5. MAC protocol specification

This clause specifies the MAC sublayer of this standard. The MAC sublayer handles all access to the physical layer and is responsible for the following tasks:

a) Generating network beacons if the device is a coordinator
b) Synchronizing to network beacons
c) Supporting VPAN association and disassociation
d) Supporting color function
e) Supporting visibility
f) Supporting dimming
g) Flicker mitigation scheme
h) Supporting visual indication of device status and channel quality
i) Supporting device security
j) Providing a reliable link between two peer MAC entities
k) Supporting mobility

Peer-to-peer, star and broadcasting capabilities, as shown in Figure 1, are provided with a single MAC frame structure. All of these diverse modes are supported via a single low complexity integrated frame structure.

Constants and attributes that are specified and maintained by the MAC sublayer are written in the text of this clause in italics. Constants have a general prefix of “a”, e.g., *aBaseSlotDuration*, and are listed in Table 59 (see 6.4.1). Attributes have a general prefix of “mac”, e.g., *macAckWaitDuration*, and are listed in Table 60 (see 6.4.2), while the security attributes are listed in Table 66 (see 7.5.1).

5.1 MAC functional description

This subclause provides a detailed description of the MAC functionality. Subclause 5.1.1 describes the following two mechanisms for channel access: contention based and contention free. Contention-based access allows devices to access the channel in a distributed fashion using an unslotted random access back-off algorithm. Contention-free access is controlled entirely by the coordinator through the use of GTSs.

The mechanisms used for starting and maintaining a VPAN are respectively described in 5.1.2 and 5.1.3. Channel scanning is used by a device to assess the current state of a channel (or channels), locate all beacons within its operating space, or locate a particular beacon with which it has lost synchronization. Before starting a new VPAN, the results of a channel scan can be used to select an appropriate logical channel, as well as a VPAN identifier that is not being used by any other VPAN in the area. Because it is still possible for the operating space of two VPANs with the same VPAN identifier to overlap, a procedure exists to detect and resolve this situation. Following a channel scan and suitable VPAN identifier selection, operation as a coordinator shall commence. Also described in the subclause is a method to allow coordinator beaconing to discover other such devices during normal operations, i.e., when not scanning.

The mechanisms to allow devices to join or leave a VPAN are defined in 5.1.4. The association procedure describes the conditions under which a device may join a VPAN and the conditions necessary for a coordinator to permit devices to join. Also described is the disassociation procedure, which can be initiated by the associated device or its coordinator.

The mechanisms to allow devices to acquire and maintain synchronization with a coordinator are described in 5.1.5. Synchronization on a beacon-enabled VPAN is described after first explaining how a coordinator generates beacon frames. Following this explanation, synchronization on a nonbeacon-enabled VPAN is described. Also described is a procedure to reestablish communication between a device and its coordinator, as it is possible that a device may lose synchronization in the case of either a beacon-enabled or a nonbeacon-enabled VPAN.

This standard has been designed so that application data transfers can be controlled by the devices on a VPAN rather than by the coordinator. The procedures the coordinator uses to handle multiple transactions while preserving this requirement are described in 5.1.6.

The mechanisms for transmitting, receiving, and acknowledging frames, including frames sent using indirect transmission, are described in 5.1.7. In addition, methods for retransmitting frames are also described.

The mechanisms for allocating and deallocating a GTS are described in 5.1.8. The deallocation process may result in the fragmentation of the GTS space, i.e., an unused slot or slots. The subclause describes a mechanism to resolve fragmentation.

The MAC sublayer uses the mechanisms described in clause 7 for all incoming and outgoing frames.

Throughout this subclause, the receipt of a frame is defined as the successful receipt of the frame by the PHY and the successful verification of the FCS by the MAC sublayer, as described in 5.2.1.9.

The mechanisms to allow devices to recover quickly in case of temporary interference using a fast link recovery process are defined in 5.1.9. The fast link recovery process also enables devices to save power by letting the infrastructure initiate the link recovery.

The mechanisms to allow devices to use multiple channels in case of limited time resources or interference are defined in 5.1.10. Multiple channel resource assignment uses information about multiple channel support and band hopping in order to support more users or improve performance.

The mechanisms to support mobility of the device under an infrastructure that supports multiple optical elements over a wide coverage area are defined in 5.1.11. The concept of a cell is introduced and the support for mobility across multiple cells supported by the infrastructure is presented.

The mechanisms to visually indicate to the user the various states using various colors are defined in 5.1.12. The various states such as device discovery (scan, association, disassociation), file transfer status, wavelength quality indication and acknowledgments can be visually indicated to the user to help with device alignment for communication.

The mechanisms to stabilize the optical color emitted by the transmitter are defined in 5.1.13. The CVD frames are used to estimate the change in color and this information can be provided as feedback to the transmitter to stabilize its color.

The mechanisms for using the visibility and dimming information in the MAC are defined in 5.1.14. Features such as an extended preamble mode for providing visibility with improved synchronization performance, dimming overrides, adjusting the MAC layer transmission schedule to accommodate dimming, association and link adaptation in the presence of dimming are provided.

5.1.1 Channel access

This subclause describes the mechanisms for accessing the physical optical channel. The standard provides a single VLC MAC frame structure that can be configured for multiple modes. The frame is composed of a variable number of slots. A slot can be defined as the minimum time needed to communicate to send the smallest data to a device and is fixed.

5.1.1.1 Superframe structure

A coordinator on a VPAN 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 inactive portion. The coordinator may enter a low-power (sleep) mode during the inactive portion.

The structure of this superframe is described by the values of *macBeaconOrder* and *macSuperframeOrder*. The MAC PIB attribute *macBeaconOrder*, describes the interval at which the coordinator shall transmit its beacon frames. The value of *macBeaconOrder*, *BO*, and the beacon interval, *BI*, are related as follows: for $0 \leq BO \leq 14$, $BI = aBaseSuperframeDuration * 2^{BO}$ optical clocks. If $BO = 15$, the coordinator shall not transmit beacon frames except when requested to do so, such as on receipt of a beacon request command. The value of *macSuperframeOrder* shall be ignored if $BO = 15$.

The MAC PIB attribute *macSuperframeOrder* describes the length of the active portion of the superframe, which includes the beacon frame. The value of *macSuperframeOrder*, *SO*, and the superframe duration, *SD*, are related as follows: for $0 \leq SO \leq BO \leq 14$, $SD = aBaseSuperframeDuration * 2^{SO}$ optical clocks. If $SO = 15$, the superframe shall not remain active after the beacon. If $BO = 15$, the superframe shall not exist (the value of *macSuperframeOrder* shall be ignored), and *macRxOnWhenIdle* shall define whether the receiver is enabled during periods of transceiver inactivity.

The active portion of each superframe shall be divided into *aNumSuperframeSlots* equally spaced slots of duration $2^{SO} * aBaseSlotDuration$ and is composed of three parts: a beacon, a CAP and a CFP. The beacon shall be transmitted, without the use of any random access, at the start of slot 0, and the CAP shall commence immediately following the beacon. The start of slot 0 is defined as the point at which the first bit of the beacon PPDU is transmitted. The CFP, if present, follows immediately after the CAP and extends to the end of the active portion of the superframe. Any allocated GTSs shall be located within the CFP.

The MAC sublayer shall ensure that the integrity of the superframe timing is maintained, e.g., compensating for clock drift error.

VPANs that wish to use the superframe structure (referred to as a beacon-enabled VPAN) shall set *macBeaconOrder* to a value between 0 and 14, both inclusive, and *macSuperframeOrder* to a value between 0 and the value of *macBeaconOrder*, both inclusive.

VPANs that do not wish to use the superframe structure (referred to as a nonbeacon-enabled VPAN) shall set both *macBeaconOrder* and *macSuperframeOrder* to 15. In this case, a coordinator shall not transmit beacons, except upon receipt of a beacon request command; all transmissions, with the exception of acknowledgment frames and any data frame that quickly follows the acknowledgment of a data request command, see 5.1.7.3, shall use an unslotted random access mechanism to access the channel. In addition, GTSs shall not be permitted.

An example of a superframe structure is shown in Figure 13. In this case, the beacon interval, *BI*, is twice as long as the active superframe duration, *SD*, and the CFP contains two GTSs.

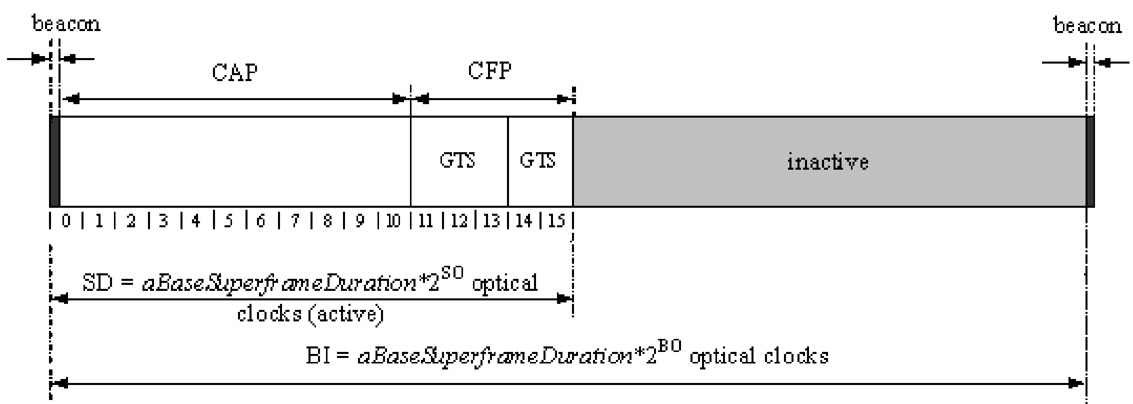

Figure 13—An example of the superframe structure

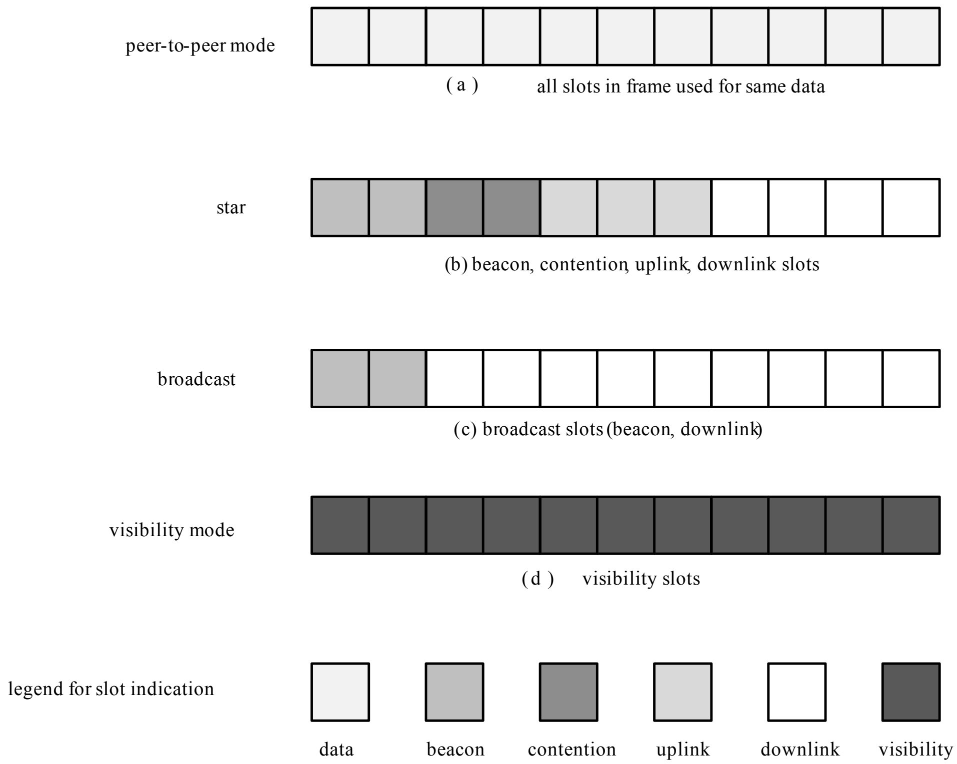

Figure 14—Example usage of frame structure for multiple topologies

Figure 14 provides an example usage of frame structure configuration for multiple topologies such as peer-to-peer, star, broadcast and visibility modes. The beacon slots are used for the beacons and the contention slots are used in the CAP period. The uplink and downlink GTS slots are used in the CFP periods. Visibility or idle patterns can be sent in the visibility slots during idle or RX modes of the infrastructure to ensure continuous output and mitigate flicker and are also used for point-and-shoot mode to ensure visibility.

5.1.1.1.1 Contention access period (CAP)

The CAP shall start immediately following the beacon and complete before the beginning of 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* optical clocks, unless additional space is needed to temporarily accommodate the increase in the beacon frame length needed to perform GTS maintenance (see 5.2.2.1.3) and shall shrink or grow dynamically to accommodate the size of the CFP.

All frames, except acknowledgment frames and any data frame that quickly follows the acknowledgment of a data request command (see 5.1.7.3) transmitted in the CAP shall use a slotted random access mechanism to access the channel. A device transmitting within the CAP shall ensure that its transaction is complete (i.e., including the reception of any acknowledgment) one IFS period (see 5.1.1.2) before the end of the CAP. If this is not possible, the device shall defer its transmission until the CAP of the following superframe.

5.1.1.1.2 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 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 of the combined GTSs. Communication between devices can take a variable number of slots. A single device or user can have access to more than a single slot for sustained data transfer in the frame, if there are slots available.

No transmissions within the CFP shall use a unslotted random access mechanism to access the channel. A device transmitting in the CFP shall ensure that its transmissions are complete one IFS period (see 5.1.1.2) before the end of its GTS.

5.1.1.1.3 Visibility support during channel access

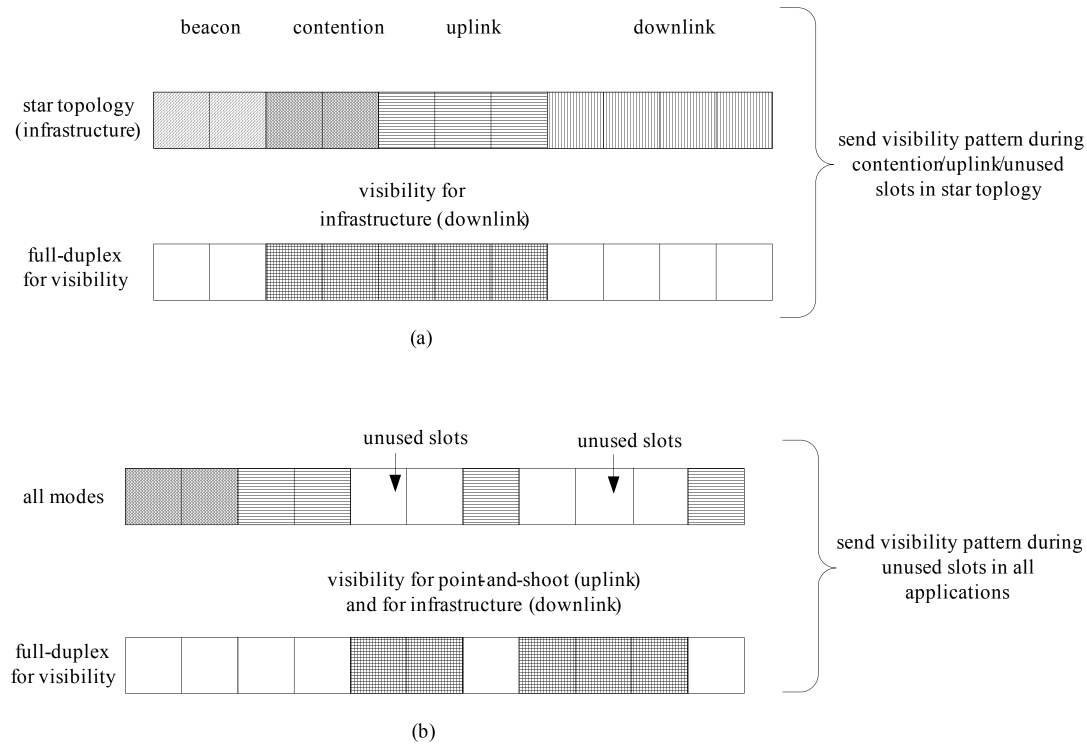

Figure 15—Usage of CVD frames during idle or RX modes of operation

The visibility slots can be used during contention and uplink slots in star topology mode and unused slots in all modes to maintain visibility, reduce flicker and keep the transmitter always "ON" for the infrastructure. This is shown in Figure 15. Visibility support is a very important distinguishing feature for VLC. One may need to transmit idle patterns during receive and idle modes. This can be done by simultaneous reception of data and the transmission of visibility or idle patterns. This is possible due to spatial separation of the light source and the receiving circuitry. As shown in Figure 15, idle patterns are sent during contention, uplink slots and unused downlink slots by the infrastructure to maintain visibility. Idle patterns are also sent during unused slots by the mobile device to help with pointing and alignment for optimal data transfer.

If the continuous visibility bit is set in the capabilities field shown in Table 16, then infrastructure devices shall provide continuous visibility.

5.1.1.2 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; if the first transmission requires an acknowledgment, the separation between the acknowledgment frame and the second transmission shall be at least an IFS period. The length of the IFS period is dependent on the size of the frame that has just been transmitted. Frames (i.e., MPDUs) of up to *aMaxSIFSFrameSize* octets in length shall be followed by a SIFS period of a duration of at least *macMinSIFSPeriod* optical clocks. Frames (i.e., MPDUs) with lengths greater than *aMaxSIFSFrameSize* octets shall be followed by a LIFS period of a duration of at least *macMinLIFSPeriod* optical clocks. Burst frames shall have an RIFS of exactly *macMinRIFSPeriod*. The IFS for the different modes are defined in 8.3.4 and the concepts are illustrated in Figure 16.

The slotted random access algorithm shall take this requirement into account for transmissions in the CAP.

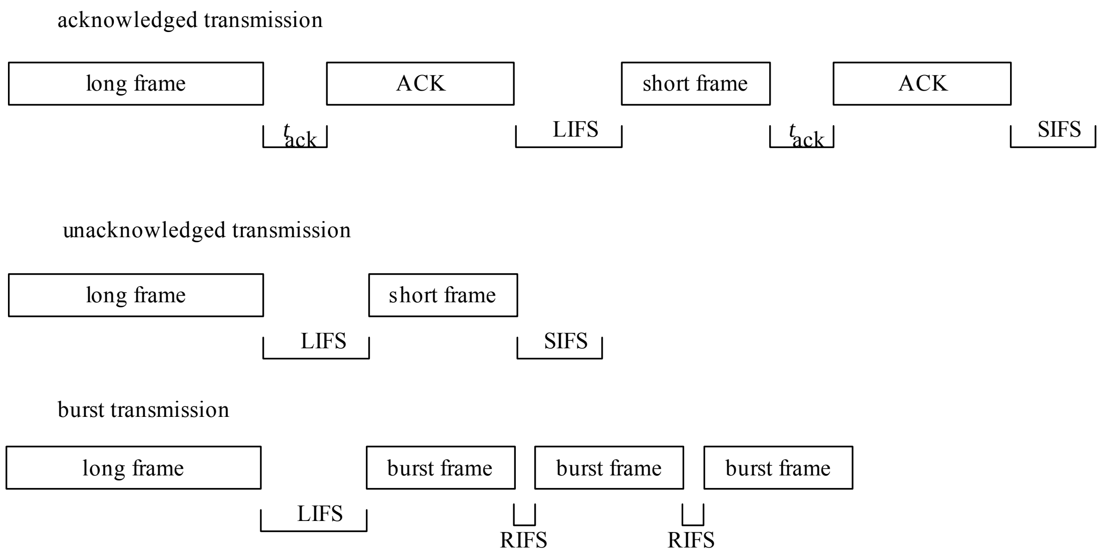

Figure 16—Interframe Spacing

5.1.1.3 Random access algorithm

The slotted random access 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 (as defined in 5.1.7.3 for timing requirements). None of the random access algorithms shall be used for the transmission of beacon frames in a beacon-enabled VPAN, acknowledgment frames, or data frames transmitted in the CFP.

If periodic beacons are being used in the VPAN, the MAC sublayer shall employ the slotted version of the random access algorithm for transmissions in the CAP of the superframe. Conversely, if periodic beacons are not being used in the VPAN or if a beacon could not be located in a beacon-enabled VPAN, the MAC sublayer shall transmit using the unslotted version of the random access algorithm. In both cases, the algorithm is implemented using units of time called backoff periods, where one backoff period shall be equal to *aUnitBackoffPeriod* optical clocks.

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

Each device shall maintain two variables for each transmission attempt: *NB* and *BE*. *NB* is the number of times the access algorithm was required to backoff while attempting the current transmission; this value shall be initialized to zero before each new transmission attempt. The variable *BE* is the backoff exponent, which is related to how many backoff periods a device shall wait before attempting to access/assess a channel. *BE* shall be initialized to the value of *macMinBE*.

Figure 17 illustrates the steps of the access algorithm. The MAC sublayer shall first initialize NB and BE for slotted random access then locate the boundary of the next backoff period. The MAC sublayer shall delay for a random number of complete backoff periods in the range 0 to $2^{BE} - 1$ [step (2)] and then request that the PHY perform a transmission or optionally a CCA. In a slotted random access system, the transmission, or CCA if active, shall start on a backoff period boundary. In an unslotted system, the transmission, or CCA if active, shall start immediately.

In a slotted random access system, the MAC sublayer shall ensure that, after the random backoff, the remaining slotted random access operations can be undertaken and the entire transaction can be transmitted before the end of the CAP. Note that any bit padding used by the supported PHY shall be considered in making this determination. If the number of backoff periods is greater than the remaining number of backoff periods in the CAP, the MAC sublayer shall pause the backoff countdown at the end of the CAP and resume it at the start of the CAP in the next superframe. If the number of backoff periods is less than or equal to the remaining number of backoff periods in the CAP, the MAC sublayer shall apply its backoff delay and then evaluate whether it can proceed. The MAC sublayer shall proceed if the remaining unslotted random access algorithm steps, the frame transmission, and any acknowledgment can be completed before the end of the CAP. If the MAC sublayer can proceed and CCA is active, it shall request that the PHY perform the CCA in the current superframe. If the MAC sublayer cannot proceed, it shall wait until the start of the CAP in the next superframe and apply a further random backoff delay before evaluating whether it can proceed again.

If CCA is active and the channel is assessed to be busy, the MAC sublayer shall increment both *NB* and *BE* by one, ensuring that *BE* shall be no more than *macMaxBE*. If the value of *NB* is less than or equal to *macMaxRABackoffs*, the access algorithm shall return to perform a random backoff as shown in Figure 17. If the value of *NB* is greater than *macMaxRABackoffs*, the access algorithm shall terminate with a channel access failure status.

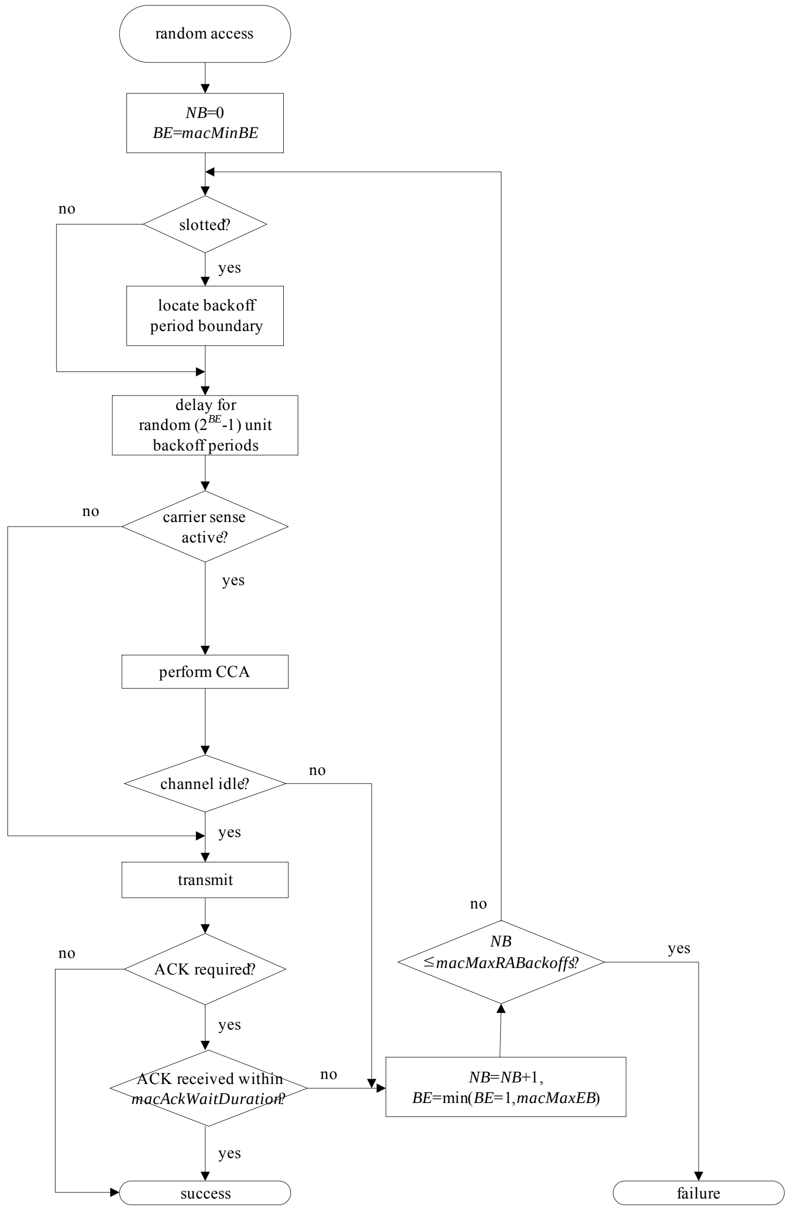

Figure 17—Random Access Flowchart

5.1.2 Starting a VPAN

5.1.2.1 Scanning through channels

All devices shall be capable of performing passive scans across a specified list of channels. In addition, a coordinator shall be able to perform active scans. A device is instructed to begin a channel scan through the MLME-SCAN.request primitive. For the duration of the scan, the device shall suspend beacon transmissions, if applicable, and shall only accept frames received over the PHY data service that are relevant to the scan being performed. Upon the conclusion of the scan, the device shall recommence beacon transmissions. The results of the scan shall be returned via the MLME-SCAN.confirm primitive.

5.1.2.1.1 Active channel scan

It is anticipated that the active channel scan is used with the peer-to-peer topology.

An active scan allows a device to locate any coordinator transmitting beacon frames within its coverage area. This could be used by a prospective VPAN coordinator to select a VPAN identifier prior to starting a new VPAN, 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. If a beacon frame is received that contains the address of the scanning device in its list of pending addresses, the scanning device shall not attempt to extract the pending data.

