IEEE P802.15
Wireless Personal Area Networks

Project	IEEE P802.15 Working Group for Wireless Personal Area Networks (WPANs)		
Title	**<Preliminary draft for 4k>**		
Date Submitted	[February 24, 2012]		
Source	[Monique Brown] [<On Ramp Wireless, Inc.>] [San Jose, California U.S.A.]	Voice: Fax: E-mail:	[+1 954 608 7521] [ ] [monique.brown@ieee.org]
Re:	15-11-0863-03-004k-lecim-dsss-draft-text 15-11-0864-04-004k-802-14-4k-fsk-phy-draft-text 15-12-0044-00-004k_Informative_annex 15-12-0882-03-004k-working-draft-for-tg4k-mac-work		
Abstract	[This document combines the work done thus far by TG4k into a single document.]		
Purpose	[This document is the first step in preparing a draft for letter ballot. Rev.1 includes the MAC text from doc. 15-12-0882-03.]		
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.		

P802.15.4k/D0.1, March, 2012

IEEE Draft Standard for Local and metropolitan area networks—

Part 15.4: Low-Rate Wireless Personal Area Networks (WPANs)

Amendment X: Physical Layer Specifications for Low Energy, Critical Infrastructure Monitoring Networks

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Abstract:
Keywords: low data rate, low power, LR-WPAN, PAN, personal area network, radio frequency, RF, wireless personal area network, WPAN

Three Park Avenue
New York, New York 10016-5997, USA

This document is an unapproved draft of a proposed IEEE Standard. As such, this document is subject to change. USE AT YOUR OWN RISK! Because this is an unapproved draft, this document must not be utilized for any conformance/compliance purposes. Permission is hereby granted for IEEE Standards Committee participants to reproduce this document for purposes of international standardization consideration. Prior to adoption of this document, in whole or in part, by another standards development organization permission must first be obtained from the IEEE Standards Activities Department (stds.ipr@ieee.org). Other entities seeking permission to reproduce this document, in whole or in part, must also obtain permission from the IEEE Standards Activities Department.

IEEE Standards Activities Department
445 Hoes Lane
Piscataway, NJ 08854, USA

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information contained in its standards.

Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied "**AS IS**."

The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon the advice of a competent professional in determining the exercise of reasonable care in any given circumstances.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position, explanation, or interpretation of the IEEE.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Comments on standards and requests for interpretations should be addressed to:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854
USA

Authorization to photocopy portions of any individual standard for internal or personal use is granted by the Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Introduction

This introduction is not part of IEEE P802.15.4k/D0.1, IEEE Draft Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (WPANs)—Amendment X: Physical Layer Specifications for Low Energy, Critical Infrastructure Monitoring Networks.

This amendment specifies ...TBD

Notice to users

Laws and regulations

Users of these documents should consult all applicable laws and regulations. Compliance with the provisions of this standard does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Copyrights

This document is copyrighted by the IEEE. It is made available for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making this document available for use and adoption by public authorities and private users, the IEEE does not waive any rights in copyright to this document.

Updating of IEEE documents

Users of IEEE standards should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect. In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE Standards Association website at http://ieeexplore.ieee.org/xpl/standards.jsp, or contact the IEEE at the address listed previously.

For more information about the IEEE Standards Association or the IEEE standards development process, visit the IEEE-SA website at http://standards.ieee.org.

Errata

Errata, if any, for this and all other standards can be accessed at the following URL: http://standards.ieee.org/reading/ieee/updates/errata/index.html. Users are encouraged to check this URL for errata periodically.

Interpretations

Current interpretations can be accessed at the following URL: http://standards.ieee.org/reading/ieee/interp/index.html.

Patents

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

At the time the draft of this standard was sent to sponsor ballot, the IEEE P802.15 Working Group had the following voting members:

Robert F. Heile, *Chair*
Rick Alfvin, *Co-Vice Chair*

Patrick W. Kinney, *Co-Vice Chair and Secretary*

James P. K. Gilb, *Working Group Technical Editor*

Patrick W. Kinney, *Task Group 4k Chair*

, *Task Group 4k Vice Chair*

Monique B. Brown, *Task Group 4k Technical Editor*

Betty Zhao, *Task Group 4k Secretary*

<insert names here>

Major contributions were received from the following individuals:

The following members of the balloting committee voted on this standard. Balloters may have voted for approval, disapproval, or abstention. *<insert names here>*

When the IEEE-SA Standards Board approved this standard on *DD MM* 201x, it had the following membership: *<insert names here>*

,

Also included are the following nonvoting IEEE-SA Standards Board liaisons: *<insert names here>*

IEEE Standard for
Local and metropolitan area networks—

Part 15.4: Low-Rate Wireless Personal Area Networks (WPANs)

Amendment X: Physical Layer Specifications for Low Energy, Critical Infrastructure Monitoring Networks

NOTE—The editing instructions contained in this amendment define how to merge the material contained therein into the existing base standard and its amendments to form the comprehensive standard.

The editing instructions are shown in ***bold italic***. Four editing instructions are used: change, delete, insert, and replace. ***Change*** is used to make corrections in existing text or tables. The editing instruction specifies the location of the change and describes what is being changed by using ~~strikethrough~~ (to remove old material) and underscore (to add new material). ***Delete*** removes existing material. ***Insert*** adds new material without disturbing the existing material. Deletions and insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. ***Replace*** is used to make changes in figures or equations by removing the existing figure or equation and replacing it with a new one. Editing instructions, change markings, and this NOTE will not be carried over into future editions because the changes will be incorporated into the base standard.

3. Definitions, acronyms, and abbreviations

3.1 Definitions

Insert the following definitions alphabetically into 3.1:

3.2 Acronyms and abbreviations

Insert the following acronyms alphabetically into 3.2:

CDMA	code division multiple access
CIC	central inventory control
CLON	co-located orthogonal network
FVS	fragment validation sequence
HWSL	hybrid wakeup sample listening
I-ACK	fragment incremental acknowledgment
LECIM	low energy, critical infrastructure monitoring
OVSF	orthogonal variable spreading factor
PBRI	pruned bit reversal interleaving
P-FSK	position-based frequency shift keying
P-GFSK	position-based Gaussian frequency shift keying
RSLN	relayed slot-link network

4. General description

4.2 Components of the IEEE 802.15.4 WPAN

Insert the following paragraph at the end of 4.2:

Low energy critical infrastructure monitoring (LECIM) networks typically are asymmetric in power consumption and capability, with a coordinator that is mains powered (or otherwise provided a substantial power source), and have energy constrained endpoints which must have minimum energy consumption.

4.3 Network topologies

4.3.1 Star network formation

Insert the following paragraphs at the end of 4.3.1:

LECIM networks primarily operate in a star topology. The coordinator is not as limited with respect to energy and available resources as endpoints devices, in which energy consumption is critical and resources may be very limited. This asymmetry is a characteristic feature of the LECIM network.

For extending networking coverage, a star network may include end points that relay MAC frames synchronously inward to the PAN coordinator or outward to a device, to form a relayed link network operating as a virtual star network.

