

IEEE P802.15

Wireless Personal Area Networks

Project	IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs)	
Title	Example what TSCH feature subclause could look like	
Date Submitted	13 September 2022	
Source	Tero Kivinen	E-mail: kivinen@iki.fi
Re:	September TG4me meeting	
Abstract	Provide example what TSCH subsection of the new IEEE Std 802.15.4-2024 would look like.	
Purpose	TG4me discussion	
Notice	This document has been prepared to assist the IEEE P802.15. It is offered as a basis for discussion and is not binding on the contributing individual(s) or organization(s). The material in this document is subject to change in form and content after further study. The contributor(s) reserve(s) the right to add, amend, or withdraw material contained herein.	
Release	The contributor acknowledges and accepts that this contribution becomes the property of IEEE and may be made publicly available by P802.15.	

5 Functional overview

5.1 TSCH Slotframes

In a TSCH PAN, the concept of the superframe is replaced with a slotframe. The slotframe also contains defined periods of communications between peers that are either CSMA-CA or guaranteed, but the slotframe automatically repeats based on the participating devices' shared notion of time. Unlike the superframe, slotframes and a device's assigned timeslot(s) within the slotframe can be initially communicated by beacon, but are typically configured by a higher layer as the device joins the network. Because all devices share common time and channel information, devices hop over the entire channel space to minimize the negative effects of multipath fading and interference and do so in a slotted way to avoid collisions, minimizing the need for retransmissions. Both of these features are desirable for operation in a harsh industrial environment.

6 MAC functional description

6.1 Generic non-beacon enabled

6.2 TSCH

6.2.1 Channel Access

6.2.1.1 TSCH CCA algorithm

When a device is operating in TSCH mode as described in 6.3.6, CCA may be used to promote coexistence with other users of the radio channel. For other devices in the same network, the start time of transmissions, *macTsTxOffset*, is closely aligned making intra-network collision avoidance using CCA ineffective. TSCH devices also use channel hopping, so there is no backoff period used when CCA prevents a transmission.

As illustrated in Figure 6-5, when a TSCH device has a packet to transmit, it shall wait for a link to the destination device. If *TschCca* was set to ON in the MLME-TSCH-MODE.request primitive, the MAC requests the PHY to perform a CCA at the designated time in the timeslot, *macTsCcaOffset*, without any backoff delays.




Figure 6-5 – TSCH CSMA-CA algorithm

6.2.1.2 TSCH CSMA-CA retransmission algorithm

Shared links (links with the linkOptions Bitmap set to shared transmission) are intentionally assigned to more than one device for transmission. This can lead to collisions and result in a transmission failure detected by not receiving an acknowledgment. To reduce the probability of repeated collisions when the packets are retransmitted, the retransmission backoff algorithm shown in Figure 6-6 shall be implemented for shared links. BE and NB are described in 6.2.5.1. The value of NB in Figure 6-5 is carried over to Figure 6-6.

When a packet is transmitted on a shared link for which an acknowledgment is expected and none is received, the transmitting device shall invoke the TSCH retransmission procedure, which includes a backoff algorithm. Subsequent retransmissions may be in either shared links or dedicated links as retransmission occurs in the next link to the destination. The device shall use an exponential backoff mechanism analogous to that described in 6.2.5.1. The MAC sublayer

shall delay for a random number in the range 0 to $(2^{\text{BE}} - 1)$ shared links (on any slotframe) before attempting a retransmission on a shared link. Retransmission on a dedicated link may occur at any time. For each successive failure on a shared link, the device shall increase the backoff exponent until the backoff exponent = *macMaxBe*.

The retransmission procedure shall follow the additional rules:

- The wait period, based on the backoff algorithm, applies only to the transmission on shared links. There is no waiting for transmission on dedicated links.
- The BE increases for each consecutive failed transmission in a shared link.
- A successful transmission resets the BE to the minimum value *macMinBe*.
- The BE does not change when a transmission is a failure in a dedicated link.

When the device is in TSCH mode the backoff is calculated in number of shared links. The values for *macMaxBe* and *macMinBe* for TSCH mode are defined in Table 8-96.




