

# IEEE 802.15.4 - Mac Sublayer - A brief summary

Luca Bedogni

September 21, 2010

## **Abstract**

This document aims to give to the reader an idea of what IEEE 802.15.4 is, and a brief introduction about the requirements and the capabilities of this standard. This is not a replacement of the standard.

The document on what this summary is based is the “IEEE Std 802.15.4-2006”  
This document is not to be intended as a replacement of the standard, but a  
summary on some peculiarities and characteristics.  
For errors, revisions and other, please mail me at *lbedogni@cs.unibo.it*.

# 1 Acronyms and Abbreviations

Acronym or abbreviation	Meaning
BE	backoff exponent
BI	beacon interval
BLE	battery life extension
BO	beacon order
BSN	beacon sequence number
CAP	contention access period
CCA	clear channel assessment
CCM*	extension of CCM
CFP	contention free period
CRC	cyclic redundancy check
CSMA-CA	carrier sense multiple access with collision avoidance
CW	contention window (length)
DSN	data sequence number
ED	energy detection
FCS	frame check sequence
FFD	full-function device
GTS	guaranteed time slot
IFS	interframe space or spacing
LIFS	long IFS
LQI	link quality indication
LR-WPAN	low-rate wireless personal area network
LSB	least significant bit
MAC	medium access control
MFR	MAC footer
MHR	MAC header
MIC	message integrity code
MLME	MAC sublayer management entity
MLME-SAP	MLME service access point
MSB	most significant bit
MPDU	MAC protocol data unit
MSDU	MAC service data unit
NB	number of backoff (periods)
PAN	personal area network
PIB	PAN information base
POS	personal operating space
RFD	reduced-function device
SD	superframe duration
SFD	start-of-frame delimiter
SHR	synchronization header
SIFS	short interframe spacing
SO	superframe order
SPDU	SSCS protocol data units
SRD	short-range device
SSCS	service-specific convergence sublayer
WLAN	wireless local area network
WPAN	wireless personal area network

## 2 Introduction

An LR-WPAN is a simple, low-cost communication network that allows wireless connectivity in applications with limited power and relaxed throughput requirements. The main objectives of an LR-WPAN are ease of installation, reliable data transfer, short-range operation, extremely low-cost, and a reasonable battery life, while maintaining a simple and flexible protocol.

Some of the characteristics of an LR-WPAN are as follows:

- Over-the-air data rates of 250kb/s, 100kb/s, 40kb/s and 20kb/s
- Star or peer-to-peer operation
- Allocated 16-bit short or 64-bit extended addresses
- Optional allocation of guaranteed time slots (GTSs)
- Carrier sense multiple access with collision avoidance (CSMA-CA) channel access
- Fully acknowledged protocol for transfer reliability
- Low power consumption
- Energy detection (ED)
- Link quality indication (LQI)
- 16 channels in the 2450 MHz band, 30 channels in the 915 MHz band, and 3 channels in the 868 MHz band

There are two different devices that can participate in an IEEE 802.15.4 network, a full-function device (FFD) and a reduced-function device (RFD).

A full function device can act as a personal area network (PAN) coordinator, a coordinator, or a device. An FFD can talk to RFDs or other FFDs, while an RFD can talk only to a FFD. An RFD is intended for applications that are extremely simple, such as light switch or a passive infrared sensor. An RFD does not have to send lot of data over the time, and can associate with only FFD at a time. Thus, an RFD can be implemented using minimal resources and memory capacity.

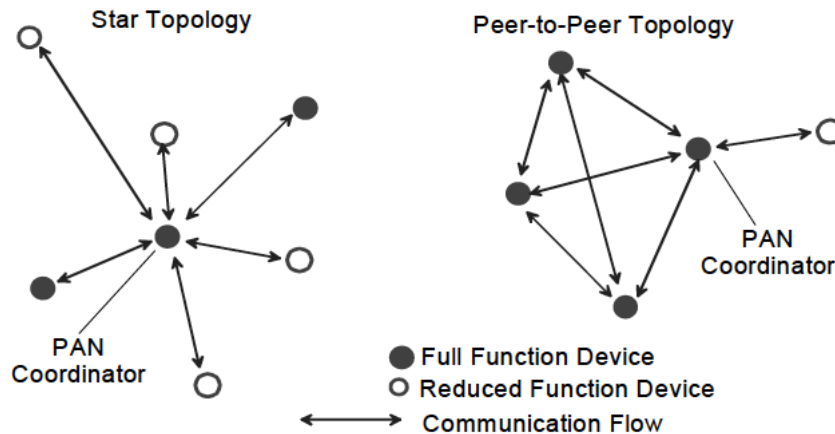
## 3 Network

### 3.1 Network components

A network must have at least one FFD, and at least one PAN coordinator. We have a WPAN when two (or more) devices within a POS communicate on the same physical channel.

### 3.2 Network topologies

An IEEE 802.15.4 LR-WPAN may operate in either of two topologies: the star topology or the peer-to-peer topology.



In the star topology, the communication is established between devices and a single central controller, called the PAN coordinator. A device typically has some associated application and is either the initiation point or the termination point for network communication. A PAN coordinator may also have specific application, but it can be used to initiate, terminate or route communication around the network. The PAN coordinator is the primary controller of the PAN. All devices operating on a network of either topology shall have unique 64-bit addresses. This address may be used for direct communication within the PAN, or a short address may be allocated by the PAN coordinator when the device associates and used instead. The PAN coordinator might often be mains powered, while the devices will most likely be battery powered. Applications that benefit from a star topology include home automation, personal computer peripherals, toys and games, and personal health care.

The peer-to-peer topology also has a PAN coordinator; however, it differs from the star topology in that any device may communicate with any other device as long as they are in range of one another. Peer-to-peer topology allows more complex network formations to be implemented, such as mesh networking topology. Applications such as industrial control and monitoring, wireless sensor networks, asset and inventory tracking, intelligent agriculture and security would benefit from such network topology. A peer-to-peer network can be ad hoc, self-organizing and self-healing. Such functions can be added at the higher layer, but are not part of this standard.

Each independent PAN selects a unique identifier. This PAN identifier allows communication between devices within a network using short address and enables transmissions between devices across independent networks. The mechanism by which identifiers are chosen is outside the scope of this standard.

The network formation is performed by the higher layer, which is not part of this standard, but 5.3.1 and 5.3.2 (both of the original document, not this) provide a brief overview on how each supported topology can be formed.