4.4 Architecture

4.4.2 MAC sublayer

Insert the following paragraph after the second paragraph of 4.4.2:

The following MAC enhancements are included to support the LECIM PHYs defined in Clause 19:

- Enhanced timing and synchronization capabilities to support synchronous and asynchronous channel access in both beacon-enabled and nonbeacon-enabled operation
- Enhanced low energy mechanisms
- MAC protocol data unit (MPDU) fragmentation to support extremely low data rates and limited PHY service data unit (PSDU) sizes
- Priority channel access
- MAC sublayer management entity (MLME) service access point (SAP) (known as MLME-SAP) and PAN information base (PIB) extensions for PHY control and configuration

4.5 Functional overview

4.5.1 Superframe structure

4.5.1.1 General

Insert the following items at the end of the list in 4.5.1.1:

- Superframe structure described in 4.5.1.2, based on beacons defined in 5.2.2.1, with an Information Element (IE) defined in 5.2.4.3.4 (Deterministic and synchronous multi-channel extension [DSME] PAN Descriptor), and priority channel access slots

- Support for priority channel access in the contention access period (CAP) of the superframe structure, as described in 4.5.1.5

Insert the following new subclause (4.5.1.5) after 4.5.1.4:

4.5.1.5 Use of superframe structure for LECIM

Priority-based contention channel access is provided in beacon-enabled mode by allocating the first two time slots of the CAP in each superframe, as shown in Figure 8. When the multi-superframe structure shown in Figure 4a is in use, the first two time slots in each CAP in the multi-superframe are allocated for high priority access.

When configured to support priority access, the priority access slots are used by devices with critical events to report, as defined by the higher layer via a PIB attribute and/or MLME or MAC common part sublayer (MCPS) data service parameters, using the currently configured clear channel assessment (CCA) mode (8.2.7). Transmission of messages with other-than-critical event priority will commence following the priority access time slots.

In a relayed slot-link network (RSLN), the PAN coordinator generates a cyclic-superframe that periodically transmits slotted-superframes, which can be combined into multi-superframes. The slotted-superframe contains a beacon slot, prioritized device slots, coordinator slots, and bidirectional device slots, as shown in Figure 4c. The prioritized device slot starts immediately following the beacon and provides an up-link to the coordinator for transmitting delay sensitive data from devices. The coordinator slot provides a down-link to devices for broadcasting frames. The bidirectional device slots in a cyclic-superframe are assigned to each device in an RSLN and provides a bidirectional link between a certain device and the PAN coordinator.

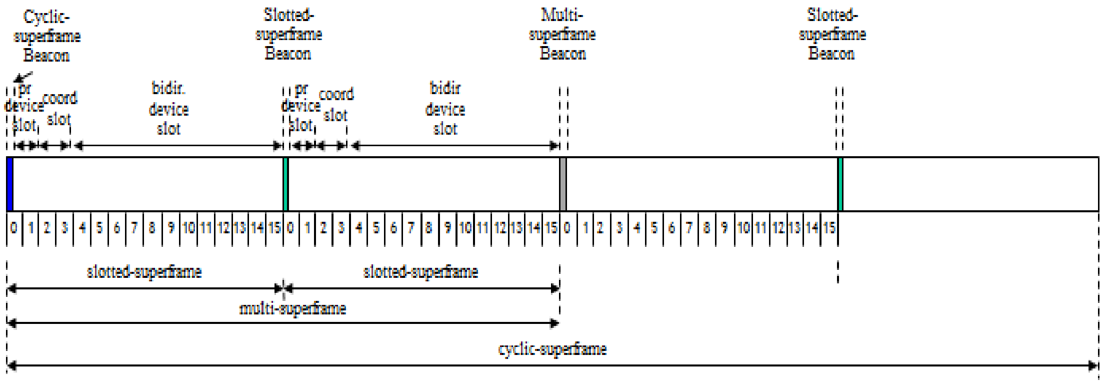

Figure 4c—An example of the cyclic-superframe structure

4.5.2 Data transfer model

4.5.2.1 Data transfer to a coordinator

Insert the following paragraph at the end of 4.5.2.1:

In an RSLN PAN, there are two methods for transferring data to a coordinator. When a device wishes to transfer delay sensitive data to a coordinator, the device transmits the data frame in the earliest prioritized device slot of a slotted-superframe. If a device fails to transfer the data in a prioritized device slot and wishes to transfer data to a coordinator, the data frame will be transmitted on the bidirectional device slot, which is allocated exclusively for the device in an RSLN.

4.5.2.2 Data transfer from a coordinator

Insert the following paragraph at the end of 4.5.2.2:

In an RSLN PAN, there are two methods for transferring data to a device. When the coordinator wants to broadcast data to devices in an RSLN, the coordinator may use the coordinator slot. When the coordinator wishes to transfer data to a device without notification, the coordinator may transmit the data frame continuously on the bidirectional device slot assigned to the device of a cyclic-superframe until the device acknowledges the successful reception of the data.

4.5.4 Improving probability of successful delivery

Insert the following new subclause (4.5.4.1a) after 4.5.4.1:

4.5.4.1a CSMA-CA used with priority channel access

When using the critical event priority access in a nonbeacon-enabled PAN where unslotted carrier sense multiple access with collision avoidance (CSMA-CA) channel access mechanism is applied, priority channel access is achieved by an use of the alternate backoff mechanism, as described in 5.1.1.4. The alternate mechanism uses a fixed backoff window instead of an exponential backoff and will, on average, provide less backoff duration for priority access than for normal access. In addition, the priority channel access continues to follow the channel, even if it is assessed to be busy, in order to gain immediate access to the channel once it is assessed to be idle.

Beacon-enabled PANs using a critical event priority access dedicate CAP time slots in the beginning of a superframe for priority channel access. Priority frames may commence in the priority slot(s) and continue through the duration of the CAP. Priority frames sent in the priority access slots utilize persistent CSMA-CA with a reduced contention window length and an alternate backoff mechanism for channel access, as described in [xref]. Priority frames sent in the non-priority slots of the CAP utilize CSMA-CA, as described in 5.1.1.4, with an alternate backoff mechanism.

4.5.4.2 ALOHA mechanism

Insert following paragraph after the last paragraph of 4.5.4.2:

When priority channel access is enabled, the alternate (i.e., fixed window) backoff mechanism is used, as described in xref. When operating in a beacon-enabled PAN, slotted ALOHA improves efficiency of channel access. When using slotted ALOHA with priority access, the first two time slots after the beacon are dedicated for priority channel access traffic. The backoff slot length is PHY dependent and should be able to accommodate, at minimum, the transmission of a single MPDU fragment.

Insert the following new subclause after 4.5.4.2:

4.5.4.2a MPDU fragmentation

With the addition of very low data rate PHY operating modes, the resulting increase in over-the-air duration of the MAC frame can lead to increased interference potential, susceptibility to channel conditions changing during the duration of a MAC frame transmission, and other effects that may reduce reliable transfer in some environments typical of LECIM applications. The long packet duration also brings a large cost for retransmission, both in terms of energy consumed and interference footprint. MPDU fragmentation can improve the probability of successful transmission and reduce the cost of retransmission. With fragmentation, each fragment is packaged into a PPDU for transmission, and this smaller PPDU has a reduced interference footprint. Also, retransmissions can be performed on a per fragment basis without needing to retransmit the entire original packet.

MPDU fragmentation operates on the complete MPDU and adapts it to the specific PHY and PHY operating mode. To reduce over-the-air overhead, some MAC header information is compressed or suppressed in the over-the-air exchange, by establishing a fragment sequence (transaction) context. The combination of the information in the fragment and the fragment sequence context provides identification of the individual fragment, the sequence to which it belongs, and where the fragment fits into the sequence. Each fragment carries an incremental validity check sequence for detecting errors. A schematic view of the fragmentation process is shown in Figure 6a.

In this standard, the term “fragment” refers to an individual MPDU fragment, the term “fragment sequence” refers to the collection of fragments transmitted that together comprise the original MPDU, “fragment number” is the position in the sequence of an individual fragment, and the “fragment sequence ID” identifies the fragment sequence.

