8. MAC Functional Description

The MAC functional description is presented in this clause. The process of starting or associating with a piconet is described in sub-clause 8.1. As a part of starting the piconet, sub-clause 8.1.3 describes the PNC selection process. If there is more than one device that is capable of performing PNC activities in a piconet, those devices can compete and win the responsibility of the PNC during the start of a piconet. When the current PNC is going away, it can choose a device that is capable to be the new PNC. This coordination handover mechanism is described in sub-clause 8.1.8.

The channel access mechanisms are described in sub-clause 8.2. There are two different mechanisms used for channel access, contention based and contention free. In the contention access period (CAP), devices can access the channel in a distributed style as described in sub-clause 8.2.2. On the other hand, the PNC controls all the accesses to the channel during the contention free period (CFP), as described in sub-clause 8.2.3. The PNC allocates channel time within CFP for each device in the piconet, based on the currently pending requests by all the devices in the piconet and the available channel time within CFP. The channel time request and allocation procedures are described in sub-clause 8.4. The required synchronization for the operation of the piconet and the channel access is described in sub-clause 8.3.

The process of stream connection, disconnection and the rest of the stream management are described in sub-clause 8.4. A stream is connected or disconnected only after tripartite communication/negotiation among the device that is originating the stream, the device that is intended receiver of the stream and the PNC. Once connected, a stream may be communicated in a peer-peer style. The fragmentation and defragmentation of the MSDUs at a device is described in sub-clause 8.5. The acknowledgement and retransmission mechanisms are described in sub-clause 8.6. A device that is transmitting a frame may request either immediate or delayed acknowledgement for the frame.

If the link between two devices is not satisfactory, either of those devices can request the PNC to provide the repeater service for that entire link. The repeater service is described in sub-clause 8.9. To overcome the problems due to overlapping piconets and interference in a given channel, the PNC may choose to move the operations of the piconet to a new channel. The process of dynamic channel selection (DCS) is described in sub-clause 8.10. The devices in the piconet can employ power save techniques to reduce their power consumption. The operation and the negotiations required for such power management are described in sub-clause 8.11. Each device in the piconet can choose transmission power based on the current channel conditions. The operation and the negotiations required for such transmit power control (TPC) are described in sub-clause 8.12.

The frame exchange sequences at the MAC layer are described in sub-clause 8.13.

In this clause, unless otherwise indicated, receiving a frame means that the PHY has successfully received a data stream over the medium and both the FCS and HCS calculations match their respective data as defined in 7.2.1.9 and 11.2.8.

8.1 Associating or starting a piconet

A device that is instructed to associate with a piconet through MLME.Associate shall try only to associate with an existing piconet and shall not attempt to start its own piconet. Similarly a device that is instructed to start a piconet through MLME.StartPiconetRequest shall try only to start its own piconet and shall not attempt to associate with an existing piconet. During the start of its own piconet, the device shall make sure there is not an already started 802.15.3 piconet in the same channel.

8.1.1 Scanning through channels

All devices shall use passive scanning to detect an started piconet. That is, devices shall be in the listen mode for a period of time in a channel and look for beacon frames from a PNC. During such scanning the device shall ignore all the received frames with a different PNID than the one the device is searching for. Hence all the following description apply only to frames received with the desired PNID.

Devices search for the piconet by always starting from the first channel and traversing through all the indexed channels available in the PHY. For the 2.4 GHz PHY the available channels are specified in 11.2.3. During such searching, if any frame is received with the desired PNID, the searching device shall stay in the channel for a minimum of a aChannelScan from the time of reception of last frame and look for a beacon from the PNC. If the beacon frame from the PNC is received , the piconet detection is considered complete.

During searching, if the device receives a PNC selection frame from another device, the searching device shall stay in the same channel until the PNC selection is complete. If the device is allowed to start its own piconet, then the device shall participate in the PNC selection process.

If all the channels are traversed and no frames with the desired PNID is received in any channel, then the device takes one of two actions depending on the MLME-SAP command that initiated the scan. If the instruction through MLME-SAP allowed the start of its own piconet, the device shall choose a channel and start the PNC selection process in that channel. If, however, the MLME-SAP command was only to associate with a piconet, then the MAC shall return a fail indication through MLME.Associate.indicate primitive.

Channel change after the piconet is started is described in sub-clause 8.10.

8.1.2 Randomization of PNID

The randomization process required to choose a unique PNID is beyond the scope of this piconet.

8.1.3 PNC selection process

The PNC selection process begins with one AC sending a PNC selection frame, defined in <TBD: 7.3.4>, in a channel where there is currently no other started piconet. The AC initiating this process shall send the PNC-selection frame for at least aCSFRameRepeat inviting other capable ACs to participate. At the end of the indicated timeout in its PNC-selection frame, the AC shall start sending a beacon to start its own piconet if there are no other participants. During the indicated timeout the AC shall broadcast the PNC-selection frame at least once in aCSFrameBroadcast. The process of PNC selection is illustrated in Figure 40.

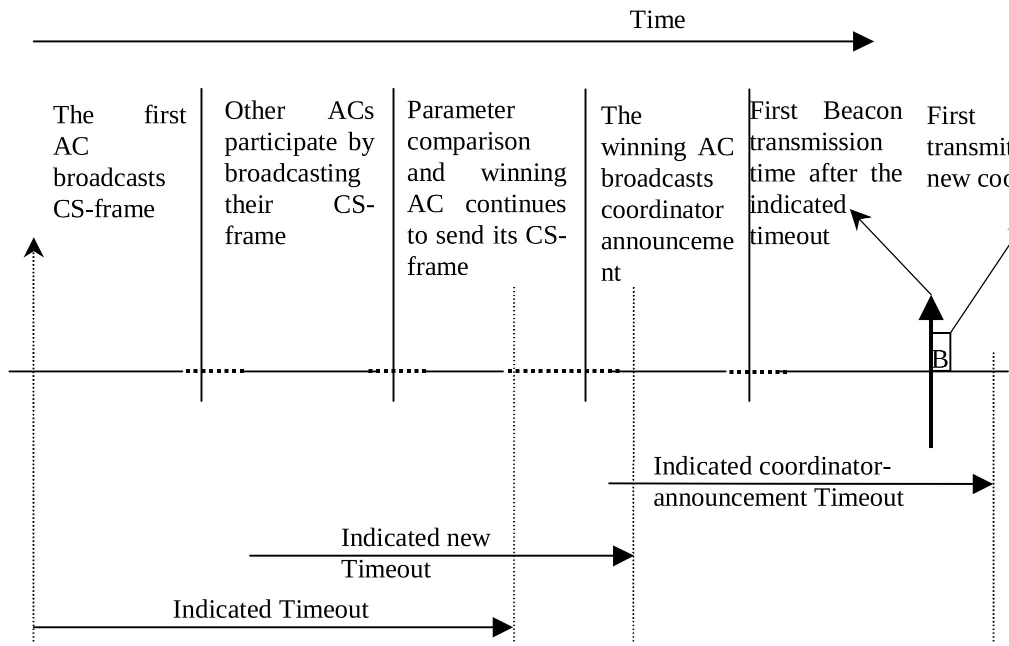

Figure 40—Illustration of coordination selection process

If more than one AC participates in the PNC seletion process, then the AC receiving a PNC-selection frame from another AC shall compare the received PNC selection parameters with its own following the evaluation defined in Table 40. If the AC finds that its parameters score higher than the received ones, the AC shall continue to broadcast the PNC-selection frame. Under this scenario, the winning AC may also send a directed PNC-selection frame to the other AC and the other AC is expected to ACK the frame if it is received correctly. If the AC finds that its parameters score lower than the ones received, the AC shall announce its pullout from the competition by sending a broadcast PNCS-frame with the action type set to AC pullout. In addition, the AC may send a directed PNCS-frame with the action type set to AC pullout announcing the pullout to the winning AC. The intended recipient shall ACK the frame. This process shall continue till all the AC’s except the winning AC announce their pullout.

When no other AC is announcing its intention to compete, the winning AC shall broadcast a PNCS-frame with action type set to PNC announcement and a timeout indicated. The winning AC may announce this more than once via the PNCS-frame before the indicated timeout. The winning AC shall take the responsibility of the PNC and start beacons before that indicated timeout.

8.1.4 Authentication

Authentication is described in clause 10.

8.1.5 Association

All the frames between the PNC and the device before the completion of the association of the device shall be exchanged only in the CAP of the superframe.

An unassociated device initiates the association process by sending an association request frame, defined in 7.3.5, during the CAP of an existing piconet. When the PNC receives a valid association request frame, it shall send an association response frame, indicating that the device has been associated and its AD-AD or that the request has been rejected with the reason for the rejection, as defined 7.3.6.

All association request frames shall be acknowledged by the PNC by sending an ACK frame. The Ack to association request frame does not mean that the device got associated. The PNC needs some time to make sure there enough resources available to support another device on the piconet and allocate an AD-AD. After a decision is made regarding the association and AD-AD, the PNC sends association response frame to indicate the acceptance/rejection of association. The time difference between sending ACK to an association request and sending an association response frame meant for the same device shall not exceed the ATP mentioned in the assocaition request frame.

The association response frame is not a directed frame and its DA is 0xFE. Hence if there is an ack required for this frame, when there are multiple devices trying to associate during the same time interval, all of them will try to ack and collide. Hence there is no ACK required for the association response frame. Instead each device trying to associate shall compare its own Device-ID with the Device-ID field in the association response frame and if there is a match, accept the AD-AD as its address for all future communicatons and consider itself associated.