Before commencing an active scan, the MAC sublayer shall store the value of *macVPANId* and then set it to 0xffff for the duration of the scan. This enables the receive filter to accept all beacons rather than just the beacons from its current VPAN, see 5.1.7.2. On completion of the scan, the MAC sublayer shall restore the value of *macVPANId* to the value stored before the scan began. An active scan over a specified set of logical channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate an active scan. For each logical channel, the device shall first switch to the channel, by setting *phyCurrentChannel* accordingly, and send a beacon request command, see 5.3.6. Upon successful transmission of the beacon request command, the device shall enable its receiver for at most [*aBaseSuperframeDuration* × (2n + 1)] optical clocks, where n is the value of the ScanDuration parameter. During this time, the device shall reject all nonbeacon frames and record the information contained in all unique beacons in a VPAN descriptor structure, see 6.3.3.1, including the channel information and the preamble code. If a beacon frame is received when *macAutoRequest* is set to TRUE, the list of VPAN descriptor structures shall be stored by the MAC sublayer until the scan is complete; at this time, the list shall be sent to the next higher layer in the VPANDescriptorList parameter of the MLME-SCAN.confirm primitive. A device shall be able to store between one and an implementation-specified maximum number of VPAN descriptors. A beacon frame shall be assumed to be unique if it contains both a VPAN identifier and a source address that has not been seen before during the scan of the current channel. If a beacon frame is received when *macAutoRequest* is set to FALSE, each recorded VPAN descriptor is sent to the next higher layer in a separate MLME-BEACON-NOTIFY. indication primitive. A received beacon frame containing one or more octets of payload shall also cause the VPAN descriptor to be sent to the next higher layer via the MLME-BEACON-NOTIFY.indication primitive. If a protected beacon frame is received, i.e., the Security Enabled subfield in the frame control field is set to one, the device shall attempt to unsecure the beacon frame using the unsecuring process described in 7.2.3. The security-related elements of the VPAN descriptor corresponding to the beacon, see 6.3.3.1, shall be set to the corresponding parameters returned by the unsecuring process. The SecurityFailure element of the VPAN descriptor shall be set to SUCCESS if the status from the unsecuring process is SUCCESS and set to one of the other status codes indicating an error in the security processing otherwise. The information from the unsecured frame shall be recorded in the VPAN descriptor even if the status from the unsecuring process indicated an error. If a coordinator of a beacon-enabled VPAN 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 VPAN receives this command, it shall transmit a single beacon frame using unslotted random access or unslotted CSMA-CA.

If *macAutoRequest* is set to TRUE, the active scan on a particular channel shall terminate when the number of beacons found equals the implementation-specified limit or the channel has been scanned for the full time, as specified in 5.1.2.1.1. If *macAutoRequest* is set to FALSE, the active scan on a particular channel shall terminate when the channel has been scanned for the full time. If a channel was not scanned for the full time, it shall be considered to be unscanned.

If *macAutoRequest* is set to TRUE, the entire scan procedure shall terminate when the number of VPAN descriptors stored equals the implementation-specified maximum or every channel in the set of available channels has been scanned. If *macAutoRequest* is set to FALSE, the entire scan procedure shall only terminate when every channel in the set of available channels has been scanned.

5.1.2.1.2 Passive channel scan

It is anticipated that the passive channel scan is used with the star or broadcast topology.

A passive scan, like an active scan, allows a device to locate any coordinator transmitting beacon frames within its coverage area. The beacon request command, however, is not transmitted. This type of scan could be used by a device prior to association. During a passive scan, the MAC sublayer shall discard all frames received over the PHY data service that are not beacon frames. If a beacon frame is received that contains the address of the scanning device in its list of pending addresses, the scanning device shall not attempt to extract the pending data.

Before commencing a passive scan, the MAC sublayer shall store the value of *macVPANId* and then set it to 0xffff for the duration of the scan. This enables the receive filter to accept all beacons rather than just the beacons from its current VPAN, see 5.1.7.2. On completion of the scan, the MAC sublayer shall restore the value of *macVPANId* to the value stored before the scan began. A passive scan over a specified set of logical channels is requested using the MLME-SCAN.request primitive with the ScanType parameter set to indicate a passive scan. For each logical channel, the device shall first switch to the channel, by setting *phyCurrentChannel* accordingly, and then enable its receiver for at most [*aBaseSuperframeDuration* × (2n + 1)] optical clocks, where n is the value of the ScanDuration parameter. During this time, the device shall reject all non-beacon frames and record the information contained in all unique beacons in a VPAN descriptor structure, see 6.3.3.1. If a beacon frame is received when *macAutoRequest* is set to TRUE, the list of VPAN descriptor structures shall be stored by the MAC sublayer until the scan is complete; at this time, the list shall be sent to the next higher layer in the VPANDescriptorList parameter of the MLME-SCAN.confirm primitive. A device shall be able to store between one and an implementation- specified maximum number of VPAN descriptors. A beacon frame shall be assumed to be unique if it contains both a VPAN identifier and a source address that has not been seen before during the scan of the current channel. If a beacon frame is received when *macAutoRequest* is set to FALSE, each recorded VPAN descriptor is sent to the next higher layer in a separate MLME-BEACON-NOTIFY. indication primitive. Once the scan is complete, the MLME-SCAN.confirm shall be issued to the next higher layer with a null VPANDescriptorList. A received beacon frame containing one or more octets of payload shall also cause the VPAN descriptor to be sent to the next higher layer via the MLME-BEACON-NOTIFY. indication primitive.

If a protected beacon frame is received (i.e., the Security Enabled subfield in the frame control field is set to one), the device shall attempt to unsecure the beacon frame using the unsecuring process described in 7.2.3.

The security-related elements of the VPAN descriptor corresponding to the beacon, as shown in 6.3.3.1, shall be set to the corresponding parameters returned by the unsecuring process. The SecurityFailure element of the VPAN descriptor shall be set to SUCCESS if the status from the unsecuring process is SUCCESS and set to one of the other status codes indicating an error in the security processing otherwise.

The information from the unsecured frame shall be recorded in the VPAN descriptor even if the status from the unsecuring process indicated an error.

If *macAutoRequest* is set to TRUE, the passive scan on a particular channel shall terminate when the number of beacons found equals the implementation specified limit or the channel has been scanned for the full time. If *macAutoRequest* is set to FALSE, the passive scan on a particular channel shall terminate when the channel has been scanned for the full time. If a channel was not scanned for the full time, it shall be considered to be unscanned.

If *macAutoRequest* is set to TRUE, the entire scan procedure shall terminate when the number of VPAN descriptors stored equals the implementation-specified maximum or every channel in the set of available channels has been scanned. If *macAutoRequest* is set to FALSE, the entire scan procedure shall terminate only when every channel in the set of available channels has been scanned.

5.1.2.2 VPAN initiation

The broadcast mode does not have any requirements for starting a VPAN. Capability exchange should occur for all bi-directional communication during device discovery. If a device supports multiple transmit color channels, it can exchange the WQI metrics for channel selection. There is no channel selection process requirement if the device supports only a single color channel. For a star topology, the coordinator establishes the VPAN by sending beacon frames. For peer-to-peer topology, a device can either send an association or active scan command to initiate communication with the peer device.

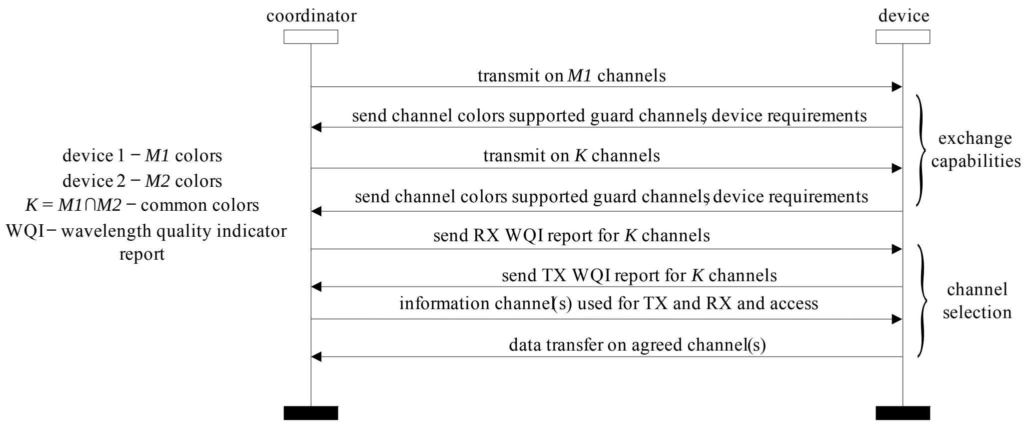

Figure 18—VPAN initiation

Let device 1 support *M1* color channels and let device 2 support *M2* color channels. Let *K* be the number of channels shared between device 1 and device 2, where $K \geq 1$ for communication. For a peer-to-peer network, the first device, which may be the device or coordinator, initiates the communications and transmits on all supported *M1* channels. If there is independent hardware for each color at the transmitter and receiver, parallel transmissions are possible as long as guard color channels are not used for any particular color choice. Each device communicates the capabilities of each device and application requirements via the MAC and PHY capabilities IE provided. The MAC also reports the number of supported aggregated channels and the associated guard colors for each channel. Next, the other device attempts to receive and synchronize on all *K* channels shared between the devices. However, it may be able to receive on only 'x' channels, where $1 \leq x \leq K$, due to interference with other light sources. The second device shall receive on at least one channel in order to communicate. The *K* channels and device capabilities are obtained from the mentioned information. Based on the interference energy from ambient light and the energy received during transmission, a WQI is calculated for all *K* channels. The second device then transmits on all *K* common channels to the first device. The second device also provides its supported channels, guard channels and application requirements as part of its capabilities information exchange. Next, the first device attempts to receive and synchronize on all *K* channels. It may receive on only 'y' channels, where $1 \leq y \leq K$, due to interference. Since VLC is very directional, it is possible that '*x*' and '*y*' may be different. For example, if first device is closer to a window, it may receive more ambient light interference than the second device. The first device calculates its RX WQI for all *K* channels as well and transmits the WQI report back to the second device

Simultaneously, the second device calculates the WQI metrics based on the received information from the first device. Channels where reception is not possible or where other piconets are known to operate by the second device will be tagged unusable with a reception WQI of 0. The second device then reports this RX WQI for all *K* channels back to the first device

The initiating device collects the information for the transmission such as the transmission and reception capabilities of the two devices, the WQI reports, the selected guard color channels for each channel and the requirements of the application. Based on this information, the first device determines a single or multiple channels for communication. The first device then reports the communication channels to the second device. Thus, at the end of this exchange, both devices have an estimate of the WQI for their transmissions that is most suitable for reception at the other end. From that point, both devices can communicate on the agreed channel or channels.

The support for WQI (wavelength quality indication) is provided in the PHY and shall be passed to the MAC via the PD-SAP interface as shown in Table 97.

For a star topology network, the coordinator acts as the initiator for device discovery and association and uses the CAP for association requests and the beacon/management frames to broadcast its association grants.

Starting a VPAN is only applicable to bi-directional communication modes and not for broadcasting.

5.1.2.3 Beacon generation

A device shall be permitted to transmit beacon frames only if *macShortAddress* is not equal to 0xffff.

A coordinator shall use the MLME-START.request primitive to begin transmitting beacons only if the BeaconOrder parameter is less than 15. The coordinator may begin beacon transmission either as the coordinator of a new VPAN or as a device on a previously established VPAN, depending upon the setting of the VPANCoordinator parameter, as shown in 6.3.11.1. The coordinator shall begin beacon transmission on a previously established VPAN only once it has successfully associated with that VPAN.

For the coordinator (i.e., the VPANCoordinator parameter is set to TRUE), the MAC sublayer shall ignore the StartTime parameter and begin beacon transmissions immediately. Setting the StartTime parameter to zero shall also cause the MAC sublayer to begin beacon transmissions immediately. If not acting as the coordinator and the StartTime parameter is nonzero, the time to begin beacon transmissions shall be calculated using the following method. The StartTime parameter, which is rounded to a backoff slot boundary, shall be added to the time, obtained from the local clock, when the MAC sublayer receives the beacon of the coordinator through which it is associated. The MAC sublayer shall then begin beacon transmissions when the current time, obtained from the local clock, equals the number of calculated optical clocks. In order for the beacon transmission time to be calculated by the MAC sublayer, the MAC sublayer shall first track the beacon of the coordinator through which it is associated. If the MLME-START.request primitive is issued with a nonzero StartTime parameter and the MAC sublayer is not currently tracking the beacon of its coordinator, the MLME shall not begin beacon transmissions but shall instead issue the MLME-START.confirm primitive with a status of TRACKING_OFF.

If a device misses between one and (*aMaxLostBeacons*–1) consecutive beacon frames from its coordinator, the device shall continue to transmit its own beacons based on both *macBeaconOrder* (see 5.1.3.5) and its local clock. If the device then receives a beacon frame from its coordinator and, therefore, does not lose synchronization, the device shall resume transmitting its own beacons based on the StartTime parameter and the incoming beacon. If a device does lose synchronization with its coordinator, the MLME of the device shall issue the MLME-SYNC-LOSS.indication primitive to the next higher layer and immediately stop transmitting its own beacons. The next higher layer may, at any time following the reception of the MLME-SYNC-LOSS.indication primitive, resume beacon transmissions by issuing a new MLME-START.request primitive.

On receipt of the MLME-START.request primitive, the MAC sublayer shall set the VPAN identifier in *macVPANId* and use this value in the Source VPAN Identifier field of the beacon frame. The address used in the Source Address field of the beacon frame shall contain the value of *aExtendedAddress* if *macShortAddress* is equal to 0xfffe or *macShortAddress* otherwise.

The time of transmission of the most recent beacon shall be recorded in *macBeaconTxTime* and shall be computed so that its value is taken at the same position in each beacon frame, the location of which is implementation specific. The position, which is specified by the *macTimeStampOffset* attribute, is the same as that used in the timestamp of the incoming beacon frame, as described in 5.1.5.1.

All beacon frames shall be transmitted at the beginning of each superframe at an interval equal to *aBaseSuperframeDuration* $* 2^n$ optical clocks, where n is the value of *macBeaconOrder* (the construction of the beacon frame is specified in 5.2.2.1).

Beacon transmissions shall be given priority over all other transmit and receive operations.

5.1.2.4 Device discovery

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

Device discovery shall be performed at 11.67 Kbps with a 200 KHz optical clock for PHY I and at 1.25 Mbps with a 3.75 MHz optical clock for PHY II. PHY III does not provide device discovery support and shall rely on device discovery using PHY II before operating in that mode. The dimmed OOK mode can be used to support dimming in the device discovery process. This mode is indicated using the MAC PIB attribute, *macUseDimmedOOKmode*, as defined in Table 60. The MAC and PHY capabilities are exchanged in the device discovery process. The clock rate support capabilities are also exchanged. Once the capabilities are exchanged, regular data transmission mode resumes for all three PHY types. Device discovery requires bi-directional communication and is not applicable for broadcasting.

5.1.2.5 Guard and aggregation color channels

The bandplan provides support for seven logical channels in the MAC. However, in order to support association without knowledge of receiver capabilities and to support unidirectional broadcasting, the VLC receiver shall support reception on the entire visible light spectrum with any type of optical light source.

Channel aggregation is used to indicate optical sources that span multiple (>1) bands in the proposed bandplan and are intentionally transmitting on multiple bands due to the choice of optical light source. Guard channels are used to indicate optical sources that unintentionally leak into other bands, whose information can be discarded at the receiver for better performance.

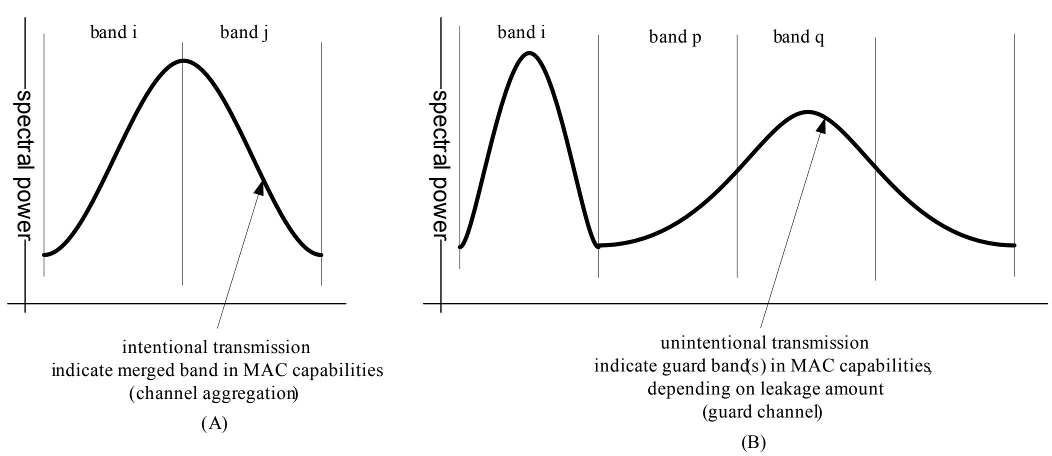

Figure 19—Concept of aggregation channel and guard channel

If multiple bands are aggregated or multiple optical sources are transmitting simultaneously, the same data shall be sent on all optical sources during the preamble and header during device discovery since it is not known what the receiver capabilities are. The details on channel aggregation and guard channel support are provided in the PHY capabilities information element of the MAC. The criterion used for defining a guard color channel or aggregated channel is based on out-of-band leakage exceeding 20 dB over maximum in-channel value. The transmitting device shall indicate channel aggregation and guard channel support using the PHY capabilities during device discovery and association for bi-directional communication modes.

5.1.3 Maintaining VPANs

In some instances a situation could occur in which two VPANs exist in the same operating space with the same VPAN identifier. If this conflict happens, the coordinator and its devices shall perform the VPAN identifier conflict resolution procedure.

This procedure is optional for a device.

5.1.3.1 Detection

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

— A beacon frame is received by the VPAN coordinator with the VPAN Coordinator subfield, see 5.2.2.1.2, set to one and the VPAN identifier equal to *macVPANId*.

— A VPAN ID conflict notification command, see 5.3.5, is received by the VPAN coordinator from an associated device on its VPAN.

A device that is associated through the VPAN coordinator (i.e., *macAssociatedVPANCoord* is set to TRUE) shall conclude that a VPAN identifier conflict is present if the following applies:

— A beacon frame is received by the device with the VPAN Coordinator subfield set to one, the VPAN identifier equal to *macVPANId*, and an address that is equal to neither *macCoordShortAddress* nor *macCoordExtendedAddress*.

5.1.3.2 Resolution

On the detection of a VPAN identifier conflict by a device, it shall generate the VPAN ID conflict notification command, defined in 5.3.5, and send it to its coordinator. Since the VPAN ID conflict notification command contains an acknowledgment request (see 5.3.3.1), the coordinator shall confirm its receipt by sending an acknowledgment frame. Once the device has received the acknowledgment frame from the coordinator, the MLME shall issue an MLME-SYNC-LOSS.indication primitive with the LossReason parameter set to VPAN_ID_CONFLICT. If the device does not receive an acknowledgment frame, the MLME shall not inform the next higher layer of the VPAN identifier conflict.

On the detection of a VPAN identifier conflict by the coordinator, the MLME shall issue an MLME-SYNC-LOSS.indication to the next higher layer with the LossReason parameter set to VPAN_ID_CONFLICT. The next higher layer of the coordinator may then perform an active scan and, using the information from the scan, select a new VPAN identifier. The algorithm for selecting a suitable VPAN identifier is out of the scope of this standard. If the next higher layer does select a new VPAN identifier, it may then issue an MLME-START.request with the CoordRealignment parameter set to TRUE in order to realign the VPAN, as described in 5.1.3.3.

5.1.3.3 Realigning a VPAN

If a coordinator receives the MLME-START.request primitive (see 6.3.11.1) with the CoordRealignment parameter set to TRUE, the coordinator shall attempt to transmit a coordinator realignment command containing the new parameters for VPANId, LogicalChannel.

When the coordinator is already transmitting beacons and the CoordRealignment parameter is set to TRUE, the next scheduled beacon shall be transmitted on the current channel using the current superframe configuration, with the frame pending subfield of the frame control field set to one. Immediately following the transmission of the beacon, the coordinator realignment command shall also be transmitted on the current channel using unslotted random access.

When the coordinator is not already transmitting beacons and the CoordRealignment parameter is set to TRUE, the coordinator realignment command shall be transmitted immediately on the current channel using unslotted random access.

If the transmission of the coordinator realignment command fails due to a channel access failure, the MLME shall notify the next higher layer by issuing the MLME-START.confirm primitive with a status of CHANNEL_ACCESS_FAILURE. The next higher layer may then choose to issue the MLME-START.request primitive again.

Upon successful transmission of the coordinator realignment command, the new superframe configuration and channel parameters shall be put into operation as described in 5.1.3.5 at the subsequent scheduled beacon, or immediately if the coordinator is not already transmitting beacons, and the MAC sublayer shall issue the MLME-START.confirm primitive with a status of SUCCESS.

5.1.3.4 Realignment in a VPAN

If a device has received the coordinator realignment command (see 5.3.7) from the coordinator through which it is associated, the MLME shall issue the MLME-SYNC-LOSS.indication primitive with the LossReason parameter set to REALIGNMENT and the VPANId, LogicalChannel, and the security-related parameters set to the respective fields in the coordinator realignment command. The next higher layer of a coordinator may then issue an MLME-START.request primitive with the CoordRealignment parameter set to TRUE. The next higher layer of a device that is not a coordinator may instead change the superframe configuration or channel parameters through use of the MLME-SET.request primitive.

5.1.3.5 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 VPAN. If *macBeaconOrder* is less than 15, the MAC sublayer will set *macSuperframeOrder* to the value of the SuperframeOrder parameter. The MAC sublayer shall also update *macVPANId* with the value of the VPANId parameter and update *phyCurrentChannel* with the values of the LogicalChannel parameters by issuing the PLME-SET.request primitive.

5.1.4 Association and disassociation

This subclause specifies the procedures for association and disassociation.

5.1.4.1 Association

A device shall attempt to associate only after having first performed a MAC sublayer reset, by issuing the MLME-RESET.request primitive with the SetDefaultPIB parameter set to TRUE, and then having completed either an active channel scan, see 5.1.2.1.1, or a passive channel scan as shown in 5.1.2.1.2. The results of the channel scan would have then been used to choose a suitable VPAN. The algorithm for selecting a suitable VPAN with which to associate from the list of VPAN descriptors returned from the channel scan procedure is out of the scope of this standard.

Following the selection of a VPAN with which to associate, the next higher layers shall request through the MLME-ASSOCIATE.request primitive that the MLME configures the following PHY and MAC PIB attributes to the values necessary for association:

— *phyCurrentChannel* shall be set equal to the LogicalChannel parameter of the MLME-ASSOCIATE.request primitive.
— *macVPANId* shall be set equal to the CoordVPANId parameter of the MLME-ASSOCIATE.request primitive.
— *macCoordExtendedAddress* or *macCoordShortAddress*, depending on which is known from the beacon frame from the coordinator through which it wishes to associate, shall be set equal to the CoordAddress parameter of the MLME-ASSOCIATE.request primitive.

A coordinator shall allow association only if *macAssociationPermit* is set to TRUE. Similarly, a device should attempt to associate only with a VPAN through a coordinator that is currently allowing association, as indicated in the results of the scanning procedure. If a coordinator with *macAssociationPermit* set to FALSE receives an association request command from a device, the command shall be ignored.

In order to optimize the association procedure on a beacon-enabled VPAN, a device may begin tracking the beacon of the coordinator through which it wishes to associate. This is achieved by the next higher layer issuing the MLME-SYNC.request primitive with the TrackBeacon parameter set to TRUE.

A device that is instructed to associate with a VPAN, through the MLME-ASSOCIATE.request primitive, shall try to associate only with an existing VPAN and shall not attempt to start its own VPAN.

The MAC sublayer of an unassociated device shall initiate the association procedure by sending an association request command, see 5.3.1, to the coordinator of an existing VPAN; if the association request command cannot be sent due to a channel access failure, the MAC sublayer shall notify the next higher layer. Because the association request command contains an acknowledgment request (see 5.3.1.1), the coordinator shall confirm its receipt by sending an acknowledgment frame.

The acknowledgment to an association request command does not mean that the device has associated. The next higher layer of the coordinator needs time to determine whether the current resources available on the VPAN are sufficient to allow another device to associate. The next higher layer should make this decision within *macResponseWaitTime* optical clocks. If the next higher layer of the coordinator finds that the device was previously associated on its VPAN, all previously obtained device-specific information should be removed. If sufficient resources are available, the next higher layer should allocate a 16-bit short address to the device, and the MAC sublayer shall generate an association response command, see 5.3.2, containing the new address and a status indicating a successful association. If sufficient resources are not available, the next higher layer of the coordinator should inform the MAC sublayer, and the MLME shall generate an association response command containing a status indicating a failure as shown in Table 11. The association response command shall be sent to the device requesting association using indirect transmission, i.e., the association response command frame shall be added to the list of pending transactions stored on the coordinator and extracted at the discretion of the device concerned using the method described in 5.1.7.3.

If the allocate address subfield of the capability information field (see 5.3.19.1.1) of the association request command is set to one, the next higher layer of the coordinator shall allocate a 16-bit address with a range depending on the addressing mode supported by the coordinator, as described in Table 2. If the Allocate Address subfield of the association request command is set to zero, the 16-bit short address shall be equal to 0xfffe. A short address of 0xfffe is a special case that indicates that the device has associated, but has not been allocated a short address by the coordinator. In this case, the device shall use only its 64-bit extended address to operate on the network.

On receipt of the acknowledgment to the association request command, the device shall wait for at most *macResponseWaitTime* optical clocks for the coordinator to make its association decision; the PIB attribute *macResponseWaitTime* is a network-topology-dependent parameter and may be set to match the specific requirements of the network that a device is trying to join. If the device is tracking the beacon, it shall attempt to extract the association response command from the coordinator whenever it is indicated in the beacon frame. If the device is not tracking the beacon, it shall attempt to extract the association response command from the coordinator after *macResponseWaitTime* optical clocks. If the device does not extract an association response command frame from the coordinator within *macResponseWaitTime* optical clocks, the MLME shall issue the MLME-ASSOCIATE.confirm primitive with a status of NO_DATA, and the association attempt shall be deemed a failure. In this case, the next higher layer shall terminate any tracking of the beacon. This is achieved by issuing the MLME-SYNC.request primitive with the TrackBeacon parameter set to FALSE.

The MLME-ASSOCIATE.response and the subsequent Association response (see 5.3.2) also contain information about what capabilities the device and the coordinator will and will not use during future communication.