### 3.3 Star network

After an FFD is activated, it can establish its own network and become the PAN coordinator. Once the PAN coordinator has chosen the PAN identifier, other de-

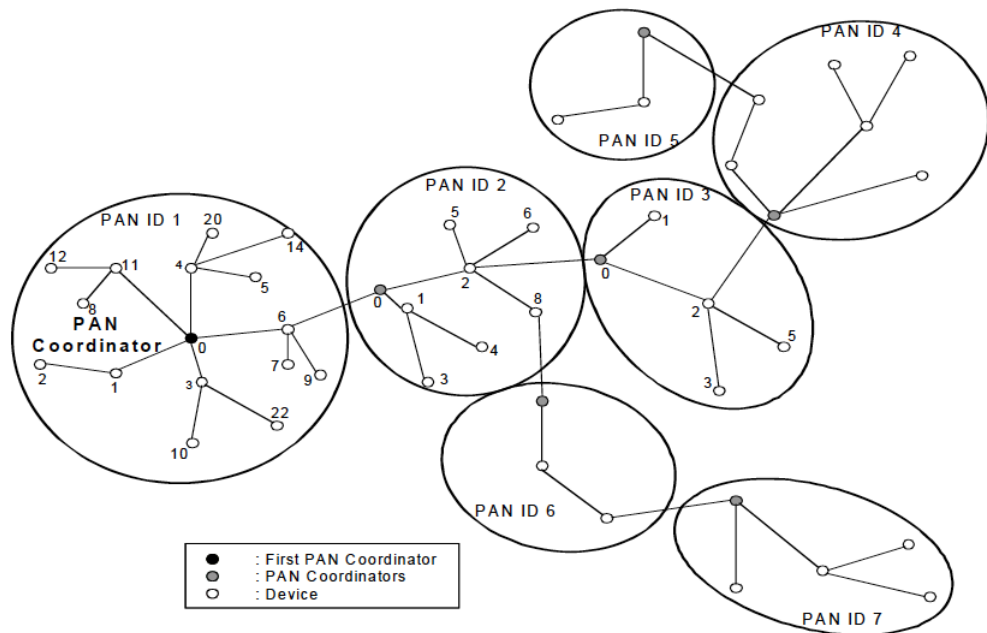
vices, both FFDs and RFDs, can join the network. The higher layer can use the procedures described in 7.5.2 and 7.5.3 (of the original document, not this) to form a star network.

### 3.4 Peer-to-peer network

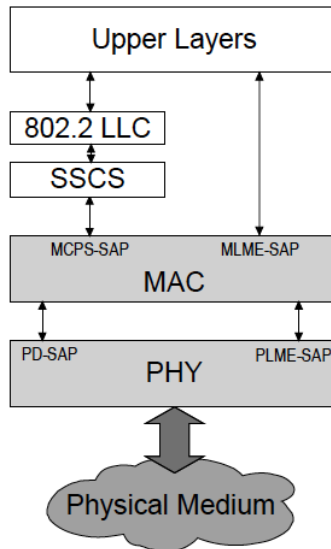
In a peer-to-peer topology, every device is able to communicate with every other device within its range. One device has to be the PAN coordinator, usually the first that has created the network. Other types of networks may form inside the peer-to-peer network, and it's possible to impose topological restriction on how to form the network.

#### 3.4.1 Cluster tree

An example of peer-to-peer network is the cluster tree. It's a special case of a peer-to-peer network in which most devices are FFDs. An RFD connects to a cluster tree network as a leaf device at the end of a branch, because RFDs do not allow other devices to associate. Every FFD can act as a coordinator, but there can be only one PAN coordinator. Every FFD that joins the network adds the PAN coordinator as parent, and the PAN coordinator adds the new device as a child. Then, the child begins to transmit periodic beacons to permit other device to join the network.



### 3.5 Architecture - MAC



The MAC sublayer provides two services: the MAC data service and the MAC management service interfacing to the MAC suayer management entity (MLME) service access point (SAP), known as MLME-SAP. The MAC data service enables the transmission and reception of MAC protocol data units (MPDUs) across the PHY service.

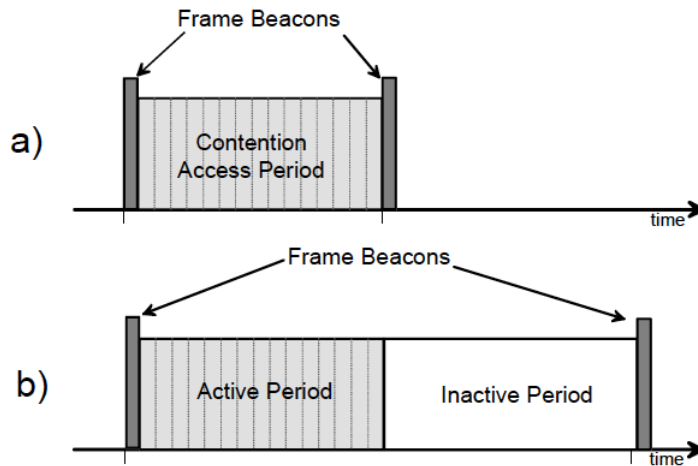
The MAC sublayer is responsible for:

- beacon management
- channel access
- GTS management
- frame validation
- acknowledged frame delivery
- association and disassociation


## 4 Functional overview

### 4.1 Superframe structure

This standards allows the optional use of a superframe structure, which format is defined by the coordinator. The superframe is bounded by network beacons, sent by the coordinator and is divided into 16 equally sized slot. Optionally, the superframe can have an active and an inactive portion, where the coordinator can enter a low-power mode. The beacon is transmitted in the first slot of each superframe. If a coordinator does not wish to use a superframe structure, it will turn off beacon transmissions.



The beacons are used to identify the PAN and to describe the nature of the superframes. Every device that wishes to communicate during the CAP competes with other devices using a slotted CSMA-CA mechanism, and must complete every transaction before the next beacon.



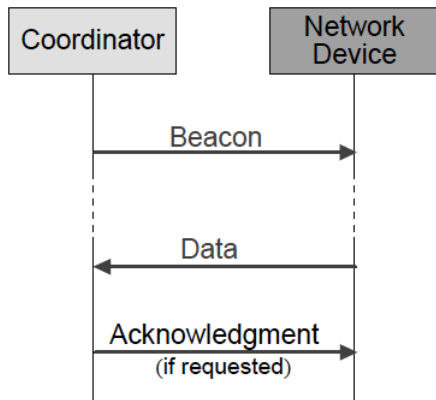
For low-latency applications or for applications that require specific data transfers and bandwidth, it is possible to allocate, at the end of the CAP, some slots (called GTSS) that allow clients to have more guarantees on the timing to transmit. This is called the contention free period, or CFP.

## 4.2 Data transfer model


There are mainly three types of data transfer: from a device to a coordinator, from a coordinator to a device, between two devices. In the star topology, only the first two are present, while in the peer-to-peer topology we can find all three of them.



#### 4.2.1 Data transfer to a coordinator




If we are in a beacon-enabled PAN, the device has to wait for the beacon from the coordinator, after that it can transmit the data needed.




If we are not in a beacon-enabled PAN, the device simply transmit the data, waiting for an optional ACK from the coordinator.

#### 4.2.2 Data transfer from a coordinator



In a beacon-enabled PAN, when the coordinator wishes to transfer data to a device, it indicates that in the beacon. Then, the device request the data from the coordinator, which send the data after an ACK is sent in answer to the request. When the transaction is finished, the message is removed from the list of the pending messages.