Since the DA of 0xFE is equivalent to broadcast, all DEVs are allowed to use this frame from PNC to update their association tables and hence keep track of the peer-DEVs in piconet, if needed.

In the absence of a valid pending frame, the PNC shall send a directed command frame with null payload to the currently associating device with the ACK policy set to Imm-ACK and the DA equal to the new AD-AD allocated to the currently associating device. This sequence of association response followed by a directed frame to the associating device confirms the reception of the association response by the newly associated device. If the ACK from the newly associated device is not received at the PNC, the PNC shall repeat the sequence of association response and the directed frame as illustrated in in Figure 41. The message sequence chart for the same is shown in Figure 42.

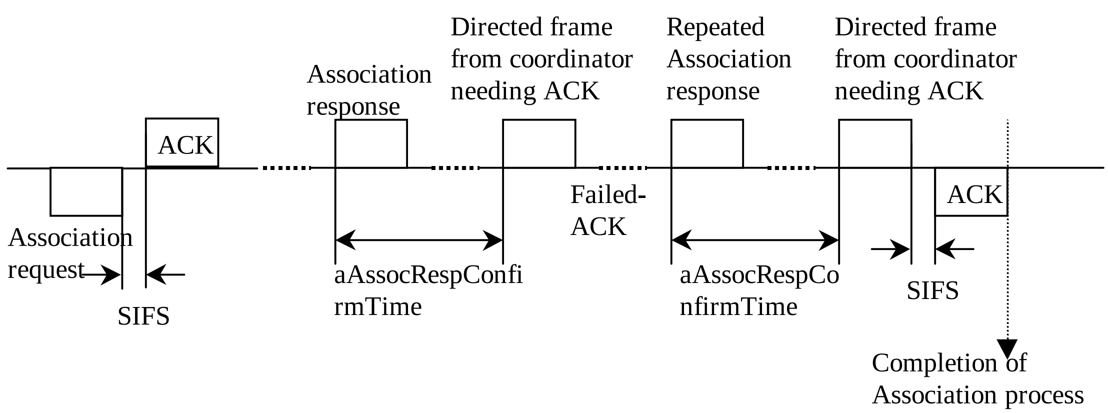

Figure 41—Illustration of association process

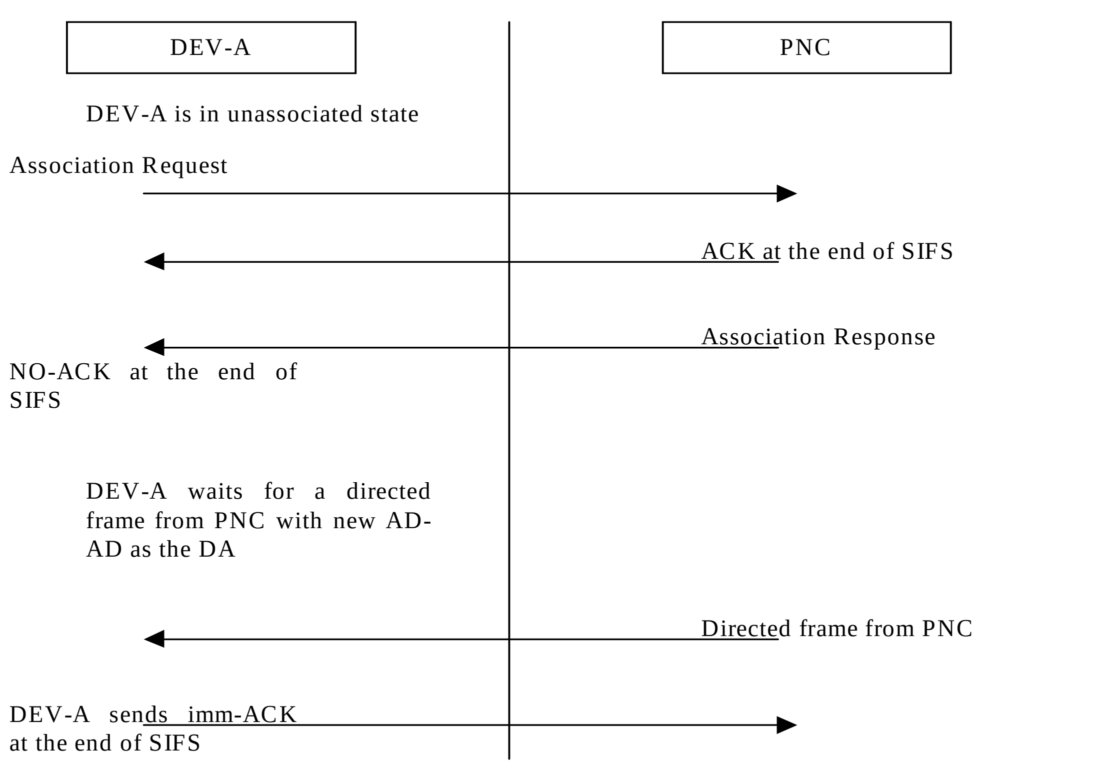

Figure 42—Message sequence chart for frame exchange during association

The addresses (AD-ADs) shall be assigned in sequence (increasing order) by PNC except when PNC wishes to use the AD-AD that was freed up when devices leave the piconet. However reallocation of the same AD-AD by PNC shall be at least aADAdressReuseTime after the disassociation of the device that was allocated the same AD-AD. The coordinator shall ensure that there is only one associated device that has been allocated a given AD-AD at any given time within the piconet. Similarly any associated device shall be allocated only one AD-AD. Only exception to this is the PNC itself. The DEV serving as PNC shall have two values of AD-AD associated with it. The AD-AD value of '0x00' shall be assigned to the coordinator fuction within the DEV and the other non-zero value of AD-AD shall be use for all the non-coordinator traffic. Hence the PNC shall be viewed as two logical operational entities within the same DEV.

8.1.6 Disassociation

When a PNC wants to disconnect a device, the PNC shall send a directed disassociation frame to that device with a reason code. Similarly when a device wants to disconnect from the piconet, the device shall send a directed disassociation frame to the PNC with a reason code.

All the disassociation frames when received correctly are acknowledged by the intended recipient through an ACK frame.

8.1.7 Device registration

An open scheme for Device registration and/or association is allowed to form adhoc piconets in which a common PNID can be sent by the PNC to the requesting device. Note that a PNC may not allow devices to register and associate using this process if the user does not wish the piconet to allow such an open registration and/orassociation.

The open scheme for device registration and association is illustrated in Figure 41. The unregistered device shall detect the presence of the piconet that it wishes to register and send an association request using an all-zero nework-ID. The PNC that allow devices to register and/or associate using this mechanism responds by sending the association response with an allocated AD-AD. The PNC shall always use its valid PNID in its frames. The device shall look for association responses from the PNC. A valid value of AD-AD in the association response containing the device's Device-ID indicates that the device is registered.

8.1.8 Coordination handover

If during the life of a piconet the PNC decides to leave the piconet, the PNC shall attempt to choose a device that is capable of being a PNC as its successor. The AC bit in the capability bit is used to decide whether a device is capable of being a PNC. Hence any device that does not wish to be the PNC shall set the AC bit to '0' in its capabilities field. The PNC shall send a coordination-handover command to its chosen device with an indication of the handover timeout. The device shall always accept the nomination and obtain the device information from the current PNC within the indicated timeout period. The new PNC shall announce its new responsibility as PNC in at least aCHRepeat of the superframes before the indicated timeout period. The new PNC shall send its first beacon at the first expected beacon transmission time after the timeout period indicated in the co-ordination handover frame. The new PNC shall begin using address of 0x00 for all (tx and rx) coordination related traffic, but it shall continue to use its previously assigned AD-AD for all non-coordinator traffic. When the coordination handover is succesful, the association of the remaining devices with the piconet is unaffected and hence they are not required to re-associate with the new-PNC. The process of coordination handover is described in Figure 43.

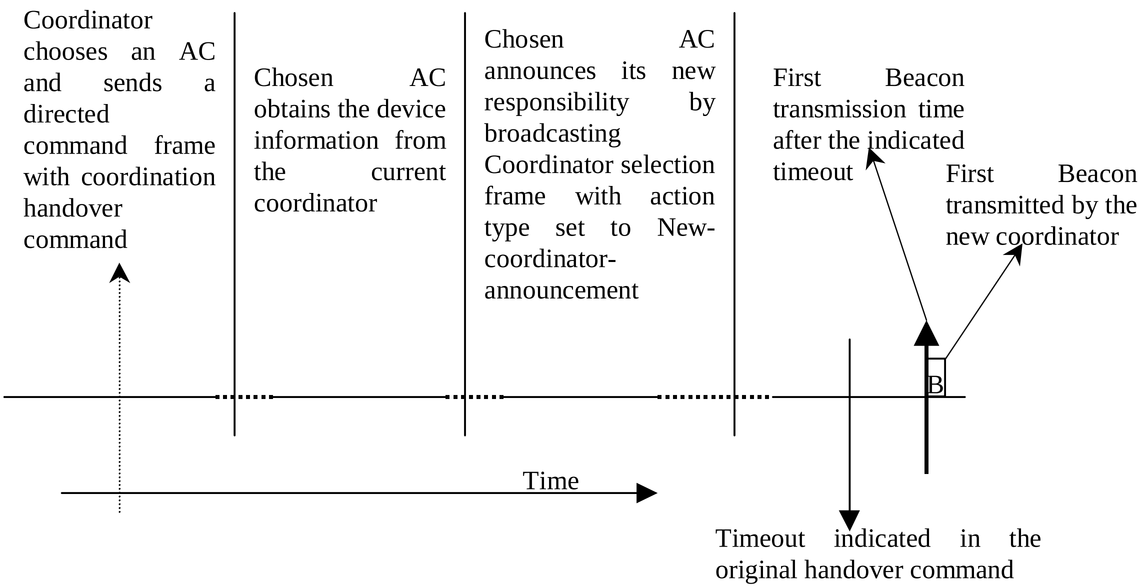

Figure 43—Illustration of coordination handover

Note that the coordinator handover need not always necessarily stop all the stream transmission. Figure 43 intends to describe the handover process and hence does not show other traffic. The current-PNC may choose long enough timeout period and hence allowing enough time for the new-PNC to obtain the device information in small quantities per supeframe. When the current-PNC chooses a small timeout the current-PNC must allocate enough time for new-PNC so that the new-PNC can obtain all the device information in that short time. However other stream-data traffic may or may not be affected by this depending on the traffic conditions within the piconet. To facilitate fast handovers the PNC shall broadcast the device information table, described in <TBD 7.5.12.2>, at least once every aBroadcastDevInfoDuration.

