

18 HDMI PAL

18.1 Introduction

This clause describes an HDMI protocol adaptation layer which preserves the HDMI content protection scheme. Details on HDMI can be found in the HDMI specification.

The 60 GHz wireless solution is placed between the HDMI source and HDMI sink as shown in Figure 177 and Figure 178.

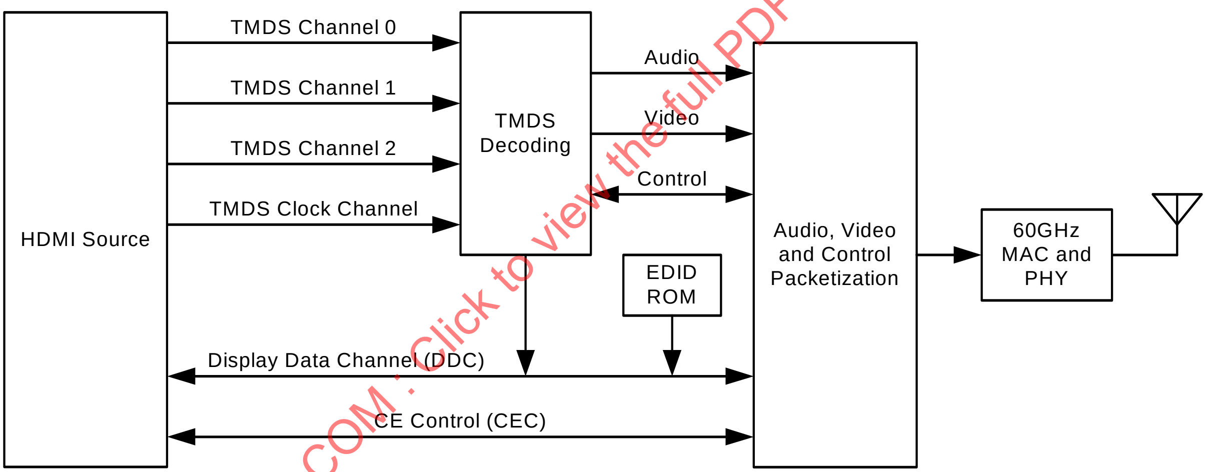

Figure 177 - Wireless HDMI transmitter

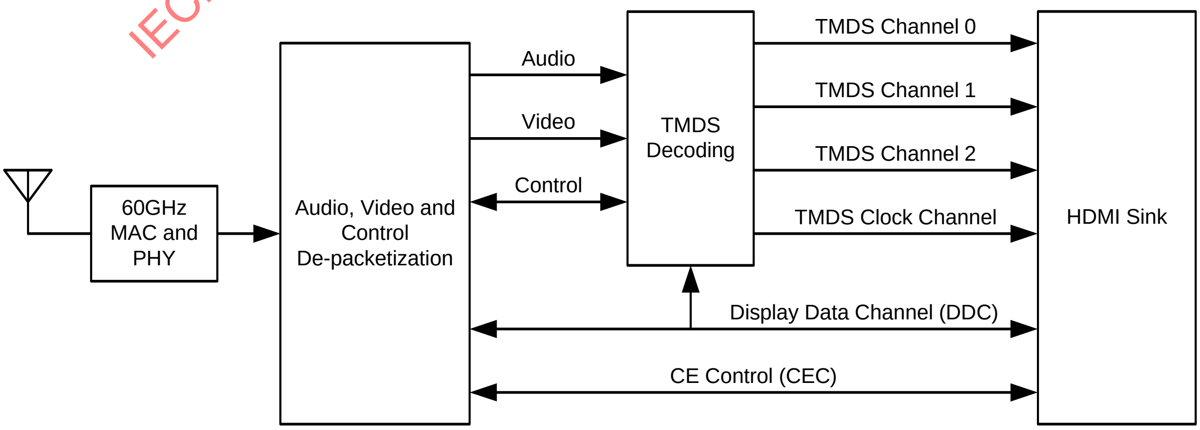

Figure 178 - Wireless HDMI receiver

18.2 HDMI transmission

Atthe transmitter, the TDMS coding is removed prior to transmission. The three data channels are then multiplexed together, along with the display data channel and the CE control. Both the Control Period and the Data Island Period contain the HSYNC and VSYNC information. If the data is video, then the pixel words are identified and the LSBs are flagged for unequal error protection (UEP). If the data is not video, then the data is flagged for equal error protection (EEP). The serial bit stream is then framed and presented to the MAC SAP.

18.2.1 Identification of video vs. data

In each video frame, a Control Period is required between any two periods that are not Control Periods. It is used to identify the type of the following Period (Video Data Period or Data Island Period).

18.2.2 TMDS removal

There are four colour depths supported: 24, 30, 36 and 48 bits per pixel. The pixel encoding method is pre-negotiated as part of the E-EDID (Extended display identification data) structure.

The 10 bit code words are identified on the TMDS channels 0 through 2. The device extracts 10 bits of data per Clock period on each TMDS pin.

18.2.3 Data type multiplexing

18.2.3.1 Video data

The video packet contains the PAL header and packet body. The leading guard band shall be stripped by the PAL and will not be part of payload. It is the receiver PAL's responsibility to insert the leading guard band into the HDMI Sink's TMDS data channel based on the PAL Header.

The layout of the packet body with "payload size = n bytes" is described below.