Figure 6-6—TSCH Retransmission backoff algorithm

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.

A device in a TSCH PAN may also use the PCA backoff method for critical event messages, as defined in 6.2.5.4.

6.2.1.3 TSCH slotframe structure

6.2.1.3.1 General

A slotframe is a collection of timeslots repeating in time. Each timeslot allows enough time for a pair of devices to exchange a frame and an acknowledgment. It is possible, although usually undesirable, to define a timeslot that is not long enough for a pair of devices to exchange a maximum length frame and an Ack frame. The number of timeslots in a given slotframe (slotframe size) determines how often each timeslot repeats, thus setting a communication schedule for nodes that use the timeslots. When a slotframe is created, it is associated with a slotframe handle (*macSlotframeHandle*) for identification. Each slotframe repeats on a cycle dependent on its length. Each timeslot is an opportunity for a device to send or receive a single frame, and optionally receive or transmit an acknowledgment to that frame. Slotframes and timeslots are configured by a higher layer.

When a timeslot has a node designated to transmit and a node designated to receive according to a *macLinkTable* schedule as per section 8.4.3.3.3, that timeslot is considered as having a link scheduled. If any of the conditions listed is absent, then there is no link scheduled for that timeslot.

6.2.1.3.2 Absolute slot number (ASN)

The total number of timeslots that has elapsed since the start of the network or an arbitrary start time determined by the PAN coordinator is called the Absolute Slot Number (ASN). It increments globally in the network every *macTsTimeslotLength*, as defined in Table 8-99. It may be beacons by devices already in a TSCH PAN, allowing new devices to synchronize. It is used globally by devices in a TSCH PAN as the frame counter (thus allowing for time-dependent security) and is used to compute the channel for any given pairwise communication as described in 6.2.6.3.

6.2.1.3.3 Links

Figure 6-9 illustrates an example of nodes communicating in a sample three-timeslot slotframe. Nodes A and B communicate during timeslot 0, nodes B and C communicate during timeslot 1, and timeslot 2 is not being used. Every three timeslots, the schedule repeats, but note that ASN

increments continuously. The pairwise assignment of a directed communication between devices for a given *macSlotFrameHandle*, as defined in 8.4.3.3.2, in a given timeslot on a given *macChannelOffset* is a link.

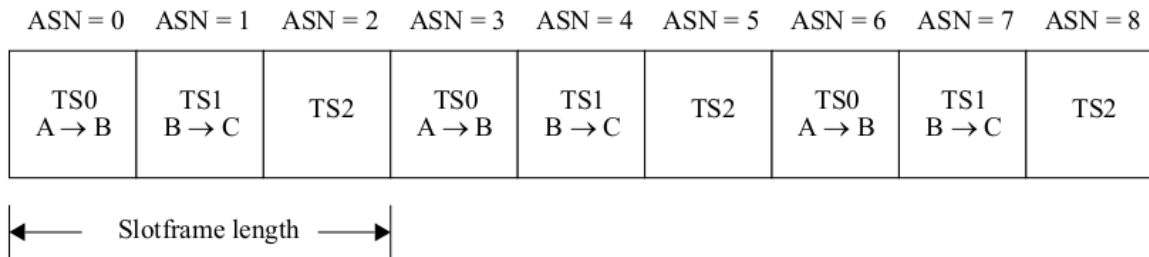


Figure 6-9 – Example of three timeslot slotframe

The physical channel, *CH*, in a link is calculated as follows:

$$CH = \text{macHoppingSequenceList} [(\text{macAsn} + \text{macChannelOffset}) \% \text{macHoppingSequenceLength}]$$

where $a\%b$ indicates a modulo b .

Use of a *macChannelOffset* allows for different channels to be used at a given *macAsn* for a given *macHoppingSequenceList*. There are *macNumberOfChannels* channel offsets that will result in a unique channel for that combination of *macAsn* and *macHoppingSequenceList*.