Because the association response command contains an acknowledgment request (see 5.3.2.1), the device requesting association shall confirm its receipt by sending an acknowledgment frame. If the Association Status field of the command indicates that the association was successful, the device shall store the address contained in the 16-bit Short Address field of the command in *macShortAddress*; communication on the VPAN using this short address shall depend on its range, as described in Table 2. If the original beacon selected for association following a scan contained the short address of the coordinator, the extended address of the coordinator, contained in the MHR of the association response command frame, shall be stored in *macCoordExtendedAddress.*

Table 2—Usage of the 16-bit short address

Value of *macShortAddress*	Description
0x0000–0xfffd	If a source address is included, the device shall use short source addressing mode for beacon and data frames and the appropriate source addressing mode specified in 5.3 for MAC command frames.
0xfffe	If a source address is included, the device shall use extended source addressing mode for beacon and data frames and the appropriate source addressing mode specified in 5.3 for MAC command frames.
0xffff	The device is not associated and, therefore, shall not perform any data frame communication. The device shall use the appropriate source addressing mode specified in 5.3 for MAC command frames.

If the Association Status field of the command indicates that the association was unsuccessful, the device shall set *macVPANId* to the default value (0xffff).

5.1.4.2 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 VPAN, 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. If TxIndirect is TRUE, the MLME of the coordinator shall send the disassociation notification command to the device using indirect transmission, i.e., the disassociation notification command frame shall be added to the list of pending transactions stored on the coordinator and extracted at the discretion of the device concerned using the method described in 5.1.7.3. If the command frame is not successfully extracted by the device, the coordinator should consider the device disassociated. Otherwise, the MLME shall send the disassociation notification command to the device directly. In this case, if the disassociation notification command cannot be sent due to a channel access failure, the MAC sublayer shall notify the next higher layer.

Because the disassociation command contains an acknowledgment request (see 5.3.3.1), 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.

If an associated device wants to leave the VPAN, the MLME of the device shall send a disassociation notification command to its coordinator. If the disassociation notification command cannot be sent due to a channel access failure, the MAC sublayer shall notify the next higher layer. Because the disassociation command contains an acknowledgment request (see 5.3.3.1), the coordinator shall confirm its receipt by sending an acknowledgment frame. However, even if the acknowledgment is not received, the device should consider itself disassociated.

If the source address contained in the disassociation notification command is equal to *macCoord-ExtendedAddress*, the device should consider itself disassociated. If the command is received by a coordinator and the source is not equal to *macCoordExtendedAddress*, it shall verify that the source address corresponds to one of its associated devices; if so, the coordinator should consider the device disassociated. If none of the conditions for disassocation are satisfied, the command shall be ignored.

An associated device shall disassociate itself by removing all references to the VPAN; the MLME shall set *macVPANId*, *macShortAddress*, *macAssociatedVPANCoord*, *macCoordShortAddress* and *macCoordExtendedAddress* to the default values. The next higher layer of a coordinator should disassociate a device by removing all references to that device.

The next higher layer of the requesting device shall be notified of the result of the disassociation procedure through the MLME-DISASSOCIATE.confirm primitive.

5.1.5 Synchronization

This subclause specifies the procedures for coordinators to generate beacon frames and for devices to synchronize with a coordinator. For VPANs supporting beacons, synchronization is performed by receiving and decoding the beacon frames. For VPANs not supporting beacons, synchronization is performed by polling the coordinator for data.

5.1.5.1 Synchronization with beacons

All devices operating on a beacon-enabled VPAN (i.e., *macBeaconOrder* < 15) shall be able to acquire beacon synchronization in order to detect any pending messages or to track the beacon. Devices shall be permitted to acquire beacon synchronization only with beacons containing the VPAN identifier specified in *macVPANId*. If *macVPANId* specifies the broadcast VPAN identifier (0xffff), a device shall not attempt to acquire beacon synchronization.

A device is instructed to attempt to acquire the beacon through the MLME-SYNC.request primitive. If tracking is specified in the MLME-SYNC.request primitive, the device shall attempt to acquire the beacon and keep track of it by regular and timely activation of its receiver. If tracking is not specified, the device shall either attempt to acquire the beacon only once or terminate the tracking after the next beacon if tracking was enabled through a previous request.

To acquire beacon synchronization, a device shall enable its receiver and search for at most [*aBaseSuperframeDuration* * (2^n + 1)] optical clocks, where *n* is the value of *macBeaconOrder*. If a beacon frame containing the current VPAN identifier of the device is not received, the MLME shall repeat this search. Once the number of missed beacons reaches *aMaxLostBeacons*, the MLME shall notify the next higher layer by issuing the MLME-SYNC-LOSS.indication primitive with a loss reason of BEACON_LOSS.

The MLME shall timestamp each received beacon frame at the same symbol boundary within each frame, the location of which is described by the *macTimeStampOffset* attribute. The position shall be the same as that used in the timestamp of the outgoing beacon frame, stored in *macBeaconTxTime*. The timestamp value shall be that of the local clock of the device at this position. The timestamp is intended to be a relative time measurement that may or may not be made absolute, at the discretion of the implementer.

If a protected beacon frame is received (i.e., the Security Enabled subfield in the frame control field is set to one), the device shall attempt to unsecure the beacon frame using the unsecuring process described in 7.2.3.

If the status from the unsecuring process is not SUCCESS, the MLME shall issue an MLME-COMM-STATUS.indication primitive with the status parameter set to the status from the unsecuring process, indicating the error.

The security-related elements of the VPAN descriptor corresponding to the beacon (see Table 38) shall be set to the corresponding parameters returned by the unsecuring process. The SecurityFailure element of the VPAN descriptor shall be set to SUCCESS if the status from the unsecuring process is SUCCESS and set to one of the other status codes indicating an error in the security processing otherwise.

If a beacon frame is received, the MLME shall discard the beacon frame if the Source Address and the Source VPAN Identifier fields of the MHR of the beacon frame do not match the coordinator source address (*macCoordShortAddress* or *macCoordExtendedAddress*, depending on the addressing mode) and the identifier of the device (*macVPANId*).

If a valid beacon frame is received and *macAutoRequest* is set to FALSE, the MLME shall indicate the beacon parameters to the next higher layer by issuing the MLME-BEACON-NOTIFY.indication primitive. If a beacon frame is received and *macAutoRequest* is set to TRUE, the MLME shall first issue the MLME-BEACON-NOTIFY.indication primitive if the beacon contains any payload. The MLME shall then compare its address with those addresses in the Address List field of the beacon frame. If the Address List field contains the 16-bit short or 64-bit extended address of the device and the source VPAN identifier matches *macVPANId*, the MLME shall follow the procedure for extracting pending data from the coordinator as shown in 5.1.7.3.

If beacon tracking is activated, the MLME shall enable its receiver at a time prior to the next expected beacon frame transmission, i.e., just before the known start of the next superframe. If the number of consecutive beacons missed by the MLME reaches *aMaxLostBeacons*, the MLME shall respond with the MLME-SYNC-LOSS.indication primitive with a loss reason of BEACON_LOST.

5.1.5.2 Synchronization without beacons

All devices operating on a nonbeacon-enabled VPAN (*macBeaconOrder* = 15) 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 as shown in 5.1.7.3.

5.1.6 Transaction handling

Transactions can be instigated from the devices themselves rather than from the coordinator. In other words, either the coordinator needs to indicate in its beacon when messages are pending for devices or the devices themselves need to poll the coordinator to determine whether they have any messages pending. Such transfers are called indirect transmissions.

The coordinator shall begin handling a transaction on receipt of an indirect transmission request either via the MCPS-DATA.request primitive or via a request from the MLME to send a MAC command instigated by a primitive from the next higher layer, such as the MLME-ASSOCIATE.response primitive as shown in 6.3.1.3. On completion of the transaction, the MAC sublayer shall indicate a status value to the next higher layer. If a request primitive instigated the indirect transmission, the corresponding confirm primitive shall be used to convey the appropriate status value. Conversely, if a response primitive instigated the indirect transmission, the MLME-COMM-STATUS.indication primitive shall be used to convey the appropriate status value. The MLME-COMM-STATUS.indication primitive can be related to its corresponding response primitive by examining the Destination Address field.

The information contained in the indirect transmission request forms a transaction, and the coordinator shall be capable of storing at least one transaction. On receipt of an indirect transmission request, if there is no capacity to store another transaction, the MAC sublayer shall indicate to the next higher layer a status of TRANSACTION_OVERFLOW in the appropriate corresponding primitive.

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 *macTransactionPersistenceTime*. If the transaction is not successfully extracted by the appropriate device within this time, the transaction information shall be discarded and the MAC sublayer shall indicate to the next higher layer a status of TRANSACTION_EXPIRED in the appropriate corresponding primitive. In order to be successfully extracted, an acknowledgment shall be received if one was requested.

If the transaction was successful, the transaction information shall be discarded, and the MAC sublayer shall indicate to the next higher layer a status of SUCCESS in the appropriate corresponding primitive.

If the coordinator transmits beacons, it shall list the addresses of the devices to which each transaction is associated in the Address List field and indicate the number of addresses in the Pending Address Specification field of the beacon frame. If the coordinator is able to store more than seven pending transactions, it shall indicate them in its beacon on a first-come-first-served basis, ensuring that the beacon frame contains at most seven addresses. For transactions requiring a GTS, the coordinator shall not add the address of the recipient to its list of pending addresses in the beacon frame. Instead it shall transmit the transaction in the GTS allocated for the device as shown in 5.1.8.3.

On a beacon-enabled VPAN, if there is a transaction pending for the broadcast address, the frame pending subfield of the frame control field in the beacon frame shall be set to one, and the pending message shall be transmitted immediately following the beacon using the unslotted random access algorithm. If there is a second message pending for the broadcast address, its transmission shall be delayed until the following superframe. Only one broadcast message shall be allowed to be sent indirectly per superframe.

On a beacon-enabled VPAN, a device that receives a beacon containing its address in the list of pending addresses shall attempt to extract the data from the coordinator. On a nonbeacon-enabled VPAN, a device shall attempt to extract the data from the coordinator on receipt of the MLME-POLL.request primitive. The procedure for extracting pending data from the coordinator is described in 5.1.7.3. If a device receives a beacon with the frame pending subfield set to one, it shall leave its receiver enabled for up to *macMaxFrameTotalWaitTime* optical clocks to receive the broadcast data frame from the coordinator.

5.1.7 Transmission, reception, and acknowledgment

This subclause describes the fundamental procedures for transmission, reception, and acknowledgment.

5.1.7.1 Transmission

Each device shall store its current DSN value in the MAC PIB attribute *macDSN* and initialize it to a random value; the algorithm for choosing a random number is out of the scope of this standard. 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; the algorithm for choosing a random number is out of the scope of this standard. 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.

It should be noted that both the DSN and BSN are 8-bit values and, therefore, have limited use to the next higher layer (e.g., in the case of the DSN, in detecting retransmitted frames).

The Source Address field, if present, shall contain the address of the device sending the frame. When a device has associated and has been allocated a 16-bit short address (i.e., *macShortAddress* is not equal to 0xfffe or 0xffff), it shall use that address in preference to its 64-bit extended address (i.e., *aExtendedAddress*) wherever possible. When a device has not yet associated to a VPAN or *macShortAddress* is equal to 0xffff, it shall use its 64-bit extended address in all communications requiring the Source Address field. If the Source Address field is not present, the originator of the frame shall be assumed to be the coordinator, and the Destination Address field shall contain the address of the recipient.

The Destination Address field, if present, shall contain the address of the intended recipient of the frame, which may be either a 16-bit short address or a 64-bit extended address. If the Destination Address field is not present, the recipient of the frame shall be assumed to be the coordinator, and the Source Address field shall contain the address of the originator.

If both destination and source addressing information is present, the MAC sublayer shall compare the destination and source VPAN identifiers. If the VPAN identifiers are identical, the VPAN ID Compression subfield of the frame control field shall be set to one, and the source VPAN identifier shall be omitted from the transmitted frame. If the VPAN identifiers are different, the VPAN ID Compression subfield of the frame control field shall be set to zero, and both Destination VPAN Identifier and Source VPAN Identifier fields shall be included in the transmitted frame. If only either the destination or the source addressing information is present, the VPAN ID Compression subfield of the frame control field shall be set to zero, and the VPAN identifier field of the single address shall be included in the transmitted frame.

If the frame is to be transmitted on a beacon-enabled VPAN, the transmitting device shall attempt to find the beacon before transmitting. If the beacon is not being tracked, as shown in 5.1.5.1, and hence the device does not know where the beacon will appear, it shall enable its receiver and search for at most [*aBaseSuperframeDuration* * (2^n + 1)] optical clocks, where *n* is the value of *macBeaconOrder,* in order to find the beacon. If the beacon is not found after this time, the device shall transmit the frame following the successful application of the unslotted version of the random access algorithm as shown in 5.1.1.3. Once the beacon has been found, either after a search or due to its being tracked, the frame shall be transmitted in the appropriate portion of the superframe. Transmissions in the CAP shall follow a successful application of the slotted version of the random access algorithm, see 5.1.1.3, and transmissions in a GTS shall not use any random access.

If the frame is to be transmitted on a nonbeacon-enabled VPAN, the frame shall be transmitted following the successful application of the unslotted version of the random access algorithm as shown in 5.1.1.3.

For either a beacon-enabled VPAN or a nonbeacon-enabled VPAN, if the transmission is direct and originates due to a primitive issued by the next higher layer and the access algorithm fails, the next higher layer shall be notified. If the transmission is indirect and the access algorithm fails, the frame shall remain in the transaction queue until it is requested again and successfully transmitted or until the transaction expires.

The device shall process the frame using the outgoing frame security procedure described in 7.2.1.

If the status from the outgoing frame security procedure is not SUCCESS, the MLME shall issue the corresponding confirm or MLME-COMM-STATUS.indication primitive with the status parameter set to the status from the outgoing frame security procedure, indicating the error.

To transmit the frame, the MAC sublayer shall first enable the transmitter by issuing the PLME-SET-TRX-STATE.request primitive with a state of TX_ON to the PHY. On receipt of the PLME-SET-TRX-STATE.confirm primitive with a status of either SUCCESS or TX_ON, the constructed frame shall then be transmitted by issuing the PD-DATA.request primitive. Finally, on receipt of the PD-DATA.confirm primitive, the MAC sublayer shall disable the transmitter by issuing the PLME-SET-TRX-STATE.request primitive with a state of RX_ON or TRX_OFF to the PHY, depending on whether the receiver is to be enabled following the transmission. In the case where the Acknowledgment Request subfield of the frame control field is set to one, the MAC sublayer shall enable the receiver immediately following the transmission of the frame by issuing the PLME-SET-TRX-STATE.request primitive with a state of RX_ON to the PHY.

5.1.7.2 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 requests 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 completion of each transceiver task, the MAC sublayer shall request that the PHY enables or disables its receiver, depending on the values of *macBeaconOrder* and *macRxOnWhenIdle*. If *macBeaconOrder* is less than 15, the value of *macRxOnWhenIdle* shall be considered relevant only during idle periods of the CAP of the incoming superframe. If *macBeaconOrder* is equal to 15, the value of *macRxOnWhenIdle* shall be considered relevant at all times.

A device with its receiver enabled will be able to receive and decode transmissions from all devices complying with this standard that are currently operating on the same channel and are in its operating space, along with interference from other sources. The MAC sublayer shall, therefore, be able to filter incoming frames and present only the frames that are of interest to the upper layers.

The MAC sublayer shall discard all received frames that do not contain a correct value in their FCS field in the MFR (see 5.2.1.9). The FCS field shall be verified on reception by recalculating the purported FCS over the MHR and MSDU of the received frame and by subsequently comparing this value with the received FCS field. The FCS field of the received frame shall be considered to be correct if these values are the same and incorrect otherwise.

The MAC sublayer shall accept only frames that satisfy all of the following filtering requirements:

— The Frame Type subfield shall not contain a reserved frame type.
— If a destination VPAN identifier is included in the frame, it shall match *macVPANId* or shall be the broadcast VPAN identifier (0xffff).
— If a short destination address is included in the frame, it shall match either *macShortAddress* or the broadcast address (0xffff). 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 VPAN identifier shall match *macVPANId* unless *macVPANId* is equal to 0xffff, in which case the beacon frame shall be accepted regardless of the source VPAN 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 coordinator and the source VPAN identifier matches *macVPANId*.

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

If the VPAN ID Compression subfield of the frame control field is set to one and both destination and source addressing information is included in the frame, the MAC sublayer shall assume that the omitted Source VPAN Identifier field is identical to the Destination VPAN Identifier field.

The device shall process the frame using the incoming frame security procedure described in 7.2.3.

If the status from the incoming frame security procedure is not SUCCESS, the MLME shall issue the corresponding confirm or MLME-COMM-STATUS.indication primitive with the status parameter set to the status from the incoming frame security procedure, indicating the error, and with the security-related parameters set to the corresponding parameters returned by the unsecuring process.

If the valid frame is a data frame, the MAC sublayer shall pass the frame to the next higher layer. This is achieved by issuing the MCPS-DATA.indication primitive containing the frame information. The security-related parameters of the MCPS-DATA.indication primitive shall be set to the corresponding parameters returned by the unsecuring process.

If the valid frame is a MAC command or beacon frame, it shall be processed by the MAC sublayer accordingly, and a corresponding confirm or indication primitive may be sent to the next higher layer. The security-related parameters of the corresponding confirm or indication primitive shall be set to the corresponding parameters returned by the unsecuring process.

5.1.7.3 Extracting pending data from a coordinator

A device on a beacon-enabled VPAN can determine whether any frames are pending for it by examining the contents of the received beacon frame, as described in 5.1.5.1. 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, see 5.3.4, to the coordinator during the CAP with the Acknowledgment Request subfield of the frame control field set to one; the only exception to this is if the beacon frame is received while performing an active or passive scan as shown in 5.1.3.1. There are two other cases for which the MLME shall send a data request command to the coordinator. The first case is when the MLME receives the MLME-POLL.request primitive. In the second case, a device may send a data request command *macResponseWaitTime* optical clocks after the acknowledgment to a request command frame, such as during the association procedure. If the data request is intended for the coordinator, the destination address information may be omitted.

If the data request command originated from an MLME-POLL.request primitive, the MLME shall perform the security process on the data request command based on the SecurityLevel, KeyIdMode, KeySource, and KeyIndex parameters of the MLME-POLL.request primitive, according to 7.2.1. Otherwise, the MLME shall perform the security process on the data request command based on the *macAutoRequestSecurityLevel*, *macAutoRequestKeyIdMode*, *macAutoRequestKeySource,* and *macAutoRequestKeyIndex* PIB attributes, according to 7.2.1.

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 (see 5.1.7.4.2), it shall set the frame pending subfield of the frame control field of 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.

On receipt of the acknowledgment frame with the frame pending subfield set to one, a device shall enable its receiver for at most *macMaxFrameTotalWaitTime* CAP optical clocks in a beacon-enabled VPAN, or in a nonbeacon-enabled VPAN, to receive the corresponding data frame from the coordinator. If there is an actual data frame pending within the coordinator for the requesting device, the coordinator shall send the frame to the device using one of the mechanisms described in this subclause. If there is no data frame pending for the requesting device, the coordinator shall send a data frame without requesting acknowledgment to the device containing a zero length payload, indicating that no data are present, using one of the mechanisms described in this subclause.

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

- Without using slotted random access, if the MAC sublayer can commence transmission of the data frame between *aTurnaroundTime-RX-TX* and (*aTurnaroundTime-RX-TX* + *aUnitBackoffPeriod*) optical clocks, on a backoff slot boundary, and there is time remaining in the CAP for the message, appropriate IFS, and acknowledgment as defined in 9.5.1. 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 slotted random access, otherwise.

If the requesting device does not receive a data frame from the coordinator within *macMaxFrameTotalWaitTime* CAP optical clocks in a beacon-enabled VPAN, or in a nonbeacon-enabled VPAN, 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.

If the frame pending subfield of the frame control field of the data frame received from the coordinator is set to one, the device still has more data pending with the coordinator. In this case it may extract the data by sending a new data request command to the coordinator.

5.1.7.4 Use of acknowledgments and retransmissions

A data or MAC command frame shall be sent with the Acknowledgment Request subfield of its frame control field set appropriately for the frame. A beacon or acknowledgment frame shall always be sent with the Acknowledgment Request subfield set to zero. Similarly, any frame that is broadcast shall be sent with its Acknowledgment Request subfield set to zero.

5.1.7.4.1 No acknowledgment

A frame transmitted with its Acknowledgment Request subfield set to zero shall not be acknowledged by its intended recipient. The originating device shall assume that the transmission of the frame was successful.

The message sequence chart in Figure 20 shows the scenario for transmitting a single frame of data from an originator to a recipient without requiring an acknowledgment. In this case, the originator transmits the data frame with the Acknowledgment Request (AR) subfield of the frame control field equal to zero.

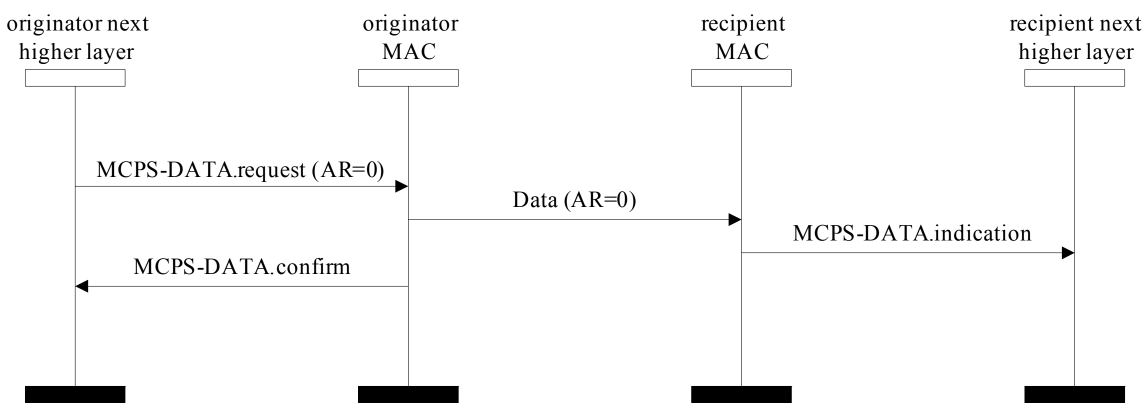

Figure 20—Successful data transmission without an acknowledgment

5.1.7.4.2 Acknowledgment

A frame transmitted with the Acknowledgment Request subfield of its frame control field set to one shall be acknowledged by the recipient. If the intended recipient correctly receives the frame, it shall generate and send an acknowledgment frame containing the same DSN from the data or MAC command frame that is being acknowledged.

The transmission of an acknowledgment frame in a nonbeacon-enabled VPAN or in the CFP shall commence *aTurnaroundTime-RX-TX* optical clocks after the last optical clock of the data or MAC command frame. The transmission of an acknowledgment frame in the CAP shall commence either *aTurnaroundTime-RX-TX* optical clocks after the reception of the last optical clock of the data or MAC command frame or at a backoff slot boundary. In the latter case, the transmission of an acknowledgment frame shall commence between *aTurnaroundTime-RX-TX* and (*aTurnaroundTime-RX-TX* + *aUnitBackoffPeriod*) optical clocks after the reception of the last optical clock of the data or MAC command frame. The constants *aTurnaroundTime-RX-TX* and *aTurnaroundTime-TX-RX* are defined in Table 99.

The message sequence chart in Figure 21 shows the scenario for transmitting a single frame of data from an originator to a recipient with an acknowledgment. In this case, the originator indicates to the recipient that it requires an acknowledgment by transmitting the data frame with the Acknowledgment Request (AR) subfield of the frame control field set to one.

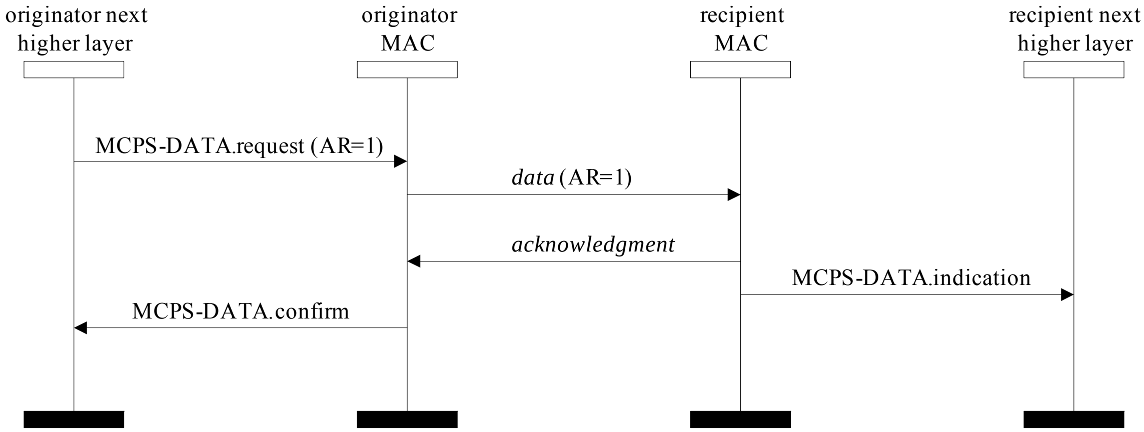

Figure 21—Successful data transmission with an acknowledgment

5.1.7.4.3 Retransmissions

A device that sends a frame with the Acknowledgment Request subfield of its frame control field set to zero shall assume that the transmission was successfully received and shall hence not perform the retransmission procedure.

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

If a single transmission attempt has failed and the transmission was indirect, the coordinator shall not retransmit the data or MAC command frame. Instead, the frame shall remain in the transaction queue of the coordinator and can only be extracted following the reception of a new data request command. If a new data request command is received, the originating device shall transmit the frame using the same DSN as was used in the original transmission.

If a single transmission attempt has failed and the transmission was direct, the device shall repeat the process of transmitting the data or MAC command frame and waiting for the acknowledgment, up to a maximum of *macMaxFrameRetries* times. The retransmitted frame shall contain the same DSN as was used in the original transmission. Each retransmission shall only be attempted if it can be completed within the same portion of the superframe, i.e., the CAP or a GTS in which the original transmission was attempted. If this timing is not possible, the retransmission shall be deferred until the same portion in the next superframe. If an acknowledgment is still not received after *macMaxFrameRetries* retransmissions, the MAC sublayer shall assume the transmission has failed and notify the next higher layer of the failure.

5.1.7.5 Transmission scenarios

Due to the imperfect nature of the wireless medium, a transmitted frame does not always reach its intended destination. Figure 22 illustrates three different data transmission scenarios:

— *Successful data transmission*. The originator MAC sublayer transmits the data frame to the recipient via the PHY data service. In waiting for an acknowledgment, the originator MAC sublayer starts a timer that will expire after *macAckWaitDuration* optical clocks. The recipient MAC sublayer receives the data frame, sends an acknowledgment back to the originator, and passes the data frame to the next higher layer. The originator MAC sublayer receives the acknowledgment from the recipient before its timer expires and then disables and resets the timer. The data transfer is now complete, and the originator MAC sublayer issues a success confirmation to the next higher layer.