8.2 Channel Access

The channel time is divided into superframes, with each superframe beginning with a beacon. The superframe is composed of three major parts: the beacon, the CAP and the CFP as shown in Figure 44. The CFP is used for asynchronous and synchronous data streams with QoS provisions while the CAP is used for non-QoS frames. During the CAP, the devices can access the channel in a distributed style using CSMA and a backoff procedure. During the CFP, the PNC controls the channel access by assigning time slots to individual devices with each time slot having a fixed start time.

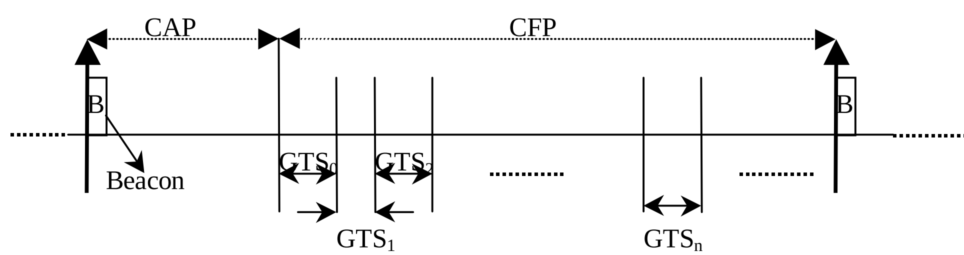

Figure 44—Superframe structure

8.2.1 Inter frame spacing (IFS)

There are two IFS that are defined. They are short inter frame space (SIFS) and retransmission inter frame space (RIFS). The relation between the IFS is, (SIFS < RIFS). The actual values of IFSs are PHY dependent. For the 2.4 GHz PHY they are listed in 11.2.6.1.

Both in the CAP and the CFP, a response frame (ACK) transmission over the medium shall start within a SIFS duration after the end of the transmission of the previous frame for which the response is intended. Similarly a SIFS duration shall be allowed between a frame that does not expect an immediate response and the next successive frame transmitted over the medium.

During the CFP, all devices shall use an RIFS for retransmissions. During the CAP, however, the retransmissions shall not use an RIFS and shall instead follow the CAP rules described in sub-clause 8.2.2. The rules for acknowledgement and retransmissions are described in sub-clause 8.6. Beacon frame transmission shall use RIFS for their transmission.

All devices shall use CCA for RIFS detection. In order to declare the detection of a RIFS the channel shall be idle for full contiguous aCCADetectionDuration after SIFS after the end of the transmission by the DEV. The use of CCA and RIFS is described in 8.6.5.

8.2.2 Contention based channel access

The basic medium access mechanism during the CAP is carrier sense multiple access with collision avoidance (CSMA/CA). To minimize collisions, a transmitting device is required to first sense that the medium is idle for a random length of time, as defined in 8.2.2.1. Only if the medium is idle after that time shall the device start its transmission. This process of waiting before transmission is termed “backoff” and is described in 8.2.2.1. The backoff procedure shall not be applicable for the transmission of the beacon that is transmitted by the coordiantor at the beginning of superframe.

During CAP a DEV is allowed to transmit one frame at a time with backoff being applied to every frame, except the imm-ACK frame, attempted during CAP. In no case shall a device extend its transmissions that started during the CAP in to CFP. Hence, once a DEV decrements its backoff counter to zero, it shall check whether there is enough time in the CAP for the transmission of current frame and SIFS. If an imm-ACK is expected for that frame, the remaining time in CAP must be large enough to accommodate the current frame, (2*SIFS) and the imm-ACK frame at the same PHY rate as the transmitted frame. If there is not enough room for this entire frame exchange sequence, then the DEV shall abort the transmission, choose a backoff value with minimum allowed backoff-window size. The backof procedure is described in the following sub-clause.

The MAC shall use the PHY-CCA.indication (see 6.8.4.10) from the PHY to detect whether the channel is busy or idle.

8.2.2.1 Backoff Procedure

Except when transmitting an immediate response frame (ACK), the following backoff procedure is performed.

A backoff value is randomly (uniformly distributed) chosen in the range of (RIFS+aRXTXTurnaroundTime to backoff_window[retry_count]), both inclusive. The retry_count shall be zero for the first transmission attempt of a frame. The table of backoff_window shall be composed of the following values:{80, 160, 320, 640} with each entry in microseconds. A backoff counter is maintained for this purpose and the backoff counter is decremented only when the medium is idle for the entire interval minus aRXTXTurnaroundTime. Whenever the channel is busy, the backoff counter is suspended. This avoids the problem of unfair channel access when a backoff counter of a device ending in the middle of a reception and hence resulting in larger backoff for that device while another device starting after the current reception choosing smaller value for backoff. When the backoff counter reaches zero, the device shall transmit its frame.

The backoff counter is also suspended outside CAP duration.

When a directed frame is transmitted and the expected immediate response (ACK) is not correctly received by the device, the retry_count is incremented and if the maximum number of retries has not been exceeded, the backoff procedure is again resumed. The device shall attempt transmission of a frame for a maximum of aMaxRetryCount times before the failure in frame transmission is reported through MLME interface.

OPEN ISSUE: the backoff_window values :{80, 160, 320, 640} are chosen with the assumption of 10microsec rx/tx turnaround time. Is this good enough?

8.2.3 Contention free channel access

The channel access during CFP employs slotted time architecture with one slot being assigned to a device. For reasons of power saving and synchronous transmissions, all slots have guaranteed latest start time. All the slot allocations within the current superframe are broadcast in the beacon. Hence if a device did not correctly receive the beacon, that device is not allowed to access the channel during CFP. The process of channel time request and allocation are described in 8.2.3.2.

8.2.3.1 Guaranteed time slots (GTS)

The PNC divides the CFP into several Guaranteed Time Slots (GTS). Each GTS is a single time slot with guaranteed start time and a guaranteed time duration reserved within the CFP. Hence a device that is allocated with a GTS is guaranteed that no other devices will compete for the channel during the indicated time and duration of the GTS. A device with a GTS may or may not make use of all the allocated time duration within the GTS. The selection of a stream for transmission during a GTS is determined locally by the device depending on the number of pending frames and the priority of the corresponding streams, see 8.4.3 for more information on priority management.

8.2.3.2 Channel time allocation (CTA) and channel time usage

The devices associated with a PNC shall send their changes in channel time requirement whenever they observe the change. Once a request for channel time is received from a station, the PNC shall remember that as the outstanding request for every superframe until, a change in request is received from the device. In addition to this the PNC shall make use of the properties of the stream provided during the stream connection process. The slot assignments within CFP are based on the current pending requests from all the devices and the currently available channel time within CFP. The slot assignments may change from superframe to superframe as felt required by the PNC. All the slot assignments are broadcast in the beacon. PNC may announce the slot assignments in directed or broadcast channel time grant command in addition to the announcement in beacon. But the additional announcements by the PNC shall not change from what was broadcast in the Beacon. The start time of all the GTSs are with reference to the start of beacon frame, whether they were announced in beacon or channel time grant command. The algorithm used to allocate the channel time and assign slots is beyond the scope of this standard.

When delayed-ack is employed, the recipient of the "stream-data" type frame is responsible for obtaining the time required for transmission of the return delayed-ack unlike in the case of imm-ack, where the sender of the "stream-data" guarantees the time required for imm-ack right after the end of current stream-data frame plus SIFS duration. Each device is expected to estimate the channel time required depending o the number (and type) of streams that they are sending and receiving, and request for that channel-time with PNC. When PNC grants one or more time-slots, the devices are responsible for using those time-slots efficiently without abusing the priority order of the streams that are pending locally for transmission. If there is a device that has not requested any time-slot (or obtained one) then the device has wished to use CAP only for sending its delayed-ack frames.

When a DEV has a frame of type other than Stream-Data type for dest-DEV, the DEV is free to send it during an allocated time slot for that (source-DEV, dest-DEV) pair or use CAP to communicate that frame.

In any superframe there may be one or more devices in the piconet that receives the Beacon in error. This may not happen to the same device all the time but may happen to different devices at different times depending upon their location and type of interference they are subjected to. If a device did not receive the CTA information, it can not access the channel during CFP. The channel time grant command gives the flexibility to the PNC to broadcast the CTA information during the supeframe in addition to the broadcast in Beacon and hence increasing the chances of all devices obtaining the allocation information. In addition this also provides the flexibility to PNC to help preserve the Qos by sending directed channel time grant command to a device that may be experiencing more than usual channel errors during certain time segments. The PNC may use the channel statistics to decide whether to send such a directed channel time grant command. Note that when channel becomes too severe for the device not to receive beacon, or the directed channel time grant command or the stream-data itself there is little help that can be provided to that device through these channel time grant commands.