6.2.1.3.4 Multiple slotframes

A given network using timeslot-based access may contain several concurrent slotframes of different sizes. Multiple slotframes may be used to define a different communication schedule for various groups of nodes or to run the entire network at different duty cycles by giving some devices many active timeslots in a slotframe, and others few or none.

A network device may participate in one or more slotframes simultaneously, and not all devices need to participate in all slotframes. By configuring a network device to participate in multiple overlapping slotframes of different sizes, it is possible to establish different communication schedules and connectivity matrices that all work at the same time.

Slotframes can be added, removed, and modified while the network is running. Even though this is the case, all slotframes are aligned to timeslot boundaries, and timeslot 0 of the first repetition of every slotframe is projected back to *macAsn* = 0, which is determined by the PAN coordinator (or other network device that starts the network). Because of this, timeslots in different slotframes are always aligned, even though the beginning and end of a particular repetition of that slotframe may not be as illustrated in Figure 6-10. When, for any given timeslot, a device has links in multiple slotframes, transmissions take precedence over receives, and lower *macSlotframeHandle* slotframes takes precedence over higher *macSlotframeHandle* slotframes.

	ASN = 0	ASN = 1	ASN = 2	ASN = 3	ASN = 4	ASN = 5	ASN = 6	ASN = 7	
Slotframe 1 5 slots	TS0	TS1	TS2	TS3	TS4	TS0	TS1	TS2	...
Slotframe 2 3 slots	TS0	TS1	TS2	TS0	TS1	TS2	TS0	TS1	...

Figure 6-10 – Multiple slotframes in the network

6.2.1.4 Channel hopping

TSCH uses slotted mode channel hopping and uses network coordination within a slotframe, via a shared hop sequence to which devices participating in the network synchronize as described in 6.2.10.

6.2.2 Starting and maintaining PANs

6.2.2.1 Scanning through channels

6.2.2.1.1 Active and passive channel scan

TSCH uses Slotted CSMA-CA for channel access for responses to an Enhanced Beacon Request command. It will respond in the next available timeslot on same channel.

6.2.2.2 Realignment in a PAN

TSCH devices shall not send a Coordinator Realignment command. TSCH devices that receive a Coordinator Realignment command shall ignore the command.

6.2.2.3 Beacon generation

See Table 6-3 for IEs to include in the Enhanced Beacon Request command for TSCH.

6.2.2.4 TSCH PAN formation

A TSCH PAN is formed when a device, referred to as an *advertising device*, advertises the presence of the network by sending Enhanced Beacon frames upon receipt of a MLME-

BEACON.request from a higher layer. In a TSCH PAN the Enhanced Beacon frames contain the following IEs:

— TSCH Synchronization IE, as described in 7.4.4.2, containing timing information so new devices can synchronize to the network.

— *Channel hopping IE*, as described in 7.4.4.31, containing channel hopping information, as described in 6.2.10.

— *TSCH Timeslot IE*, as described in 7.4.4.4, containing timeslot information describing when to expect a frame to be transmitted and when to send an acknowledgment.

— *TSCH Slotframe and Link IE*, as described in 7.4.4.3, containing initial link and slotframe information so new devices know when to listen for transmissions from the advertising device and when they can transmit to the advertising device.

Enhanced Beacon frames in TSCH mode shall not be encrypted, but may be authenticated (i.e., security level 1, 2 or 3).

NOTE—If Enhance Beacon frames were encrypted, the TSCH Synchronization IE used to transmit the ASN to joining devices would be encrypted as well. A joining device or a device that has lost synchronization with the network would be unable to decrypt the Enhanced Beacon frame as the current ASN is required to generate the nonce. Thus, these devices would be unable to join the network or resynchronize with the network.

The device wishing to join a TSCH network begins passively (preferred) or actively scanning for the network as the result of receiving an MLME-SCAN.request from a higher layer. Once the listening device has heard a valid Enhanced Beacon, it generates an MLME-BEACON-NOTIFY.indication to a higher layer. The higher layer may wait for additional MLME-BEACON-NOTIFY.indication primitives before selecting a TSCH network based upon the value of the Join Metric field in the TSCH Synchronization IE. The higher layer may initialize the slotframe and links contained in the Enhanced Beacon from the preferred TSCH network and switch the device into TSCH mode with a MLME-TSCH-MODE.request.