— *Lost data frame*. The originator MAC sublayer transmits the data frame to the recipient via the PHY data service. In waiting for an acknowledgment, the originator MAC sublayer starts a timer that will expire after *macAckWaitDuration* optical clocks. The recipient MAC sublayer does not receive the data frame and so does not respond with an acknowledgment. The timer of the originator MAC sublayer expires before an acknowledgment is received; therefore, the data transfer has failed. If the transmission was direct, the originator retransmits the data, and this entire sequence may be repeated up to a maximum of *macMaxFrameRetries* times; if a data transfer attempt fails a total of (1 + *macMaxFrameRetries*) times, the originator MAC sublayer will issue a failure confirmation to the next higher layer. If the transmission was indirect, the data frame will remain in the transaction queue until either another request for the data is received and correctly acknowledged or until *macTransactionPersistenceTime* is reached. If *macTransactionPersistenceTime* is reached, the transaction information will be discarded, and the MAC sublayer will issue a failure confirmation to the next higher layer.

— *Lost acknowledgment frame*. The originator MAC sublayer transmits the data frame to the recipient via the PHY data service. In waiting for an acknowledgment, the originator MAC sublayer starts a timer that will expire after *macAckWaitDuration* optical clocks. The recipient MAC sublayer receives the data frame, sends an acknowledgment back to the originator, and passes the data frame to the next higher layer. The originator MAC sublayer does not receive the acknowledgment frame, and its timer expires. Therefore, the data transfer has failed. If the transmission was direct, the originator retransmits the data, and this entire sequence may be repeated up to a maximum of *macMaxFrameRetries* times. If a data transfer attempt fails a total of (1 + *macMaxFrameRetries*) times, the originator MAC sublayer will issue a failure confirmation to the next higher layer. If the transmission was indirect, the data frame will remain in the transaction queue either until another request for the data is received and correctly acknowledged or until *macTransactionPersistenceTime* is reached. If *macTransactionPersistenceTime* is reached, the transaction information will be discarded, and the MAC sublayer will issue a failure confirmation to the next higher layer.

5.1.8 GTS allocation and management

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

A GTS shall be allocated before use, with the coordinator deciding whether to allocate a GTS based on the requirements of the GTS request and the current available capacity in the superframe. GTSs shall be allocated on a first-come-first-served basis, and all GTSs shall be placed contiguously at the end of the superframe and after the CAP. Each GTS shall be deallocated when the GTS is no longer required, and a GTS can be deallocated at any time at the discretion of the coordinator or by the device that originally requested the GTS. A device that has been allocated a GTS may also operate in the CAP.

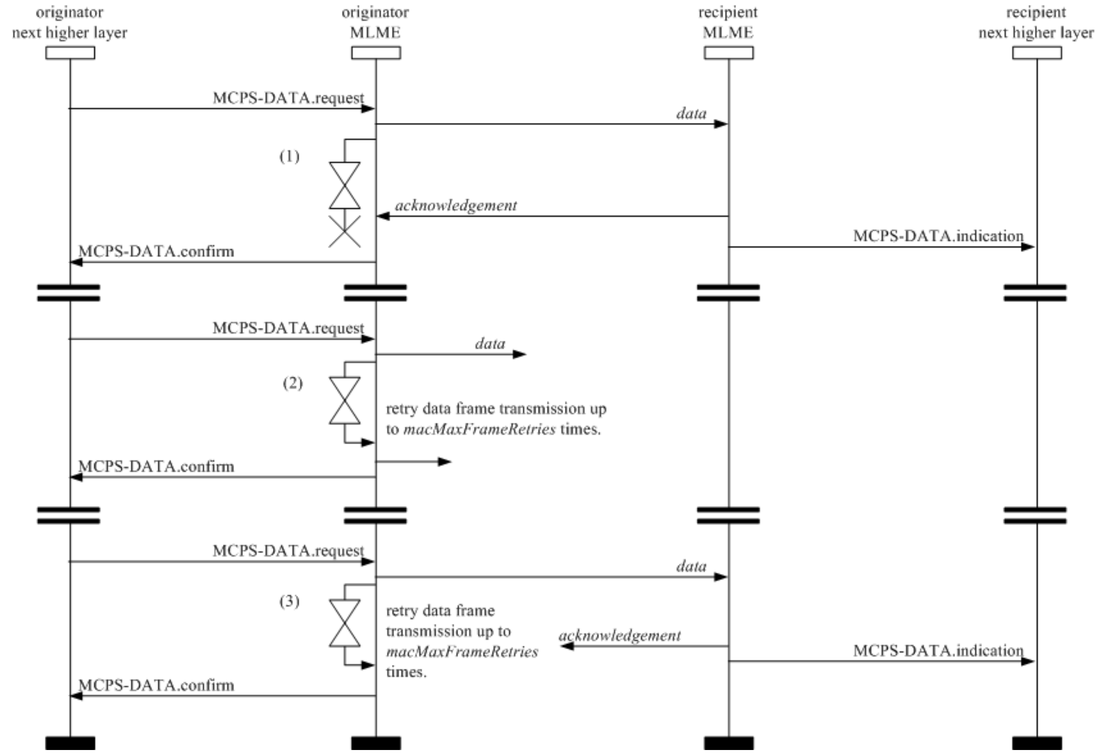

Figure 22—Transmission scenarios, using direct transmission, for frame reliability

A data frame transmitted in an allocated GTS shall use only short addressing.

The management of GTSs shall be undertaken by the coordinator only. To facilitate GTS management, the coordinator shall be able to store all the information necessary to manage seven GTSs. For each GTS, the coordinator shall be able to store its starting slot, length, direction, and associated device address.

The GTS direction, which is relative to the data flow from the device that owns the GTS, is specified as either transmit or receive. The device address and direction shall, therefore, uniquely identify each GTS. Each device may request one transmit GTS and/or one receive GTS. For each allocated GTS, the device shall be able to store its starting slot, length, and direction. If a device has been allocated a receive GTS, it shall enable its receiver for the entirety of the GTS. In the same way, the coordinator shall enable its receiver for the entirety of the GTS if a device has been allocated a transmit GTS. If a data frame is received during a receive GTS and an acknowledgment is requested, the device shall transmit the acknowledgment frame as usual. Similarly, a device shall be able to receive an acknowledgment frame during a transmit GTS.

A device shall attempt to allocate and use a GTS only if it is currently tracking the beacons. The MLME is instructed to track beacons by issuing the MLME-SYNC.request primitive with the TrackBeacon parameter set to TRUE. If a device loses synchronization with the coordinator, all its GTS allocations shall be lost.

The use of GTSs is optional.

5.1.8.1 CAP maintenance

The 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 accommodation of the temporary increase in the beacon frame length needed to perform GTS maintenance. If preventative action becomes necessary, the action chosen is left up to the implementation, but may include one or more of the following:

— Limiting the number of pending addresses included in the beacon.
— Not including a payload field in the beacon frame.
— Deallocating one or more of the GTSs.

5.1.8.2 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.

To request the allocation of a new GTS, the MLME shall send the GTS request command, see 5.3.13, to the coordinator. The Characteristics Type subfield of the GTS Characteristics field of the request shall be set to one (GTS allocation), and the length and direction subfields shall be set according to the desired characteristics of the required GTS. Because the GTS request command contains an acknowledgment request (see 5.3.3.1), the coordinator shall confirm its receipt by sending an acknowledgment frame.

On receipt of a GTS request command indicating a GTS allocation request, the coordinator shall first check if there is available capacity in the current superframe, based on the remaining length of the CAP and the desired length of the requested 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 coordinator provided there is sufficient bandwidth available. The coordinator shall make this decision within *aGTSDescPersistenceTime* superframes.

On receipt of the acknowledgment to the GTS request command, the device shall continue to track beacons and wait for at most *aGTSDescPersistenceTime* superframes. If no GTS descriptor for the device appears in the beacon within this time, the MLME of the device shall notify the next higher layer of the failure. This notification is achieved when the MLME issues the MLME-GTS.confirm primitive, see 6.3.5.3, with a status of NO_DATA.

When the 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. If the GTS was allocated successfully, the coordinator shall set the start slot in the GTS descriptor to the superframe slot at which the GTS begins and the length in the GTS descriptor to the length of the GTS. In addition, the coordinator shall notify the next higher layer of the new GTS. This notification is achieved when the MLME of the coordinator issues the MLME-GTS.indication primitive, see 6.3.5.2, with the characteristics of the allocated GTS. If there was not sufficient capacity to allocate the requested GTS, the start slot shall be set to zero and the length to the largest GTS length that can currently be supported. The coordinator shall then include this GTS descriptor in its beacon and update the GTS Specification field of the beacon frame accordingly. The coordinator shall also update the Final CAP Slot subfield of the Superframe Specification field of the beacon frame, indicating the final superframe slot utilized by the decreased CAP. The GTS descriptor shall remain in the beacon frame for *aGTSDescPersistenceTime* superframes, after which it shall be removed automatically. The coordinator shall be allowed to reduce its CAP below *aMinCAPLength* to accommodate the temporary increase in the beacon frame length due to the inclusion of the GTS descriptor.

On receipt of a beacon frame containing a GTS descriptor corresponding to *macShortAddress*, the device shall process the descriptor. The MLME of the device shall then notify the next higher layer of whether the GTS allocation request was successful. This notification is achieved when the MLME issues the MLME-GTS.confirm primitive with a status of SUCCESS (if the start slot in the GTS descriptor was greater than zero) or DENIED (if the start slot was equal to zero or if the length did not match the requested length).

5.1.8.3 GTS usage

When the MAC sublayer of a device that is not the coordinator receives an MCPS-DATA.request primitive, see 6.2.1, with the TxOptions parameter indicating a GTS transmission, it shall determine whether it has a valid transmit GTS. If a valid GTS is found, the MAC sublayer shall transmit the data during the GTS, i.e., between its starting slot and its starting slot plus its length. At this time, the MAC sublayer shall transmit the MPDU immediately without using any random access, provided the requested transaction can be completed before the end of the GTS. If the requested transaction cannot be completed before the end of the current GTS, the MAC sublayer shall defer the transmission until the specified GTS in the next superframe. Note that the MAC shall allow for the PHY overhead in making this determination.

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.

When the MAC sublayer of the coordinator receives an MCPS-DATA.request primitive with the TxOptions parameter indicating a GTS transmission, it shall determine whether it has a valid receive GTS corresponding to the device with the requested destination address. If a valid GTS is found, the coordinator shall defer the transmission until the start of the receive GTS. In this case, the address of the device with the message requiring a GTS transmission shall not be added to the list of pending addresses in the beacon frame as shown in 5.1.6. At the start of the receive GTS, the MAC sublayer shall transmit the data without using any random access, provided the requested transaction can be completed before the end of the GTS. If the requested transaction cannot be completed before the end of the current GTS, the MAC sublayer shall defer the transmission until the specified GTS in the next superframe.

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

Before commencing transmission in a GTS, each device shall ensure that the data transmission, the acknowledgment, if requested, and the IFS, suitable to the size of the data frame, can be completed before the end of the 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 GTSs deallocated.

5.1.8.4 GTS deallocation

A device is instructed to request the deallocation of an existing GTS through the MLME-GTS.request primitive specified in 6.3.5.1, using the characteristics of the GTS it wishes to deallocate. From this point onward, the GTS to be deallocated shall not be used by the device, and its stored characteristics shall be reset.

To request the deallocation of an existing GTS, the MLME shall send the GTS request command, specified in 5.3.13, to the coordinator. The Characteristics Type subfield of the GTS Characteristics field of the request shall be set to zero (i.e., GTS deallocation), and the length and direction subfields shall be set according to the characteristics of the GTS to deallocate. Because the GTS request command contains an acknowledgment request, specified in 5.3.3.1, the coordinator shall confirm its receipt by sending an acknowledgment frame. On receipt of the acknowledgment to the GTS request command, the MLME shall notify the next higher layer of the deallocation. This notification is achieved when the MLME issues the MLME-GTS.confirm primitive, see 6.3.5.3, with a status of SUCCESS and a GTSCharacteristics parameter with its Characteristics Type subfield set to zero. If the GTS request command is not received correctly by the coordinator, it shall determine that the device has stopped using its GTS by the procedure described in 5.1.8.6.

On receipt of a GTS request command with the Characteristics Type subfield of the GTS Characteristics field set to zero (GTS deallocation), the coordinator shall attempt to deallocate the GTS. If the GTS characteristics contained in the GTS request command do not match the characteristics of a known GTS, the coordinator shall ignore the request. If the GTS characteristics contained in the GTS request command match the characteristics of a known GTS, the MLME of the coordinator shall deallocate the specified GTS and notify the next higher layer of the change. This notification is achieved when the MLME issues the MLME-GTS.indication primitive, see 6.3.5.2, with a GTSCharacteristics parameter containing the characteristics of the deallocated GTS and a Characteristics Type subfield set to zero. The coordinator shall also update the Final CAP Slot subfield of the Superframe Specification field of the beacon frame, indicating the final superframe slot utilized by the increased CAP. It shall not add a descriptor to the beacon frame to describe the deallocation.

GTS deallocation may be initiated by the coordinator due to a deallocation request from the next higher layer, the expiration of the GTS (see 5.1.8.6), or maintenance required to maintain the minimum CAP length, *aMinCAPLength* (see 5.1.8.1).

When a GTS deallocation is initiated by the next higher layer of the coordinator, the MLME shall receive the MLME-GTS.request primitive with the GTS Characteristics field of the request set to zero (i.e. GTS deallocation) and the length and direction subfields set according to the characteristics of the GTS to deallocate.

When a GTS deallocation is initiated by the coordinator either due to the GTS expiring or due to CAP maintenance, the MLME shall notify the next higher layer of the change. This notification is achieved when the MLME issues the MLME-GTS.indication primitive with a GTSCharacteristics parameter containing the characteristics of the deallocated GTS and a Characteristics Type subfield set to zero.

In the case of any deallocation initiated by coordinator, the coordinator shall deallocate the GTS and add a GTS descriptor into its beacon frame corresponding to the deallocated GTS, but with its starting slot set to zero. The descriptor shall remain in the beacon frame for *aGTSDescPersistenceTime* superframes. The coordinator shall be allowed to reduce its CAP below *aMinCAPLength* to accommodate the temporary increase in the beacon frame length due to the inclusion of the GTS descriptor.

On receipt of a beacon frame containing a GTS descriptor corresponding to *macShortAddress* and a start slot equal to zero, the device shall immediately stop using the GTS. The MLME of the device shall then notify the next higher layer of the deallocation. This notification is achieved when the MLME issues the MLME-GTS.indication primitive with a GTSCharacteristics parameter containing the characteristics of the deallocated GTS and a Characteristics Type subfield set to zero.

5.1.8.5 GTS reallocation

The deallocation of a GTS may result in the superframe becoming fragmented. For example, Figure 23 shows three stages of a superframe with allocated GTSs. In stage 1, three GTSs are allocated starting at slots 14, 10, and 8, respectively. If GTS 2 is now deallocated (stage 2), there will be a gap in the superframe during which nothing can happen. To solve this, GTS 3 will have to be shifted to fill the gap, thus increasing the size of the CAP (stage 3).

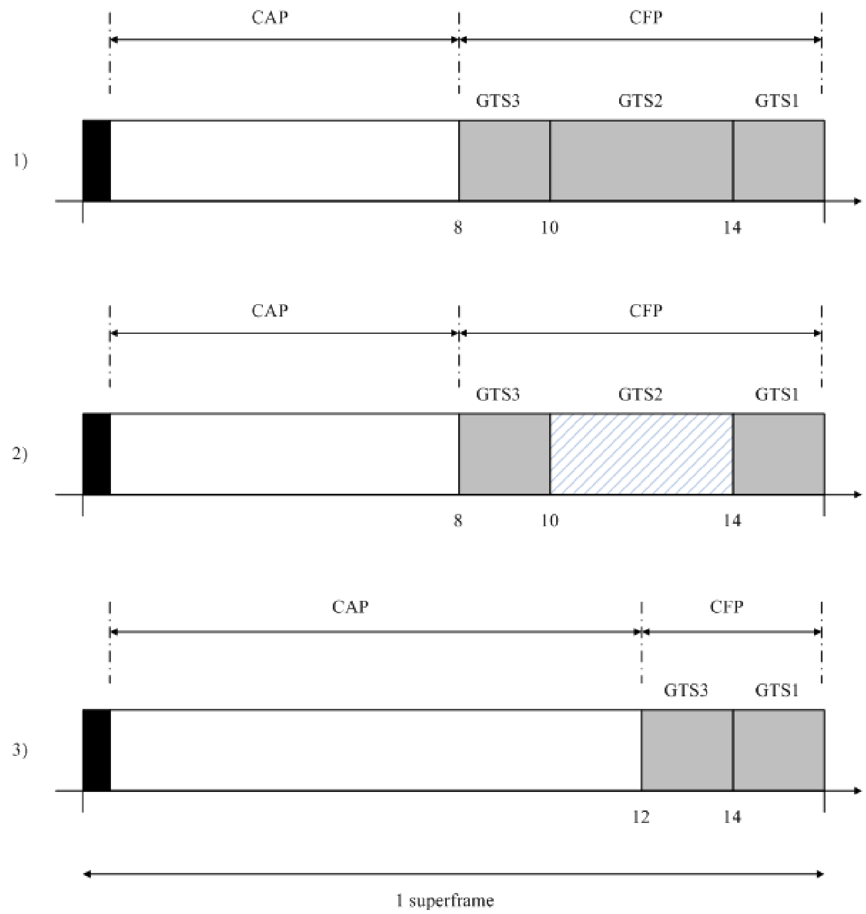

Figure 23—CFP defragmentation on GTS deallocations

The 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.

When a GTS is deallocated by the coordinator, it shall add a GTS descriptor into its beacon frame indicating that the GTS has been deallocated. If the deallocation is initiated by a device, the coordinator shall not add a GTS descriptor into its beacon frame to indicate the deallocation. For each device with an allocated GTS having a starting slot lower than the GTS being deallocated, the coordinator shall update the GTS with the new starting slot and add a GTS descriptor to its beacon corresponding to this adjusted GTS. The new starting slot is computed so that no space is left between this GTS and either the end of the CFP, if the GTS appears at the end of the CFP, or the start of the next GTS in the CFP.

In situations where multiple reallocations occur at the same time, the coordinator may choose to perform the reallocation in stages. The coordinator shall keep each GTS descriptor in its beacon for *aGTSDescPersistenceTime* superframes.

On receipt of a beacon frame containing a GTS descriptor corresponding to *macShortAddress* and a direction and length corresponding to one of its GTSs, the device shall adjust the starting slot of the GTS corresponding to the GTS descriptor and start using it immediately.

In cases where it is necessary for the coordinator to include a GTS descriptor in its beacon, it shall be allowed to reduce its CAP below *aMinCAPLength* to accommodate the temporary increase in the beacon frame length. After *aGTSDescPersistenceTime* superframes, the coordinator shall remove the GTS descriptor from the beacon.

5.1.8.6 GTS expiration

The MLME of the 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 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, where *n* is defined below.
— For receive GTSs, the MLME of the 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, where *n* is defined below. If the data frames sent in the GTS do not require acknowledgment frames, the MLME of the coordinator will not be able to detect whether a device is using its receive GTS. However, the coordinator is capable of deallocating the GTS at any time.

The value of *n* is defined as follows:

$n = 2^{(8-macBeaconOrder)}$ $\quad 0 \leq macBeaconOrder \leq 8$

$n = 1$ $\quad 9 \leq macBeaconOrder \leq 14$

5.1.9 Fast link recovery

In the star topology, a fast link recovery process may be triggered at the device end during communication. The trigger may be initiated when the device does not receive ACKs for a number of times given by the MAC PIB attribute *macNumAcks,* as defined in Table 60. In the fast link recovery process, the device may decide on its own to stop sending data. The device may also send the fast link recovery (FLR) signal repeatedly (within the allocated resource) to the coordinator if the device is connected to mains power. Upon receiving the FLR signal, the coordinator shall send a FLR response to the device. The communication resumes after the device receives the response. If there is bi-directional data transfer during communication, the device may wait after stopping sending data. If the device does not receive any FLR response signal within a timer given by the MAC PIB attribute *macLinkTimeOut*, the device may assume the link is broken and may disassociate.

The FLR signal and response are defined in 5.3.11. The FLR signal and response shall be sent at the lowest data rate corresponding to the currently negotiated optical rate.

Figure 24 shows an example of the process of device stopping sending data based on the retransmission count.

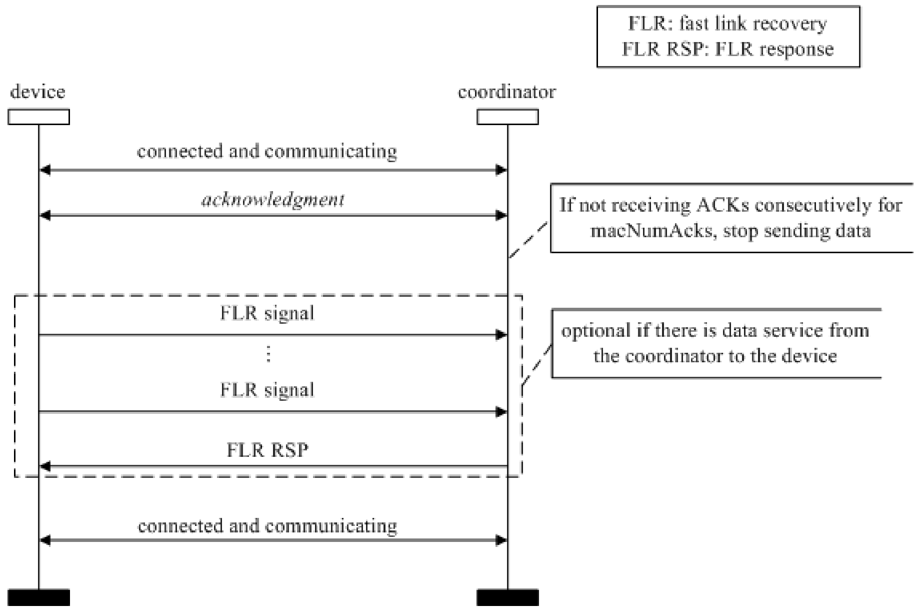

Figure 24—An example of the process of device stopping data transmission based on the retransmission count, and triggering FLR.

In the star topology, a fast link recovery may also be triggered by the coordinator. The trigger may be initiated when the coordinator does not receive contiguous ACKs for a number of times given by the MAC PIB attribute *macNumAcks*. In the fast link recovery process, the coordinator may stop sending data to the device. The coordinator then sends fast link recovery (FLR) signal repeatedly to the device. The coordinator may hold the uplink grant allocated to the device. Upon receiving FLR signal, the device shall send a FLR response to the coordinator. The communication resumes after the device receives the response.

Figure 25 shows an example of the process of the coordinator stopping sending data based on the retransmission count.

Figure 25—An example of the process of the coordinator stopping sending data based on the retransmission count, and triggering FLR.

In peer-to-peer VLC, the devices may let each other know their battery life. If the conditions to trigger the fast link recovery process are satisfied, the device may further compare its own battery life with the battery life of its peer (the one it is communicating). If the battery life of the device is shorter than its peer's, then the device stops sending data, and waits. If the battery life of the device is longer than its peer's, then the device stops sending data, and initiates the fast link recovery process.

When the fast link recovery is triggered, and if the device has spare wavelength bands, some or all of the spare bands may also start sending fast link recovery signals to recover the link. The device then shall choose a band which gets the fast link recovery response to continue the communication.

The address field of MHR in FLR signal and response may include the address or the identifier of the color bands.

Figure 26 shows a flowchart of the process for color band assisted fast link recovery.

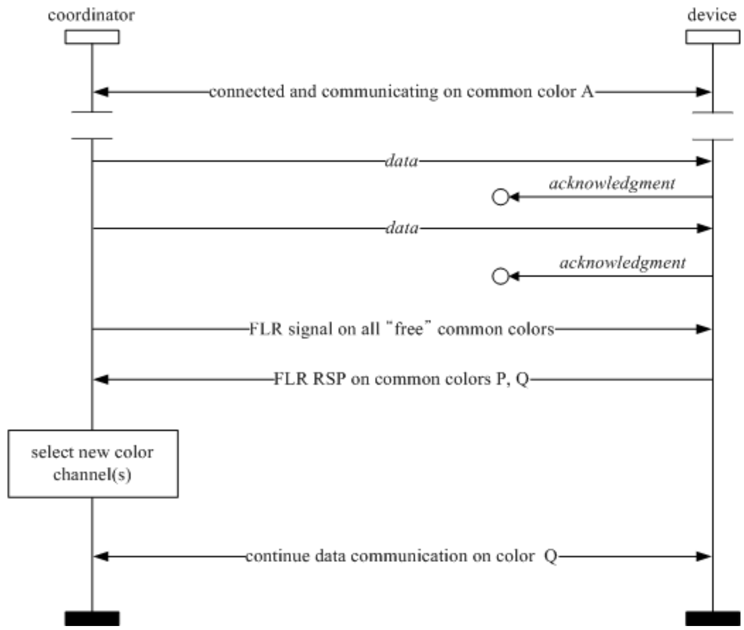

Figure 26—Flowchart of process for color band assisted fast link recovery

When the fast link recovery is triggered, if the device has other communication directions/angles (e.g., a light with multiple LEDs with different angles) some or all of the other angles may also start sending fast link recovery signaling to recover the link. The device then shall choose an angle which gets the fast link recovery response to continue the communication. The process of fast link recovery on other directions/ angles is done successively (i.e. one direction after another). The direction is indicated in the link recovery mechanism provided by the command frame structure.

The address field of MHR in FLR signal and response may include the address or the identifier of the angles or directions.

Figure 27 shows a flowchart of the process for multiple angle assisted fast link recovery.