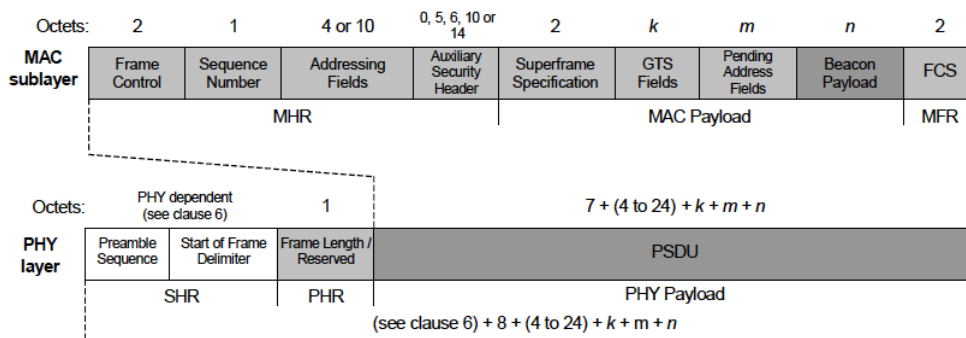
In a non beacon-enabled PAN, when a coordinator wishes to transfer data to a client, it simply waits for a client requesting the data. If that happens, the coordinator acknowledge the request, following with the data to be transmitted. If it has no data, it signal this fact to the client requesting data.

### 4.3 Frame structure

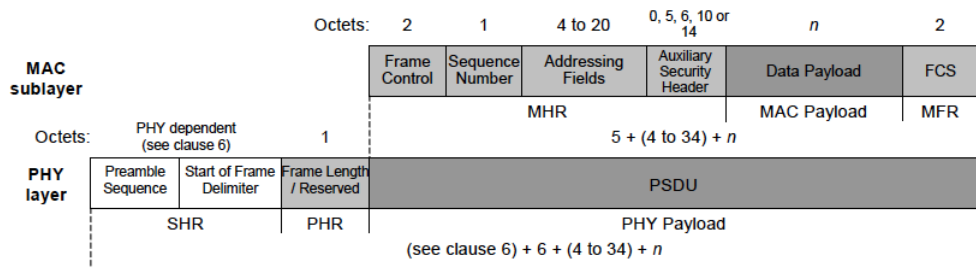
There are four different frame structures:

- Beacon frame
- Data frame
- ACK frame
- MAC command frame

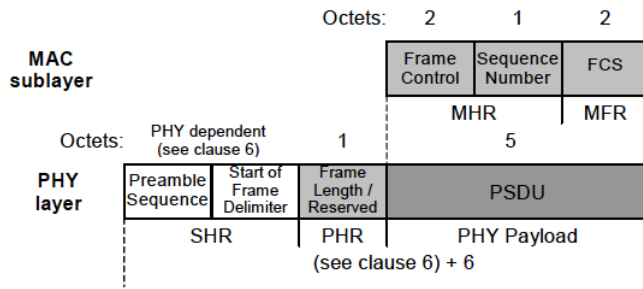
#### 4.3.1 Beacon frame



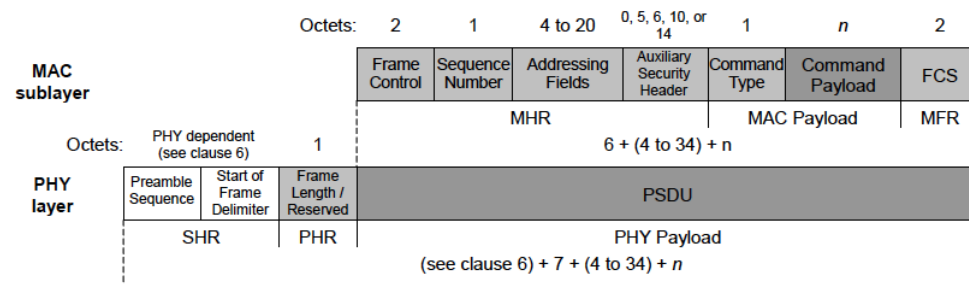
#### 4.3.2 Data frame



### 4.3.3 ACK frame



### 4.3.4 MAC command frame



## 4.4 Improving probability of successful delivery

There are mainly three methods to improve the probability of successful data transmission, and they are:

- CSMA-CA
- Frame ACK
- Data verification

### 4.4.1 CSMA-CA

Nonbeacon-enabled PANs use an unslotted CSMA-CA channel access mechanism. Every time a device has to send data, it waits for a random period, and then check the channel status: if its idle, it transmits the data, if its busy, then the device waits for another random period before trying to transmit again. ACKs are sent without CSMA-CA.

Beacon-enabled PANs, otherwise, use a slotted CSMA-CA mechanism, where the backoff slots are aligned with the start of the beacon transmission. Each time a device has to send data frames during the CAP, it locates the boundary of the next backoff slot and then waits for a random number of backoff slots. If the channel is busy, following the random backoff, the device waits for another random number of backoff slots before trying to access the channel again. ACKs and beacon frames are sent without CSMA-CA.

#### 4.4.2 Frame ACK

When the ACK is not required, the sender assumes the transmission was successful. Otherwise, if a frame is to be expected, then the sender retry the transmission until a positive ACK is received.

#### 4.4.3 Data verification

In order to detect bit errors, an FCS mechanism employing a 16-bit International Telecommunication Union-Telecommunication Standardization Sector (ITU-T) cyclic redundancy check (CRC) is used to detect errors in every frame.

### 4.5 Power consumption considerations

This standard is designed to be low-cost in terms of power consumption. Most of the devices using this standard will be battery power, with no practical ability to change batteries in short periods.

The programmer has the choice to balance between high performances, having the devices that listens to the channel frequently, or conservative, having the devices to be often in sleep state, and waking up only if there is some data to be transmitted.

### 4.6 Security

Devices for which this standard was designed have low power and limited capabilities in terms of computer power, but security is a major concern in wireless ad hoc networks.

Many of the architectural aspects of security are left for the higher layers, leaving in this standard only a symmetric-key cryptographic system, using keys that are provided by the higher layers.

Security mainly wishes to achieve:

- *Data confidentiality*: assurance that data is delivered only to those intended by the sender.
- *Data authenticity*: assurance that the data reaches the destination with data and sender intact.
- *Replay protection*: assurance that duplicates are detected.

### 4.7 Primitives

This standard is based on the concept of primitives, which can be found more detailed in ISO/IEC 8802.2. The basic idea is that a layer provides some functions to the higher layer, and these function are built on the sublayer.

A primitive can be one of four types:

- *Request*: this type indicate a request to initiate a service.
- *Indication*: this type is used internally to signal some occurred event.
- *Response*: this type is used to reply to an indication.
- *Confirm*: this type is used to reply to a request.

## 5 MAC sublayer

The MAC sublayer is responsible for the following tasks:

- Generating network beacons if the device is a coordinator
- Synchronizing to network beacons
- Supporting PAN association and disassociation
- Supporting device security
- Employing the CSMA-CA mechanism for channel access
- Handling and maintaining the GTS mechanism
- Providing a reliable link between two peer MAC entities

The MAC sublayer provides an interface between the SSCS and the PHY. It includes a management entity called the MLME. The MAC layer provides two services, accessed through two SAPs:

- The MAC data service, accessed through the MAC common part sublayer (MCPS) data SAP (MCPS-SAP)
- The MAC management service, accessed through the MLME-SAP

These two services provide the interface between the SSCS and the PHY, via the PD-SAP and PLME-SAP interfaces.

*Missing figure 29*

### 5.1 MAC data service

The MCPS-SAP support the transport of SSCS protocol data units (SPDUs) between peer SSCS entities.