At this point the device is synchronized to the network and may optionally send an Association Request command. If the device uses association, it may request a short address. The sequence of messages exchanged to synchronize a device to the networks is shown in Figure 6-21, and the process of synchronization is described in 6.5.4.

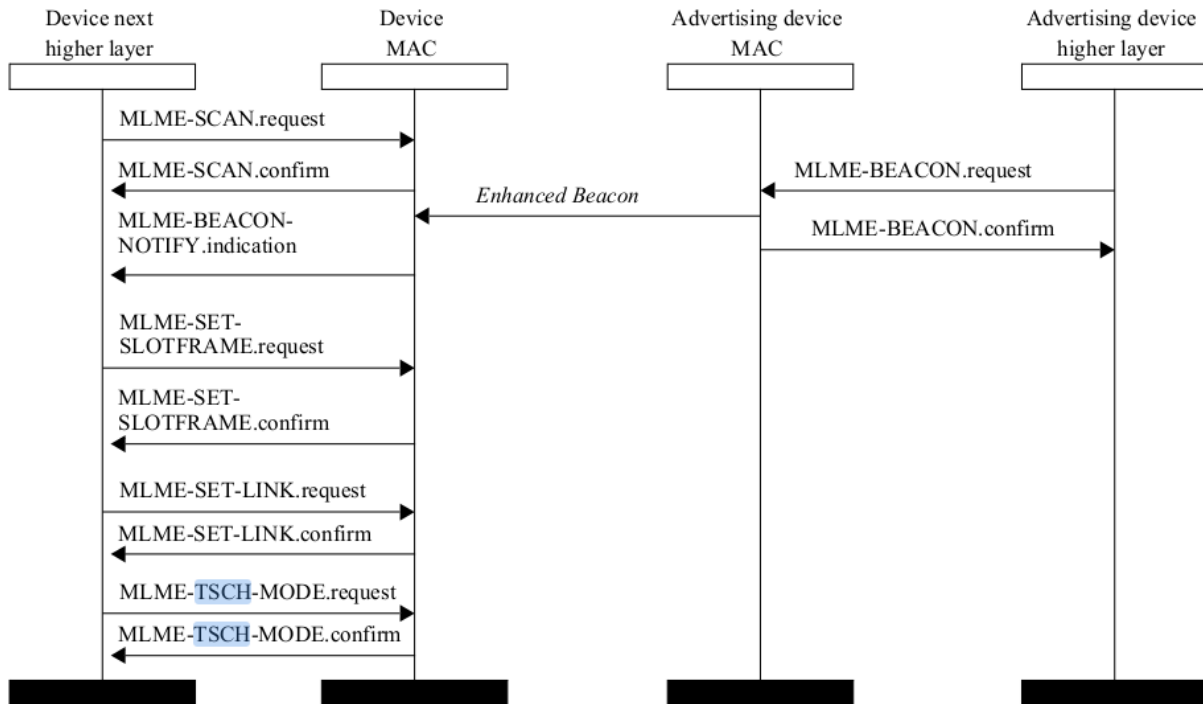


Figure 6-21 – Message sequence chart for TSCH procedure to find an advertising device

Typically at this point the device will go through a procedure to allocate additional communication resources (slotframes and links) to the joining device. This procedure may include a security handshake to mutually authenticate the joining device, configure encryption keys, and configure routing information. The mechanism and rules for setting up these additional communication links needs to be defined in a higher layer standard.

Once synchronized and configured by a higher layer to do so, all FFDs that are already part of the network may send Enhanced Beacon frames announcing the presence of the network. The advertising rate and content is configured by a higher layer as appropriate to the density of devices, the desired rate of network formation, and the energy devoted to network formation.

After joining, the device may receive additional slotframes and links from a higher layer management entity or peer. Likewise, the device may be instructed to remove certain slotframes and links obtained from the Enhanced Beacon.

6.2.2.5 Association and diassociation

6.2.2.5.1 Association

Association is optional for devices operating in TSCH mode.