In no case shall a device extend its transmissions that started during an allocated time slot beyond the end of that time slot. Hence, the source-DEV shall check whether there is enough time in the time slot for the transmission of current frame and SIFS. If an imm-ACK is expected for that frame, the remaining time in the time slot must be large enough to accommodate the current frame, (2*SIFS) and the imm-ACK frame at the same PHY rate as the transmitted frame. If there is not enough room for this entire frame exchange sequence, then the DEV shall abort the transmission and let the remaining duration of the time slot unused.

PNC may compute more than one superframe slot allocations at a time and keep them repeating over time till the situation change. The allocation by the PNC may not exactly match the duration of time requested by a DEV. The PNC may set the Slot-dur in CTA-element to indicate the rejection of the channel time request. But the rejection shall be a directed channel time grant command and shall not be in the Beacon frame. When a DEV sees a time slot allocated to it but not exactly what it had requested, the DEV is free to request again while using whatever is currently allocated. However the allocation shall be atleast the minimum requested (by DEV) in the channel time request command except when the request is rejected.

8.3 Synchronization

All devices within a single piconet shall be synchronized to a common clock using the time synchronization function (TSF). The beacon sent at the beginning of every superframe contains the information necessary to time-synchronize the device. See 7.4.2 for the definition of the TSF parameters sent in the beacon.

8.3.1 TSF timer and time synchronization

The TSF is a one microsecond resolution clock that is maintained at all the devices with the TSF at the PNC being the reference clock. The PNC shall initialize its TSF timer at the time of the initiation of the piconet.

The PNC sends the value of its TSF timer as the time stamp in each of its beacons to synchronize the other devices in the piconet. The time stamp value shall be equal to the value of PNC's local TSF timer value at the time of transmission of first bit of the Time-stamp field. A device shall always accept the TSF value in the beacons from its associated PNC and compare that as the reference time with its local TSF timer. Note that since the frame checking is done after receiving the entire frame, the device is required to take a copy of the local TSF timer at the end of beacon frame and adjust it appropriately. The adjustment is required to reflect the local receive time at first bit of the Time-stamp field in the beacon frame. This adjusted value is compared with the time-stamp field in the received beacon. If such an adjusted value of the local TSF timer value is different from the reference time stamp received in the beacon from its PNC, the device shall set its local timer to the received time stamp value advanced appropriately to reflect the end of beacon reception time.

The TSF timer is 64 bit wide and hence it rolls over one clock after it reaches (2**64)-1.

8.3.2 TSF time accuracy

All device implementations shall maintain the accuracy of the TSF timer to be ±0.01%

8.3.3 Beacon generation

The PNC shall send a beacon at the beginning of each superframe. The slot-start-time in all the CTA blocks in the beacon is offset from the start of the superframe and hence the start of the beacon frame.

8.3.4 Beacon reception

All of the devices that are associated shall use the beacon start time and the time slot allocations contained in beacon to start their transmissions. The auperframe duration and the aMaxCFPDuration in the beacon are used to accurately mark the begining and the end of the CFP and the CAP.

8.3.5 Acquiring synchronization

All devices acquire synchronization through beacons from the PNC. Unassociated devices that wish to associate with the piconet shall use passive scanning to collect beacons and use the information within the beacons for synchronization. A device shall hear at least one beacon before sending its request for association.

If an associated device did not receive a beacon from its PNC, it shall use the same superframe duration indicated in the last correctly received beacon from its PNC. If a device did not receive a beacon for more than aConnectionTimeOut period, then the device shall stop all of its transmissions and waits for a beacon before starting the association process again.

8.4 Stream management

Stream management involving stream connection, stream disconnection and priority management is described in this sub-clause. Each DEV must support at least one stream connection. A DEV may support more than one stream depending on the application for which they are built to support. A PNC must support as many as it desires to source and sink with some additional streams to support the repeater service.

8.4.1 Stream connection

A stream shall be connected only after tripartite communication/negotiation among the device that is originating the stream, device that is the intended receiver of the stream and the PNC. For a broadcast stream, the involvement of intended receiver is precluded. Once connected, a stream can be communicated in a peer-peer style.

Either the sending device or the intended recipient device for the new stream may send a stream management command with the request for stream connection. The process of stream connection is illustrated in the Figure 45. In all stream management communications from the PNC to the other involved device, the PNC appropriately changes the value of the direction field to imply the same direction of the stream as originally requested.

The values for direction, security, stream type and priority shall be non-negotiable and are decided by the device that is sending the stream connection request. These values shall not be changed anytime after the first transmission of the command frame containing the request for that stream.

All the bandwidth and latency related requirements of the stream shall be negotiated between the sender of the stream and the PNC. The PNC decision on the values of the stream QoS parameters that are supported in the piconet shall be final.

The retransmission window shall be decided between the device that is originating the stream and the device that is the intended receiver of the stream for a directed stream. If the two devices suggest different sizes for the retransmission window, smaller of the two shall be adopted.

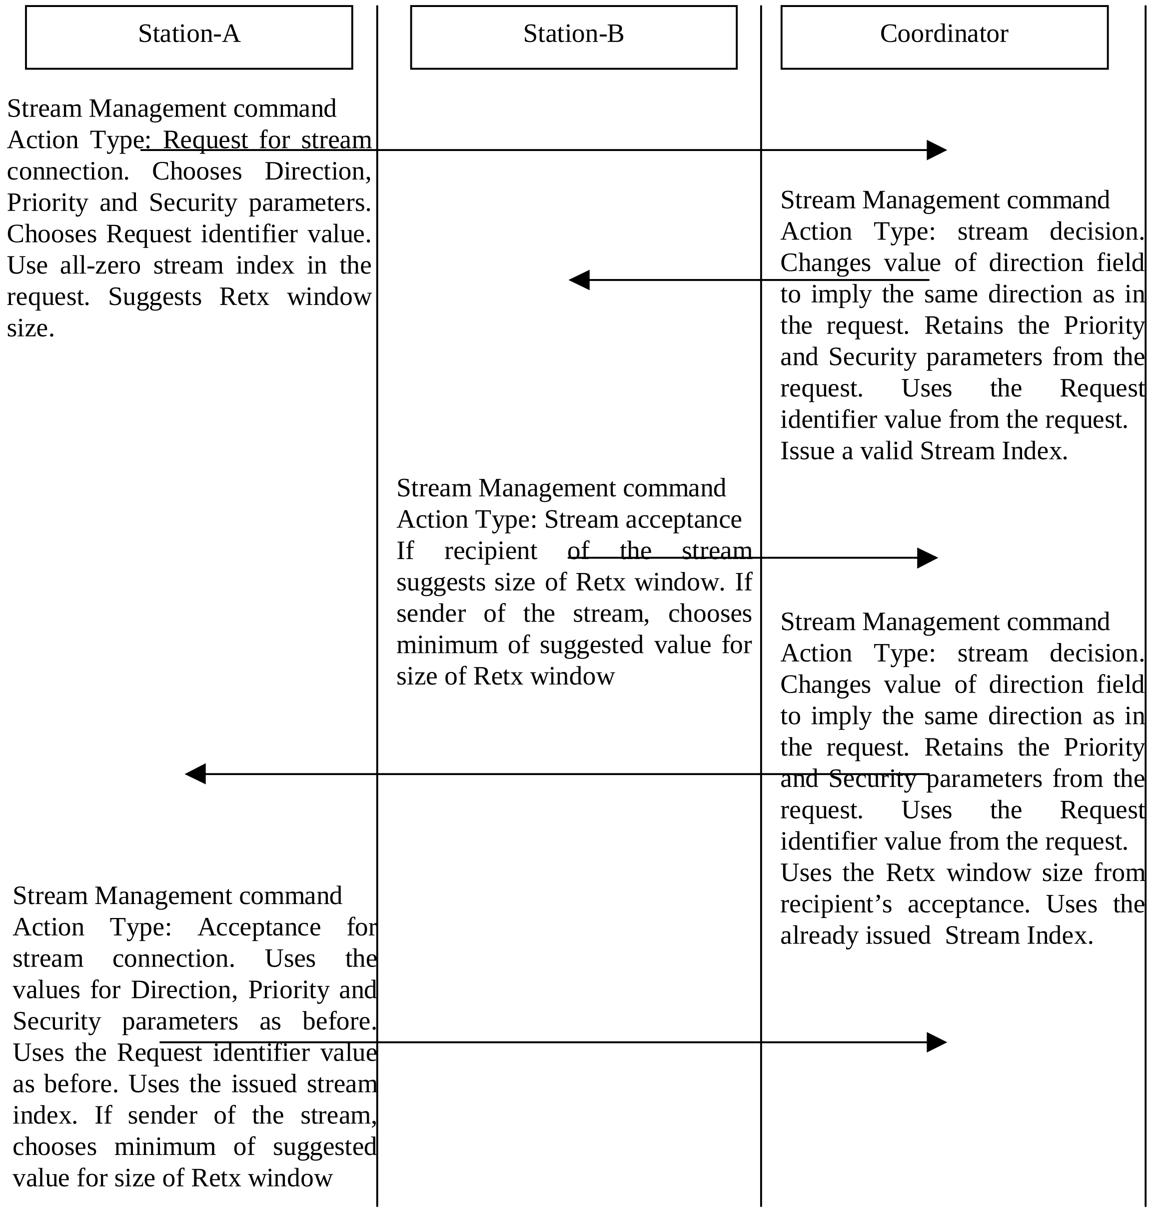

Figure 45—Stream connection process with sender of the stream initiating the stream connection process

In some cases it may be necessary for the PNC to initiate the stream connection process, even though the stream is supposed to be exchanged between two other devices in the same piconet. In this case the PNC sends the request for stream connection to the expected originator of the stream. The originator of the stream may change the QoS parameters in the command. The originator of the stream follows the same sequence of transactions described before for the connection of stream. This is illustrated in Figure 46.