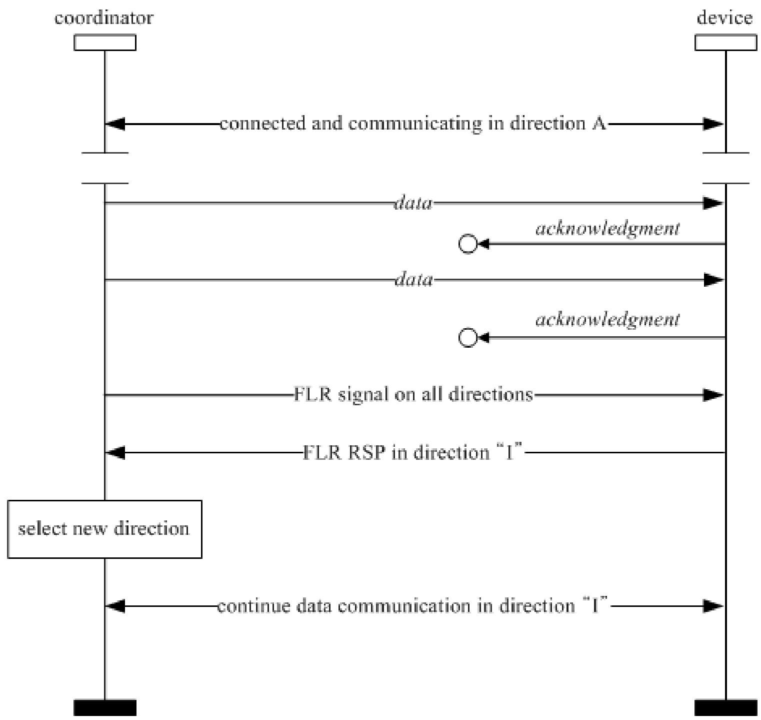

Figure 27—Flowchart showing process of multiple angles assisted fast link recovery

5.1.10 Multiple channel resource assignment

5.1.10.1 Multiple channel information

When the coordinator does not have time slot resources to assign for new user, the coordinator should extend the resource by using multiple bands. Figure 28 shows one example of multiple band usage.

Figure 29 describes the procedure of multiple band usage when the multiple band function is needed. When device 2 tries to initially access the coordinator for communication and no time slot is available but other bands are available for device 2, the coordinator can assign another band except the default band. Capability exchange should occur for all bi-directional communication during device discovery (see 5.1.2.4). If multiple bands are used, the coordinator should transmit to the device the "Src_multi_info" in the MAC command payload field which is defined in Table 3 to the device. Then the device 2 shall respond to the coordinator using the "Des_multi_info" which is defined in Table 3, informing the device of available multiple bands of the device.

If the coordinator does not support multiple bands, because the coordinator has a single band light source, or does not want to use multiple bands, the coordinator should transmit Src_multi_info set with code '0000000' as shown in Annex D.

If the device also cannot support multiple bands due to hardware limitations, such as a single band light source or an interference situation, or does not want to use multiple bands, the device should respond with Des_multi_info set to code '0000000' as shown in Annex D.

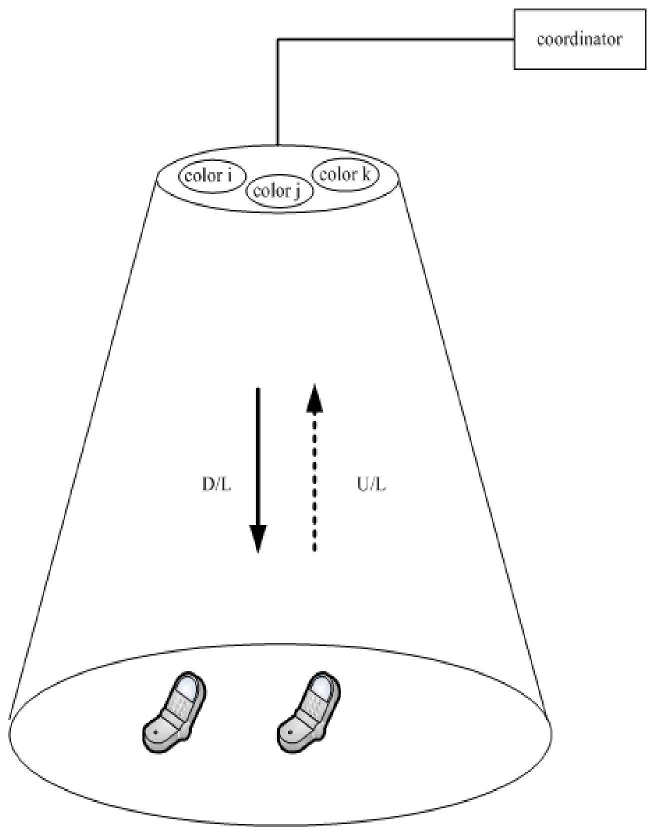

Figure 28—Example of multiple channel usage

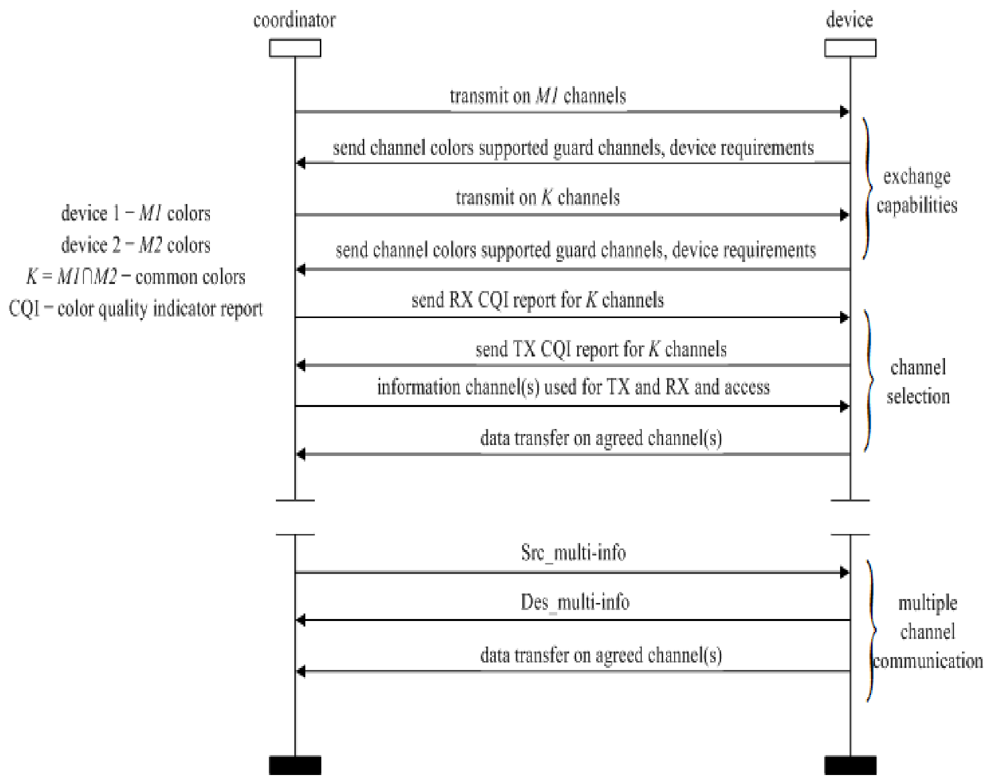

Figure 29—MSC for multi-band information

Table 3—Command frame payload for multiple bands

MAC command frame payload	Bits	Usage/Description	Down/Up Link
Src_multi_info	b0...b6	Bit map that indicates the available channels to the coordinator ex: 0000000: No multiple channel mode ex: 0000001: using channel "Band 7" ex: 0000101: using channel "Band 5" and "Band 7"	D/L
Des_multi_info	b0...b6	Bit map that indicates the available channels to the mobile device ex: 0000000: No multiple channel mode ex: 0000001: using channel "Band 7" ex: 0000101: using channel "Band 5" and "Band 7"	U/L

5.1.10.2 Band hopping for interference avoidance

A single coordinator can service multiple cells.

If interference is being experienced from an adjacent light then hopping can be used to mitigate the interference. When spatial reuse due to direction optics is not present, and when the VLC communications system uses the same time slot between the adjacent light sources or cells with multiple band communication, and when multiple bands are supported by the PHY, band hopping can be used. In order to avoid interference and increase system capacity, pre-assigned hopping patterns should be adopted.

The hopping pattern should be assigned to the device and then the device should operate and hop based on the assigned hopping pattern. The coordinator shall transmit to the device the 'H_pattern' using the MAC command frame payload which is defined in Table 4.

If the VLC system does not use multiple bands (Src_multi_info is set to code '0000000'), then the hopping function is not supported. The hopping patterns shall be structured so as not to change the visual perception of the light. For example, the patterns could hop between RGB in the proper time averaged portion so as to appear white.

Table 4—Command frame payload for channel hopping

MAC command frame payload	Bit	Usage/Description	Down/Up Link
H_pattern	b0, b1, b2, b3, b4	Band hopping information	D/L

5.1.11 VLC cell design and mobility support

There may be a need to support link switching due to physical movement or interference. Mobility can be of two types: physical and logical. Physical mobility occurs when the VLC device M1 changes its position due to the movement within the coverage area of infrastructure I1 while logical mobility occurs when the device M1 changes its communication link from a link with infrastructure I2 to one with infrastructure I3 due to interference or deliberate channel switching, as shown in Figure 30.

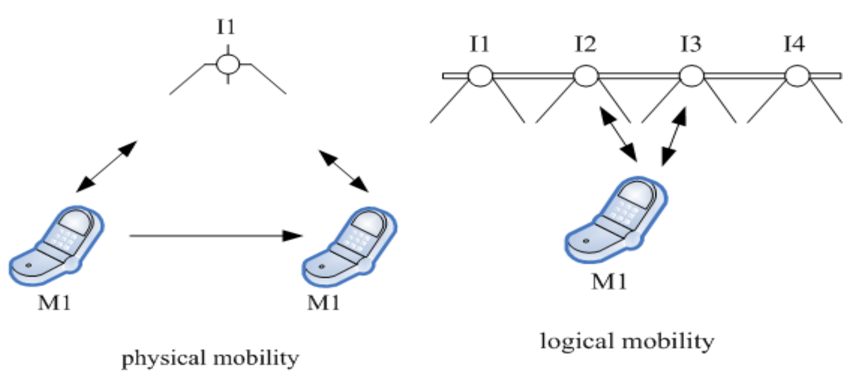

Figure 30—Physical and logical mobility

A coordinator DME can separate the optical media into multiple cells for supporting applications such as location based services.

5.1.11.1 Mobility using boundary information

A single coordinator can support mobility of the device through multiple cells using the PHY switch, controlled by the DME, as shown in Figure 31. Each optical element in a cell is denoted by *cell_ID(i,j)* where *j* is the index of the element in the i^{th} cell. The size and the position of the cell in the optical media shown in Figure 3 can be variable and can be programmed by the DME. The actual size and position determination for the cell by the coordinator DME is not defined in the standard. If device 1 moves to the next cell, for example, from *cell_ID(i,j)* to *cell_ID(i+1,j)*, the coordinator can detect the mobility of the device using the uplink signal (i.e. acknowledgment frame).

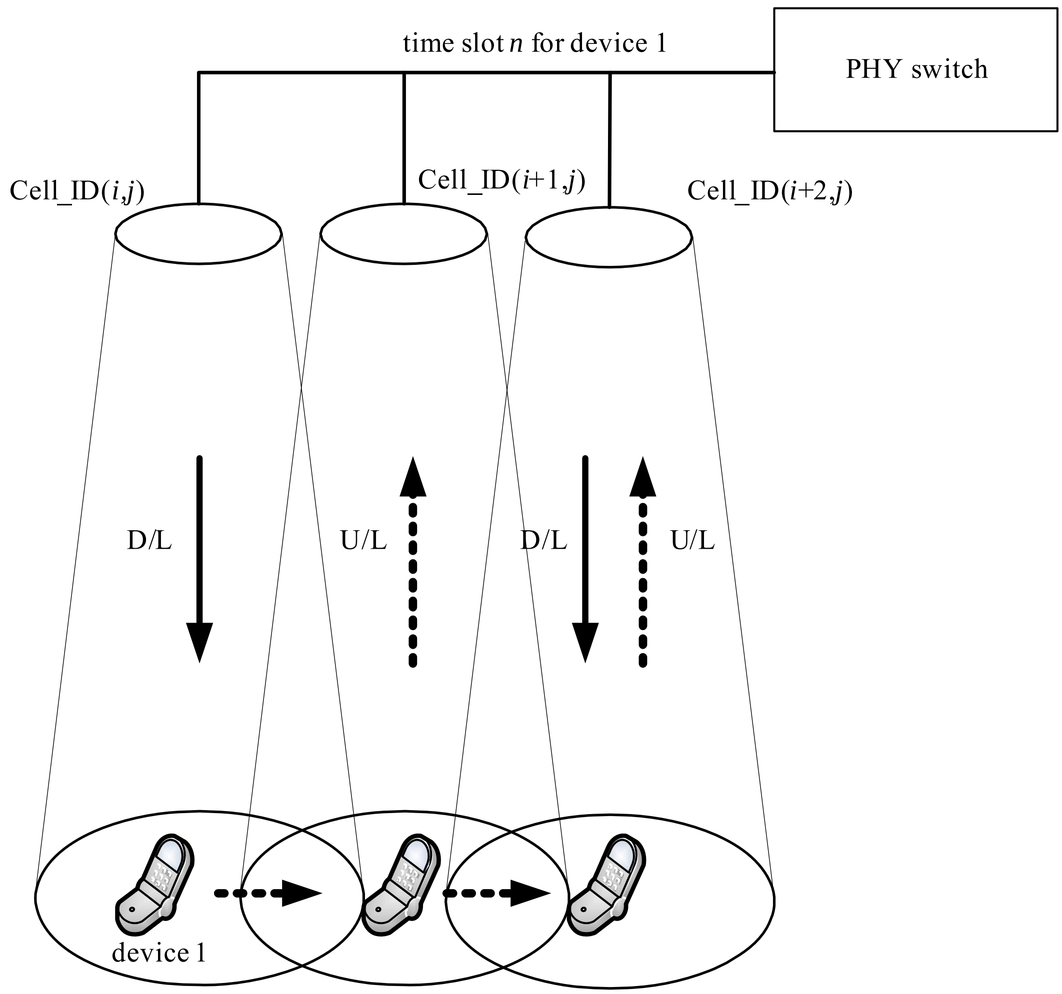

Figure 31—Cell configuration for VLC mobility

Figure 32 shows the mobility support for a device through multiple cells. When device 1 moves out from *Cell_ID(i,j)* to *Cell_ID(i+1,j)*, the coordinator may not receive the uplink transmission (for example, acknowledgment frame or CVD frame) from *Cell_ID(i,j)*. The coordinator may then search for the device through the adjacent cells such as *Cell_ID(i+1,j)* and *Cell_ID(i-1,j)* during the same time slots assigned to device 1 in the superframe. The other devices in *cell_ID(i,j)* will continue communication in the same cell. The co-ordinator may also expand the cell size in order to provide coverage for mobility of the device. The coordinator can decide on the new cell selection for the device on receiving the uplink transmission from device 1.Thus, if the coordinator can resume communication with the device in *cell_ID(i+1,j)*, the coordinator DME may set the PHY switch to use *cell_ID(i+1,j)* for device 1 during the time slots allocated for device 1 and then switch back to *cell_ID(i,j)* to service any existing devices in *cell_ID(i,j)* in the remaining time slots. The searching process can be terminated if the device is not found within the link timeout period, defined in MAC PIB attribute *macLinkTimeOut* in Table 60, and the device can then be considered to be disassociated from the coordinator.

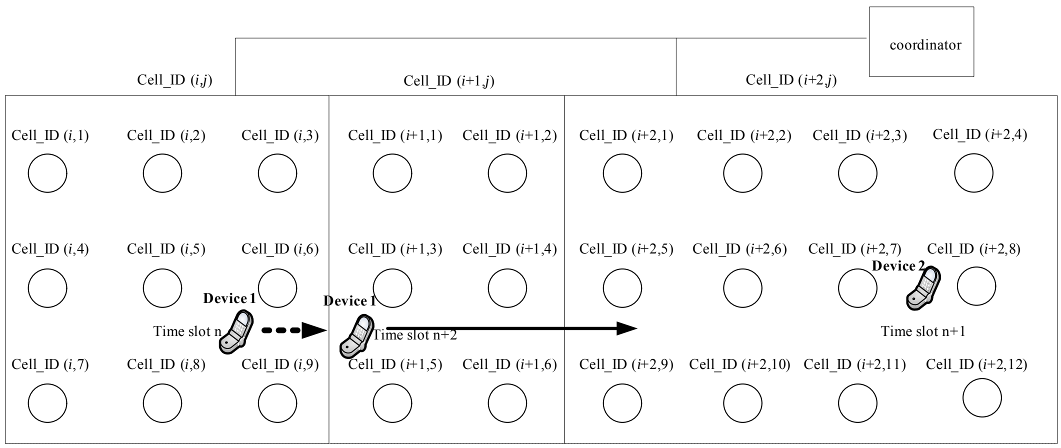

Figure 32—Mobility support for a device through multiple cells

5.1.11.2 Cell configuration during superframe

In order to support access for new devices through the entire superframe, the entire optical media shall be configured to a single macro cell during the beacon and CAP periods. Once devices are discovered and associated, the cell sizes and positions can be determined and the cell structure can be applied to the individual device(s) for communication, as shown in Figure 33.

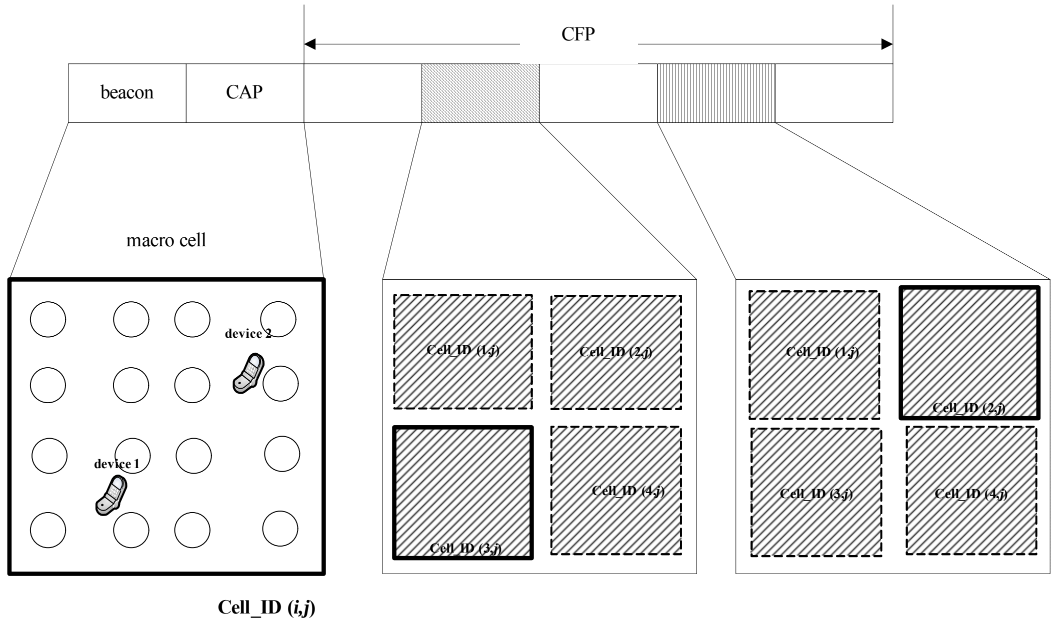

Figure 33—Superframe configuration for mobility support

5.1.11.3 Cell size and location search procedure

Once a device is associated with a coordinator using the beacon and CAP, the coordinator may establish the size and location of the cell in order to service the new device in the CFP with a smaller cell size. In order to determine the size and location of the cell, the coordinator first sets the *cellSearchEn* bit in the superFrame specification field of the beacon frame as defined in Figure 49. If the *cellSearchEn* bit is set, the *cellSearchLength* is transmitted as an additional field in the beacon frame, as shown in Figure 46. If the *cellSearchEn* bit is set, the coordinator readjusts its superframe GTS allocation to ensure the first *cellSearchLength* slots of the CFP are allocated for cell size and location search.

The first *cellSearchLength* slots are used as visibility slots by the coordinator and the devices. During the first *cellSearchLength* slots, the coordinator sequentially cycles through the *cellSearchLength* cells and transmits CVD frames in all the cells. Figure 34 shows an example of the sequential search for 4 cells. CS1 to CS4 are the 4 cell search slots that are made available for searching via setting the *cellSearchLength* to 4 and setting the *cellSearchEn* bit in the beacon frame.

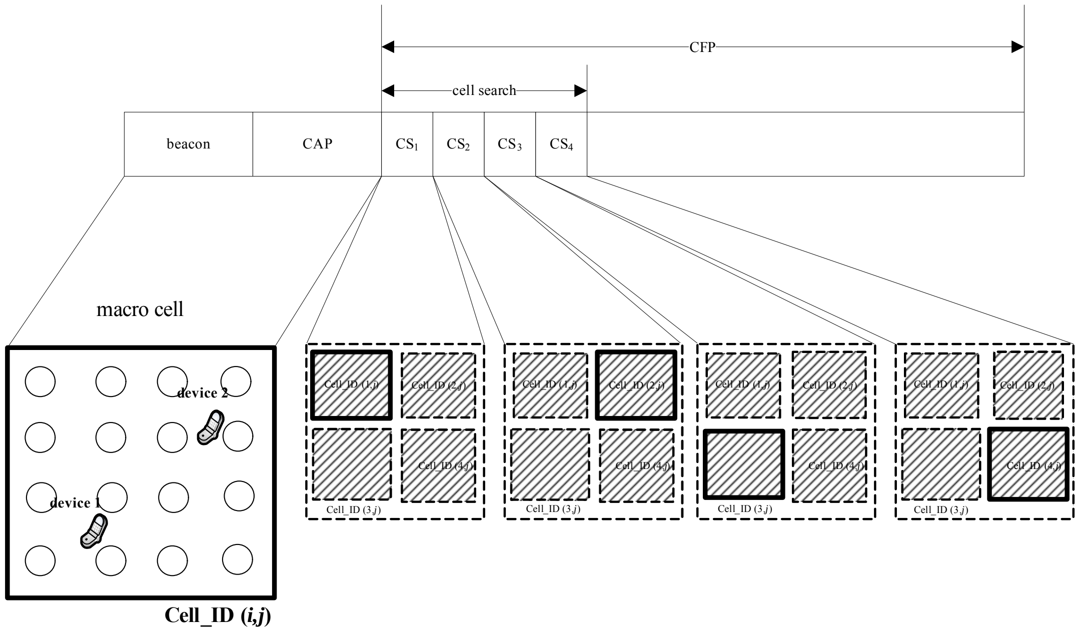

Figure 34—Cell size and location search procedure

If a device receives a beacon with the *cellSearchEn* bit set to 1, the device shall also continuously transmit CVD frames during the *cellSearchEn* slots while also monitoring the CVD frame reception from the coordinator. The device shall note the WQI during each of the *cellSearchLength* slots and shall report this information back to the coordinator using the mobility notification command frame, as described in 5.3.12.

The coordinator makes the determination of the cell sizes and location based on the information from the mobility notification command and its own reception of the CVD frames from the device during the cell search slots.

5.1.12 Color function support

The CVD frame, using various colors, can be used to display various statuses of a device. The colors mapped for each status of the devices are based on the phyColorFunction (see Table 100). The colors chosen for different statuses are left to the discretion of the implementer. Multiple statuses may choose the same color, depending on the number of colors supported by the device. The use of color function through the CVD frame has the potential to change the color of the emitted light.

5.1.12.1 CVD frame usage for MAC state indication

The CVD frames are used between state changes to provide visual information to the user regarding the communication status. The MLME primitives for association (see 6.3.1.1), scan (see 6.3.8.1) and disassociation (see 6.3.2.1) are used to support this functionality. The corresponding colors, as described in Table 5 can be used to display various states of a device. The MAC PIB attributes, *macDuringASSOCColor*, *macDuringDISASSOCColor*, and *macDuringSCANColor* as shown in Table 60, are used for the color assignment of the CVD frame when the CVD frame is sent to indicate the MAC state during the association, disassociation, or scan process.

Table 5—Color table for MAC state indication

State	Color choice	Color resolution range
scan	Color “A”	0-255
association	Color “B”	0-255
disassociation	Color “C”	0-255

For example, the device sends an association request to the coordinator (see Figure 35) and indicates this to the user with a chosen color. This information about the color choice is communicated using the MLME-ASSOCIATE.request primitive as in 6.3.1.1.

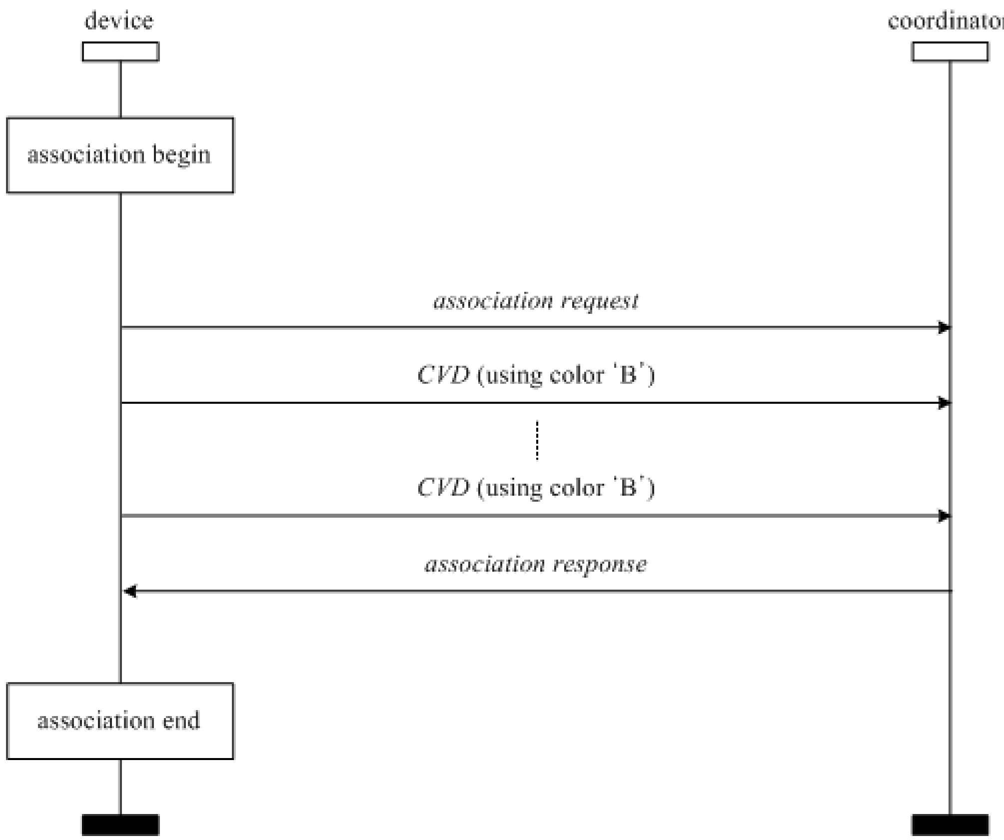

Figure 35—MSC when color function for association indication is invoked

5.1.12.2 CVD frame usage for acknowledgment indication

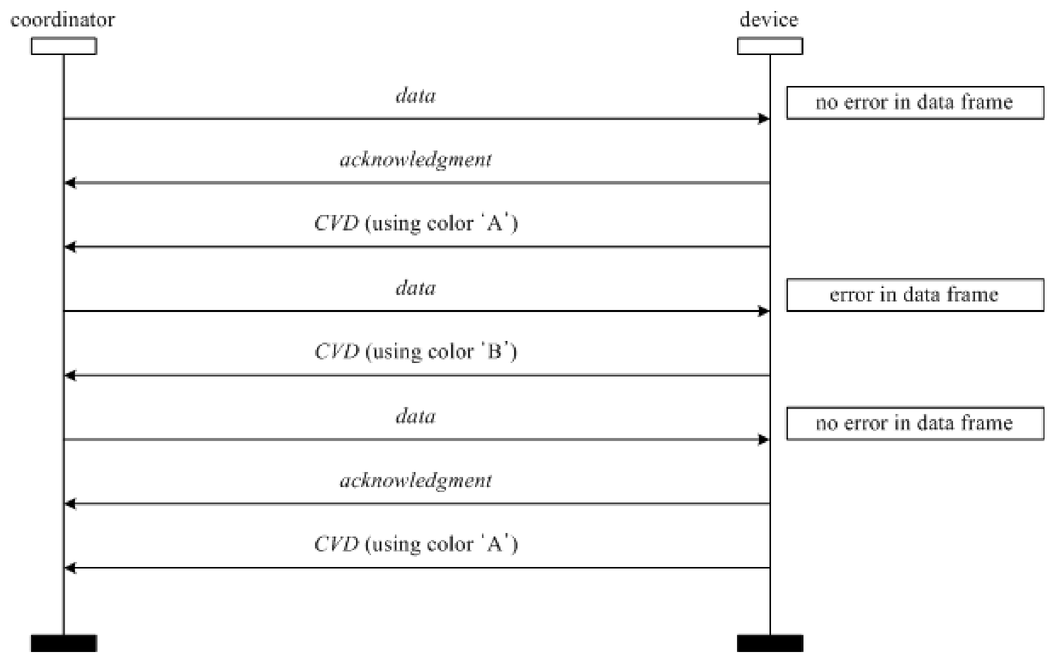

Figure 36—CVD frame usage for acknowledgment indication

Figure 36 shows an example of how the user can infer whether a receiver successfully receives data or not. According to this figure, the device sends a CVD frame after the ACK frame has been sent. The CVD frame can indicate that the received data has errors or is error-free, based on the choice of colors. The MAC PIB attribute, *macColorReceived* as shown in Table 60, is used for the color assignment of the CVD frame when the ACK frame is sent and the color function for the ACK state indication is achieved by the CVD frame. The MAC PIB attribute, *macColorNotReceived* as shown in Table 60, is used for the color assignment of the CVD frame when the ACK frame is not sent but the color function for the non-ACK state indication is achieved by the CVD frame.