6.2.2.5.2 Disassociation

For devices using the optional TSCH mode, additional disassociation behavior is required. A TSCH device shall disassociate from the PAN if it receives a Disassociation Notification command from either the PAN coordinator, or all of its time source neighbors, as defined in 6.5.4.

Upon determining that it should disassociate from the PAN, the device shall transmit Disassociate Notification commands to all its neighbors on any available link for `macDisconnectTime` timeslots, after which it should clear all synchronization information and leave the PAN.

6.2.3 Synchronization

6.2.3.1 General

For devices using the optional TSCH mode, initial synchronization is performed by the use of Enhanced Beacon frames, and synchronization is maintained by slotted communication with other devices in the PAN.

In a TSCH PAN, all communication happens in timeslots as described in 6.5.4.2. To remain synchronized, the devices should have the same notion of when each timeslot begins and ends, within $\pm macTsRxWait/2$. In a typical TSCH PAN, time propagates outwards from the PAN coordinator. A device shall periodically synchronize its network time to at least one other network device, which is a time source neighbor. A time source neighbor is another device for which the `macLinkTimekeeping` is TRUE in the `macLinkTable` of that link. The device may also provide its network time to one or more network devices via the ASN. A higher layer may add or change time source neighbors at any time.

Note that a device sending Enhanced Beacons to advertise a TSCH PAN should set the Timekeeping bit in the Link Option field, as described in 7.4.4.3, for the joining devices' receive link so that joining devices can maintain time synchronization until additional time source neighbors are configured by a higher layer.

A network device may have more than one neighbor as its time source. In such cases, the device shall synchronize its time to all of the neighbors that are acting as its time source, synchronizing to the relative drift of all its time source neighbors.

Figure 6-29 shows an example of time propagation in a TSCH PAN. The arrows indicate the direction of time distribution. In this example, the PAN coordinator acts as the time source for the entire network. Device 1 synchronizes to the PAN coordinator only, and is the time source for device 3. Device 2 synchronizes its time to both 1 and the PAN coordinator, and device 4 synchronizes to the PAN coordinator, device 2, and device 3.

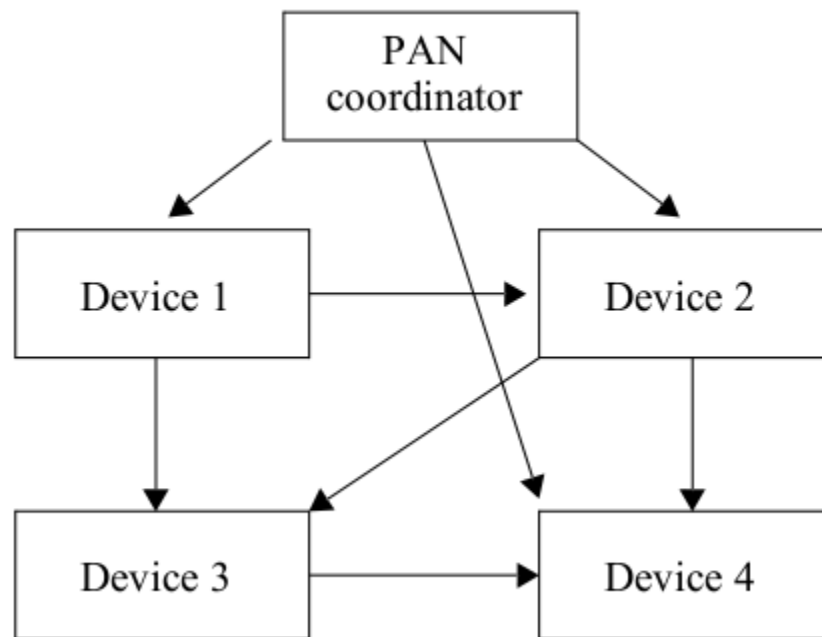


Figure 6-29—Example of possible time propagation in a TSCH PAN

Synchronization is possible whenever a device exchanges a frame with a time source neighbor. This can either come from receiving an acknowledgment with time correction information or from the arrival time of a frame from the time source neighbor. Both methods are described in 6.5.4.3.