The primitives supported by MCPS-SAP are:

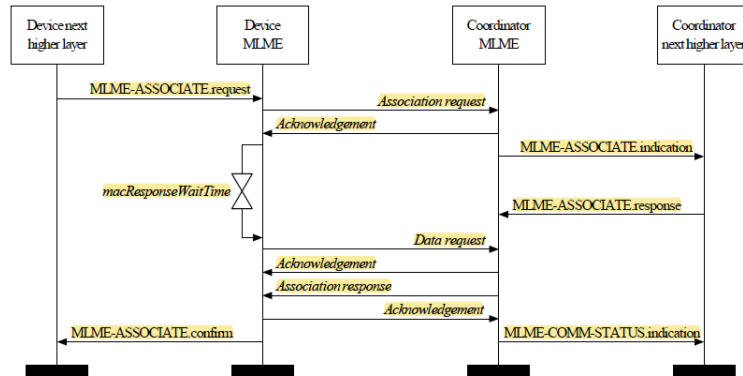
- **MCPS-DATA.request**: request the transfer of a data SPDU.
- **MCPS-DATA.confirm**: reports the result of a request to transfer an SPDU.
- **MCPS-DATA.indication**: indicates the transfer of a data SPDU from the MAC sublayer to the local SSCS entity.
- **MCPS-PURGE.request**: allows the next higher layer to purge an MSDU from the transaction queue. This primitive is optional for an RFD.
- **MCPS-PURGE.confirm**: allows the MAC sublayer to notify the next higher layer of the success of its request to purge an MSDU from the transaction queue. This primitive is optional for an RFD.

### 5.2 MAC management service

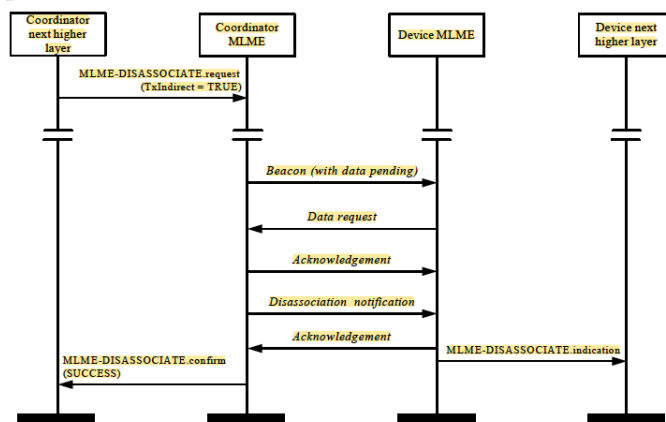
The MLME-SAP allows the transport of management commands between the next higher layer and the MLME.

- **MLME-ASSOCIATE.request**: allows a device to request an association with a coordinator.
- **MLME-ASSOCIATE.indication**: is used to indicate the reception of an association request command. This primitive is optional for an RFD.

- **MLME-ASSOCIATE.response**: is used to initiate a response to an MLME-ASSOCIATION.indication primitive. This primitive is optional for an RFD.
- **MLME-ASSOCIATE.confirm**: is used to inform the next higher layer of the initiating device whether its request to associate was successful or unsuccessful.



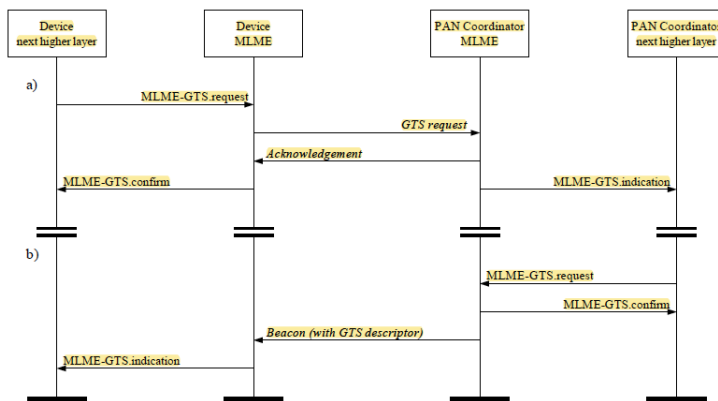
- **MLME-DISASSOCIATE.request**: is used by an associated device to notify the coordinator of its intent to leave the PAN. It is also used by a coordinator to instruct an associated device to leave the PAN.
- **MLME-DISASSOCIATE.indication**: is used to indicate the reception of a disassociation notification command.
- **MLME-DISASSOCIATE.confirm**: report the results of an MLME-DISASSOCIATION.request primitive.



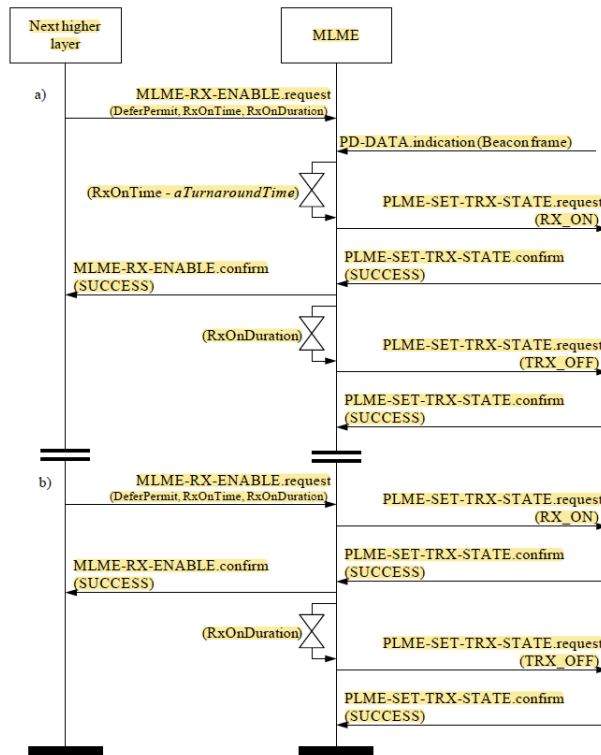
- **MLME-BEACON-NOTIFY.indication**: is used to send parameters contained within a beacon frame received by the MAC sublayer to the next higher layer. The primitive also sends a measure of the LQI and the time the beacon frame was received.
- **MLME-GET.request**: requests information about a given PIB attribute.
- **MLME-GET.confirm**: reports the result of an information request from the PIB.
- **MLME-GTS.request**: allows a device to send a request to the PAN coordinator to allocate new GTS or to deallocate an existing GTS. This primitive

is also used by the PAN coordinator to initiate a GTS deallocation. This primitive is optional for both an RFD and an FFD.

- **MLME-GTS.confirm**: reports the result of a request to allocate a new GTS or deallocate an existing GTS. This primitive is optional for both an RFD and an FFD.
- **MLME-GTS.indication**: indicates that a GTS has been allocated or that a previously allocated GTS has been deallocated. This primitive is optional for both an RFD and an FFD.
- *Missing figure 34*

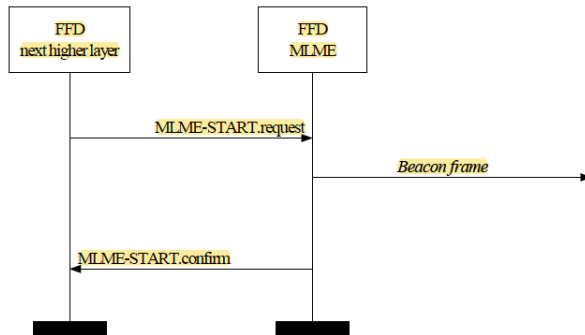


- **MLME-ORPHAN.indication**: allows the MLME of a coordinator to notify the next higher layer of the presence of an orphaned device. This primitive is optional for an RFD.
- **MLME-ORPHAN.response**: allows the next higher layer of a coordinator to respond to the MLME-ORPHAN.indication primitive. This primitive is optional for an RFD.
- *Missing figure 36*
- **MLME-RESET.request**: allows the next higher layer to request that the MLME performs a reset operation.
- **MLME-RESET.confirm**: reports the result of the reset operation.
- **MLME-RX-ENABLE.request**: allows the next higher layer to request that the receiver is either enabled for a finite period of time or disabled. This primitive is optional for both an RFD and an FFD.
- **MLME-RX-ENABLE.confirm**: reports the result of the attempt to enable or disable the receiver. This primitive is optional for both an RFD and an FFD.