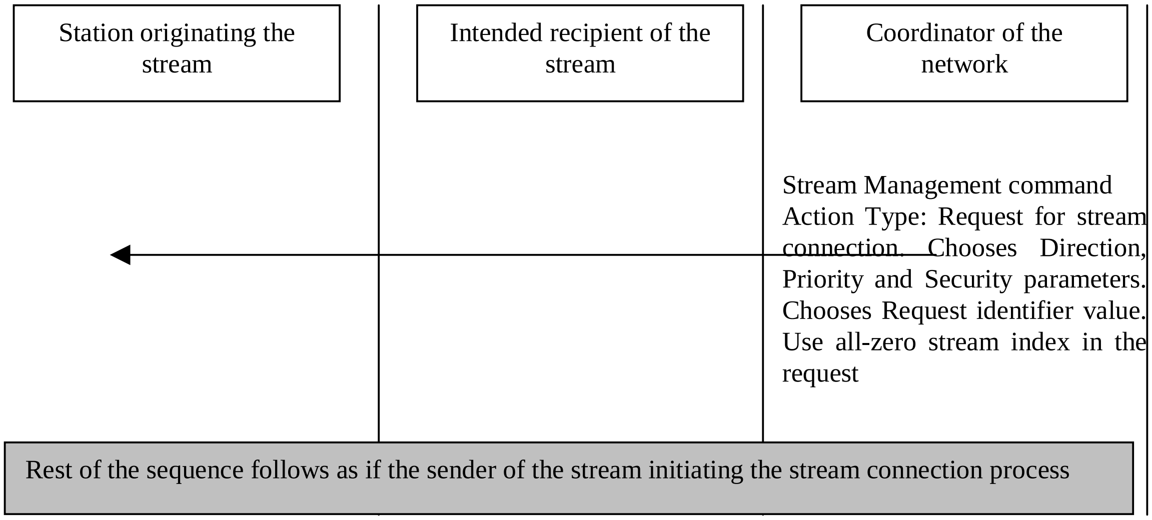

Figure 46—Stream connection process with the PNC initiating the stream connection process

If the stream is a broadcast stream, the command transactions during the stream connection are only between the sender of the stream and the PNC.

Devices may use the stream ID of all-zero in data frames that do not need a stream connection.

8.4.2 Stream disconnection

During any of the transactions described in any of the scenarios described for stream connection in 8.4.1, a device (or the PNC) may reject the stream connection. The frame used for either rejecting or disconnecting a stream shall be a stream management command with the action type set to rejection/disconnection as specified in 7.5.9. If the device wishes to reject the stream before it is connected, it shall send the appropriate frame to the PNC. In that case, the PNC shall communicate the disconnection information to the other device via a similar frame. If the PNC wishes to reject the connection, it shall send the disconnection frames to the devices involved in the stream connection process.

If a stream is connected and if any of the three devices involved, sender, recipient or PNC, wishes to disconnect the stream, it shall send the stream management commanddescribed above. If the device initiating such a transaction is not the PNC, it sends this command to the PNC. The PNC then communicates the disconnection to the other involved device via the stream management command. If the PNC wishes to disconnect the stream, it sends the stream management command to the other devices involved in the connection.

The device sending the command to disconnect a stream shall consider the immediate acknowledgement received for that command frame as the acceptance by the other device.

8.4.3 Priority management

This sub-clause provides guidelines for priority management. The actual algorithms required for priority management is beyond the scope of this standard.

The PNC shall always provide preference for higher priority streams to be connected in the piconet. The PNC shall remember the priority of the streams already connected and use that information to allocate the required channel time in the order of stream priority.

The devices shall make local decisions on how the streams are transmitted in their allocated time slots. However the devices shall transmit an higher priority stream before lower priority stream data, except for isochronous streams. The frames that belong to isochronous streams and the streams meant for power saving devices may have fixed transmission time and hence they should be treated differently so as for them to be transmitted only during those fixed times.

8.5 Fragmentation and defragmentation

Fragmentation is performed at the transmitting DEV on each MSDU or MCDU whose size is greater than fragmentation threshold, aFragThreshold bytes. The transmitting DEV is free to choose different values for aFragThreshold parameter for different streams and change them as deemed fit. However, all the fragments in an MSDU/MCDU shall be of the same size except the last fragment that can be smaller than the rest of the fragments in the same MSDU/MCDU. Once fragmented and attempted a transmission, the MSDU/MCDU shall not be refragmented for any reason and all the retransmissions shall obey the original fragmentation threshold used on the MSDU/MCDU.

The first fragment of all MSDU/MCDUs shall have the Frag-start bit in the frame control field set to ‘1’. The first fragment shall have Frag-end bit in the frame control field set to ‘0’ unless it is also the last fragment of the MSDU/MCDU. The last fragment of all MSDU/MCDUs shall have the Frag-end bit in the frame control field set to ‘1’. The last fragment shall have Frag-start bit in the frame control field set to ‘0’ unless it is also the first fragment of the MSDU/MCDU. The rest of the fragments of the same MSDU/MCDU in between the first and the last fragment shall have both Frag-start and Frag-end bit in the frame control field set to ‘0’. The frame consisting of an MSDU/MCDU that has only one fragment shall have both Frag-start and Frag-end bit in the frame control field set to ‘1’.

Defragmentation of MSDU/MCDU is the reassembly of the received fragments into an MSDU/MCDU before delivering the data unit to next higher processing layer. The defragmentation for MSDU and MCDU are handled differently. The receiving DEV can process, and if needed respond to, all of the commands in a received MCDU fragment for which the command payload was completely received. For the last command in a MCDU fragment, the payload may continue to the next fragment unless the fragment has the frag-end bit set to ‘1’. Hence the receiving device shall wait for the next MCDU fragment from the same sender to complete processing of that command.

Defragmentation of MSDUs are performed completely before delivering to the rx data interface of MAC. The MAC is responsible for obtaining all the fragments of an MSDU before defragmentation. If for some reason, the receive DEV could not obtain all the fragments of an MSDU and the transmitting DEV indicated that it has discarded the frame, the receiving DEV shall discard all the fragments of the MSDU and indicate a receive failure to the rx interface of the MAC. The transmitting DEV can indicate that it has discarded the frame through any one of three different ways. If the stream does not need retransmissions at all, the Ack-policy employed, and indicatedin FC field, is no-ack. Hence the receive DEV must through away an MSDU if one of its fragments is missing and the stream has no-ack as the ack-policy. If the stream has immediate-ack or piggy-back-ack as the policy, reception of new frame with a different sequence number is an indication to the receive DEV that the transmitting DEV will not attempt retransmission on any of the frames in the same stream with older sequence numbers. If the stream has employed delayed-ack as the policy, the transmitting DEV shall inform the receive DEV to flush all its retransmission requests upto a sequence numebr using “Retransmission Sequence Resync” command described in <TBD: 7.5.12.2>.

There is no theoretical limit for the size of MSDU/MCDU with this fragmentation mechanism. However for practical reasons of limited resources at the receive DEV, the largest size of MSDU/MCDU that is accepted for fragmentation is limited to aMaxTransferUnitSize.

8.6 Acknowledgement and retransmission

There are two types of acknowledgements, immediate acknowledgment (Imm-ACK) and delayed acknowledgment (Del-ACK).

8.6.1 No acknowledgment

A transmitted frame with an ack policy set to indicate no-ack shall not be acknowledged by the intended recipient(s). The transmitting DEV assumes that the frame is succesful for all its local management and proceeds to the next frame inline for transmission.

8.6.2 Immediate acknowledgement frame

A directed frame that expects an immediate response frame (Imm-ACK) shall have the ACK-policy in that directed frame set to indicate the same, as defined in <TBD 7.2.1.1.2>. If the intended recipient of a directed frame correctly receives the frame, it shall start the transmission of the response frame within SIFS duration after the end of transmission of the directed frame.

8.6.3 Delayed acknowledgement

Delayed acknowledgement shall be applicable only for directed Stream Data frames. The intended recipient of the directed stream data frames is allowed to group the acknowledgement indications into retransmission request command described in 7.5.8.1. The negotiated retransmission window for the stream and the available transmission opportunities for the intended recipient of the stream govern the frequency of the retransmission request command. The intended recipient of the stream is free to send the delayed-ack command anytime before the expiry of the retransmission window. However the reciepient of the stream shall send delayed-ack command at the expiry of the retransmission window. In addition to this the source device for the stream can solicit the delayed-ack frame by setting the “Del-Ack-Request” bit in the FC of any frame of a stream that is employing delayed-ack policy.