In order to ensure that it remains synchronized with the TSCH PAN, a network device shall ensure that it communicates with each of its timekeeping neighbors, as defined in 8.2.19.7.

6.2.3.2 Timeslot communication

During a timeslot in a slotframe, one node typically sends a frame, and another sends back an Enh-Ack frame containing the Time Correction IE, as described in 7.4.2.7, if it successfully receives that frame. A positive acknowledgment indicates that the receiver has successfully received the frame and has taken ownership of it for further routing. A negative acknowledgment indicates that the receiver cannot accept the frame at this time, but has received it with no errors. The Time Correction IE includes timing information used by nodes to maintain network synchronization. Frames sent to a unicast node address require that a link-layer acknowledgment be sent in response during the same timeslot as shown in Figure 6-30. If an acknowledgment is requested and not received within the time-out period, retransmission of the frame waits until the next assigned transmit timeslot to that address occurs.

As shown in Figure 6-30, the timeslot starts at time $T = 0$ from the transmitting device’s perspective. The transmitter waits $macTsCcaOffset$ and then performs CCA (if active). At $macTsTxOffset$, the device begins transmitting the frame. If an acknowledgment is expected, the device waits $macTsRxAckDelay$ and then enables the receiver to await the acknowledgment. If the acknowledgment does not arrive within the expected time, the device may idle the radio and consider the transmission a failure. If no acknowledgment is expected, the transmitter may idle the radio after sending the frame.

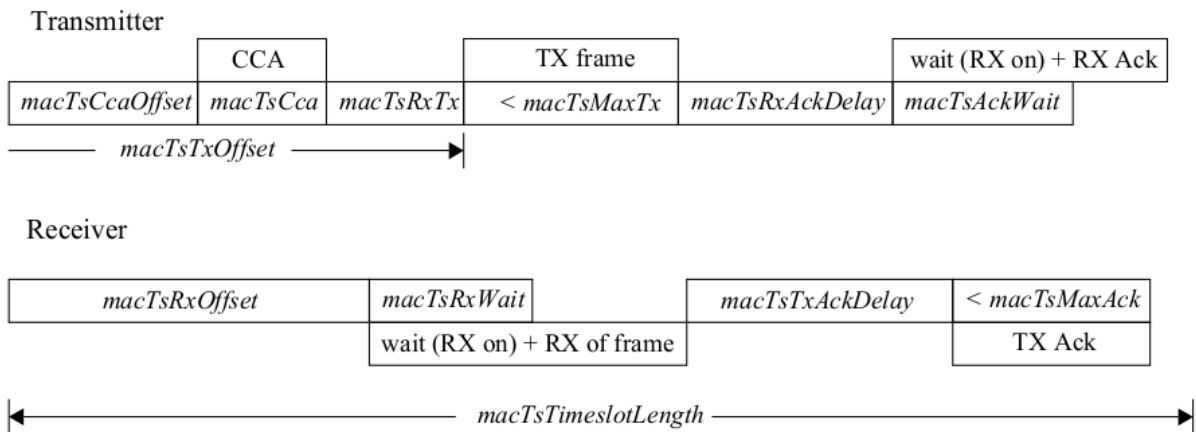


Figure 6-30—Timeslot diagram of acknowledged transmission

On the receiver’s side, at its estimate of $T = 0$ it waits $macTsRxOffset$ and then goes into receive mode for $macTsRxWait$. If the frame has not started by that time, it may idle the receiver. Otherwise, once the frame has been received, the receiver waits $macTsTxAckDelay$ and then sends an acknowledgment.

The transmitter or receiver may resynchronize clocks as described in 6.5.4.3.

6.2.3.3 Node synchronization

Device-to-device synchronization is necessary to maintain connection with neighbors in a slotframe-based network. There are two methods for a device to synchronize to the network, acknowledgment-based and frame-based. Originator in this context is the device sending a frame, and receiver is the device receiving that frame and sending back an acknowledgment as is appropriate. Since timestamps are required to maintain synchronization in a TSCH PAN, all devices shall have $macTimestampSupported = TRUE$.