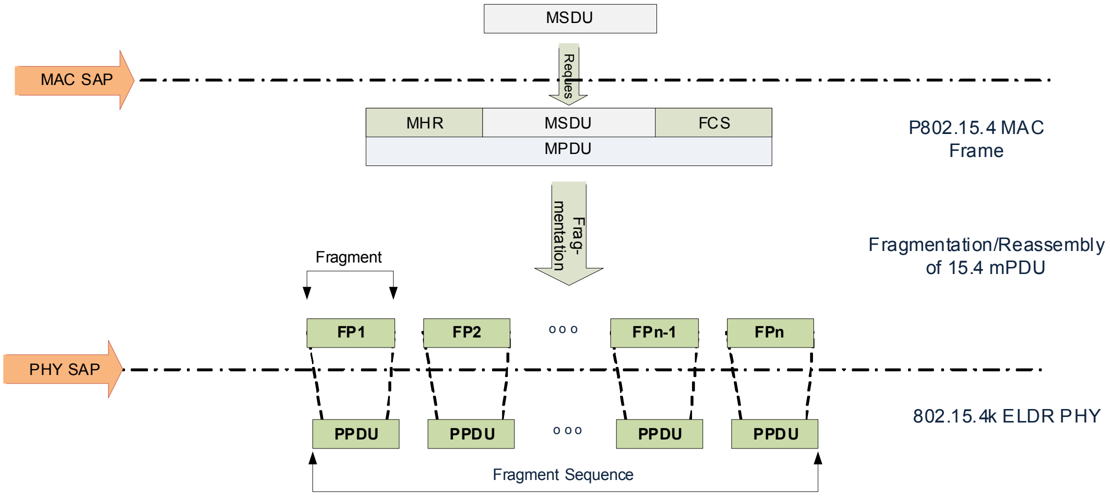

Figure 6a—Schematic view of MPDU fragmentation

Each fragment may be individually acknowledged and retransmitted. Retransmission of only the missed fragments can reduce air time and improve reliability. The complete MPDU transaction may be acknowledged.

4.5.4.3 Frame acknowledgment

Insert the following new subclauses (4.5.4.3.1, 4.5.4.3.2) after 4.5.4.3:

4.5.4.3.1 Fragment incremental acknowledgment (I-ACK)

The incremental acknowledgement (I-ACK) is used during the fragment sequence transfer to determine which fragments have been received successfully and which need to be retransmitted. An I-ACK may aggregate the status of one or more fragments. The number of fragment status reports grouped into an I-ACK is controlled by the higher layer. The format of the I-ACK is given in 5.1.6.4.5.

4.5.4.3.2 MPDU completion acknowledgment

The MPDU acknowledgement mechanism may be used to report the status of a fragment sequence transaction upon reconstruction of the MPDU at the recipient. The reassembly and validation process may require processing time in the MAC sublayer or higher layers prior to transmitting the final acknowledgement. A method is provided to coordinate the acknowledgement. Because fragment failures may occur due to conditions on a specific frequency channel, transmitting the acknowledgement and subsequent retransmissions on a different channel may also be desirable. A coordination mechanism is provided to support this capability, as described in xref. Means are also provided to include feedback to the initiator of the transaction, such as link quality information (xref to IE definitions), which is made available to the higher layer and may be used for adjusting fragmentation parameters or PHY configuration based on performance.

4.5.4.4 Data verification

Insert the following paragraphs at the end of 4.5.4.4:

To accommodate individual fragment acknowledgement, a fragment validation sequence (FVS) is included with each fragment. The recipient uses the FVS and fragment number to determine which fragments of the sequence have been received correctly and which are missing. The I-ACK reports the status of one or more received fragments. The FVS is described in (xref).

The reassembled MPDU also carries a frame check sequence (FCS). The MAC may apply this FCS as a validity check of the reassembled MPDU.

Insert the following new subclause (4.5.4.6) after 4.5.4.5:

4.5.4.6 Multiple grades of synchronous channel access

The times of occurrence of events are often crucial for the observer, and maintaining synchronous channels can support accurate time-stamping for measuring events. The synchronous channel access helps to distribute the data transfers in time scale and can provide multiple grades of channel access. In an RSLN PAN, three grades of synchronous channel access are provided: grade 0 for transmitting delay sensitive data, grade 1 for the reliable transmission of data, and grade 2 for the best effort data transmission.

For grade 0 channel access, a device first searches the earliest prioritized device slot. If the device fails to transmit the data in the prioritized device slot, the device will continue trying to transmit the data in either a bidirectional device slot or in another prioritized device slot, whichever comes first. A device using grade 1 channel access waits for the primary bidirectional device slot in the cyclic-superframe and transmits the data. If the device fails to transmit the data in the primary bidirectional device slot, the device will keep searching supplementary bidirectional device slots for the duration of the cyclic-superframe or will search the coming cyclic-superframe for an opportunity to transmit the data.

Add some detail about grade 2.

4.5.5 Power consumption considerations

Insert the following new subclause (4.5.5.3) after 4.5.5.2:

4.5.5.3 Low energy extension of networking coverage by synchronous relaying

In a star network, the coverage of networking will be limited by the transmission range of the device. For low powered devices, transmit power may be limited to increasing the life span of the device within the network. Compared to the energy-constrained end point device, the abundantly powered coordinator can have greater responsibility to extend the coverage of the star network with no burden to a device while preserving the topology. In an RSLN PAN, a cyclic-superframe repeater provides synchronous relaying of the frames inward or outward between the PAN coordinator and end device, in order to extend the coverage of a star network.

5. MAC protocol

5.1 MAC functional description

5.1.1 Channel access

5.1.1.1 Superframe structure

Insert the following paragraphs after the last paragraph of 5.1.1.1:

When priority access is enabled and the superframe structure shown in Figure 8 is in use, the first two time slots in the CAP shall be dedicated for priority channel access. When priority access is enabled and the multi-superframe structure shown in Figure 34g is in use, the first two time slots in each CAP shall be dedicated for priority channel access. See 5.1.1.4 for more information.

The superframe structure used for RSLN applications is described in 5.1.1.8.

5.1.1.4 CSMA-CA algorithm

Insert the following new subclauses (5.1.1.4.5, 5.1.1.4.6) after 5.1.1.4.4:

5.1.1.4.5 CSMA with priority channel access

This subclause describes the alternate backoff procedure used to support priority channel access for transmission of a critical event priority message. This backoff procedure shall be used when the CCA returns channel busy and priority access is enabled.

When operating a LECIM PHY in a nonbeacon-enabled PAN using unslotted CSMA-CA, the critical event priority transmission may be initiated at any time. During transmission of a priority message, when the CCA returns a status of channel busy, the alternate backoff procedure shall be used: the backoff exponent *BE* remains constant for subsequent retransmissions. The first transmission attempt shall set *BE* to the value of *macMinBE*–1 (with a default value of *macMinBE* = 2). In addition, the priority channel access follows a persistent CSMA mechanism, where a device continues to monitor the channel and decrements the value of unit backoff periods any time the channel is sensed idle for a duration of a backoff slot, in order to gain access to the channel as soon as possible.

In a beacon-enabled PAN, a critical event priority message transmission may be initiated in any part of the CAP. When transmission is initiated in the priority time slots, and the CCA returns channel busy, the alternate backoff mechanism shall be used as follows: *BE* remains constant (tentatively two, or *macMinBE*) for retransmissions. The first transmission attempt shall set *BE* to the value of *macMinBE*–1.