8.6.4 Implied acknowledgement

Implied acknowledgement shall be applicable only for directed Stream Data frames sent during a time slot during CFP, with the source-address (SA) in that frame being the address of DEV that has been allocated that time slot. This ack policy shall not be used during CAP. The intended recipient of the directed stream data frames is allowed to send another Stream-data type frame in response to a directed frame received with acthe k policy set to implied-ack. The start of transmission of the response stream-data type frame shall start at the end of SIFS, like imm-ACK frame transmission, and end before the end of duration indicated in the received frame. The ack-policy field in the response stream-data frame shall take any value except the implied-ack itself. Hence the responding DEV may expect an imm-ACK to the stream-data type frame that is being sent as implied-ack response to the original received stream-data type frame. Since the ack-policy in the response stream-data type frame can not be set to expect an implied-ack, the chaining of implied-acks is not permitted. The frame exchange sequence involving implied-ack is shown in Figure 47. If the responding DEV could not fit its SIFS + response frame + SIFS (+ imm_ack + SIFS, if imm-ack is expected) within the duration indiated in the original stream-data type frame, then the responding DEV shall simply send an imm-ACK frame starting at the end of SIFS at the same PHY rate as the original received stream-data type frame. This is illustrated in Figure 48.

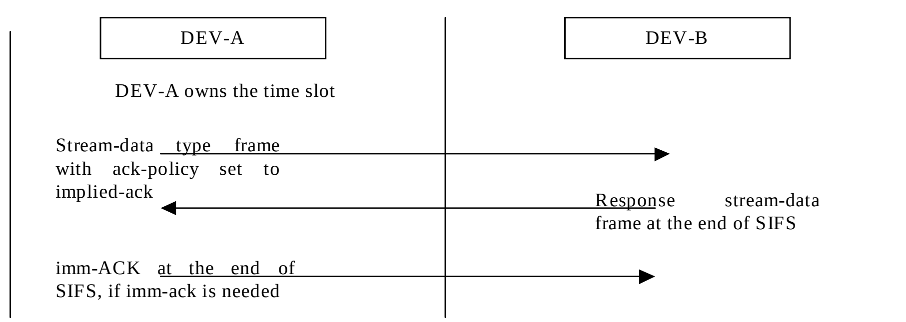

Figure 47—DEV-B responding with a stream-data frame for a frame received from DEV-A that has ack-policy set to implied-ack

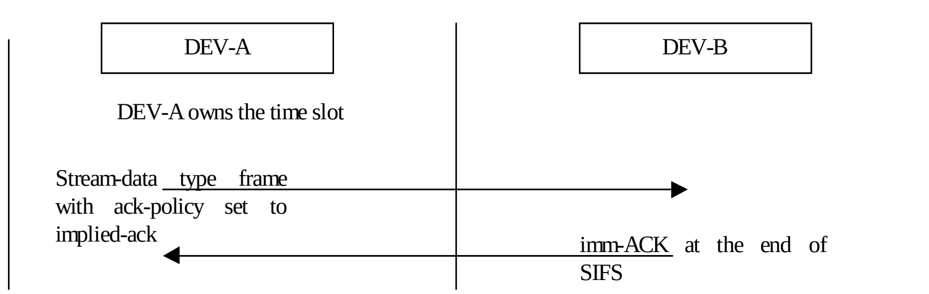

Figure 48—DEV-B sending an imm-ack as response to frame received from DEV-A that has ack-policy set to implied-ack

8.6.5 Retransmissions

During the CAP the retransmissions shall follow backoff rules as specified in sub-clause 8.2.2.

During the time slots within CFP when either imm-ack or the implied-ack is expected, the source-DEV shall measure CCA, between the end of SIFS and the end of RIFS (aCCADetectionDuration), to detect the start of the response frame. When no CCA is detected the source-DEV shall start the retransmission of the frame (or new transmission) at the end of RIFS as long as there is enough channel time remaining in the time slot for the entire frame exchange. When an imm-ACK is expected and the source-DEV detects CCA, as described above, it shall wait for the duration of imm-ACK frame plus SIFS before attempting another transmission. When an implied-ack is expected and the source-DEV detects CCA, the source DEV shall wait for the minimum of the duration that was indicated in its immediately previous transmission or the CCA indicating that the channel is idle. The retransmission of the frame or the transmission of a new frame shall start only after SIFS duration of the above events indicating the channel is available for transmission for the source-DEV of the time slot.

When delayed acknowledgement is used for a stream, the device transmitting the stream data may reject retransmission requests for the frames beyond the negotiated retransmission window for the stream. When retransmissions are rejected, the device transmitting the stream shall send retransmission sequence resync command, as defined in <TBD 7.5.8.2>, to the recipient of the stream in order to synchronize the retransmission requests.

A collision during the transmission of a directed frame during the CAP is detected by the absence of the immediate response (ACK) for that frame. Hence only the broadcast frames and the frames transmitted during burst transfer during the CAP may have ack-policy bits set to require delayed-ACK or no-acknowledgement.

Any frame can be attempted atmost aMaxRetransmissionLimit number of times before the transmitting DEV gives up on that frame and discards it. However a DEV might choose to transmit a “Stream-data” type frame for less number of times as some data streams have short life time.

8.6.6 Duplicate detection

Since the device sending the stream data may not correctly receive an ACK, duplicate frames may be sent even though the intended recipient has already received and acknowledged the frame. Hence all the devices are expected to detect such multiple receptions and indicate the stream-data frames to the higher layers only once. The stream ID and the sequence number field are used to detect multiple receptions of the same frame. Duplicate detection is also required by the PNC when it is providing repeater service. The PNC is expected to detect multiple receptions of a frame and not send the repeated frame over the medium to the final recipient.

8.7 Peer discovery

Each DEV associated with a PNC can use device information request command and obtian the information about its peers in the piconet. In addition the DEV can use probe information command to obtain other information required for peer-peer comunication (example TPC/DCS).

To accommodate the peer-discovery and peer-peer communication it is required to specify the response time for probe information command. A DEV that receives a Probe information command with any information request bit set shall respond with a Probe information command with its response within aProbeResponseDelay. If an appropriate Response to request in probe information command is not received by the requesting DEV within aProbeResponseDelay, the requesting DEV can send another probe information command with the same request. The time between these Probe information frames with the same request shall increase by a factor of 2 (i.e. exponential increasing). If the command frame containing the probe information command is sent during CAP, the incurred backoff channel time will be in addition to the delay described above.

8.8 Multirate support

In any approved IEEE standard for PHY for PANs there may be more than one rate specified. In each PHY there will be one mandatory base rate specified for the purposes described in this sub-clause. The base rate and other supported data rates for the 2.4 GHz PHY, they are defined in 11.3. A DEV shall send a frame with one of the supported data rates to a destination device only when the destination device is known to support that rate. The supported rates and the mandatory rate are PHY dependent. For the 2.4 GHz PHY, the supported rates are specified in 11.3. Each DEV in a piconet shall use probe information command, defined in <TBD: 7.5.2> to obtain the supported rates from the peer-DEV(s) that it is interested in communicating with. Similarly each DEV shall periodically use channel status request command, defined <TBD 7.5.5> to obtain the channel status information from the peer-DEV(s) that it is interested and decide the PHY rate to be used in transmissions to that peer-DEV(s).

All broadcast frames regardless of their type shall be sent at the base rate that is supported by the piconet as all devices must be able to receive these frames. The Allowed PHY rates for each of the different types of frames are listed in Table 45.

Table 45—Frame data rate limits

Frame type	PHY Data rate limitation
All group addressed frames (including Beacon)	Base rate
Immediate Acknowledgement (ACK)	Same rate as the rx-frame that is being ACK'ed
PNC selection	Base rate
Association Request	Base rate
Association Response	Base rate
Disassociation Request	Base rate
Directed Command frame	Any rate supported by both the source and destination
Directed Stream-Data	Any rate supported by both the source and destination

8.9 Repeater service

If the link between the two devices in a piconet is not satisfactory, either of those devices may request that the PNC to provide repeater service for that link. The PNC may grant the service if there is enough channel time available for the PNC to repeat the entire link while preserving the quality of service that is currently expected by the streams in that link. The PNC shall communicate the repeater service grant command to both the devices. The sequence of events for the establishment of the repeater service is illustrated in Figure 49. During this process, if either of the devices does not wish to continue, that device shall send the repeater service reject command, defined in 7.5.3.3, to the PNC. The PNC shall also send the repeater service reject command to the other device, if it is already involved.

Once the repeater service is granted between two devices, the sending device shall continue to send its frames as before. The PNC shall collect all the frames exchanged between the two devices and repeat them in the time slot of the PNC (or CAP) with PNC repeater field set in the MAC header of the repeated frames as specified in 7.2.1.1.8. The PNC shall decide the type of acknowledgement used on the repeated frames.

Once the repeater service has been established, the PNC may not be able to provide the best service on this link all of the time. If the service is not satisfactory, it is up to the devices to reject the service by sending the repeater service reject command to the PNC. If the PNC receives the repeater service reject command, it shall send the repeater service rejection command to the other involved device. Also, if the PNC feels it can no longer provide the repeater service for this link, the PNC may also send the repeater service rejection command to both the devices to reject the service.