Name	*Offset*	*Size*	*Content*
PAL Header	0	32 bits	Header data
Pixel stream on TMDS 0	PAL Header+0	8 bits	Pixel data
Pixel stream on TMDS 1	PAL Header+1	8 bits	Pixel data
Pixel stream on TMDS 2	PAL Header+2	8 bits	Pixel data
.			
.			
Pixel stream on TMDS 0	PAL Header+n-3	8 bits	Pixel data
Pixel stream on TMDS 1	PAL Header+n-2	8 bits	Pixel data
Pixel stream on TMDS 2	PAL Header+n-1	8 bits	Pixel data

Figure 179 - Video packet body layout

18.2.3.2 Data island

The data island packet contains the PAL header and packet body. The leading guard band and trailing guard band shall be stripped by the PAL and will not be part of payload. It is the receiver PAL's responsibility to insert the leading guard band and trailing guard band into the HDMI Sink's TMDS data channel base on the PAL Header.

The minimum size of payload size = 32 TMDS periods (1 data packet * 32 Data TMDS period each). The maximum payload size = 576 TMDS periods (18 data packets * 32 Data TMDS period each).

1 Data packet = 32 TMDS periods = 30 bits * 32 = 960 bits = 120 bytes

The layout of the packet body with "payload size = n bytes" is described below.

Name	*Offset*	*Size*	*Content*
PAL Header	0	32 bits	Header data
Data packet clock 0 on TMDS 0	PAL Header+0	8 bits	MSB 8 bits of TMDS 0
Data packet clock 0 on TMDS 1	PAL Header+1	8 bits	LSB 2 bits of TMDS 0 + MSB 6 bits of TMDS 1
Data packet clock 0 on TMDS 2	PAL Header+2	8 bits	LSB 4 bits of TMDS 1 + MSB 4 bits of TMDS 2
Data packet clock 1 on TMDS 0	PAL Header+3	8 bits	LSB 6 bits of TMDS 2 + MSB 2 bits of TMDS 0
Data packet clock 1 on TMDS 1	PAL Header+4	8 bits	LSB 8 bits of TMDS 0
Data packet clock 1 on TMDS 2	PAL Header+5	8 bits	MSB 8 bits of TMDS 1
.			
.			
Data packet clock k on TMDS 2	PAL Header+n-1	8 bits	LSB 8 bits of TMDS 2

Figure 180 - Data packet body layout

18.2.3.3 Control

The control period contains a 4 bytes PAL header and packet body (control only). The preamble shall be stripped by the PAL and will not be part of payload. It is the receiver PAL's responsibility to insert the Preamble into the HDMI Sink's TMDS data channel base on the PAL Header.

The layout of the packet body with "payload size = n bytes" is described below.

Minimum Control Packet Payload Size = 4 TMDS periods = 15 bytes

Name	*Offset*	*Size*	*Content*
PAL Header	0	32 bits	Header data
Control packet clock 0 on TMDS 0	PAL Header+0	8 bits	MSB 8 bits of TMDS 0
Control packet clock 0 on TMDS 1	PAL Header+1	8 bits	LSB 2 bits of TMDS 0 + MSB 6 bits of TMDS 1
Control packet clock 0 on TMDS 2	PAL Header+2	8 bits	LSB 4 bits of TMDS 1 + MSB 4 bits of TMDS 2
Control packet clock 1 on TMDS 0	PAL Header+3	8 bits	LSB 6 bits of TMDS 2 + MSB 2 bits of TMDS 0
Control packet clock 1 on TMDS 1	PAL Header+4	8 bits	LSB 8 bits of TMDS 0
Control packet clock 1 on TMDS 2	PAL Header+5	8 bits	MSB 8 bits of TMDS 1
.			
.			
Control packet clock k on TMDS 2	PAL Header+n-1	8 bits	LSB 8 bits of TMDS 2

Figure 181 - Control packet body layout

18.2.3.4 DDC/CEC packet

The DDC/CEC packet contains the PAL header and packet body. The packet body includes the 6 bytes DDC/CEC message header, followed by the variable size message body. The format of the DDC and CEC message body can be find in the HDMI specification. The layout of the packet body is described below.

Name	*Offset*	*Size*	*Content*
PAL Header	0	32 bits	Header data
Message Header (Type)	PAL Header+0	8 bits	EDID Protocol: 1 CEC Protocol: 2
Message Header (RdWrt)	PAL Header+1	8 bits	Read: 0 Write: 1
Message Header (Address)	PAL Header+2	8 bits	Address
Message Header (SegPtr)	PAL Header+3	8 bits	Segment Pointer
Message Header (MsgLen)	PAL Header+4	16 bits	Message Length: 1 - 65535
Message body byte 0	PAL Header+Message Header+0		
Message body byte 1	PAL Header+Message Header+1	8 bits	
.			
.			
Message body byte n	PAL Header+Message Header+n-1	8 bits	

Figure 182 - DDC/CEC packet body layout

18.3 HDMI reception

Atthe receiver, the packet is received via the MAC SAP. The PAL header is read to indicate if the packet is video, audio, control or DDC/CEC. If the packet is video then the TDMS coding is applied prior to passing the video data onto the HDMI receiver. The three data channels are demultiplexed apart, along with the display data channel and the CE control.

18.3.1 TMDS encoding

For video packets the PAL needs to apply the TMDS coding prior to passing it to the HDMI sink, as described in the HDMI specification.

18.3.2 Packet demultiplexing

The packet body is routed, as indicated, to the appropriate port of the HDMI sink.

18.3.2.1 Video

Video packets are demultiplexed and distributed to TMDS data channels 0 through 2. The receiver PAL shall receive and examine the PAL header for packet type. If the type of packet is video, then a leading guard band shall be insert into TMDS data channel before the video stream.

After the leading guard band, the PAL shall encode the 8 bits video data into 10 bits format before mapping to the TMDS data channel. The 2-stage TMDS encoding process shall be performed before mapping into the TMDS data channel, starting with TMDS data channel 0 for the first video data byte in the video packet, as shown in the figure below.