When a critical event priority transmission is initiated within the CAP in a time slot that is not a priority access time slot, the primary CSMA-CA, as defined in 5.1.1.4, with the above alternate backoff mechanism shall be used.

5.1.1.4.6 LECIM ALOHA priority channel access

When critical event priority channel access is in use with CCA Mode 4 (ALOHA), priority channel access is achieved by using an alternate backoff mechanism. A backoff period is defined as *macLECIMAlohaBackoffSlot* durations. A *macLECIMAlohaBackoffSlot* duration is a both a PHY and deployment-dependent parameter. It shall be sufficiently long in order to accommodate the transmission of a single MPDU fragment with associated interframe spacing (IFS) periods and any ACK frames. The backoff window size shall stay constant during retransmissions

In beacon-enabled PANs, slotted ALOHA is applied for more efficient channel access. When critical event priority channel access is in use, the slot length shall be equal to *macLECIMAlohaBackoffSlot* duration. In addition, the first two time slots after the beacon transmission are dedicated for priority channel access traffic. Priority frames may be transmitted in the entire CAP portion of the superframe.

5.1.1.7 LE-Functional description

Change the first paragraph of 5.1.1.7 as indicated:

This subclause specifies functionalities of devices supporting the following PIB attributes:

- *macCSLPeriod*
- *macRITPeriod*
- *macCSLMaxPeriod*
- *macHWSLMaxPeriod*
- *macHWSLPeriod*
- *macIRITEnabled*
- *macLowEnergySuperframeSupported*
- *macLowEnergySuperframeSyncInterval*

5.1.1.7.1 LE-Contention access period (LE-CAP)

Change the first paragraph of 5.1.1.7.1 as indicated:

When *macCSLPeriod* is non-zero, CSL is deployed in CAP, and HWSL is deployed in CAP when *macHWSLPeriod* is non-zero. CSL behavior is defined in 5.1.11.1, and HWSL behavior is defined in 5.1.11.3. The *macRITPeriod* shall be set to zero in a beacon-enabled PAN.

5.1.1.7.4 LE-Scan

Change the first paragraph of 5.1.1.7.4 as indicated:

When *macCSLPeriod* is non-zero, CSL is deployed in channel scans. When *macCSLMaxPeriod* is non-zero, each coordinator broadcasts beacon frames with wakeup frame sequence. When *macHWSLPeriod* is non-zero, each endpoint device deploys HWSL in channel scans. When *macHWSLMaxPeriod* is non-zero, each coordinator sends a wakeup sequence. Both cases ~~This~~ allow~~s~~ devices to perform channel scans with low duty cycles.

Insert the following new subclauses (5.1.1.8–5.1.1.8.5) after 5.1.1.7.4:

5.1.1.8 RSLN slot-link structure

5.1.1.8.1 General

A relayed slot-link network has slot-links between the PAN coordinator and each device in the network. A slot-link is the pairwise assignment of a directed communication between the PAN coordinator and a device(s) in a given time slot. The PAN coordinator generates a sequence of time slots and repeats the sequences to form a cyclic-superframe, as shown in Figure 11i. Time slots in a cyclic-superframe may be a 1-to-1 link (i.e., a link between the PAN coordinator and a single device) or a 1-to-*n* link (i.e., a link between the PAN coordinator and *n* devices).

The cyclic-superframe provides slot-links to devices, the slotted-superframe, in time scale. The slotted-superframe consists of a beacon slot, a prioritized device slot, a coordinator slot, and bidirectional device slots.

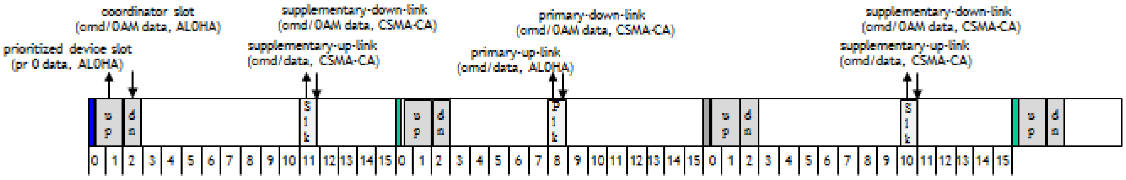

Figure 11i—Example of slot-links in a cyclic-superframe

5.1.1.8.2 Beacon slot

The beacon slot provides a link for transmitting a beacon from the PAN coordinator to devices. The beacon slot is reserved for the RSLN PAN coordinator, and the start of every slotted-superframe is indicated with the transmission of a beacon.

The beacon provides information, such as the structure of the cyclic-superframe and global time information, to the RSLN PAN.

5.1.1.8.3 Prioritized device slot

The prioritized device slot provides a link for transmitting delay sensitive data from a device to the PAN coordinator. The number of the prioritized device slots is defined as *macNumPrioritizedDeviceSlot*.

A device shall use the slotted ALOHA mechanism to access the prioritized device slot-link.

5.1.1.8.4 Coordinator slot

The coordinator slot provides a link for transmitting data from the PAN coordinator to devices. The number of the coordinator slot is defined as *macNumCoordSlot*.

The PAN coordinator shall use the slotted ALOHA mechanism to access the coordinator slot-link.

5.1.1.8.5 Bidirectional device slot

Each bidirectional device slot provides a link for transmitting data either from a device to the PAN coordinator or from the PAN coordinator to a device. Bidirectional device slot-links are assigned to all the devices in an RSLN PAN. If the number of bidirectional device slots in a cyclic-superframe is larger than the number of devices in the RSLN PAN, each device may be assigned a preemptive bidirectional device slot-link. Alternatively, some devices may share a bidirectional device slot-link to the PAN coordinator.

The channel access mechanism of a bidirectional slot-link depends upon the direction of transmission. On the access of the bidirectional device slot-link, the device gives priority in use. A device transmits at the start of the assigned bidirectional device slot without sensing the medium. Each time the PAN coordinator wishes to transmit data on a bidirectional device slot-link assigned to a certain device, it waits a random number of backoff periods at the start of the assigned bidirectional device slot. If the slot-link is found to be idle, the PAN coordinator begins transmitting.

One primary bidirectional device slot and multiple supplementary bidirectional device slots are allocated to each device in an RSLN PAN. The supplementary bidirectional device slot provides additional slots for initial transmissions or for retransmitting a frame which failed to transmit in the primary bidirectional device slot. A device shall use a slotted CSMA-CA mechanism when accessing a supplementary bidirectional device slot-link.

5.1.2 Starting and maintaining PANs

New methods to support LECIM go here, additions to scan for example.

Insert the following new subclause (5.1.2.7) after 5.1.2.6:

5.1.2.7 RSLN Pan formation

An RSLN PAN is formed when an MLME-RSLN-START.request primitive, as defined in 6.2.x.x, is issued to the MLME of the PAN coordinator requesting that an enhanced beacon be transmitted in every beacon slot of a cyclic-superframe. Each enhanced beacon contains:

- Cyclic-superframe specification (*BO*, *SO*, *MO*, number of prioritized device slot, number of coordinator slot)
- Time synchronization specification (global clock timestamp, slotted-superframe ID)
- Synchronous relaying specification (current depth of relaying, beacon information about neighbor tiers' repeaters)
- Indirect data transmission information

Need an introductory sentence for this figure

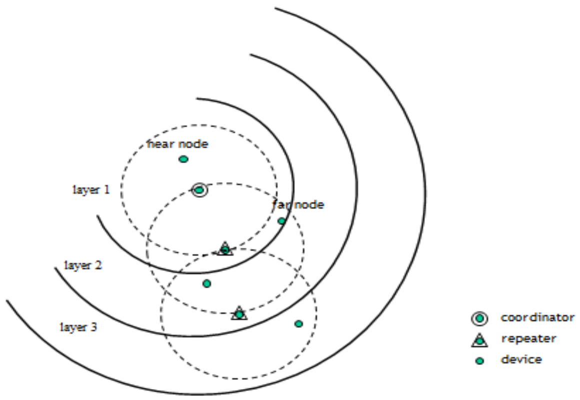

Figure 16b—Relayed slot-link network (RSLN) PAN

5.1.3 Association and disassociation

Insert the following new subclauses (5.1.3.4, 5.1.3.5) after 5.1.3.3:

5.1.3.4 Link context association when MPDU fragmentation is used

Describe the link context setup for mapping full address to coordinator and device short addresses.