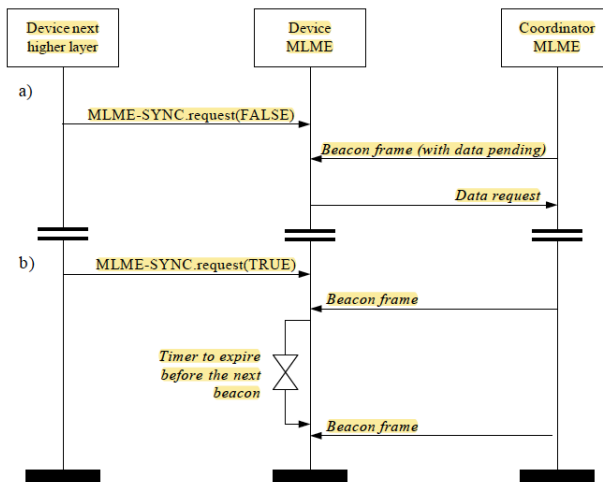


- **MLME-SCAN.request**: is used to initiate a channel scan over a given list of channels. A device can use a channel scan to measure the energy on the channel, search for the coordinator with which it associated, or search for all coordinators transmitting beacon frames within the POS of the scanning device.
- **MLME-SCAN.confirm**: reports the result of the channel scan request.
- *Missing images 79, 82, 83, 86*
- **MLME-COMM-STATUS.indication**: allows the MLME to indicate a communication status.
- **MLME-SET.request**: attempts to write the given value to the indicated PIB attribute.
- **MLME-SET.confirm**: reports the result of an attempt to write a value to a PIB attribute.
- **MLME-START.request**: allows the PAN coordinator to initiate a new PAN or to begin using a new superframe configuration. This primitive may also be used by a device already associated with an existing PAN to begin using a new superframe configuration. This primitive is optional for an RFD.
- **MLME-START.confirm**: reports the results of the attempt to start using a new superframe configuration. This primitive is optional for an RFD.





- **MLME-SYNC.request**: requests to synchronize with the coordinator by acquiring and, if specified, tracking its beacons. This primitive is optional for both an RFD and an FFD.



- **MLME-SYNC-LOSS.indication**: indicates the loss of synchronization with a coordinator.
- **MLME-POLL.request**: prompts the device to request data from the coordinator.
- **MLME-POLL.confirm**: reports the result of a request to poll the coordinator for data.
- *Missing figure 40*

### 5.3 MAC frame formats

Each MAC frame consist of the following elements:

- A MHR, which comprises frame control, sequence number, address information and security-related information.
- A MAC payload, of variable length, which contains information specific to the frame type. Acknowledgment frames do not contain a payload.
- A MFR, which contains a FCS.

<b>Octets:</b> 2	1	0/2	0/2/8	0/2	0/2/8	0/5/6/10/ 14	variable	2
Frame Control	Sequence Number	Destination PAN Identifier	Destination Address	Source PAN Identifier	Source Address	Auxiliary Security Header	Frame Payload	FCS
		Addressing fields						
MHR							MAC Payload	MFR

This is the basic MAC frame, including the MHR, the payload and the MFR.

#### 5.4 MAC command frame

<b>Octets:</b> 2	1	(see 7.2.2.4.1)	0/5/6/10/14	1	variable	2
Frame Control	Sequence Number	Addressing fields	Auxiliary Security Header	Command Frame Identifier	Command Payload	FCS
MHR				MAC Payload		MFR

An FFD shall be capable of transmitting and receiving all command frame types, with the exception of the GTS request command. For an RFD the requirements are smaller instead.

- 0x01 - **Association request** - RFD transmit only. Allows a device to request an association with a PAN through the PAN coordinator or a coordinator.
- 0x02 - **Association response** - RFD receive only. Allows the PAN coordinator or a coordinator to communicate the results of an association attempt back to the device requesting association.
- 0x03 - **Disassociation notification** - RFD transmit and receive. Allows a PAN coordinator, a coordinator or an associated device to send the disassociation command.
- 0x04 - **Data request** - RFD transmit only. Allows an associated device to request data to a PAN coordinator or a coordinator.
- 0x05 - **PAN ID conflict notification** - RFD transmit only. Allows an associated device to send a message to the PAN coordinator when a PAN identifier conflict is detected.
- 0x06 - **Orphan notification** - RFD transmit only. This command is used by an associated device that has lost synchronization with its coordinator.
- 0x07 - **Beacon request** - RFD optional. Is used by a device to locate all coordinators within its POS during an active scan.
- 0x08 - **Coordinator realignment** - RFD receive only. This command is sent by a PAN coordinator or a coordinator following the request of an orphan notification command from a device that is recognized to be on its PAN or when any of its PAN configuration attributes change due to the receipt of an MLME-START.request primitive.
- 0x09 - **GTS request command** - RFD optional. Is used by an associated device that is requesting the allocation of a new GTS or the deallocation of an existing GTS from the PAN coordinator. Only devices that have a 16-bit short address less than 0xffff shall send this command.

## 5.5 MAC constants

*TODO - insert the constants from table 85 of the original document.*

*TODO - insert the MAC PIB attributes from table 86 of the original document.*

## 6 MAC functional description

There are two mechanism to access the channel, contention based and contention free. Contention based access allows devices to access the channel in a distributed fashion using a CSMA-CA backoff algorithm. Contention-free access is controlled entirely by PAN coordinator using GTSs.

Before starting a new PAN, the results of a channel scan can be used to select an appropriate logical channel and channel page, as well as a PAN identifier that is not being used by ant other PAN in the area.

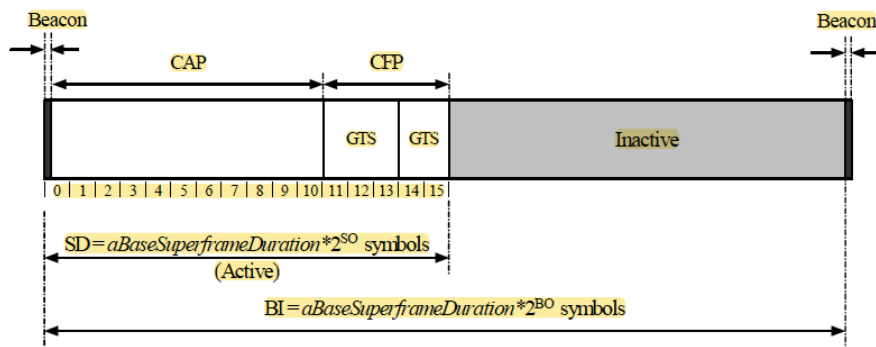
This standard is designed to allow devices to control the data transfer rather than the coordinator.