Acknowledged communication provides a basic method of time synchronization through the exchange of Data frames and Ack frames. The algorithm involves the receiver calculating the delta between the expected time of frame arrival and its actual arrival, and providing that

information to the transmitting node in the subsequent acknowledgment.

The acknowledgment-based synchronization algorithm can be described as follows:

- Originator sends a frame, timing the start symbol to be sent at $macTsTxOffset$ according to its clock, which would correspond to $macTsRxOffset + macTsRxWait/2$ in the receiver's clock if both clocks were perfectly synchronized.
- Receiver records the timestamp when it receives the start symbol of the frame, where frame is as defined in 3.1.
- Receiver calculates a time correction = $macTsRxOffset + macTsRxWait/2$ arrival timestamp.
- Receiver sends back the time correction in the IE field in the corresponding Enh-Ack frame (assuming the incoming frame passes validation).
- Originator receives the acknowledgment. If the receiver node is a time source neighbor, the originator adjusts its own clock by incorporating the difference into an average of the drift to all its time source neighbors. The averaging method is implementation dependent. If the receiver is not a clock source, the time correction is ignored.

In frame-based synchronization a node may synchronize its own network clock if it receives a frame from a time source neighbor. The receiver calculates the delta between expected time of frame arrival and its actual arrival time and use that information to adjust its own clock.

The frame-based synchronization algorithm can be described as follows:

- Receiver records the timestamp when it receives the start symbol of the frame.
- The receiver calculates a time correction = $macTsRxOffset + macTsRxWait/2$ arrival timestamp.
- If the originator was a time source neighbor, the receiver adjusts its own clock by incorporating the information from all of its time source neighbors in an implementation-dependent manner. If the originator is not a time source neighbor, the time correction shall be ignored.

Figure 6-31 illustrates both time synchronization mechanisms. In both cases, the receiver calculates its time adjustment to either send back to the transmitting device or to use locally.




Figure 6-31—Time synchronization

6.2.3.4 Transaction handling

TSCH mode indirect behavior is further described in 7.2.2.4.

6.2.4 Transmission

If security is enabled in TSCH mode, then only one frame shall be sent in a given slot to prevent the same nonce being used for more than one frame. The ASN is used in the nonce in TSCH mode.

6.2.5 Reception

6.2.5.1 Acknowledgement

The Time Correction IE shall be used in all Enh-Ack frames if `macTschEnabled` is TRUE. When returning Time Correction IE, as described in 7.4.2.7, in the Enh-Ack frame, the receiving device may indicate a negative acknowledgment to indicate that the frame successfully passed FCS check, but that the MAC discarded the frame.

When in TSCH mode, incoming frames are acknowledged using the Enh-Ack frame as described in 7.3.3. Security of the Enh-Ack frame shall match that of the incoming frame. When operating in the TSCH mode, the Enh-Ack frame is sent at the time specified by the `macTimeslotTemplate`.

6.2.5.2 Retransmissions

If a single transmission attempt failed and the device is operating in TSCH mode, the retransmission process is defined in 6.2.5.3.

When not using TSCH mode and a frame with the Security Enabled field set to one is

retransmitted, the frame shall be retransmitted without changes and without passing through the outgoing frame security procedure, as defined in 9.2.2.

When using TSCH mode, and a frame with the Security Enabled field set to one is retransmitted, the frame shall follow the outgoing frame security procedure, as defined in 9.2.2.

NOTE—In TSCH mode, the security processing needs to be performed again because the ASN is used in the nonce and the retransmitted frame is sent in a slot with a different ASN.

6.2.5.3 Transmission timing restrictions

In TSCH PANs certain times may be allocated to devices for transmission. For TSCH the ATI is:

— TSCH timeslot: Start and end time defined the *macSlotframeHandle*, *macTimeslot*, and *macTsTimeslotLength*

6.2.5.4 Guard time

In a TSCH PAN the synchronizing event is a Timing Correction IE received from a timekeeping neighbor.

6.2.6 Special features