5.1.3.5 RSLN repeater association

The next higher layer should attempt to associate as a repeater after having completed a channel scan. The goal of the next higher layer should be to associate to the PAN either through the PAN coordinator itself or through the repeater in the list of RSLN PAN descriptors that is closest to the PAN coordinator.

Following the selection of an inner coordinator through which to associate, the next higher layer should request through the MLME-RSLN-ASSOCIATE.request primitive, as described in 6.2.x.x, that the MLME configures the following values for repeater association to the RSLN PAN:

— RSLN PAN information (*phyCurrentChannel*, *phyCurrentPage*, *macPANId*)

— inner coordinator information (*macCoordExtendedAddress* or *macCoordShortAddress*)

— repeater information (current depth of relaying, slotted-superframe ID)

The PAN coordinator shall allow association only if the relayed beacon slot is available.

Need an introductory sentence for this figure

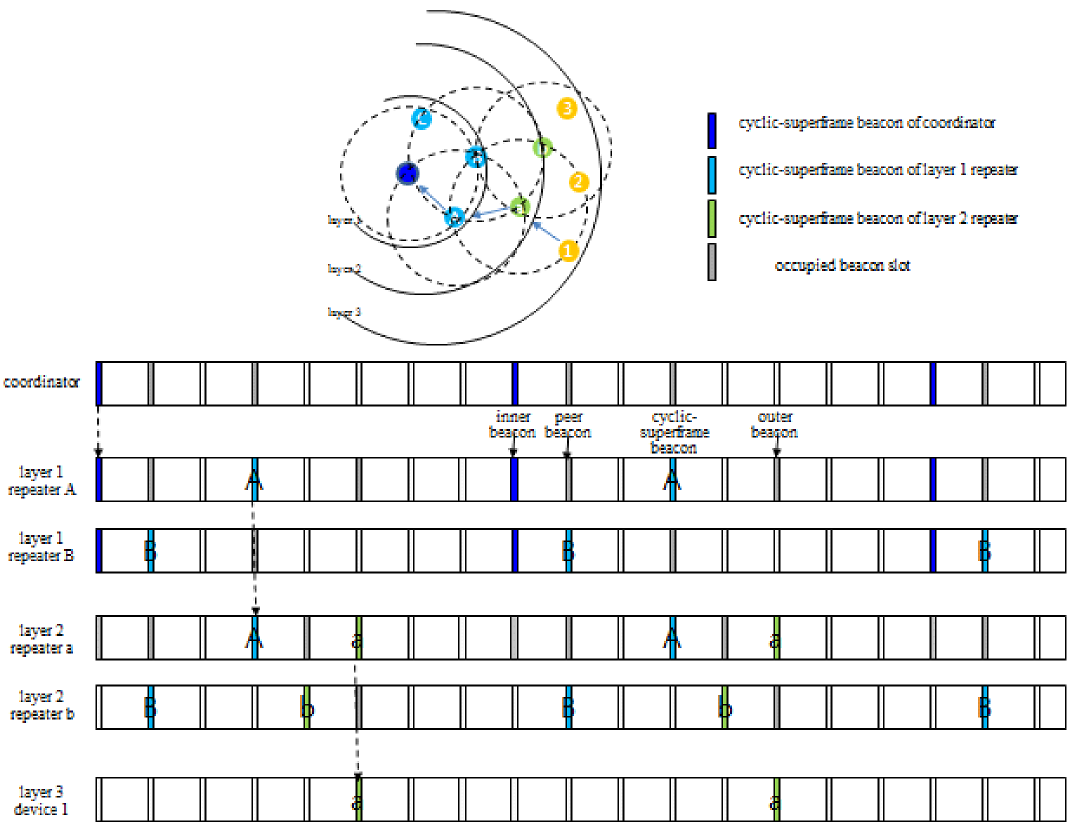

Figure 19a—Synchronous relaying of cyclic-superframe in an RSLN PAN

5.1.4 Synchronization

Probably some additional considerations for LECIM for both beacon and non beacon cases.

Insert the following new subclauses (5.1.4.4–5.1.4.4.1) after 5.1.4.3:

5.1.4.4 LECIM synchronization

Alternately we may just add a separate section, or fold in to beacon or non-beacon cases as appropriate.

5.1.4.4.1 RSLN synchronization

TBD

5.1.6 Transmission, reception, and acknowledgment

Expect all sub-clauses will have some changes related to MPDU fragmentation

5.1.6.1 Transmission

5.1.6.2 Reception and rejection

Expect some additional filtering for MPDU fragmentation will be required based on transaction ID, sequence # or something like that.

5.1.6.4 Use of acknowledgments and retransmissions

Insert the following new subclauses (5.1.6.4.5–5.1.6.4.6) after 5.1.6.4.5:

5.1.6.4.5 Incremental fragment acknowledgment

5.1.6.4.6 Incremental fragment retransmission

5.1.6.5 Promiscuous mode

5.1.6.6 Transmission scenarios

Insert the following new subclause (5.1.6.7) after 5.1.6.6:

5.1.6.7 Synchronous relaying

Each repeater relays the slot-link outward or inward. The selection of relayed slot-link depends on the direction of relaying and the type of the slot-link.

Only the beacon received in the beacon slot of an inner repeater or the PAN coordinator shall be relayed to the outward beacon slot of the inner repeater or the PAN coordinator. The relaying of the cyclic-superframe beacon slot-link of the PAN coordinator to the cyclic-superframe beacon slot-link of the repeater is synchronized by letting a cyclic-superframe beacon of the PAN coordinator relayed on the cyclic-superframe beacon slot of the repeater. The outward relaying beacon slot of the inner repeater is determined by the relative beacon slot distance between the cyclic-superframe beacon slot of the repeater and the cyclic-superframe beacon slot of the inner repeater.

Outward relaying of the coordinator slot is synchronized with the relaying of the cyclic-superframe beacon slot.

The frames received on the bidirectional device slot of the inner repeater are relayed to the bidirectional device slot of the next cyclic-superframe.

When relaying the beacon or command frames outward, the repeater updates the time synchronization specification and the synchronous relaying specification in the frames, if applicable.