### 6.1 Channel access

#### 6.1.1 Superframe structure

A coordinator on a PAN can optionally bound its channel time using a superframe structure. A superframe is bounded by the transmission of a beacon frame and can have an active portion and an incative portion. The coordinator may enter a low-power (sleep) mode during the incative portion.

The active portion of each superframe shall be divided into  $aNumSuperframeSlots$  equally spaced slots of duration  $2exp2 * SOaBaseSlotDuration$  and is composed of three parts: a beacon, a CAP and a CFP. The beacon shall be transmitted, without the use of CSMA, at the start of slot 0, and the CAP shall commence immediatly following the beacon. The start of slot 0 is defined at the point at which the first symbol of the beacon PPDU is transmitted. The CFP, if present, follows immediatly after the CAP and extends to the end of the active portion of the superframe. Any allocated GTSs shall be located within the CFP.



#### 6.1.2 Contention Access Period - CAP

The CAP shall start immediatly following the beacon and complete before the CFP on a superframe slot boundary. If the CFP is zero length, the CAP shall complete

at the end of the active portion of the superframe. The CAP shall be at least  $aMinCAPLength$  symbols, unless additional space is needed to temporarily accommodate the increase in the beacon frame length needed to perform GTS maintenance, and shall shrink or grow dynamically to accommodate the size of the CFP.

All frame, except acknowledgment frames and any data frame that quickly follows the acknowledgment of a data request command, transmitted in the CAP shall use a slotted CSMA-CA mechanism to access the channel. A device transmitting within the CAP shall ensure that its transaction is complete (including ACKs) one IFS period before the end of the CAP. If this is not possible, the device shall defer its transmission until the CAP of the following superframe.

MAC command frame shall always be transmitted in the CAP.

### 6.1.3 Contention Free Period - CFP

The CFP shall start on a slot boundary immediately following the CAP and it shall complete before the end of the active portion of the superframe. If any GTSs have been allocated by the PAN coordinator, they shall be located within the CFP and occupy contiguous slots. The CFP shall therefore grow or shrink depending on the total length of all the combined GTSs.


No transmissions within the CFP shall use a CSMA-CA mechanism to access the channel. A device transmitting in the CFP shall ensure its transmission are complete one IFS period before the end of its GTS.

### 6.1.4 Interframe spacing - IFS


The MAC sublayer needs a finite amount of time to process data received by the PHY. To allow for this, two successive frames transmitted from a device shall be separated by at least an IFS period, including ACKs.

The length of the IFS period is dependent on the size of the frame that has just been transmitted. Frames (MPDUs) of up to  $aMaxSIFSFrameSize$  octets in length shall be followed by a SIFS period of duration of at least  $macMinSIFSPeriod$  symbols. Frames (MPDUs) with lengths greater than  $aMaxSIFSFrameSize$  octets shall be followed by a LIFS period of a duration of at least  $macMinLIFSPeriod$ .

#### Acknowledged transmission



#### Unacknowledged transmission



Where  $aTurnaroundTime \leq t_{ack} \leq (aTurnaroundTime + aUnitBackoffPeriod)$

## 6.2 CSMA-CA algorithm

The CSMA-CA algorithm shall be used before the transmission of data or MAC command frames transmitted within the CAP, unless the frame can be quickly

transmitted following the acknowledgment of a data request command. The CSMA-CA algorithm shall not be used for the transmission of beacon frames in a beacon-enabled PAN, acknowledgment frames or data frames transmitted in the CFP.

If periodic beacons are being used in the PAN, the MAC sublayer shall employ the slotted version of the CSMA-CA algorithm for transmission in the CAP of the superframe. Conversely, if periodic beacons are not being used in the PAN or if a beacon could not be located in a beacon-enabled PAN, the MAC sublayer shall transmit using the unslotted version of the CSMA-CA algorithm. In both case, the algorithm is implemented using units of time called backoff periods, where one backoff shall be equal to  $aUnitBackoffPeriod$  symbols.

In slotted CSMA-CA, the backoff period boundaries of every device in the PAN shall be aligned with the superframe slot boundaries of the PAN coordinator, i.e., the start of the first backoff period of each device is aligned with the start of the beacon transmission. In slotted CSMA-CA, the MAC sublayer shall ensure that the PHY commences all of its transmission on the boundary of a backoff period. In unslotted CSMA-CA, the backoff periods of one device are not related in time to the backoff period of any other device in the PAN.

Each device shall maintain three variables for each transmission attempt: **NB**, **CW** and **BE**. **NB** is the number of times the CSMA-CA algorithm was required to backoff while attempting the current transmission; this value shall be initialized to zero before each new transmission attempt. **CW** is the contention window length, defining the number of backoff period that need to be clear of channel activity before the transmission can commence; this value shall be initialized to two before each transmission attempt and reset to two each time the channel is assessed to be busy. The **CW** variable is only used for slotted CSMA-CA. **BE** is the backoff exponent, which is related to how many backoff periods a device shall wait before attempting to assess a channel. In unslotted systems, or slotted systems with the received BLE subfield set to zero, **BE** shall be initialized to the value of  $macMinBE$ . In slotted systems with the received BLE subfield set to one, this value shall be initialized to the lesser of two and the value of  $macMinBE$ . Note that if  $macMinBE$  is set to zero, collision avoidance will be disabled during the first iteration of this algorithm.

*Missing figure 69*

## 6.3 Starting and maintaining PANs

All devices shall be capable of performing active and passive scans across a specified list of channels. In addition, an FFD shall be able to perform ED and active scans.

### 6.3.1 ED channel scan

An ED scan allows a device to obtain a measure of the peak energy in each requested channel. This could be used by a prospective PAN coordinator to select a channel on which to operate prior to starting a new PAN. During an ED scan, the MAC sublayer shall discard all frames received over the PHY data service.

The ED scan shall terminate when either the number of channel ED measurements stored equals the implementation-specified maximum or energy has been measured on each of the specified logical channels.

### 6.3.2 Active channel scan

An active scan allows a device to locate any coordinator transmitting beacon frame within its POS. This could be used by a prospective PAN coordinator to select a PAN identifier prior to starting a new PAN, or it could be used by a device prior to association.

During an active scan, the MAC sublayer shall discard all frames received over the PHY data service that are not beacon frames.

Before commencing an active scan, the MAC sublayer shall store the value of *macPANId* and then set it to *0xffff* for the duration of the scan.

If a coordinator of a beacon-enabled PAN receives the beacon request command, it shall ignore the command and continue transmitting its periodic beacons as usual. If a coordinator of a nonbeacon-enabled PAN receives this command, it shall transmit a single beacon frame using unslotted CSMA-CA.

### 6.3.3 Passive channel scan

A passive scan, like an active scan, allows a device to locate any coordinator transmitting beacon frames within its POS. The beacon request command, however, is not transmitted. This type of scan could be used by a device prior to association.

### 6.3.4 Orphan channel scan

An orphan scan allows a device to attempt to relocate its coordinator following a loss of synchronization. During an orphan scan, the MAC sublayer shall discard all frames received over the PHY data service that are not coordinator realignment command frames.