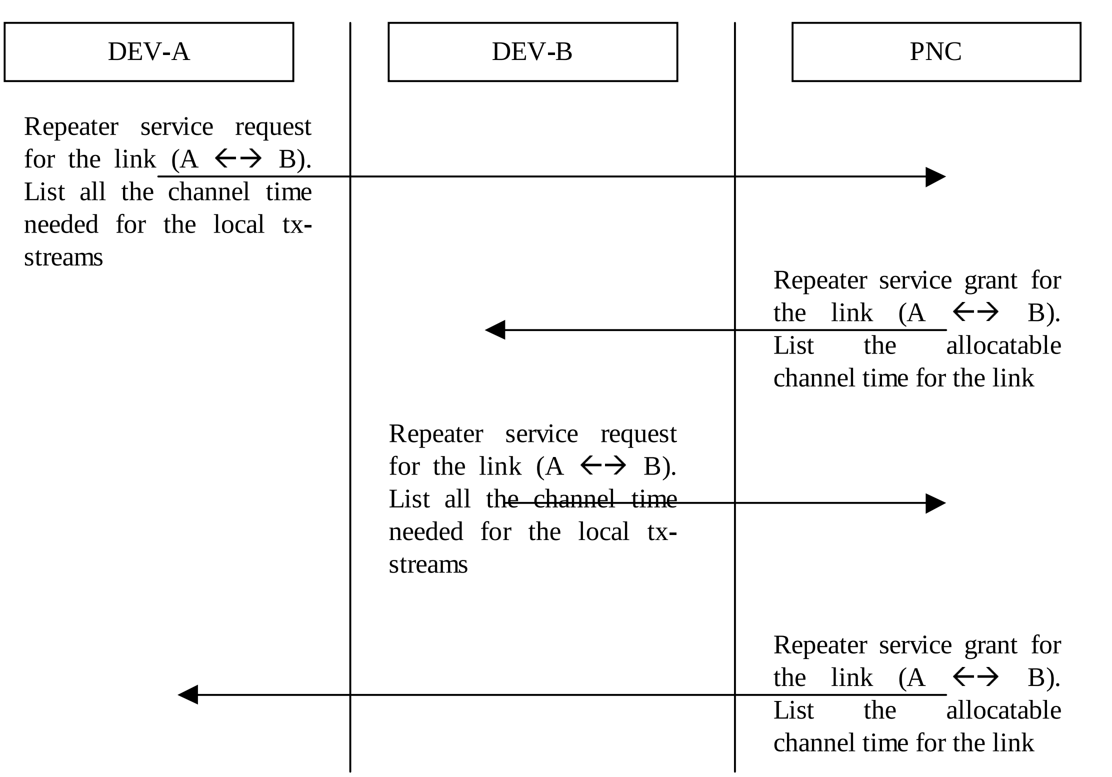

Figure 49—Illustration of establishment of repeater-service

8.10 Dynamic channel selection

The PNC initiates dynamic channel selection if the current QoS requirements cannot be satisfied due to adverse conditions of the channel. The PNC collects the channel status from the devices to arrive at this decision. The PNC shall send channel status request command, defined in 7.5.5, to request that the devices provide their channel status in return. The algorithm required to use the channel status information and decide change of channel is beyond the scope of this standard.

If the decision is made by the PNC to change the channel, the PNC shall announce to the piconet to keep quiet by broadcasting a remain-quiet command, as defined in 7.5.4, with the timeout period. Within that indicated timeout the PNC may change to one or more other channels to check if one of the other channels is better than the current channel and then return back to the current channel. If no beacons are received after this timeout period, the devices shall assume they are disconnected and start the association process.

If the PNC returns to the current channel within the indicated timeout, the PNC shall send a beacon to cancel the remain-quiet state of the piconet. Once the piconet is resumed the PNC may send a beacon with the channel change element indicating the new channel and the timeout for changing channel. The devices that received the beacon with channel change element shall change the channel to the indicated new channel within the indicated timeout duration and wait for beacons in the new channel. Whenever another channel is not available, the PNC may decide to stay in the same channel and hence may not send channel change element in its beacon following the cancellation of remain-quiet state of the piconet.

The dynamic channel selection process is illustrated in Figure 50.

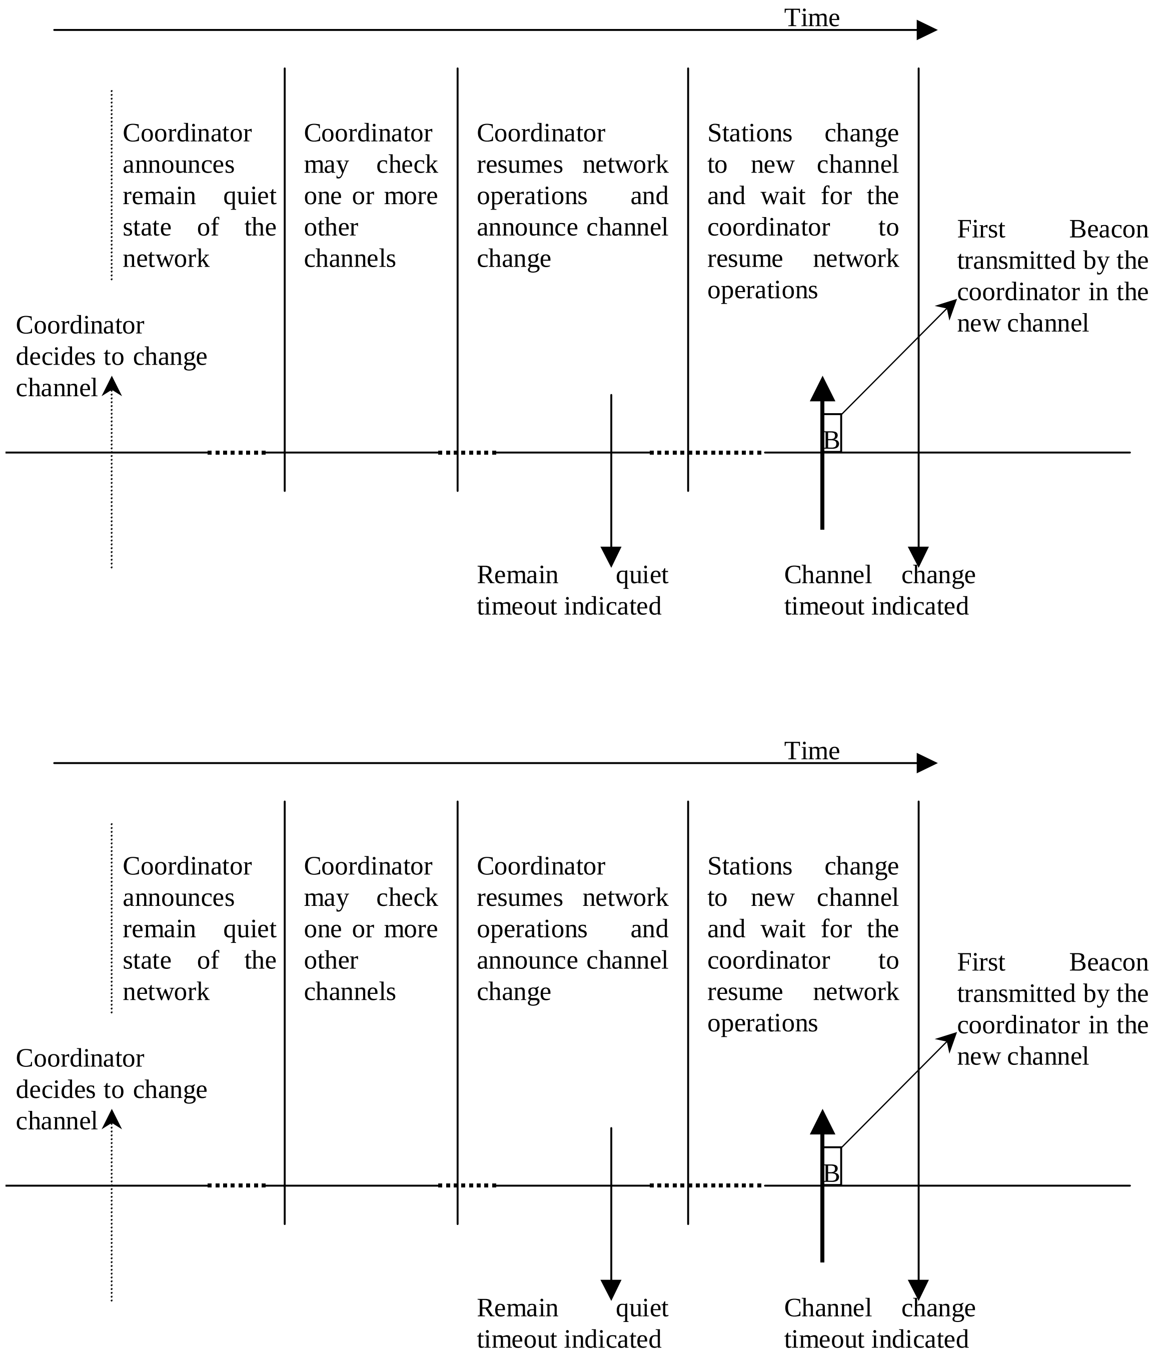

Figure 50—Illustration of dynamic channel selection process

8.11 Power management

There are three states in power management (PM) at MAC, (a) Active state (b) Snooze state and (c) Sleep state. The PS bit in capability field shall be set to 1 if the DEV is planning to use sleep state during an association state.

The DEV shall remain awake as long it is in active state.

In Snooze state, the DEV shall be awake at the beacon tx time and for the entire CAP. During CFP the DEV uses the channel time allocation received in Beacon or channel time grant command to decide when to be awake. Each DEV shall be awake at all GTSs in which its address is listed as either the source or the source DEV-address. All DEVs, except those in sleep state, shall be awake at all GTSs for which the destination DEV-address is Broadcast DEV-address. The rest of the time during the superframe the DEV is free to snooze and not transmit or receive any frame on the channel. Each DEV is free to use snooze state and save power in any superframe only if it has correctly received beacon frame for that superframe.