Need an introductory sentence for this figure

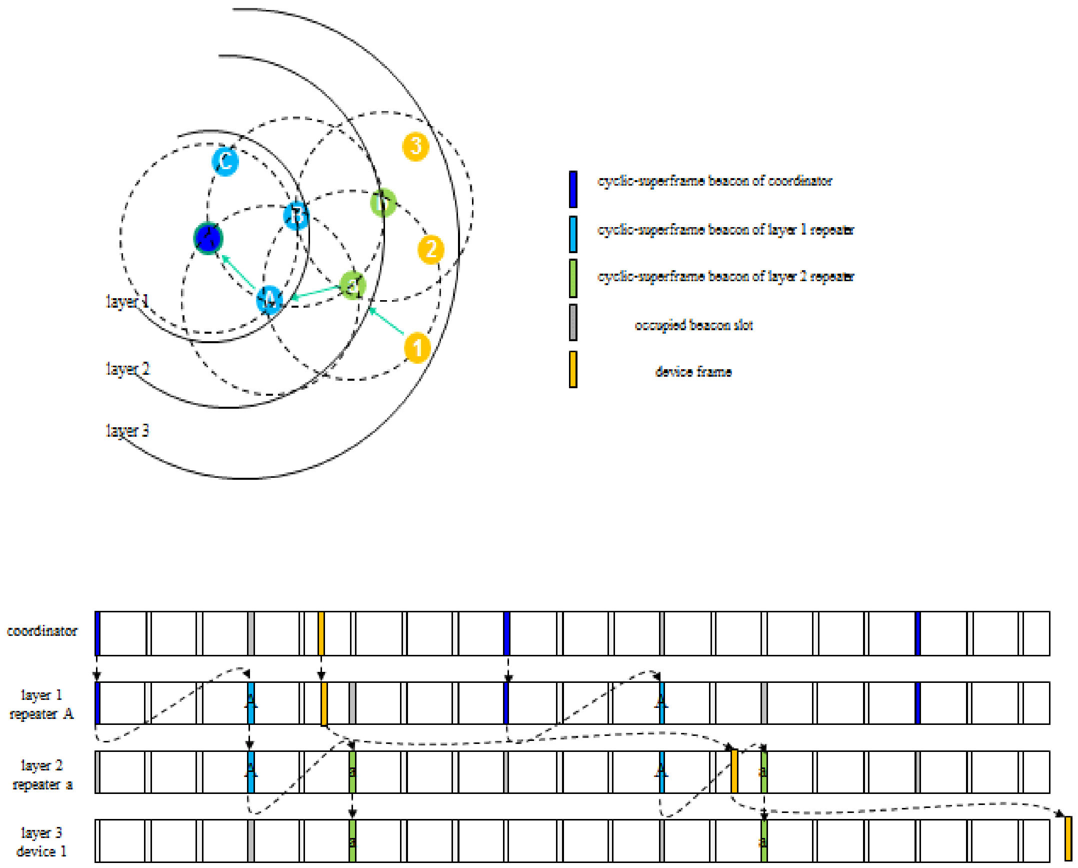

Figure 29a—Outward synchronous relaying in an RSLN PAN

The prioritized device slot of the outer repeater is relayed to the earliest available prioritized device slot of the upcoming slotted-superframe.

The frames received on the bidirectional device slot of the outer repeater are relayed to the bidirectional device slot of the next cyclic-superframe.

Need an introductory sentence for this figure

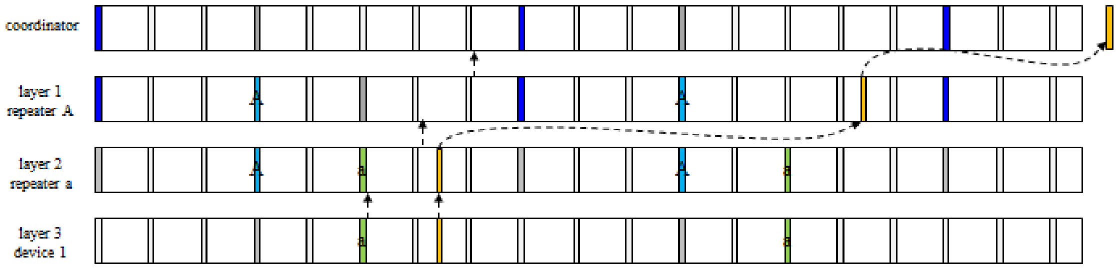

Figure 29b—Inward synchronous relaying in an RSLN PAN

5.1.7 GTS allocation and management

TBD if additions are needed.

5.1.11 LE-transmission, reception, and acknowledgment

Insert the following new subclauses (5.1.11.3–5.1.11.4.2) after 5.1.11.2.4:

5.1.11.3 LECIM alternate/hybrid LE scheme

5.1.11.3.1 General

The alternate/hybrid LE mode is active when *macLEenabled* is TRUE while CSL and RIT are disabled, as indicated by *macCSLPeriod* and *macRITPeriod* both being set to zero.

The basic LECIM hybrid LE mode is illustrated in Figure 34sa.

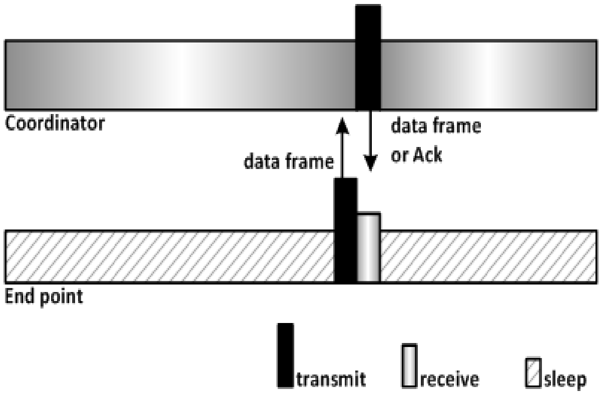

Figure 34sa—Basic LECIM LE mode operations

5.1.11.3.2 LECIM LE transmission

In LECIM networks, transmissions are mainly transmitted from an endpoint device to a coordinator. As described in Clause 4, the power of the coordinator is not as limited as that of the endpoint device when operating in LECIM LE mode. Therefore, the coordinator shall keep listening to the channel, except when it has a data frame to send or needs to send beacon frames when *macLowEnergySuperframeSupported* is TRUE.

An endpoint device shall keep sleeping for the normal time, unless it has a data frame to send. Then the endpoint device shall enable its transmitter and send the data frame.

When *macLowEnergySuperframeSupported* is TRUE, an endpoint device shall send data frames using either slotted ALOHA or slotted CSMA-CA. Otherwise, the endpoint device shall send data frames using either unslotted ALOHA or unslotted CSMA-CA.

If the coordinator has a data frame to send to an endpoint device, the coordinator shall wait until it receives a data frame from that endpoint device and then send its own data frame as an acknowledgment for the received data frame. If the coordinator has more than one data frame to send to the same endpoint device, it shall indicate the additional data frames by setting the Frame Pending field of the Frame Control field to one. If the coordinator does not have a data frame to send to the endpoint device, the coordinator shall send an acknowledgment frame in response to the received data frame.

After sending the data frame to the coordinator, the endpoint device shall wait for *macAckWaitDuration*. If an acknowledgment frame containing the same DSN as the original transmission is received within *macAckWaitDuration*, or a new data frame is received from the coordinator within *macAckWaitDuration*, the transmission is considered successful. Otherwise, the device shall conclude that the transmission has failed, and the device shall retransmit the data frame up to a maximum of *macMaxFrameRetries* times.

If the endpoint device received a data frame from the coordinator, it shall follow the acknowledgment procedure defined in this standard. The Frame Pending field of the Frame Control field in the received data frame shall determine whether the receiver is to be kept on or turned off following the reception of the data frame.

5.1.11.3.3 Hybrid wakeup sample listening (HWSL)

The hybrid wakeup sample listening (HWSL) mode guarantees timely transmission from a coordinator to an endpoint device(s). The HWSL mode shall be enabled when the PIB attribute *macHWSLEnabled* is set to

TRUE. If the value of the PIB attribute *macHWSLEnabled* is TRUE, the values of PIB attributes *macCSLPeriod* and *macRITPeriod* shall be ignored.