5.1.12.3 CVD frame usage for channel quality indication

Table 6 describes how the user can infer the quality of the data transmission or the communication quality through the CVD frame. The communication quality may be obtained by various metrics. For example, FER statistics can be averaged over multiple frames. The *ppduLinkQuality* of 9.3.3 (PD-DATA.indication) can also be used for this purpose. This information can help provide misalignment indication to the user. Different colors can be used to indicate different states of misalignment. The choice of the colors and the FER range is left to the implementer and is out of scope of the standard.

Table 6—Color table for channel quality indication

Color of CVD frame	Channel quality
Color “A”	current FER < FER #1
Color “B”	FER #1 ≤ FER < FER #2
Color “C”	current FER ≥ FER #2

5.1.12.4 CVD frame usage for file-transfer status indication

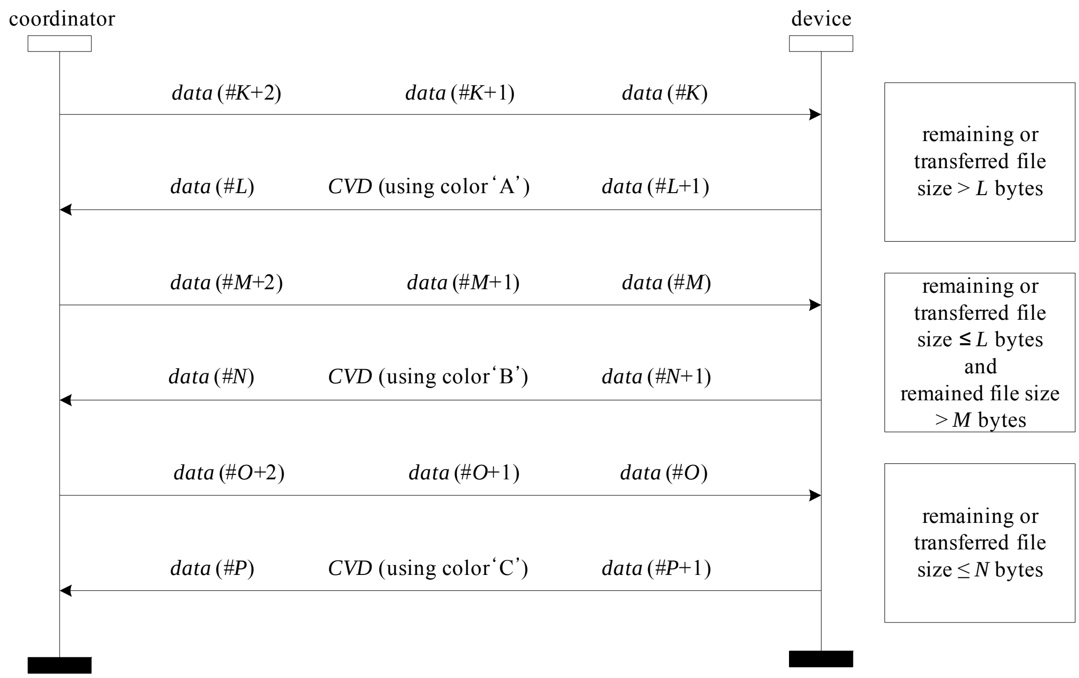

Figure 37—Example of MSC for CVD frame usage for file-transfer status indication

Figure 37 shows an example of how the user can infer the remaining or transferred file size through the color of the CVD frame. As shown in the example of Figure 37, the coordinator transfers files to the device. Different stages of the file transfer process can be represented with different choices of colors. In order to use this indication, the device needs to know the total file size to be transmitted. The remaining file size can be obtained by subtracting the transferred file size from the total file size. The MAC PIB attribute, *macCFAppColor* as shown in Table 60, is used for the color assignment of the CVD frame when the CVD frame is sent to indicate the application-dependent information, such as the file-transfer status.

5.1.12.5 Generic color assignment mechanism

The color function can be used beyond the applications as described from 5.1.12.1 to 5.1.12.4. The colors to support the various color functions shall be chosen from the *phyColorFunction* PHY PIB attribute as shown in Table 100, using the MLME-SET.request and PLME-SET.request primitives available to the DME shown in Figure 3.

5.1.13 Color stabilization

When a device joins a network (administrated by a coordinator) it advertises its capability of color stabilization in CSK links as shown in Table 16. We assume that at least one link is functioning as a CSK bidirectional link. Otherwise, no color-stabilization functionality is invoked in the network. Also, for the sake of simplicity, we assume that only the device will be requested to send color-stabilization updates.

The device and the coordinator go through the steps of association as in 5.1.4. Upon the issuance of a MLME-ASSOCIATE.request the device sends an Association request, among other things advertising its capability for Rx-side CSK-color stabilization.[1] Upon reception of this request, the coordinator MLME creates an MLME-ASSOCIATE.indication to the next higher layer in the coordinator. There, a decision is made, whether and where color stabilization will be invoked. If the link to be established is a duplex CSK link, the coordinator can also choose to stabilize the color of the device Tx. (As already mentioned, we are describing the case of color stabilization of the coordinator, but the other possible cases can be inferred from the description in a straight-forward manner). After this decision has been made, the pertinent capability negotiation response field in the MLME-ASSOCIATE.response is set according to Table 32 and the pertinent information is then translated by the coordinator MLME into the MAC association response message. Upon reception of this message, the device MLME creates the MLME-ASSOCIATE.confirm and sends it to the next higher layer in the device for further processing.

When the coordinator starts sending CVD frames to the device (identified by the pertinent MAC header as shown in 5.1.12), the device sends color stabilization information back to the coordinator. The MAC command frame used for this can be found in 5.3.17. After a time set in the variable *macColorStabilizationTimer*, as shown in Table 60, the current information is sent again from the device to the coordinator. If the coordinator wants to change the time between two such updates, it can send a color-stabilization-timer notification command (see 5.3.16) to the device, upon which the device MLME sets the pertinent timer, which is not further described in this standard.

Upon dissociation, the *macColorStabilization* variable is set back to its default value '00'.

5.1.14 Visibility and dimming support

The standard supports visibility for the following purposes:

a) Alignment (Device discovery, negotiation, connection)
b) Visible guiding for user alignment
c) Infrastructure continuous light output
d) Blinking for unexpected interference, disconnection warnings

5.1.14.1 Visibility pattern

The MAC passes the visibility pattern requirement to the PHY layer via the PLME interface using the *phyDim* PIB attribute as shown in Table 100. Sending an idle pattern is a mandatory requirement for infrastructure during idle or receive operation to ensure continuous illumination. Sending an idle pattern is optional for the mobile device.

5.1.14.2 Extended preamble mode for visibility

The MAC provides an extended preamble mode for visibility. The advantage of this mode is to provide additional time for synchronization while simultaneously providing visibility.

[1] If not otherwise mentioned, an acknowledgment shall be sent after the reception of each message.

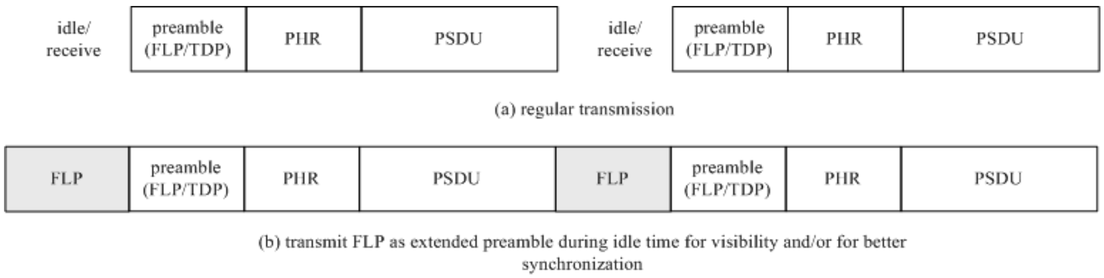

Figure 38—Extended preamble mode provided by the MAC

The MAC uses the knowledge of the idle time and may increase the number of preamble repetitions during the frame transmission to cover the idle time period. The extended preamble is made continuous to the existing preamble of the next frame transmission. There is a possibility that the idle time may not be an integral multiple of the preamble length. In such cases, it is acceptable to transmit a fraction of the preamble (the latter part) in order to maintain visibility. This fraction of the preamble can be called a truncated preamble.

The MAC can choose to either transmit a idle pattern or an extended preamble in the idle mode during regular operation. The choice is made by the DME and is indicated to the PHY via PLME access to the PHY PIB attribute *phyUseExtendedMode* (see Table 100).

The FLP part of the preamble sequence (1010...) shall be used in the extended preamble mode, as shown in Figure 39. Since idle time is not an integral multiple of the preamble, only a fraction of the preamble pattern such as '010' can be sent to complete the idle time.

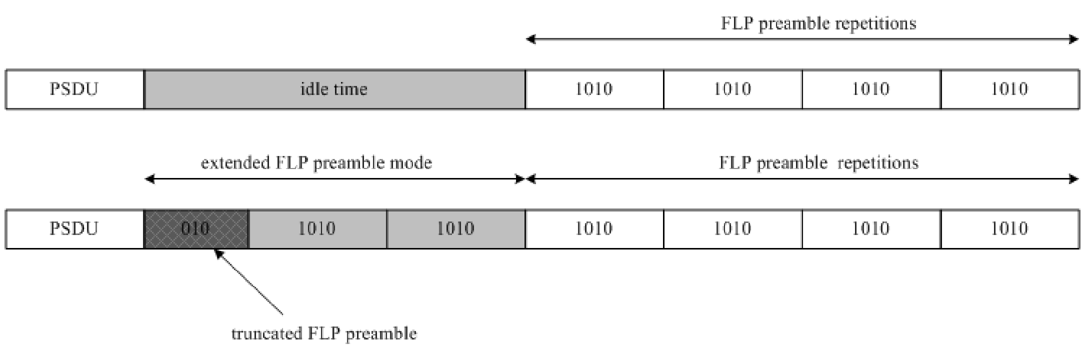

Figure 39—Truncated preamble in extended preamble mode for utilizing idle time for visibility

Note: the dimming requirements shall be met, even during the preambles. The implementer needs to be mindful of this when doing repetitions of the preamble.

5.1.14.3 Transmitting visibility pattern during uplink for star topology mode

For the star topology mode, assuming the visibility pattern is sent "in-band" as described in the modulation domain (see 4.3), the point-and-shoot visibility signal from the mobile device cannot be transmitted continuously since multiple users could be pointing to the infrastructure fixed coordinator. This makes the visibility signal difficult to attain due to the low duty cycle. Hence, the knowledge of idle periods (unused slots) is transmitted by the beacons and the mobile device uses the idle periods for transmitting the visibility pattern to the fixed coordinator. All mobile devices talking to a coordinator can share the empty slots for the CVD frame transmission during uplink.

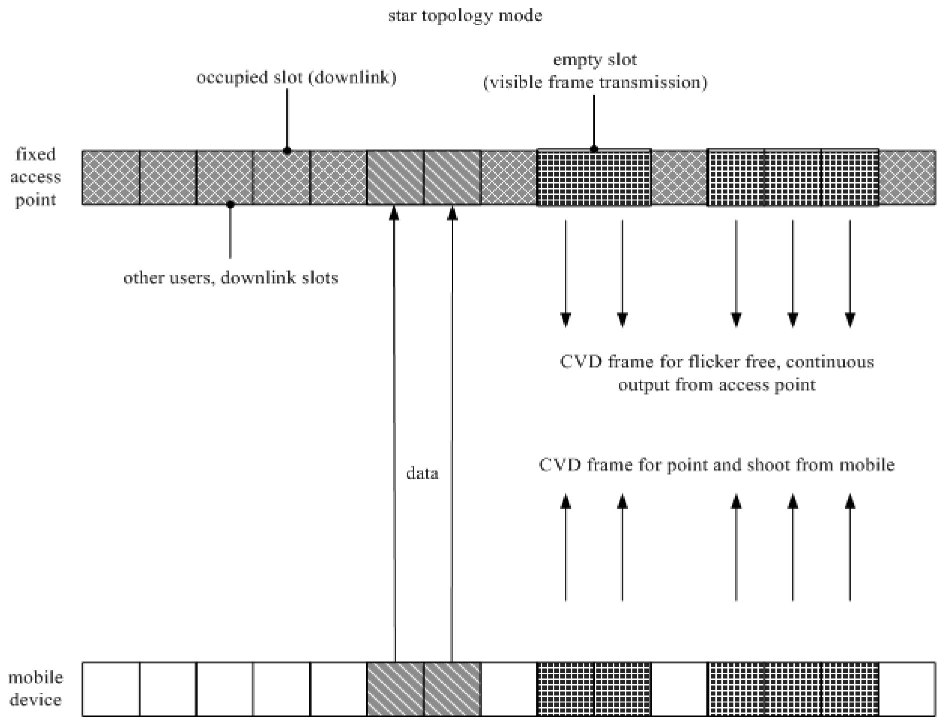

Figure 40—Usage of CVD frames during star topology operation

5.1.14.4 Dimming override capability

This standard supports bypassing the dimmer functionality during VLC operation. The dimmer control can be set to maximum brightness to facilitate VLC communication. As soon as the VLC communication is completed, the dimmer regains control of the optical source driver and resumes normal operation.

A dimmer override capability request signal is added to the MLME SAP and provided to the external dimmer interface, using the MAC PIB attribute, *macDimOverrideRequest*, as shown in Table 60. This dimmer override request attribute shall be set to '1' during VLC operation and shall be set to '0' after the communication has been completed. The dimmer circuit can decide whether to accept or reject this request. The response to this dimmer override request signal by the external dimmer circuit is out-of-scope of this standard. The MLME-GET (see 6.3.4) and MLME-SET(see 6.3.10) primitives are used to read and write PIB attributes for dimming.

5.1.14.5 PWM signal override

A PWM signal override request signal is added to the MLME SAP, using the MAC PIB attribute, *macDimPWMOverrideRequest*, as defined in Table 60 and provided to the external dimmer interface. This PWM override request attribute shall be set to '1' to inform the dimmer circuit that the VLC PHY will be responsible for dimming and to disable any PWM circuit present in the dimmer. The duty cycle for dimming is then driven by modulation mode provided by the VLC PHY (such as VPPM). The response to this PWM override request signal by the external dimmer circuit is out-of-scope of this standard. The MLME-GET (see 6.3.4) and MLME-SET (see 6.3.10) primitives are used to read and write PIB attributes for dimming.

5.1.14.6 MAC layer transmission adjustment for dimming

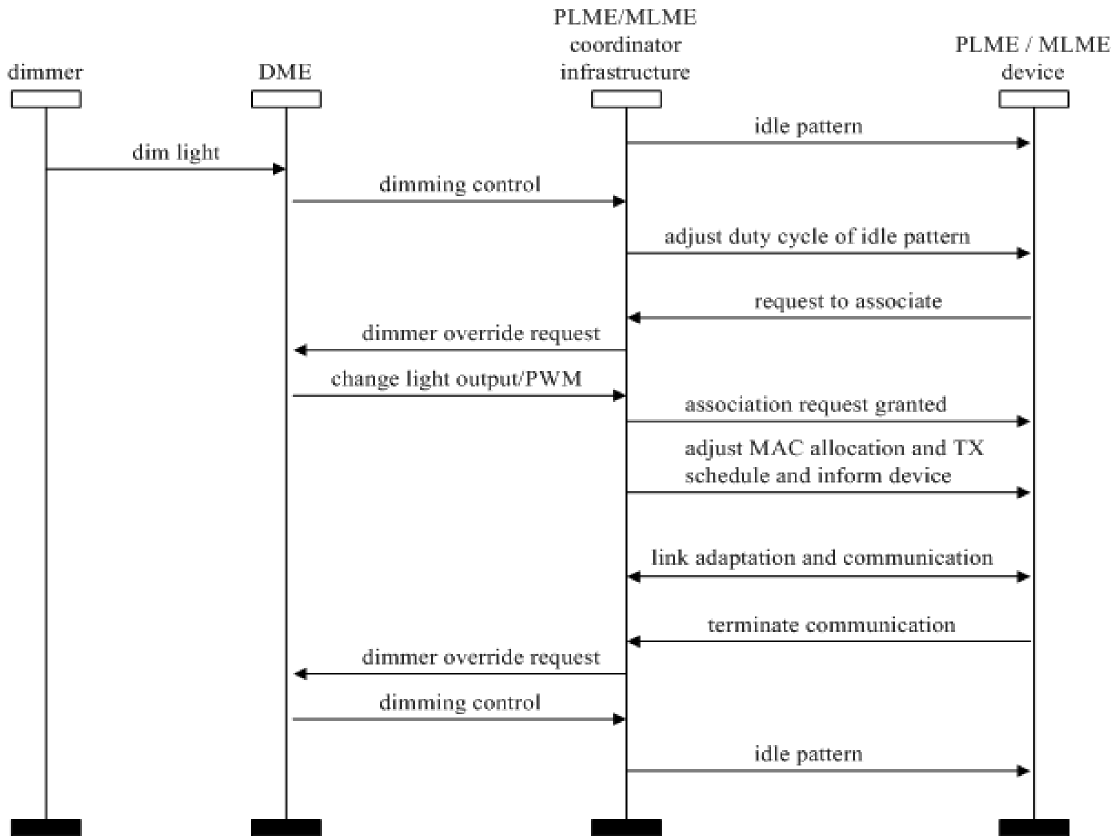

Figure 41—MSC for dimming

Referring to Figure 41, the infrastructure MAC adjusts the data transmission to match the duty cycle requirements from the dimmer.

5.1.14.7 Device discovery and association in the presence of dimming and visibility

The visibility pattern can help with device discovery when the idle pattern or the data has been modified because of the PHY and MAC layer modulation changes to support dimming. Based on the dimming pattern change and duty cycle, the VLC device may choose to associate with a different coordinator that is currently not being dimmed or has a higher duty cycle (more illumination). The visibility pattern is uncoded as shown in Figure 59. The header for the CVD frame is sent at the lowest data rate corresponding to the currently negotiated clock rate. Figure 42 shows an example of using the visibility pattern as a signal to establish the best connectivity to an infrastructure device.

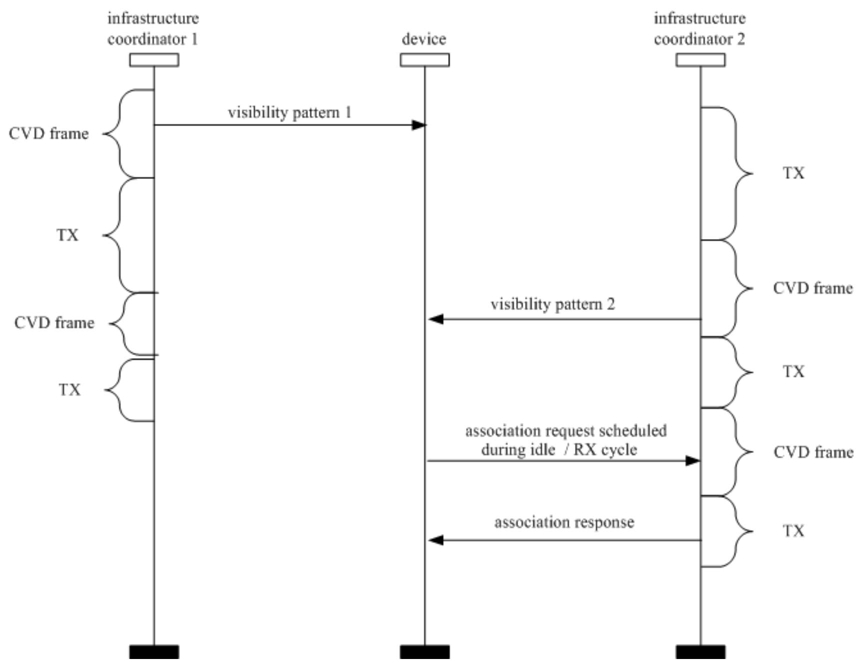

Figure 42—Example of using the visibility pattern to establish best connectivity to an infrastructure device

5.1.14.8 Link adaptation for dimming support

Dimming requirements of the infrastructure should be notified to VLC RX device, so that the VLC receiver can adapt to the dimming pattern of the data when VPPM is used. The infrastructure coordinator may receive an external dimming request. A dimming adaptation timer is used that delays the time between the dimming request and the actual dimming of the light source. With this knowledge of an incoming dimming, the link between the devices can be adapted to work at a new (lower) data rate (if dimmed) without requiring the link to be interrupted or possible link failure. This dimming adaptation is indicated and supported by the MAC dimming notification command frame in 5.3.10. Figure 43 shows an example of delay dimming and adapt resources for uninterrupted link.

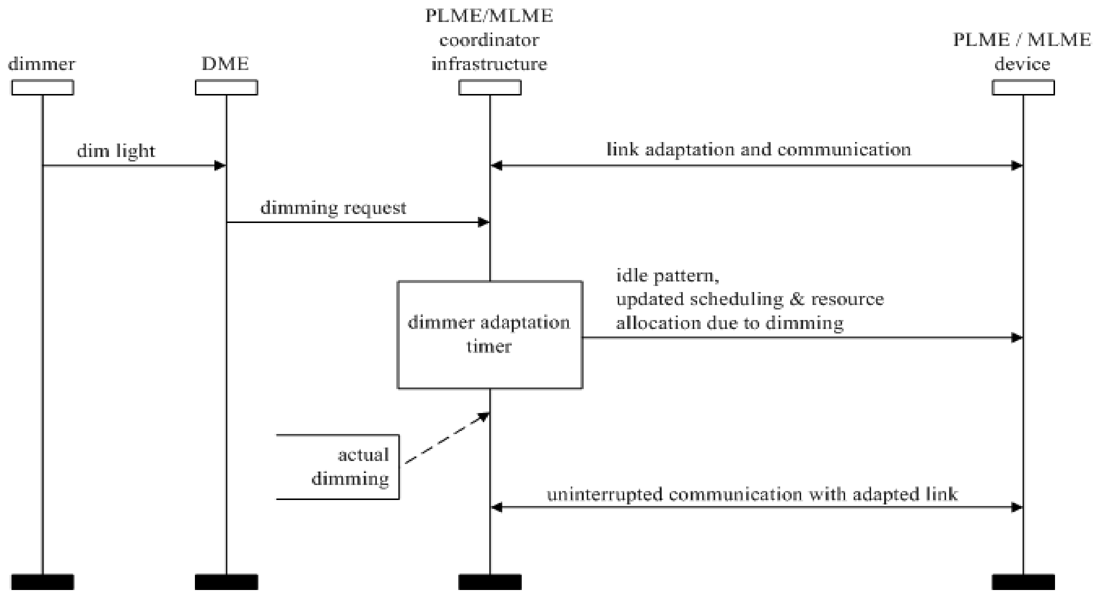

Figure 43—Usage of MAC layer to delay dimming and adapt resources for uninterrupted link

5.2 MAC frame formats

This subclause specifies the format of the MAC frame (MPDU). Each MAC frame consists of the following basic components:

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

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

For every MAC frame, all reserved bits shall be ignored upon receipt.

5.2.1 General MAC frame format

The MAC frame format is composed of a MHR, a MSDU, and a MFR. The fields of the MHR appear in a fixed order; however, the addressing fields may not be included in all frames. The general MAC frame shall be formatted as illustrated in Figure 44.

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

Figure 44—General MAC frame format

5.2.1.1 Frame control field

The frame control field is 2 octets in length and contains information defining the frame type, addressing fields, and other control flags. The frame control field shall be formatted as illustrated in Figure 45. Reserved bits are set to zero on transmission and ignored on reception.

Bits: 0-1	2-5	6–8	9	10	11	12–13	14–15
Frame Version	Reserved	Frame Type	Security Enabled	Frame Pending	Ack Request	Dest Addressing Mode	Source Addressing Mode

Figure 45—Format of the frame control field

5.2.1.1.1 Frame version subfield

The Frame Version subfield specifies the version number corresponding to the frame. This subfield shall be set to 0b00 to indicate a frame compatible with IEEE Std 802.15.7. All other subfield values shall be reserved for future use.

5.2.1.1.2 Frame type subfield

The Frame Type subfield shall be set to one of the nonreserved values listed in Table 7.