In Sleep state, the DEV may not receive or be able to transmit for several superframes at a time. The PNC is required to buffer all the frames to the sleeping DEV and all the broadcast frames received during the sleep state of the DEV and make them available to the DEV when it is awake. Hence the DEV shall inform the same to PNC its intention to go to sleep state just before going into sleep state. For this purpose, the DEV shall use “Sleep State Request” command and wait for PNC to respond. Only after the reception of “Sleep state permit” command from PNC, shall the DEV is allowed to shut itself off for a maximum sleep time duration as indicated by the PNC in “Sleep state permit” command. Note that the DEV must awake sufficient time before the expiration of that maximum sleep time in order to inform the PNC that the DEV is awake. If the DEV did not successfully indicate its awake state to PNC within the expiration of that maximum sleep time, the PNC shall disassociate the DEV from its piconet. If the DEV receives “Sleep state reject” command from the PNC instead of “Sleep state permit” command, the DEV shall not go to sleep state and wait for atleast one more beacon interal before attempting to send another “Sleep state request” command to PNC. But the device is free to use snooze state following the rules to enter snooze state described above.

When the DEV wakes up from sleep, it shall remain awake until atleast one beacon is correctly received. The DEV shall then inform the PNC that it is awake using “Active state Indication” command and wait for the PNC to deliver all the frames that are buffered for this DEV.

The PNC shall set the repeater bit to 1 in all the frames it is relaying to a DEV that is awake from sleep state. Since the PNC is buffering broadcast frames and transmitting them when the DEV is awake, it is possible for any DEV in the piconet to receive the same broadcast frames multiple times with or without repeater bit set. Hence each DEV shall use duplication detection mechanism described in 8.6.6 to reject all the multiply received broadcast frames.

The sequence chart in Figure 51 show the sequence of frames/commands exchanged between DEV and PNC for sleep-state management.

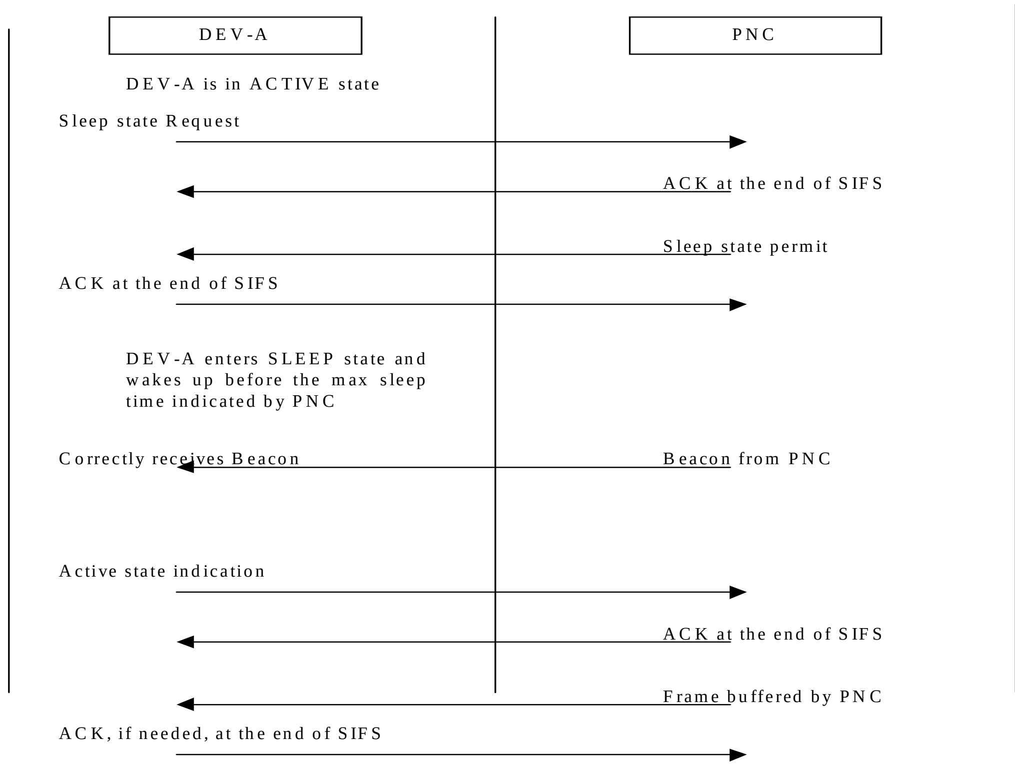

Figure 51—Message sequence chart for sleep state management at a DEV

8.12 Transmit power control

NOTE: Two modifications is needed in the PHY for TPC. In order to enforce consistency and interoperability for TPC, standardized transmitter power levels are required. Following the guidelines of the ERC, the TX power levels are specified as Effective Isotropic Radiated Power (EIRP). A table of the defined TX power level settings needs to be defined. These settings can be defined in 1 dB steps, in order to allow appropriate EIRP settings for all frequencies and regulatory domains. It is not required, nor is it expected, that all power levels will be implemented in a given DEV. The only requirement is that the device is capable of TX power levels that will allow compliance.

Under this scheme there are two basic mechanisms that can be used for TPC, a fixed maximum power level for all DEVs within a piconet or multiple power settings that are controlled by PNC. These are described below.

8.12.1 Fixed Maximum Transmitter Power

Fixed maximum transmitter power is the simplest TPC method. With this method, the appropriate maximum power level for the channel is determined by the PNC. The PNC then conveys this information to the DEVs within the Beacon frames with a TPC element. All DEVs within the piconet follow the maximum power level. Hence a DEV can simply set its TX power level to that received in TPC element of Beacon from its PNC and ignore all thechange of power level requests from its peers.

8.12.2 Adjustable Transmitter Power

Adjustable transmitter power based TPC is a flexible way to achieve transmit power control and at the same time reduce overall interference levels. With adjustable TPC, each DEV can request the peer-DEV it is communicating with to set its transmitter power to a certain level.

1. An DEV-1 estimates that its receive power from the DEV-2 is 6 dB higher than necessary.

2. The DEV-1 sends a Probe Information command with a TPC element to the DEV-2 with the “information request” bit corresponding to the TPC element set so that DEV-2 is being requested to send its current setting.

3. The DEV-2 replies with a Probe Information command that contains a TPC element providing the tx power level that is currently used during transmissions to the DEV-1

4. The DEV-1 sends another probe information command with the requested new setting that is always less than the max limit indicated by the PNC in the Beacon frame.

5. DEV-2 sets the tx power level to the minimum of max-limit indicated by the PNC in the Beacon frame and the requested power level. If the chosen power level is not implemented by DEV-2, it chooses the closest implemented tx-power level and sets that as the tx power level for all its transmissions to DEV-1.

8.13 Frame exchange rules

This clause defines the general frame exchange rules while the specific order of frames in a given frame exchange sequence is defined in an earlier clause describing the mechanism for which those frames are used. A list of the rules are specified in the tables below.

Table 46—Frame exchange limitations during CFP

Frame sequence	Number of frames in the sequence	Usage description
Association request - ACK	2	DEV requesting to associate with PNC
Association response - ACK	2	PNC allows the DEV to associate
Authentication request <TBD>- ACK	2	DEV requesting to be authenticated
PNC selection	1	PNC selection: During this time there is no CFP running as there is no PNC for the piconet that is being established
Group addressed frame	1	Broadcast frame transmission: Only if the source-DEV does not have any time slot allocated for non-Isochronous stream
Directed Command frame -ACK	2	Command frame transmission: Only if the source-DEV does not have any time slot allocated for non-Isochronous stream

Table 47—Frame exchange limitations during CFP

Frame sequence	Number of frames in the sequence	Usage description
Beacon frame	1	At the start of superframe
Group addressed frame	1	Broadcast frame transmission: Only in the time slot allocated for source-DEV with dest-DA of the time slot being the broadcast address
Association response - ACK	2	Only in the time slot allocated for PNC with dest-DA of the time slot being the broadcast address
Directed frame (requiring no-ack or delayed ack)	1	Only in the time slot allocated for the source-DEV with dest-DA of the time slot being the same as that of the transmitted frame
Directed frame (requiring imm-ack) -ACK	2	Only in the time slot allocated for the source-DEV with dest-DA of the time slot being the same as that of the transmitted frame
Directed frame (requiring implied-ack) - Directed frame (requiring no-ack or delayed-ack)	2	Only in the time slot allocated for the source-DEV with dest-DA of the time slot being the same as that of the first transmitted frame
Directed frame (requiring implied-ack) - Directed frame (requiring imm-ack) - ACK	3	Only in the time slot allocated for the source-DEV with dest-DA of the time slot being the same as that of the first transmitted frame

8.14 MAC layer parameters

The parameters that define some of the MAC characteristics are given in Table 48

Table 48—MAC layer parameters

Parameter	Value
aChannelScan	TBD ms
aCSFrameBroadcast	20 Kus
aADAdressReuseTime	1024 Kus
aRXTXTurnaroundTime	should be PHY clause
aCCADetectionDuration	should be PHY clause
aMaxBurstDuration	
aAssocRespConfirmTime	5 Kus
aAssocTimeoutPeriod	min 250 msec max two seconds
aProbeResponseDelay	8 Kus
aFragThreshold	DEV chooses this value
aMaxTransferUnitSize	8196 bytes
aMaxRetransmissionLimit	16
aBroadcastDevInfoDuration	1024 Kus