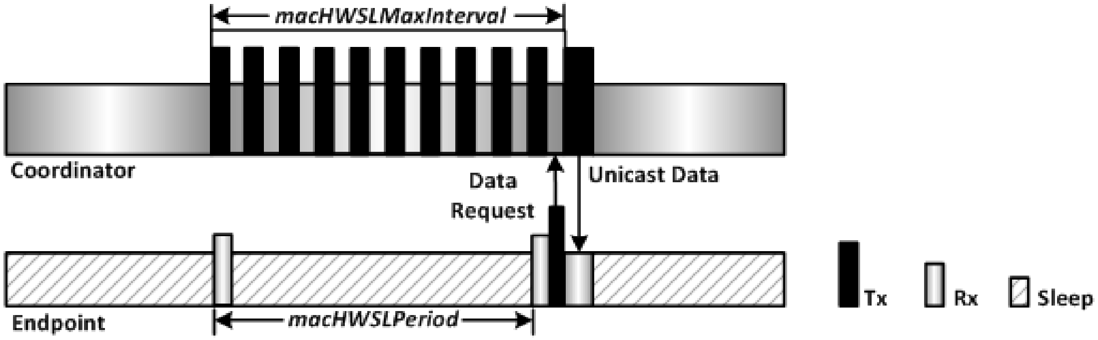

Figure 34sb—Unicast transmission in HWSL mode

As described in 5.1.11.3.2, for daily transmission from the coordinator to an endpoint device(s), the coordinator shall transmit the data the endpoint device until received data frames from the corresponding endpoint device. In some cases, the latency will be very long, HWSL mode is used for the emergency data frame from the coordinator to the endpoint device, and support broadcast data frame from the coordinator.

A coordinator operating in HWSL mode shall listen to the channel continuously. If the coordinator has an emergency data frame to send, the transmission of the payload frame shall be preceded with a sequence of HWSL wakeup frames.

The HWSL wakeup sequence consists of a sequence of HWSL wakeup frames, and the interval between two consecutive HWSL wakeup frames is defined by the PIB attribute *macHWSLWakeupInterval*. The coordinator shall listen to the channel in between wakeup frame transmissions. The maximum length of an HWSL wakeup sequence is *macHWSLMaxPeriod*.

An endpoint device performs a channel sample every *macHWSLPeriod* time. If the channel sample does not detect any HWSL wakeup frames from the coordinator, the endpoint device shall disable the receiver until the next channel sample time.

If the coordinator has a unicast frame to send, the destination address of the HWSL wakeup frame shall be set to the address of the corresponding endpoint device. On receipt of the unicast HWSL wakeup frame by the endpoint device through channel sampling, the endpoint device shall first check the destination address. If the destination address matches that of the endpoint device, the endpoint device shall request that the higher layer stop periodic channel sampling. The endpoint device shall send an HWSL data request frame to the coordinator and wait for a period of *macDataWaitDuration* for incoming unicast data frame.

If the coordinator received an HWSL data request frame from the corresponding endpoint device after sending an unicast HWSL wakeup frame, the coordinator shall stop sending the HWSL wakeup sequence and send the corresponding unicast data frame to the endpoint immediately. Following that, the coordinator shall wait for a period of *macAckWaitDuration* for the acknowledgment from the endpoint device.

On receipt of the incoming unicast data frame, the endpoint device shall send a corresponding acknowledgment to the coordinator.

If the next higher layer of the coordinator has multiple frames to transmit to the same endpoint device, the coordinator shall set the Frame Pending field of the Frame Control field to one in all but the last frame.

An HWSL unicast transmission is performed via the following steps by the MAC sublayer of the coordinator:

a) Perform CSMA-CA to acquire the channel

b) If the previously acknowledged unicast data frame had the Frame Pending field of the Frame Control field set to one and *macHWSLFramePendingWaitTime* has not been reached (defined in Table 52j), go to Step d.

c) For the duration of the wakeup sequence length, transmit the HWSL wakeup frames according to the interval *macHWSLWakeupInterval*.

d) If the coordinator has a pending unicast data frame to send, set the Frame Pending field of the Frame Control field to one, then transmit the unicast data frame.

e) Wait for up to *macAckWaitDuration* symbol time for the acknowledgment frame if the Acknowledgment Request field in the unicast data frame was set to one.

f) If the acknowledgment frame is received, go to Step g. Otherwise, start the retransmission process.

g) If the coordinator has pending unicast data to send, go to Step b. Otherwise, exit HWSL mode and keep listening to the channel.

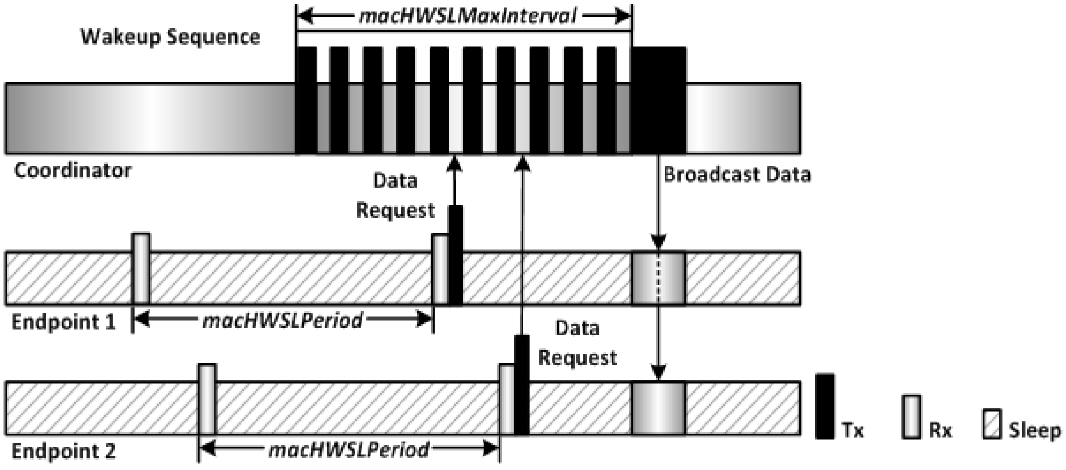

Figure 34sc—Broadcast transmission in HWSL mode

If the coordinator has a broadcast frame to send, the destination address of the HWSL wakeup frame shall be set to the broadcast address, and include the remaining time of the broadcast data frame transmission.

An endpoint device receiving the broadcast HWSL wakeup frame through channel sampling shall request that the higher layer stop the periodic channel sampling. The endpoint device shall then send an HWSL data request frame to the coordinator and return to sleep for the remaining portion of time indicated by the broadcast HWSL wakeup frame. The endpoint device shall then turn on its receiver and wait for the corresponding broadcast data frame.

If the coordinator received an HWSL data request frame from the corresponding endpoint device after sending a broadcast HWSL wakeup frame, the coordinator shall keep sending the HWSL wakeup sequence until it has received HWSL data request frames from all the endpoint devices or until *macHWSLMaxPeriod* has expired. The coordinator shall send the corresponding broadcast data frame in the designed time.