Table 7—Values of the Frame Type subfield

Frame type value b_2 b_1 b_0	Description
000	Beacon
001	Data

Table 7—Values of the Frame Type subfield

Frame type value b_2 b_1 b_0	Description
010	Acknowledgment
011	Command
100	CVD
101–111	Reserved

5.2.1.1.3 Security enabled subfield

The Security Enabled subfield is 1 bit in length, and it shall be set to one if the frame is protected by the MAC sublayer and shall be set to zero otherwise. The Auxiliary Security Header field of the MHR shall be present only if the Security Enabled subfield is set to one.

5.2.1.1.4 Frame pending subfield

The frame pending subfield is 1 bit in length and shall be set to one if the device sending the frame has more data for the recipient. This subfield shall be set to zero otherwise (see 5.1.7.3).

The frame pending subfield shall be used only in beacon frames or frames transmitted either during the CAP by devices operating on a beacon-enabled VPAN or at any time by devices operating on a nonbeacon-enabled VPAN.

At all other times, it shall be set to zero on transmission and ignored on reception.

5.2.1.1.5 Acknowledgment request subfield

The Acknowledgment Request subfield is 1 bit in length and specifies whether an acknowledgment is required from the recipient device on receipt of a data or MAC command frame. If this subfield is set to one, the recipient device shall send an acknowledgment frame only if, upon reception, the frame passes the third level of filtering as shown in 5.1.7.2. If this subfield is set to zero, the recipient device shall not send an acknowledgment frame.

5.2.1.1.6 Destination addressing mode subfield

The Destination Addressing Mode subfield shall be set to one of the nonreserved values listed in Table 8.

If this subfield is equal to zero and the Frame Type subfield does not specify that this frame is an acknowledgment or beacon frame, the Source Addressing Mode subfield shall be nonzero, implying that the frame is directed to the VLC coordinator with the VPAN identifier as specified in the Source VPAN Identifier field. If this subfield is equal to 01, the Source Addressing Mode subfield shall be equal to 01, implying that the frame is a broadcast frame, and no source or destination address fields are present in the frame.

5.2.1.1.7 Source addressing mode subfield

The Source Addressing Mode subfield shall be set to one of the nonreserved values listed in Table 8.

Table 8—Possible values of the Destination Addressing Mode and Source Addressing Mode subfields

Addressing mode value b_1 b_0	Description
00	VPAN identifier and address fields are not present.
01	No address field (broadcast only mode with no address fields present). Addresses with all ones of 16 or 64 bits are defined as broadcast.
10	Address field contains a 16-bit short address.
11	Address field contains a 64-bit extended address.

If this subfield is equal to zero and the Frame Type subfield does not specify that this frame is an acknowledgment frame, the Destination Addressing Mode subfield shall be nonzero, implying that the frame has originated from the coordinator with the VPAN identifier as specified in the Destination VPAN Identifier field.

If this subfield is equal to 01, the Source Addressing Mode subfield shall be equal to 01, implying that the frame is a broadcast frame, and no source or destination address fields are present in the frame.

5.2.1.2 Sequence number field

The Sequence Number field is 1 octet in length and specifies the sequence identifier for the frame.

For a beacon frame, the Sequence Number field shall specify a BSN. For a data, acknowledgment, or MAC command frame, the Sequence Number field shall specify a DSN that is used to match an acknowledgment frame to the data or MAC command frame.

5.2.1.3 Destination VPAN identifier field

The Destination VPAN Identifier field, when present, is 2 octets in length and specifies the unique VPAN identifier of the intended recipient of the frame. A value of 0xffff in this field shall represent the broadcast VPAN identifier, which shall be accepted as a valid VPAN identifier by all devices currently listening to the channel.

This field shall be included in the MAC frame only if the Destination Addressing Mode subfield of the frame control field is 10 or 11.

5.2.1.4 Destination address field

The Destination Address field, when present, is either 2 octets or 8 octets in length, according to the value specified in the Destination Addressing Mode subfield of the frame control field, see 5.2.1.1.6, and specifies the address of the intended recipient of the frame. A 16-bit value of 0xffff in this field shall represent the broadcast short address, which shall be accepted as a valid 16-bit short address by all devices currently listening to the channel.

This field shall be included in the MAC frame only if the Destination Addressing Mode subfield of the frame control field is nonzero.

5.2.1.5 Source VPAN identifier field

The Source VPAN Identifier field, when present, is 2 octets in length and specifies the unique VPAN identifier of the originator of the frame. This field shall be included in the MAC frame only if the Source Addressing Mode and VPAN ID Compression subfields of the frame control field are nonzero and equal to zero, respectively.

The VPAN identifier of a device is initially determined during association on a VPAN, but may change following a VPAN identifier conflict resolution as discussed in 5.1.3.

5.2.1.6 Source address field

The Source Address field, when present, is either 2 octets or 8 octets in length, according to the value specified in the Source Addressing Mode subfield of the frame control field, as shown in 5.2.1.1.7, and specifies the address of the originator of the frame. This field shall be included in the MAC frame only if the Source Addressing Mode subfield of the frame control field is 10 or 11.

5.2.1.7 Auxiliary security header field

The Auxiliary Security Header field has a variable length and specifies information required for security processing, including how the frame is actually protected (security level) and which keying material from the MAC security PIB is used (see 7.5.1). This field shall be present only if the Security Enabled subfield is set to one. For details on formatting, see 7.4.

5.2.1.8 Frame payload field

The Frame Payload field has a variable length and contains information specific to individual frame types. If the Security Enabled subfield is set to one in the frame control field, the frame payload is protected as defined by the security suite selected for that frame.

5.2.1.9 FCS field

The FCS field is 2 octets in length and is explained in Annex C. The FCS is calculated over the MHR and MSDU parts of the frame. The FCS shall be only generated for payloads greater than zero bytes.

5.2.2 Format of individual frame types

Five frame types are defined: beacon, data, acknowledgment, command and CVD. These frame types are discussed in 5.2.2.1 through 5.2.2.4.3.

5.2.2.1 Beacon frame format

The beacon frame shall be formatted as illustrated in Figure 46.

The GTS fields shall be formatted as illustrated in Figure 47, and the pending address fields shall be formatted as illustrated in Figure 48.

The order of the fields of the beacon frame shall conform to the order of the general MAC frame as illustrated in Figure 44.

5.2.2.1.1 Beacon frame MHR fields

The MHR for a beacon frame shall contain the frame control field, the Sequence Number field, the Source VPAN Identifier field, and the Source Address field.

Octets: 2	1	4/10	0/5/6/10/14	2	variable	variable	0/1	variable	2
Frame Control	Sequence Number	Addressing fields	Auxiliary Security Header	Superframe Spec	GTS fields (Figure 47)	Pending address fields (Figure 48)	cellSearch Length	Beacon Payload	FCS
MHR				MSDU					MFR

Figure 46—Beacon frame format

Octets: 1	0/1	variable
GTS Specification	GTS Directions	GTS List

Figure 47—Format of the GTS information fields

Octets: 1	variable
Pending Address Specification	Address List

Figure 48—Format of the pending address information fields

In the frame control field, the Frame Type subfield shall contain the value that indicates a beacon frame, as shown in Table 7, and the Source Addressing Mode subfield shall be set as appropriate for the address of the coordinator transmitting the beacon frame. If protection is used for the beacon, the Security Enabled subfield shall be set to one. If a broadcast data or command frame is pending, the frame pending subfield shall be set to one. All other subfields shall be set to zero by the sender and ignored on reception.

The Sequence Number field shall contain the current value of *macBSN*.

The addressing fields shall comprise only the source address fields. The Source VPAN Identifier and Source Address fields shall contain the VPAN identifier and address, respectively, of the device transmitting the beacon.

The Auxiliary Security Header field, if present, shall contain the information required for security processing of the beacon frame, as specified in 5.2.1.7.

5.2.2.1.2 Superframe specification field

The Superframe Specification field shall be formatted as illustrated in Figure 49.

Bits: 0–3	4-7	8-11	12	13	14	15
Beacon Order	Superframe Order	Final CAP Slot	Reserved	VPAN Coordinator	Association Permit	cellSearchEn

Figure 49—Format of the Superframe Specification field

The Beacon Order subfield shall specify the transmission interval of the beacon. Refer to 5.1.1.1 for an explanation of the relationship between the beacon order and the beacon interval.

The Superframe Order subfield shall specify the length of time during which the superframe is active (i.e., receiver enabled), including the beacon frame transmission time. Refer to 5.1.1.1 for an explanation of the relationship between the superframe order and the superframe duration.

The Final CAP Slot subfield specifies the final superframe slot utilized by the CAP. The duration of the CAP, as implied by this subfield, shall be greater than or equal to the value specified by *aMinCAPLength*. However, an exception is allowed for the accommodation of the temporary increase in the beacon frame length needed to perform GTS maintenance, as in 5.2.2.1.3.

The VPAN Coordinator subfield shall be set to one if the beacon frame is being transmitted by the coordinator. Otherwise, the VPAN Coordinator subfield shall be set to zero.

The Association Permit subfield shall be set to one if *macAssociationPermit* is set to TRUE (i.e., the coordinator is accepting association to the VPAN). The association permit bit shall be set to zero if the coordinator is currently not accepting association requests on its network.

If the cellSearchEn bit is set, the cellSearchLength is transmitted as an additional field in the beacon frame, as shown in Figure 46.

5.2.2.1.3 GTS specification field

The GTS Specification field shall be formatted as illustrated in Figure 50.

Bits: 0-2	3-6	7
GTS Descriptor Count	Reserved	GTS Permit

Figure 50—Format of the GTS Specification field

The GTS Descriptor Count subfield specifies the number of 3-octet GTS descriptors contained in the GTS List field of the beacon frame. If the value of this subfield is greater than zero, the size of the CAP shall be allowed to dip below *aMinCAPLength* to accommodate the temporary increase in the beacon frame length caused by the inclusion of the subfield. If the value of this subfield is zero, the GTS Directions field and GTS List field of the beacon frame are not present.

The GTS Permit subfield shall be set to one if *macGTSPermit* is equal to TRUE (i.e., the coordinator is accepting GTS requests). Otherwise, the GTS Permit field shall be set to zero.

5.2.2.1.4 GTS directions field

The GTS Directions field shall be formatted as illustrated in Figure 51.

Bits: 0-6	7
GTS Directions Mask	Reserved

Figure 51—Format of the GTS Directions field

The GTS Directions Mask subfield contains a mask identifying the directions of the GTSs in the superframe. The lowest bit in the mask corresponds to the direction of the first GTS contained in the GTS List field of the beacon frame, with the remainder appearing in the order that they appear in the list. Each bit shall be set to one if the GTS is a receive-only GTS or to zero if the GTS is a transmit-only GTS. GTS direction is defined relative to the direction of the data frame transmission by the device.

5.2.2.1.5 GTS list field

The size of the GTS List field is defined by the values specified in the GTS Specification field of the beacon frame and contains the list of GTS descriptors that represents the GTSs that are being maintained. The maximum number of GTS descriptors shall be limited to seven.

Each GTS descriptor shall be formatted as illustrated in Figure 52.

Bits: 0-15	16-19	20-23
Device Short Address	GTS Starting Slot	GTS Length

Figure 52—Format of the GTS descriptor

The Device Short Address subfield shall contain the short address of the device for which the GTS descriptor is intended.

The GTS Starting Slot subfield contains the superframe slot at which the GTS is to begin.

The GTS Length subfield contains the number of contiguous superframe slots over which the GTS is active.

5.2.2.1.6 Pending address specification field

The Pending Address Specification field shall be formatted as illustrated in Figure 53.

Bits: 0-2	3	4-6	7
Number of Short Addresses Pending	Reserved	Number of Extended Addresses Pending	Reserved

Figure 53—Format of the Pending Address Specification field

The Number of Short Addresses Pending subfield indicates the number of 16-bit short addresses contained in the Address List field of the beacon frame.

The Number of Extended Addresses Pending subfield indicates the number of 64-bit extended addresses contained in the Address List field of the beacon frame.

5.2.2.1.7 Address list field

The size of the Address List field is determined by the values specified in the Pending Address Specification field of the beacon frame and contains the list of addresses of the devices that currently have messages pending with the coordinator. The address list shall not contain the broadcast short address 0xffff.

The maximum number of addresses pending shall be limited to seven and may comprise both short and extended addresses. All pending short addresses shall appear first in the list followed by any extended addresses. If the coordinator is able to store more than seven transactions, it shall indicate them in its beacon on a first-come-first-served basis, ensuring that the beacon frame contains at most seven addresses.

5.2.2.1.8 Beacon payload field

The Beacon Payload field is an optional sequence of up to *aMaxBeaconPayloadLength* octets specified to be transmitted in the beacon frame by the next higher layer. The set of octets contained in *macBeaconPayload* shall be copied into this field.

5.2.2.2 Data frame format

The data frame shall be formatted as illustrated in Figure 54.

Octets: 2	1	(as defined in 5.2.2.2.1)	0/5/6/10/14	variable	2
frame control	Sequence Number	Addressing fields	Auxiliary Security Header	Data Payload	FCS
MHR				MSDU	MFR

Figure 54—Data frame format

The order of the fields of the data frame shall conform to the order of the general MAC frame as illustrated in Figure 44.

5.2.2.2.1 Data frame MHR fields

The MHR for a data frame shall contain the frame control field, the Sequence Number field, the destination VPAN identifier/address fields, and/or the source VPAN identifier/address fields.

In the frame control field, the Frame Type subfield shall contain the value that indicates a data frame, as shown in Table 7. If protection is used for the data, the Security Enabled subfield shall be set to one. All other subfields shall be set appropriately according to the intended use of the data frame. All reserved subfields shall be set to zero by the sender and ignored on reception.

The Sequence Number field shall contain the current value of *macDSN*.

The addressing fields shall comprise the destination address fields and/or the source address fields, dependent on the settings in the frame control field.

The Auxiliary Security Header field, if present, shall contain the information required for security processing of the data frame, as specified in 5.2.1.7.

5.2.2.2.2 Data payload field

The payload of a data frame shall contain the sequence of octets that the next higher layer has requested the MAC sublayer to transmit. The data type field is 1 byte and is explained in Table 9.

Table 9—Data Payload Field

Bits 0-1	Bits 2-7	variable
00 - Single 01 - Packed 10 - Burst 11 - Reserved	Number of PPDUs per data frame	Data payload

The data type field mentions the format used for sending the data - single, packed or burst. It also mentions the number of PPDUs that are associated for this data frame.

The payload of a data frame shall contain the sequence of octets that the next higher layer has requested the MAC sub layer to transmit.

5.2.2.3 Acknowledgment frame format

The acknowledgment frame shall be formatted as illustrated in Figure 55.

Octets: 2	1	variable	2
frame control	sequence number	B-ACK frame payload (optional)	FCS
MHR		MSDU	MFR

Figure 55—Acknowledgment frame format

The order of the fields of the acknowledgment frame shall conform to the order of the general MAC frame as illustrated in Figure 44. The sequence number is defined in 5.2.2.1.1.

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 5.2.2.3. The B-ACK frame payload is defined in Figure 56.

octets: 2	1	1	2	0-n
buffer size	frame count	reserved	sequence control	frame bitmap

Figure 56—B-ACK frame payload

The Buffer Size field specifies the maximum number of octets in the sum of the frame payloads of all frames in the next B-ACK sequence. The Frame Count field specifies the maximum number of frames in the next B-ACK sequence. The Sequence Control and frame bitmap fields together specify an acknowledgment window of MSDU fragments and their reception status. The Sequence Control field specifies the Sequence Number and Fragment Number that start the acknowledgment window.

bits: b15-b14	b13-b3	b2-b0
reserved	sequence number	fragment number

Figure 57—B-ACK frame bitmap

The frame bitmap field varies in length. A zero-length frame bitmap field indicates an acknowledgment 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 acknowledgment 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 acknowledgment 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 block ACK is applicable to the packed data type. The bitmap and sequence number is repeated for every frame in the burst mode (multiple frames)

The order of the fields of the acknowledgment frame shall conform to the order of the general MAC frame as illustrated.

The MHR for an acknowledgment frame shall contain only the frame control field and the Sequence Number field.

In the frame control field, the Frame Type subfield shall contain the value that indicates an acknowledgment frame, as shown in Table 7. If the acknowledgment frame is being sent in response to a received data request command, the device sending the acknowledgment frame shall determine whether it has data pending for the recipient. If the device can determine this before sending the acknowledgment frame (see 5.1.7.4.2), it shall set the frame pending subfield according to whether there is pending data. Otherwise, the frame pending subfield shall be set to one. If the acknowledgment frame is being sent in response to either a data frame or another type of MAC command frame, the device shall set the frame pending subfield to zero. All other subfields, except the security enabled subfield, shall be set to zero by the sender and ignored on reception.

The Sequence Number field shall contain the value of the sequence number received in the frame for which the acknowledgment is to be sent.

5.2.2.4 Command frame format

The command frame shall be formatted as illustrated in Figure 58.

The order of the fields of the MAC command frame shall conform to the order of the general MAC frame as illustrated in Figure 44.

5.2.2.4.1 MAC command frame MHR fields

The MHR for a MAC command frame shall contain the frame control field, the Sequence Number field, the destination VPAN identifier/address fields and/or the source VPAN identifier/address fields.

Octets: 2	1	(as defined in 5.2.2.4.1)	0/5/6/10/14	1	variable	2
Frame Control	Sequence Number	Addressing fields	Auxiliary Security Header	Command Frame Identifier	Command Payload	FCS
MHR				MSDU		MFR

Figure 58—Command frame format

In the frame control field, the Frame Type subfield shall contain the value that indicates a MAC command frame, as shown in Table 7. If the frame is to be secured, the Security Enabled subfield of the frame control field shall be set to one and the frame secured according to the process described in 7.5.4. Otherwise the Security Enabled subfield of the frame control field shall be set to zero. All other subfields shall be set appropriately according to the intended use of the MAC command frame. All reserved subfields shall be set to zero by the sender and ignored on reception.

The Sequence Number field shall contain the current value of *macDSN*.

The addressing fields shall comprise the destination address fields and/or the source address fields, dependent on the settings in the frame control field.

The Auxiliary Security Header field, if present, shall contain the information required for security processing of the MAC command frame, as specified in 5.2.1.7.

5.2.2.4.2 Command frame Identifier field

The Command Frame Identifier field identifies the MAC command being used. This field shall be set to one of the nonreserved values listed in Table 10.

5.2.2.4.3 Command payload field

The Command Payload field contains the MAC command itself. The formats of the individual commands are described in 5.3.

5.2.2.5 CVD frame format

Octets: 2	Octet: 2	variable
Frame control	FCS	Visibility pattern
MHR	MFR	

Figure 59—CVD frame

The structure of the CVD frame is as shown in Figure 59. The CVD frame is used to visually provide information on the communication status, such as misalignment between the two devices, transmission direction or sending data status; the data transmission quality; the transferred file size and remaining file size. The visibility pattern has no error protection. The length of the visibility pattern shall be set in the PHY header and the FCS shall not include the visibility pattern of the CVD frame. The FCS is only calculated over the frame control field (MHR) using the CRC described in Annex C. The visibility pattern will be generated based on the dimming level requirements and is described in 8.5.1.2. The CVD frame is used by the infrastructure to maintain visibility at all times and by the mobile device for point-and-shoot. The CVD frame can also be used for color stabilization for PHY III as explained in 8.5.4. It should be noted that the CVD frame is not used for communicating the dimming level; rather, the dimming notification command is used for this function as described in 5.3.10.

The CVD frame is sent at the currently negotiated optical clock.

5.3 MAC Command frames

The command frames defined by the MAC sublayer are listed in Table 10. A coordinator shall be capable of transmitting and receiving all command frame types, with the exception of the GTS request command, while the requirements for a device are indicated by an “X” in Table 10. A P2P device functioning as a coordinator shall be capable of transmitting and receiving all supported command frames in a device. MAC commands shall only be transmitted in the CAP for beacon-enabled VPANs or at any time for nonbeacon-enabled VPANs.

How the MLME shall construct the individual commands for transmission is detailed in 5.3.1 through 5.3.18. MAC command reception shall abide by the procedure described in 5.1.7.2.

Table 10—Command frames

Command frame identifier	Command name	device		P2P coordinator		Subclause
		Tx	Rx	Tx	Rx	
0x01	Association request	X		X	X	5.3.1
0x02	Association response		X	X	X	5.3.2
0x03	Disassociation notification	X	X	X	X	5.3.3
0x04	Data request	X		X	X	5.3.4
0x05	VPAN ID conflict notification	X		X	X	5.3.5
0x06	Beacon request					5.3.6
0x07	Coordinator realignment		X	X	X	5.3.7
0x08	GTS request					5.3.8
0x09	Blinking notification					5.3.9
0x0a	Dimming notification		X	X	X	5.3.10
0x0b	Fast link recovery					5.3.11
0x0c	Mobility notification					5.3.12
0x0d	GTS Response					5.3.13
0x0e	Clock rate change notification		X	X	X	5.3.14
0x0f	Multiple channel assignment					5.3.15
0x10	Band hopping					5.1.10.2

Table 10—Command frames *(continued)*

Command frame identifier	Command name	device		P2P coordinator		Subclause
		Tx	Rx	Tx	Rx	
0x11	Color stabilization timer notification	X	X			5.3.16
0x12	Color stabilization information	X	X			5.3.17
0x13	CVD disable					5.3.18
0x14	Information element	X	X	X	X	5.3.19
0x15–0xff	Reserved					—

5.3.1 Association request command

The association request command allows a device to request association with a VPAN through the coordinator. This command shall only be sent by an unassociated device that wishes to associate with a VPAN. A device shall only associate with a VPAN through the coordinator as determined through the scan procedure.

All devices shall be capable of transmitting this command, although a device is not required to be capable of receiving it.

The association request command shall be formatted as illustrated in Figure 60.

octets: (see 5.2.2.4)	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	Capability Information

Figure 60—Association request command format

5.3.1.1 MHR fields

The Source Addressing Mode subfield of the frame control field shall be set to three (64-bit extended addressing). The Destination Addressing Mode subfield shall be set to the same mode as indicated in the beacon frame to which the association request command refers.

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

The Destination VPAN Identifier field shall contain the identifier of the VPAN to which to associate. The Destination Address field shall contain the address from the beacon frame that was transmitted by the coordinator to which the association request command is being sent. The Source VPAN Identifier field shall contain the broadcast VPAN identifier (i.e., 0xffff). The Source Address field shall contain the value of *aExtendedAddress*.

5.3.2 Association response command

The association response command allows the coordinator or a coordinator to communicate the results of an association attempt back to the device requesting association.

This command shall only be sent by the coordinator or a coordinator to a device that is currently trying to associate.

All devices shall be capable of receiving this command, although a device is not required to be capable of transmitting it.

The association response command shall be formatted as illustrated in Figure 61.

octets: (see 5.2.2.4)	1	2	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	Short Address	Association Status	Capability negotiation response

Figure 61—Association response command format

The capability-negotiation response is the same as that of the color-stabilization scheme in Table 20.

5.3.2.1 MHR fields

The Destination Addressing Mode and Source Addressing Mode subfields of the frame control field shall each be set to three (i.e., 64-bit extended addressing).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

The VPAN ID Compression subfield of the frame control field shall be set to one. In accordance with this value of the VPAN ID Compression subfield, the Destination VPAN Identifier field shall contain the value of *macVPANId*, while the Source VPAN Identifier field shall be omitted. The Destination Address field shall contain the extended address of the device requesting association. The Source Address field shall contain the value of *aExtendedAddress*.

5.3.2.2 Short address field

If the coordinator was not able to associate this device to its VPAN, the Short Address field shall be set to 0xffff, and the Association Status field shall contain the reason for the failure. If the coordinator was able to associate the device to its VPAN, this field shall contain the short address that the device may use in its communications on the VPAN until it is disassociated.

A Short Address field value equal to 0xfffe shall indicate that the device has been successfully associated with a VPAN, but has not been allocated a short address. In this case, the device shall communicate on the VPAN using only its 64-bit extended address.

5.3.2.3 Association status field

The Association Status field shall contain one of the nonreserved values listed in Table 11.

Table 11—Valid values of the Association Status field

Association status	Description
0x00	Association successful.
0x01	VPAN at capacity.
0x02	VPAN access denied.
0x03–0x7f	Reserved.
0x80–0xff	Reserved for MAC primitive enumeration values.

5.3.2.4 Capability negotiation response field

The capability negotiation response field describes if and where (device and/or coordinator) color stabilization is performed. All allowed settings are shown in Table 12.

Table 12—Capability negotiation response field

Bits: 0-1	2-7
00: No color stabilization 01: Color-stabilization information to be sent from device to coordinator upon reception of CVD frames 10: Color-stabilization information to be sent from coordinator to device upon reception of CVD frames 11: Color-stabilization information to be sent from device to coordinator and from coordinator to device when either receives CVD frames	Reserved

5.3.3 Disassociation notification command

The VLC coordinator or an associated device may send the disassociate notification command.

All devices shall implement this command.

The disassociation notification command shall be formatted as illustrated in Figure 62.

octets: (see 5.2.2.4)	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	Disassociation Reason

Figure 62—Disassociation notification command format

5.3.3.1 MHR fields

The Destination Addressing Mode subfield of the frame control field shall be set according to the addressing mode specified by the corresponding primitive. The Source Addressing Mode subfield shall be set to three (i.e., 64-bit extended addressing).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

The VPAN ID Compression subfield of the frame control field shall be set to one. In accordance with this value of the VPAN ID Compression subfield, the Destination VPAN Identifier field shall contain the value of *macVPANId*, while the Source VPAN Identifier field shall be omitted. If the coordinator wants an associated device to leave the VPAN, then the Destination Address field shall contain the address of the device being removed from the VPAN. If an associated device wants to leave the VPAN, then the Destination Address field shall contain the value of either *macCoordShortAddress*, if the Destination Addressing Mode subfield is equal to two, or *macCoordExtendedAddress*, if the Destination Addressing Mode subfield is equal to three. The Source Address field shall contain the value of *aExtendedAddress*.

5.3.3.2 Disassociation reason field

The Disassociation Reason field shall contain one of the nonreserved values listed in Table 13.

Table 13—Valid disassociation reason codes

Disassociate reason	Description
0x00	Reserved.
0x01	The coordinator wishes the device to leave the VPAN.
0x02	The device wishes to leave the VPAN.
0x03	Device can not support communications for the requested dimming value
0x04f–0x7f	Reserved.
0x80–0xff	Reserved for MAC primitive enumeration values.

5.3.4 Data request command

The data request command is sent by a device to request data from the coordinator.

There are three cases for which this command is sent. On a beacon-enabled VPAN, this command shall be sent by a device when *macAutoRequest* is equal to TRUE and a beacon frame indicating that data are pending for that device is received from its coordinator. The coordinator indicates pending data in its beacon frame by adding the address of the recipient of the data to the Address List field. This command shall also be sent when instructed to do so by the next higher layer on reception of the MLME-POLL.request primitive. In addition, a device may send this command to the coordinator *macResponseWaitTime* optical clocks after the acknowledgment to an association request command.

All devices shall be capable of transmitting this command, although a device is not required to be capable of receiving it.