## 6.4 PAN identifier conflict resolution

In some instances a situation could occur in which two PANs exist in the same POS with the same PAN identifier. If this conflict happens, the PAN coordinator and its devices shall perform the PAN identifier conflict resolution procedure. This procedure is optional for an RFD.

### 6.4.1 Detection

The PAN coordinator shall conclude that a PAN identifier conflict is present if either of the following apply:

- A beacon frame is received by the PAN coordinator with the PAN coordinator subfield set to one and the PAN identifier equal to *macPANId*.
- A PAN ID conflict notification command is received by the PAN coordinator from a device associated with it on its PAN.

Instead, a device that is associated through the PAN coordinator shall conclude that a PAN identifier conflict is present if the following applies:

- A beacon frame is received by the device with the PAN coordinator subfield set to one, the PAN identifier equal to *macPANId*, and an address that is equal to neither *macCoordShortAddress* nor *macCoordExtendedAddress*.

A device that is associated through a coordinator that is not the PAN coordinator shall not be capable of detecting a PAN identifier conflict.

#### 6.4.2 Resolution

On the detection of a PAN identifier conflict by a device, it shall generate the PAN ID conflict notification command and send it to its PAN coordinator.

On the detection of a PAN identifier conflict by the PAN coordinator, the MLME shall issue an MLME-SYNC-LOSS.indication to the next higher layer with the *LossReason* parameter set to PAN\_ID\_CONFLICT. The next higher layer of the PAN coordinator may then perform an active scan and, using the information from the scan, select a new PAN identifier.

### 6.5 Starting and realigning a PAN

#### 6.5.1 Starting a PAN

A PAN should be started by an FFD only after having first performed a MAC sublayer reset, by issuing the MLME-RESET.request primitive with the *SetDefaultPIB* parameter set to TRUE, an active channel scan, and a suitable PAN identifier selection.

#### 6.5.2 Realigning a PAN

If a coordinator receives the MLME-START.request primitive with the *CoordRealignment* parameter set to TRUE, the coordinator shall attempt to transmit a coordinator realignment command containing the new parameters for PANId, LogicalChannel and, if present, ChannelPage.

#### 6.5.3 Realignment in a PAN

If a device has received the coordinator realignment command from the coordinator through it is associated and the MLME was not carrying out an orphan scan, the MLME shall issue the MLME-SYNC-LOSS.indication primitive with the *LossReason* parameter set to REALIGNMENT and the PANId, LogicalChannel, ChannelPage and the security-relate parameters set to the respective fields in the coordinator realignment command.

#### 6.5.4 Updating superframe configuration and channel PIB attributes

To update the superframe configuration and channel attributes, the MLME shall assign values from the MLME-START.request primitive parameters to the appropriate PIB attributes. The MLME shall set *macBeaconOrder* to the value of the BeaconOrder parameter. If *macBeaconOrder* is equal to 15, the MLME will also set *macSuperframeOrder* to 15. In this case, this primitive configures a nonbeacon-enabled PAN.

## 6.6 Device discovery

The PAN coordinator or a coordinator indicates its presence on a PAN to other devices by transmitting beacon frames. This allows the other devices to perform device discovery.

A coordinator that is not the PAN coordinator shall begin transmitting beacon frames only when it has successfully associated with a PAN.

## 6.7 Association and disassociation

### 6.7.1 Association

A device shall attempt to associate only after having first performed a MAC sub-layer reset, by issuing the MLME-RESET.request primitive with the *SetDefaultPIB* parameter set to TRUE, and then having completed either an active channel scan or a passive channel scan.

### 6.7.2 Disassociation

The disassociation procedure is initiated by the next higher layer by issuing the MLME-DISASSOCIATE.request primitive to the MLME. When a coordinator wants one of its associated devices to leave the PAN, the MLME of the coordinator shall send the disassociation notification command in the manner specified by the *TxIndirect* parameter of the MLME-DISASSOCIATE.request primitive previously sent by the next higher layer.

Because the disassociation command contains an acknowledgment request, the receiving device shall confirm its receipt by sending an acknowledgment frame. If the direct or indirect transmission fails, the coordinator should consider the device disassociated. The same applies if it is the device that wants to disassociate.

## 6.8 Synchronization

For PANS supporting beacons, synchronization is performed by receiving and decoding the beacon frames. For PANs not supporting beacons, synchronization is performed by polling the coordinator for data.

### 6.8.1 Synchronization with beacons

All devices operating on a beacon-enabled PAN shall be able to acquire beacon synchronization in order to detect any pending message or to track the beacon. Devices shall be permitted to acquire beacon synchronization only with beacons containing the PAN identifier specified in *macPANId*. If *macPANId* specifies the broadcast PAN identifier (*0xffff*), a device shall not attempt to acquire beacon synchronization.

### 6.8.2 Synchronization without beacons

All devices operating on a nonbeacon-enabled PAN shall be able to poll the coordinator for data at the discretion of the next higher layer. A device is instructed to poll the coordinator when the MLME receives the MLME-POLL.request primitive.



On receipt of this primitive, the MLME shall follow the procedure for extracting pending data from the coordinator.

### 6.8.3 Orphaned device realignment

If the next higher layer receives repeated communication failures following its request to transmit data, it may conclude that it has been orphaned. Then, it may instruct the MLME to either perform the orphaned device realignment procedure, or to reset the MAC sublayer and then perform the association procedure.

## 6.9 Transaction handling

Because this standard favors very low cost devices that, in general, will be battery powered, transactions can be instigated from the devices themselves rather than from the coordinator. In other words, either the coordinator needs to indicate in its beacons when messages are pending for devices or the devices themselves need to poll the coordinator to determine whether they have any message pending. Such transfer are called indirect transmissions.

If the coordinator is capable of storing more than one transaction, it shall ensure that all the transactions for the same device are sent in the order in which they arrived at the MAC sublayer. For each transaction sent, if another exists for the same device, the MAC sublayer shall set its Frame Pending subfield to one, indicating the additional pending data.

Each transaction shall persist in the coordinator for at most *macTransaction-PersistenceTime*.

## 6.10 Transmission, reception and acknowledgment

### 6.10.1 Transmission

Each device shall store its current DSN value in the MAC PIB attribute *macDSN* and initialize it to a random value. Each time a data or a MAC command frame is generated, the MAC sublayer shall copy the value of *macDSN* into the sequence Number field of the MHR of the outgoing frame and then increment it by one. Each device shall generate exactly one DSN regardless of the number of unique devices with which it wishes to communicate. The value of *macDSN* shall be permitted to roll over.

Each coordinator shall store its current BSN value in the MAC PIB attribute *macBSN* and initialize it to a random value. Each time a beacon frame is generated, the MAC sublayer shall copy the value of *macBSN* into the Sequence Number field of the MHR of the outgoing frame and then increment it by one. The value of *macBSN* shall be permitted to roll over.

Both the DSN and the BSN are 8-bit values.

When a device has been associated and has assigned a 16-bit short address, it should use that for every transmission. When it has not been associated, it should use its 64-bit extended address.

If the Destination Address field is not present, the recipient of the frame shall be assumed to be the PAN coordinator.