The data request command shall be formatted as illustrated in Figure 63.

octets: (see 5.2.2.4)	1
MHR fields	Command Frame Identifier (as defined in Table 10)

Figure 63—Data request command format

If the data request command is being sent in response to the receipt of a beacon frame indicating that data are pending for that device, the Destination Addressing Mode subfield of the frame control field may be set to zero (i.e., destination addressing information not present) if the beacon frame indicated in its Superframe Specification field (see 5.2.2.1.2) that it originated from the coordinator (see 5.2.1.1.6) or set otherwise according to the coordinator to which the data request command is directed. If the destination addressing information is to be included, the Destination Addressing Mode subfield shall be set according to the value of *macCoordShortAddress*. If *macCoordShortAddress* is equal to 0xfffe, extended addressing shall be used: the Destination Addressing Mode subfield shall be set to three, and the Destination Address field shall contain the value of *macCoordExtendedAddress*. Otherwise, short addressing shall be used: the Destination Addressing Mode subfield shall be set to two, and the Destination Address field shall contain the value of *macCoordShortAddress*.

If the data request command is being sent in response to the receipt of a beacon frame indicating that data are pending for that device, the Source Addressing Mode subfield shall be set according to the addressing mode used for the pending address. If the Source Addressing Mode subfield is set to two, short addressing shall be used: the Source Address field shall contain the value of *macShortAddress*. Otherwise, extended addressing shall be used: the Source Addressing Mode subfield shall be set to three, and the Source Address field shall contain the value of *aExtendedAddress*.

If the data request command is triggered by the reception of an MLME-POLL.request primitive from the next higher layer, then the destination addressing information shall be the same as that contained in the primitive. The Source Addressing Mode subfield shall be set according to the value of *macShortAddress*. If *macShortAddress* is less than 0xfffe, short addressing shall be used. Extended addressing shall be used otherwise.

If the data request command is being sent following the acknowledgment to an association request command frame, the Destination Addressing Mode subfield of the frame control field shall be set according to the coordinator to which the data request command is directed. If *macCoordShortAddress* is equal to 0xfffe, extended addressing shall be used. Short addressing shall be used otherwise. The Source Addressing Mode subfield shall be set to use extended addressing.

If the Destination Addressing Mode subfield is set to zero (i.e., destination addressing information not present), the VPAN ID Compression subfield of the frame control field shall be set to zero and the source VPAN identifier shall contain the value of *macVPANId*. Otherwise, the VPAN ID Compression subfield shall be set to one. In this case and in accordance with the VPAN ID Compression subfield, the Destination VPAN Identifier field shall contain the value of *macVPANId*, while the Source VPAN Identifier field shall be omitted.

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

5.3.5 VPAN ID conflict notification command

The VPAN ID conflict notification command is sent by a device to the coordinator when a VPAN identifier conflict is detected.

All devices shall be capable of transmitting this command, although a device is not required to be capable of receiving it.

The VPAN ID conflict notification command shall be formatted as illustrated in Figure 64.

octets: (see 5.2.2.4)	**1**
MHR fields	Command Frame Identifier (as defined in Table 10)

Figure 64—VPAN ID conflict notification command format

The Destination Addressing Mode and Source Addressing Mode subfields of the frame control field shall both be set to three (i.e., 64-bit extended addressing).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

The VPAN ID Compression subfield of the frame control field shall be set to one. In accordance with this value of the VPAN ID Compression subfield, the Destination VPAN Identifier field shall contain the value of *macVPANId*, while the Source VPAN Identifier field shall be omitted. The Destination Address field shall contain the value of *macCoordExtendedAddress*. The Source Address field shall contain the value of *aExtendedAddress*.

5.3.6 Beacon request command

The beacon request command is used by a device to locate all coordinators within its operating space during an active scan.

This command is optional for a device.

The beacon request command shall be formatted as illustrated in Figure 65.

octets: 7	**1**
MHR fields	Command Frame Identifier (as defined in Table 10)

Figure 65—Beacon request command format

The Destination Addressing Mode subfield of the frame control field shall be set to two (i.e., 16-bit short addressing), and the Source Addressing Mode subfield shall be set to zero (i.e., source addressing information not present).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception. The Acknowledgment Request subfield and Security Enabled subfield shall also be set to zero.

The Destination VPAN Identifier field shall contain the broadcast VPAN identifier (i.e., 0xffff). The Destination Address field shall contain the broadcast short address (i.e., 0xffff).

5.3.7 Coordinator realignment command

The coordinator realignment command is sent by the coordinator or a coordinator when any of its VPAN configuration attributes change due to the receipt of an MLME-START.request primitive.

If this command is sent when any VPAN configuration attributes (i.e., VPAN identifier, short address, or logical channel) change, it is broadcast to the VPAN.

All devices shall be capable of receiving this command, although a device is not required to be capable of transmitting it.

The coordinator realignment command shall be formatted as illustrated in Figure 66.

octets: 17/18/23/24	1	2	2	1	2
MHR fields	Command Frame Identifier (as defined in Table 10)	VPAN Identifier	Coordinator Short Address	Logical Channel	Short Address

Figure 66—Coordinator realignment command format

5.3.7.1 MHR fields

The Destination Addressing Mode subfield of the frame control field shall be set to two (e.g., 16-bit short addressing) if it is to be broadcast to the VPAN. The Source Addressing Mode subfield of the frame control field shall be set to three (e.g., 64-bit extended addressing).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception.

The Acknowledgment Request subfield of the frame control field shall be set to zero if the command is to be broadcast to the VPAN.

The Destination VPAN Identifier field shall contain the broadcast VPAN identifier (e.g., 0xffff). The Destination Address field shall contain the broadcast short address (e.g., 0xffff). The Source VPAN Identifier field shall contain the value of *macVPANId*, and the Source Address field shall contain the value of *aExtendedAddress*.

5.3.7.2 VPAN identifier field

The VPAN Identifier field shall contain the VPAN identifier that the coordinator intends to use for all future communications.

5.3.7.3 Coordinator short address field

The Coordinator Short Address field shall contain the value of *macShortAddress*.

5.3.7.4 Logical channel field

The Logical Channel field shall contain the logical channel that the coordinator intends to use for all future communications.

5.3.7.5 Short address field

If the coordinator realignment command is broadcast to the VPAN, the Short Address field shall be set to 0xffff and ignored on reception.

5.3.8 GTS request command

The GTS request command is used by an associated device that is requesting the allocation of a new GTS or the deallocation of an existing GTS from the coordinator. Only devices that have a 16-bit short address less than 0xfffe shall send this command.

This command is optional.

The GTS request command shall be formatted as illustrated in Figure 67.

octets: 7	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	GTS Characteristics

Figure 67—GTS request command format

5.3.8.1 MHR fields

The Destination Addressing Mode subfield of the frame control field shall be set to zero (e.g., destination addressing information not present), and the Source Addressing Mode subfield shall be set to two (e.g., 16-bit short addressing).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

The Source VPAN Identifier field shall contain the value of *macVPANId*, and the Source Address field shall contain the value of *macShortAddress*.

5.3.8.2 GTS characteristics field

The GTS Characteristics field shall be formatted as illustrated in Figure 68.

bits: 0–3	4	5	6–7
GTS Length	GTS Direction	Characteristics Type	Reserved

Figure 68—GTS Characteristics field format

The GTS Length subfield shall contain the number of superframe slots being requested for the GTS.

The GTS Direction subfield shall be set to one if the GTS is to be a receive-only GTS. Conversely, this subfield shall be set to zero if the GTS is to be a transmit-only GTS. GTS direction is defined relative to the direction of data frame transmissions by the device.

The Characteristics Type subfield shall be set to one if the characteristics refers to a GTS allocation or zero if the characteristics refers to a GTS deallocation.

5.3.9 Blinking notification command

The blinking notification command is sent by a coordinator when the device is no longer responding. A reason for this might be the misalignment between the device TX and the coordinator RX (limited FOV of receiver, low device TX power, mobility of the device, etc.). In such cases, the device can change the visibility indication from continuous visibility for point-and-shoot to blinking indication. The device can then change from point-and-shoot mode to blinking mode in order to indicate to the user that the uplink to the coordinator is disconnected. This indication can be applied to both P2MP and P2P modes of operation.

octets: 7	**1**	**1**
MHR fields	Command Frame Identifier (see Table 10)	Blinking frequency

Figure 69—Blinking Notification Command

The blinking notification bit shall be set when the MAC PIB attribute, *macUseBlinkingNotification* and *macBlinkingNotificationFrequency*, as defined in Table 60 indicates the blinking notification usage.

To support the blinking notification, the frequency shall be chosen from the *phyBlinkingNotificationFrequency* PHY PIB attribute as shown in Table 100, using the MLME-SET.request and PLME-SET.request primitives.

This feature can help to align the link and is only intended for mobile devices.

5.3.9.1 Blinking frequency

The frequency subfield shall contain the frequency for blinking.

Bits: 0-3	**4-7**
Frequency	Reserved

Figure 70—Blinking frequency field format

5.3.10 Dimming notification command

The DME indicates the dimming level to the MAC using the MAC PIB attribute, *macDim*, as defined in Table 60. The dimming notification command is used to communicate the dimming level set by the *macDim* PIB attribute to the receiver. The dimming notification command shall be sent at the lowest data rate corresponding to the currently negotiated optical clock rate. The symbol shape information for VPPM is derived using the algorithm of Figure 114 after the dimming level is obtained.

octets: 7	1	2	2
MHR fields	Command frame identifier (see Table 10)	Dimming level	Dimmer adaptation timer (see 5.1.14.8)

Figure 71—Dimming notification command

The dimming level is two bytes long and contains a value between 0 and 1000, where 0 represents 0% visibility and 1000 represents 100% visibility. The dimming levels are defined with a resolution of 0.1%. The dimmer adaptation timer provides a resolution of 0-16383 MAC clock cycles. The dimming notification command transmits the dimmer level from the TX to the RX along with the dimmer adaptation timer information. VPPM by default uses only 50% duty cycle, so if dimming is supported using VPPM as in subclause 8.5.2.3, the VPPM pulse shape is obtained using the dimmer notification command in conjunction with the algorithm shown in Figure 114. Before dimming is supported using VPPM, the dimming notification command needs to be sent by the MAC to the receiver.

5.3.11 Fast link recovery command

Fast link recovery command is used for the device or coordinator to send the fast link recovery (FLR) signal and the fast link recovery response (FLR RSP), to help the link recovery.

Fast link recovery signal and response use the fast link recovery command format. The fast link recovery command shall be formatted as illustrated in Figure 72.

octets: (as defined in 5.2.2.4)	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	FLR field
		FLR field explanation bit 0: ‘0’ indicating it is FLR signal, ‘1’ indicating it is FLR RSP bits 1-3: index of FLR signal direction, if bit 0 is ‘0’. received FLR signal direction index if bit 0 is ‘1’. bits 4-7: reserved

Figure 72—Fast link recovery command

The FLR signal and the FLR response (RSP) are differentiated by the first bit (bit 0) of the FLR field in the fast link recovery command frame. The device can indicate the index of FLR signal direction by using bits 1 to 3 of the FLR field in the command frame. If the device receives the FLR signal and needs to send FLR RSP, it repeats the received FLR signal direction index by using bits 1 to 3 of the FLR field in the command frame. If the device is uni-direction, it uses ‘000’ as the index of the direction.

The usage of the FLR is presented in 5.1.9.

5.3.12 Mobility notification command

The mobility notification command is shown in Figure 73. The concept of VLC cell mobility is defined in 5.1.11.

octets: 7	1	variable
MHR fields	Command Frame Identifier (Table 10)	cellSearchQuality (see 5.1.11.3)

Figure 73—Mobility notification command

The results from the cell search are provided in the mobility notification command as shown in Figure 73. The WQI values (in octets) obtained for the current channel during the cell search procedure defined in 5.1.11.3 shall be included in the command frame. The number of octets sent shall be equal to cell-SearchLength, as defined in 5.1.11.3.

5.3.13 GTS response command

The optional GTS.response primitive is generated in response to a GTS.request primitive. When used, the GTS response command shall be formatted as illustrated in Figure 74.

octets: 7	1	1	1
MHR fields	Command Frame Identifier as defined in Table 10)	GTS Characteristics	GTS Starting Slot (see 5.2.2.1.5)

Figure 74—GTS response command format

5.3.13.1 MHR fields

The Destination Addressing Mode subfield of the frame control field shall be set to zero (e.g., destination addressing information not present), and the Source Addressing Mode subfield shall be set to two (e.g., 16-bit short addressing).

The frame pending subfield of the frame control field shall be set to zero and ignored upon reception, and the Acknowledgment Request subfield shall be set to one.

The Source VPAN Identifier field shall contain the value of *macVPANId*, and the Source Address field shall contain the value of *macShortAddress*.

5.3.13.2 GTS characteristics field

The GTS Characteristics field shall be formatted as illustrated in Figure 75.

The GTS Length subfield shall contain the number of superframe slots being requested for the GTS.

The GTS Direction subfield shall be set to one if the GTS is to be a receive-only GTS. Conversely, this subfield shall be set to zero if the GTS is to be a transmit-only GTS. GTS direction is defined relative to the direction of data frame transmissions by the device.

bits: 0–3	4	5	6–7
GTS Length	GTS Direction	Characteristics Type	Reserved

Figure 75—GTS Characteristics field format

The Characteristics Type subfield shall be set to one if the characteristics refers to a GTS allocation or zero if the characteristics refers to a GTS deallocation.

5.3.14 Clock rate change notification command

The command format for the clock rate change notification is as shown in Figure 76. This clock rate change notification is sent at the current clock rate negotiated between the devices. All future transmissions from the current device to the receiving device will occur at this new clock rate.

Octets: 7	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	New clock rate for future TX

Figure 76—Clock rate change notification format

The MCS ID from Table 83 shall be used to indicate the optical clock rate. Any MCS ID corresponding to the chosen future clock rate can be used. The 6 LSBs shall be set to the MCS ID corresponding to the future clock rate. The other bits are set to 0 and reserved for future use.

5.3.15 Multiple channel assignment command

Octets: 7	1	1
MHR fields	Command Frame Identifier (as defined in Table 10)	Multiple Channels

Figure 77—Multiple channel assignment command format

Multiple channels should be used in the VLC system when time slot resources are not enough to cover all the currents users. These channels should be assigned based on the band-plan in Table 76. Refer to Table 3 for the contents of the Multiple Channels field.

5.3.16 Color stabilization timer notification command

The color stabilization timer notification command shall be formatted as illustrated in Figure 78. This command is used to inform a device or coordinator about the minimum time between two color-stabilization updates (upon reception of CVD frames).

The color stabilization timer field has the same format as the *macColorStabilizationTimer* (see Table 60).

Octets: 7	1	2	4
MHR fields	Command frame identifier (see Table 10)	Short address	Color stabilization timer

Figure 78—Color stabilization timer notification command format

5.3.17 Color stabilization information command

The color stabilization information command shall be formatted as illustrated in Figure 79. This command is used for relaying the color-stabilization updates (upon reception of CVD frames) back to the pertinent CSK transmitter (see Figure 116).

Octets: 7	1	2	Color stabilization information		
			2	2	2
MHR fields	Command frame identifier (see Table 10)	Short address	band *i*	band *j*	band *k*

Figure 79—Color stabilization information command format

The color stabilization information per band is 2 octets long and is used by the color-stabilization module (see Figure 117). It consists of the received signal levels in each of the three CSK bands. Two octets are used for each of the bands. A linear scale is used for each band, where the highest value corresponds to the maximum receive signal and the lowest value to the minimum receive signal. These fields are sent LSB first.

5.3.18 CVD disable command

Octets: 7	1	1
MHR fields	Command frame identifier (see Table 10)	CVD disable

Figure 80—CVD disable command format

The CVD frame can be transmitted depending on bi-directional, multicasting and broadcasting capabilities. A device shall not transmit a CVD frame after the device has received a frame from an associated device that has the “CVD usage option’ bit set to ’0’ as defined in Table 14. A device may resume sending CVD frames after it has received a frame from associated devices that have the “CVD usage option” bit set to ’1’. When the coordinator transmits and receives data with a device, if another device transmits an in-band CVD frame, interference may occur in the link between the coordinator and device. Out-of-band idle patterns may be used to maintain visibility when interference is seen in devices due to use of the CVD frame. In this case, the coordinator may indicate the transmission of “CVD usage option” with the CVD usage option bit set to ’0’. The CVD frame should be used prudently so as to cause minimal interference and prolong battery life. In many cases a light source is used for illumination which takes precedence over the use for communication.

The CVD usage option subfield is 1 bit in length and shall be set to ’1’ if the device is sending a CVD frame. This subfield shall be set to‘0’ otherwise.

Table 14—CVD disable field

Command frame payload	Bit	Usage/Description
CVD usage option	b0	logic 1 indicates that the device shall transmit the CVD frame logic 0 indicates that the device shall not transmit the CVD frame and may use out-of-band idle patterns if visibility needs to be maintained
	b1-b7	reserved

5.3.19 Information element command

The format of an individual information element (IE) is shown in Figure 81. The first octet is the Element ID and the second octet is the Length (Ln) of the payload of the IE in octets. The following Ln octets are the payload for the IE. Unless otherwise specified, these elements may appear in any order in the frames that are allowed to include more than one of these elements.

octets: 1	1	Ln
Element ID	Length (=Ln)	IE payload

Figure 81—Information element format

The information elements defined in this standard are listed in Table 15.

Table 15—Information Elements

Element ID hex value	Element	Subclause
0x01	Capabilities	5.3.19.1
0x02	Wavelength quality indication	5.3.19.2

When the information elements are used, they shall be added at the end of command frame format. Multiple information elements can be part of a single command frame. IE's can be added to any command frame.

5.3.19.1 Capabilities IE

The capabilities IE is used to convey device MAC and PHY capabilities to peer devices. The capabilities IE, as shown in Figure 82, consists of two fields: the capability information field, refer to Table 16, which indicates general capabilities of the device; and the aggregation bitmap field, which is specified in Figure 82.

5.3.19.1.1 Capability information Field

The capability information field is illustrated in Table 16.

octets: 8	variable: 8*n	variable: 8*n
Capability Information Field	Aggregation bitmap field	Guard bitmap field

Figure 82—Capabilities IE

Table 16—Capability information field

	Bit position	Function
MAC layer capabilities	0	Power source
	1-2	Battery information
	3	Security capability
	4	Co-ordinator capability
	5	Traffic support
	6-8	Topology support
	9-10	Device type
	11	Beacon support
	12	Dimming support
	13	Continuous visibility transmission (for infrastructure)
	14	CVD support
	15-23	Reserved
PHY layer capabilities	24	PHY I support
	25	PHY II support
	26	PHY III support
	27-28	Color stabilization capability
	29-31	Max supported TX clock
	32-34	Max supported RX clock
	35	Explicit clock notification request
	36	CCA support
	37-39	Reserved
Physical device capabilities	40-42	Number of optical sources
	43-45	Multiple direction support
	46-55	Number of cells supported (n)
Band capabilities	56-63	Bands used for PHY III (any 3 bits of the bits set to 1 can be used)

The power source subfield is 1 bit in length and shall be set to one if the device is receiving power from the alternating current mains. Otherwise, the power source subfield shall be set to zero.

The battery information subfield, shown in Table 17, is set to reserved (11) if the power source is set to 1.

Table 17—Battery Indication

Bits (b2 b1)	Battery indication
00	unknown
11	< 50% (low battery)
10	≥ 50% (sufficient battery)
11	reserved

The security capability subfield is 1 bit in length and shall be set to one if the device is capable of sending and receiving cryptographically protected MAC frames; it shall be set to zero otherwise.

The co-ordinator capability subfield is 1 bit in length and shall be set to 1 if the device is capable of functioning as a co-ordinator; it shall be set to zero otherwise.

The traffic support capability subfield is 1 bit in length. It shall be set to 0 if the device is only capable of broadcasting (unidirectional) communication. Otherwise, it shall be set to 1.

The topology support capability subfield can support multiple topologies via the bit maps of Table 18. .

Table 18—Topology support capability

Bits (b8 b7 b6)	Topology indication
b8	P2MP
b7	P2P
b6	broadcast

The device type capability subfield is set according to Table 19. This information is provided to assist upper layers.

Table 19—Device Type Capability

Bits (b10 b9)	Device Capability
00	infrastructure
01	mobile
10	vehicle
11	unknown/reserved

The beacon support capability subfield is one bit in length. It shall be set to 1 if the device is capable of sending beacons. Otherwise, it shall be set to 0.

The dimming support in MAC capability subfield is one bit in length. It shall be set to 1 if the device is capable of supporting dimming in the MAC using duty cycling and idle patterns. Otherwise, it shall be set to 0. A device shall honor all dimming requests. If the dimming support bit is not set then the device shall not attempt to communicate when a dimming request is received and shall comply with the dimming request even if the device must disassociate from the network as discussed in 5.1.4.2. Even if the device supports dimming but is unable to communicate during dimming, it shall set the *macDimDataFailureIndication* MAC PIB attribute as mentioned in Table 60, but shall still comply with the dimming request at the expense of loss of communication.

The continuous visibility transmission subfield is one bit in length. It shall be set to 1 if the device will be continuously transmitting to maintain illumination. Otherwise, it shall be to 0.

The CVD support subfield is one bit in length. It shall be set to 1 if the device is capable of transmitting various colors; otherwise, it shall be set to 0.

The PHY I support subfield is one bit in length. It shall be set to 1 if the device supports PHY I.

The PHY II support subfield is one bit in length. It shall be set to 1 if the device supports PHY II.

The PHY III support subfield is one bit in length. It shall be set to 1 if the device supports PHY III.

The color stabilization capability subfield describes if and where (device and/or coordinator) color stabilization is performed. All allowed settings are shown in Table 20.

Table 20—Color stabilization capability

Bits (b28 b27)	Color stabilization scheme
00	No color stabilization
01	Color-stabilization information to be sent from device to coordinator upon reception of CVD frames
10	Color-stabilization information to be sent from coordinator to device upon reception of CVD frames
11	Color-stabilization information to be sent from device to coordinator and from coordinator to device when either receives CVD frames

The max supported TX clock subfield and max RX clock subfields follow the usage as indicated in Table 21. Support for 200 kHz is mandatory for PHY I and support of 3.75 MHz is mandatory for PHY II. Support for 12 MHz is mandatory for PHY III and shall be indicated using bits ’100’ as in Table 21.

The explicit clock notification subfield is 1 bit in length. The subfield shall be set to 1 if the receiving device needs an explicit clock change notification from the transmitter before any change of clock frequency.

If CCA is supported then the CCA Support bit is set to 1, otherwise the bit is set to 0.

The number of optical sources subfield indicates the number of optical sources in the transmitter of the device that have distinct frequency responses.

Table 21—Maximum supported optical clock frequency

Bits (b31 b30 b29)	Description
000	200 kHz
001	≤ 400 kHz
010	≤ 3.75 MHz
011	≤ 7.5 MHz
100	≤ 15 MHz
101	≤ 30 MHz
110	≤ 60 MHz
111	≤ 120 MHz

The multiple direction support subfield indicates the number of distinct directions supported by the device transmitter supported by the multiple optical sources. This is used for fast link recovery as defined in 5.3.11.

The number of cells *n* indicates the maximum number of cells supported in the device. The number of cells supported shall not be more than 1023.

In regards to the bands used for PHY III, bit 7 is reserved and bits 0-6 map to the bits corresponding to the bandplan. Only 3 bits shall be set to indicate PHY III usage. If the device supports more colors and wants to change the PHY III usage, it needs to send the capabilities information again with the new bitmap.

5.3.19.1.2 Aggregation and guard channel

The aggregation and guard channels are used to support any visible light optical source for VLC that may have variable spectral widths and center frequencies. The aggregation and guard bitmap for a single optical source type is as shown in Figure 83. The bit map is variable in length. The length of the aggregation and guard bit maps are 'n' octets each, where 'n' is the number of optical source types. The aggregation and guard channel bit usage are defined by an 8-bit bitmap for every optical source type supported by the transmitter of the device. The 8-bit bitmap is indexed by the bandplan identification number. The bit position 'm' is set to a '1' for band 'm' if that band is used by the optical source. The reserved bit in Figure 83 shall be set to 0.

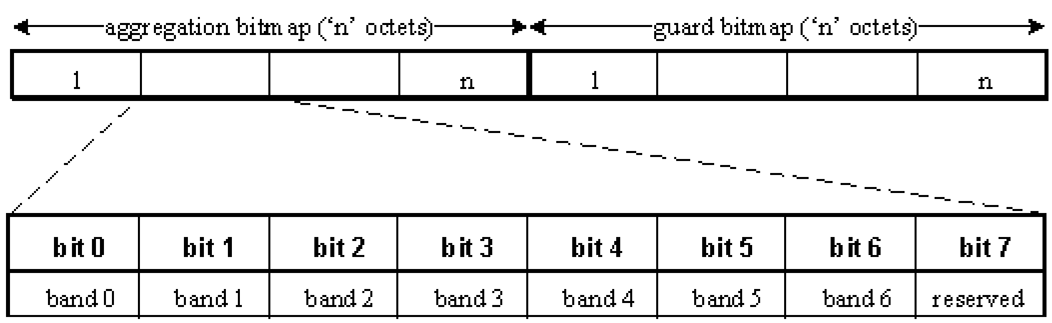

Figure 83—Aggregation and guard bitmap per optical source type

For example,

if band 1 and band 2 need to be aggregated (assuming a blue LED) the aggregation bit-map is indicated as 0110000 and the guard bit-map is indicated as 0000000.

if band 1 is being used but there is leakage in bands 3, 4, 5 (assuming a white LED which is realized via a blue LED with yellow phosphor), the aggregation bit map is indicated as 0100000 and the guard bit-map is indicated as 0001110.

5.3.19.2 Wavelength quality indication (WQI) IE

WQI is communicated to another device using the WQI Information Element. The WQI value to be sent in the Information Element may be an average value across a number of packets, and WQI value sets for a number of band plan ID's can be reported using the WQI information element as shown in Table 22. The wavelength quality indication IE is 7 octets in length and the WQI information is provided for all band plan IDs. If a band plan ID is not supported, WQI of 0 shall be reported.

Table 22—Wavelength quality indication IE

Band plan code	WQI value
0x00	0x00 to 0xff
0x01	0x00 to 0xff
0x02	0x00 to 0xff
0x03	0x00 to 0xff
0x04	0x00 to 0xff
0x05	0x00 to 0xff
0x06	0x00 to 0xff

6. MAC sublayer service specification

6.1 Overview

The MAC sublayer provides an interface between the SSCS, DME and the PHY. The MAC sublayer conceptually includes a management entity called the MLME. This entity provides the service interfaces through which layer management functions may be invoked. The MLME is also responsible for maintaining a database of managed objects pertaining to the MAC sublayer. This database is referred to as the MAC sublayer PIB.