### 6.10.2 Reception and rejection

Each device may choose whether the MAC sublayer is to enable its receiver during idle periods. During these idle periods, the MAC sublayer shall still service transceiver task request from the next higher layer. A transceiver task shall be defined as a transmission request with acknowledgment reception, if required, or a reception request. On a completion of each transceiver task, the MAC sublayer shall request that the PHY enables or disables its receiver.

For the first level of filtering by the MAC sublayer, it shall discard all received frames that do not contain a correct value in their FCS field in the MFR.

The second level of filtering shall be dependent on whether the MAC sublayer is currently operating in promiscuous mode. In promiscuous mode, the MAC sublayer shall pass all frames received after the first filter directly to the upper layers without applying any more filtering or processing. The MAC sublayer shall be in promiscuous mode if *macPromiscuousMode* is set to TRUE.

If the MAC sublayer is not in promiscuous mode, it shall accept only frames that satisfy all of the following third-level filtering requirements:

- The Frame Type subfield shall not contain a reserved frame type
- The Frame Version subfield shall not contain a reserved value
- If a destination PAN identifier is included in the frame, it shall match *macPANId* or shall be the broadcast PAN identifier
- If a short destination address is included in the frame, it shall match either *macShortAddress* or the broadcast address. Otherwise, if an extended destination address is included in the frame, it shall match *aExtendedAddress*
- If the frame type indicates that the frame is a beacon frame, the source PAN identifier shall match *macPANId* unless *macPANId* is equal to 0xffff, in which case the beacon frame shall be accepted regardless of the source PAN identifier
- If only source addressing fields are included in a data or MAC command frame, the frame shall be accepted only if the device is the PAN coordinator and the source PAN identifier matches *macPANId*

### 6.10.3 Extracting pending data from a coordinator

A device on a beacon-enabled PAN can determine whether any frames are pending for it by examining the content of the received beacon frame. If the address of the device is contained in the Address List field of the beacon frame and *macAutoRequest* is TRUE, the MLME of the device shall send a data request command to the coordinator during the CAP with the Acknowledgment Request subfield of the Frame Control field set to one.

On successfully receiving a data request command, the coordinator shall send an acknowledgment frame, thus confirming its receipt. If the coordinator has enough time to determine whether the device has a frame pending before sending the acknowledgment frame accordingly to indicate whether a frame is actually pending for the device. If this is not possible, the coordinator shall set the Frame Pending subfield of the acknowledgment frame to one.

On receipt of the acknowledgment frame with the Frame Pending subfield set to zero, the device shall conclude that there are no data pending at the coordinator.

The data frame following the acknowledgment of the data request command shall be transmitted using one of the following mechanism:

- Without using CSMA-CA, if the MAC sublayer can commence transmission of the data frame between  $aTurnaroundTime$  and  $(aTurnaroundTime + aUnitBackoffPeriod)$  symbols, on a backoff slot boundary, and there is time remaining in the CAP for the message, appropriate IFSm and acknowledgment. If a requested acknowledgment frame is not received following this data frame, the process shall begin anew following the receipt of a new data request command.
- Using CSMA-CA otherwise.

If the requesting device does not receive a data frame from the coordinator within  $macMaxFrameTotalWaitTime$  CAP symbols in a beacon-enabled PAN, or symbols in a nonbeacon-enabled PAN, or if the requesting device receives a data frame from the coordinator with a zero length payload, it shall conclude that there are no data pending at the coordinator. If the requesting device does receive a data frame from the coordinator, it shall send an acknowledgment frame, if requested, thus confirming receipt.

## 6.11 GTS allocation and management

A GTS allows a device to operate on the channel within a portion of the superframe that is dedicated exclusively to that device. A GTS shall be allocated only by the PAN coordinator, and it shall be used only for communications between the PAN coordinator and a device associated with the PAN through the PAN coordinator. A single GTS may extend over one or more superframe slots. The PAN coordinator may allocate up to seven GTSs at the same time, provided there is sufficient capacity in the superframe.

## 6.12 CAP maintenance

The PAN coordinator shall preserve the minimum CAP length of  $aMinCAPLength$  and take preventative action if the minimum CAP is not satisfied. However, an exception shall be allowed for the accomodation of the temporary increase in the beacon frame length needed to perform GTS maintenance.

### 6.12.1 GTS allocation

A device is instructed to request the allocation of a new GTS through the MLME-GTS.request primitive, with GTS characteristics set according to the requirements of the intended application.

For allocating a GTS, the superframe shall have available capacity if the maximum number of GTSs has not been reached and allocating a GTS of the desired length would not reduce the length of the CAP to less than  $aMinCAPLength$ . GTSs shall be allocated on a first-come-first-served basis by the PAN coordinator provided there is sufficient bandwidth available.

When the PAN coordinator determines whether capacity is available for the requested GTS, it shall generate a GTS descriptor with the requested specifications and the 16-bit short address of the requesting device.

### 6.12.2 GTS usage

When the MAC sublayer of a device that is not the PAN coordinator receives an MCPS-DATA.request primitive with the TxOptions parameter indicating a GTS transmissions, it shall determine whether it has a valid transmit GTS.

If the device has any receive GTSs, the MAC sublayer of the device shall ensure that the receiver is enabled at a time prior to the start of the GTS and for the duration of the GTS, as indicated by its starting slot and its length.

For all allocated transmit GTSs, the MAC sublayer of the PAN coordinator shall ensure that its receiver is enabled at a time prior to the start and for the duration of each GTS.

If a device misses the beacon at the beginning of a superframe, it shall not use its GTSs until it receives a subsequent beacon correctly. If a loss of synchronization occurs due to the loss of the beacon, the device shall consider all of its GTS deallocated.

### 6.12.3 GTS deallocation

GTS deallocation may be initiated by the PAN coordinator due to deallocation request from the next higher layer, the expiration of the GTS, or maintenance required to maintain the minimum CAP length, *aMinCAPLength*.

### 6.12.4 GTS reallocation

The deallocation of a GTS may result in the superframe becoming fragmented. *Missing figure 73*. If there are 3 GTSs, and the one in the middle is deallocated, then there will be a gap in the superframe during which nothing can happen.

The PAN coordinator shall ensure that any gaps occurring in the CFP, appearing due to the deallocation of a GTS, are removed to maximize the length of the CAP.

### 6.12.5 GTS expiration

The MLME of the PAN coordinator shall attempt to detect when a device has stopped using a GTS using the following rules:

- For a transmit GTS, the MLME of the PAN coordinator shall assume that a device is no longer using its GTS if a data frame is not received from the device in the GTS at least every  $2^n$  superframes
- For receive GTSs, the MLME of the PAN coordinator shall assume that a device is no longer using its GTS if an acknowledgment frame is not received from the device at least every  $2^n$  superframes. If the data frames sent in the GTS do not require acknowledgment frames, the MLME of the PAN coordinator will not be able to detect whether a device is using its receive GTS. However, the PAN coordinator is capable of deallocating the GTS at any time.

The value of n is:

- $n = 2^{(8 - \text{macBeaconOrder})} - 1$  — — —  $0 \leq \text{macBeaconOrder} \leq 8$
- $n = 1$  — — — — — — — — — —  $9 \leq \text{macBeaconOrder} \leq 14$

### 6.13 Frame security

The MAC sublayer is responsible for providing security services on specified incoming and outgoing frames when requested to do so by the higher layers. The services assured are:

- Data confidentiality
- Data authenticity
- Replay protection