5.1.11.4 Implicit receiver initiated transmission (I-RIT)

5.1.11.4.1 General

The implicit receiver initiated transmission (I-RIT) is an alternative low energy MAC for nonbeacon-enabled PANs. I-RIT is designed to be used for end devices, such as sensors, that primarily transmit information to a coordinator but have no way of determining when they should make use of conventional RIT. Instead of transmitting an RIT data request, when an end device has I-RIT enabled, the device turns its receiver on for a known period of time, at a known interval after each transmission, so that the end device makes itself available to receive information from the coordinator. I-RIT mode is turned on when PIB attribute *macIRITPeriod* is non-zero and is turned off when *macIRITPeriod* is zero. The values of *macCSLPeriod* (in coordinated sample listening) and *macRITPeriod* shall be set to zero when the value of *macIRITPeriod* is non-zero. Transmission and reception in I-RIT mode is illustrated in Figure 34sd.

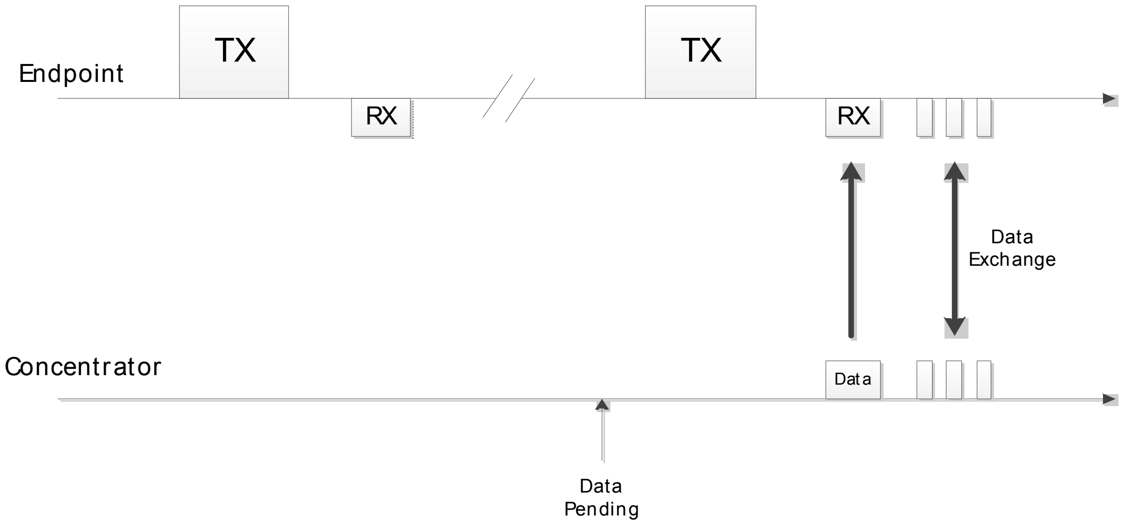

Figure 34sd—I-RIT transmission

5.1.11.4.2 I-RIT data request transmission and reception

In I-RIT mode, a device turns on its receiver *macIRITPeriod* symbol periods after the last bit of its transmitted frame for a period of *macIRITListenDuration* symbols in order to listen for an incoming frame. Then the device goes back to idle state until the next frame is transmitted.

5.1.12 Asynchronous multi-channel adaptation (AMCA)

5.2 MAC frame formats

Expect to add new Information Elements to 5.2.4, which will require adding to table 4b and new subclauses starting at 5.2.4.23.

5.2.1 General MAC frame format

5.2.1.1 Frame Control field

5.2.1.1.3 Frame Pending field

Change the third paragraph of 5.2.1.1.3 as indicated:

When operating in Low Energy (LE) CSL mode or HWSL mode, the frame pending bit may be set to one to indicate that the transmitting device has back-to-back frames to send to the same recipient and expects the recipient to keep the radio on until the frame pending bit is reset to zero.

5.2.2 Format of individual frame types

5.2.2.1 Beacon frame format

5.2.2.1.1a Information Elements (IEs) field

Change Table 3b (the entire table is not shown) as indicated:

Table 3b—EBR IEs per enabled attribute

Attribute request identifier	PIB attribute	IE type	IEs to include
3	*macLEenabled*	Header	LE CSL, or LE RIT, HWSL LE (5.2.4.7, 5.2.4.8, 5.2.4.8a)

5.2.4 Information element (IE)

5.2.4.2 Header information elements

Insert the following two rows at the end of Table 4b:

Table 4b—Element IDs, Header IEs

Element ID	Content length	Name	Description
TBD	4	HWSL LE	5.2.4.8a
TBD		MPDU Fragment Sequence Context Description	5.2.4.23

Editing Note: Element ID values are to be assigned by the 802.15 Numbering Authority (aka Gilb)

Insert the following new subclause (5.2.4.8a) after 5.2.4.8:

5.2.4.8a HWSL IE

The structure of the HWSL IE is illustrated in Figure 48ua.

Octets: 2	**2**
HWSL Phase	HWSL Remain Time

Figure 48ua—HWSL IE

The HWSL Phase field specifies the time remaining in the HWSL wakeup sequence. The range of the value of this field is 0x0000–0xffff, and the unit is 10 symbol durations.

The HWSL Remain Time specifies the remaining time of the incoming data frame. The range of the value of this field is 0x0000–0xffff, and the unit is 10 symbol durations.

Insert the following new subclauses (5.2.4.23–5.2.4.23.7) after 5.2.4.22:

5.2.4.23 MPDU Fragment Sequence Context Description IE

The MPDU Fragment Sequence Context IE contains a description of an MPDU being fragmented and associates this information with a unique fragmentation transaction ID. The transaction ID is transmitted with each fragment to identify it as part of the MPDU described by the IE. The IE is illustrated in Figure 48ub.

Octets: 2			**1**	**2**		**variable**	**variable**
Bits: 10	**5**	**1**		**10**	**6**		
Transaction ID	I-ACK Interval	Reserved	Fragment Size	PDSU Size	Addressing Information	Addressing field	PHY-dependent Parameters

Figure 48ub—MPDU Fragment Context Description IE

5.2.4.23.1 Transaction ID field

The Transaction ID field contains a value that is locally unique in the PAN and identifies the fragment sequence. It associates the context information with each fragment in the transaction. The specific method for generating the transaction ID is implementation dependent and should assure that the current value is different from the preceding value.

5.2.4.23.2 I-ACK Interval field

The I-ACK Interval field indicates the I-ACK policy to be employed. For values from one to the maximum number of fragments, an I-ACK is generated by the receiving device after it has detected a fragment cell with the fragment number greater than or equal to the [(fragment number of last I-ACK) + (I-ACK interval)].

5.2.4.23.3 Fragment Size field

TBD

5.2.4.23.4 MPDU Size field

TBD

5.2.4.23.5 Addressing Information field

The Addressing Information field describes the context of the addressing fields that follow. The fragment sequence description may contain any combination of source PAN ID, destination PAN ID, source address, and destination address in any of the allowable addressing modes defined by this standard. Figure 48uc illustrates the format of this field.

Bit: 0	**1**	**2–3**	**4–5**	**6**
Source PAN ID Present	Destination PAN ID Present	Source Address Mode	Destination Address Mode	Reserved

Figure 48uc—Addressing Information field format

The Source and Destination PAN ID Present fields shall be set respectively if a source and/or destination PAN ID is included in the Addressing field. The Source Address Mode field indicates the presence and format of a source address included in the Addressing field; the Source Address Mode field shall be set to one of the values given in Table 3. The Destination Address Mode field shall indicate the presence and format of a destination address included in the Addressing field, and the Destination Address Mode field shall be set to one of the values given in Table 3.

The setting of the Addressing Information field shall be determined by the PAN ID and addressing mode fields of the MPDU being fragmented.

5.2.4.23.6 Addressing field

The Addressing field contains source and/or destination addressing information associated with the MPDU being fragmented. The format is illustrated in Figure 48ud.

Octets: 0/16	**0/16**	**0/8/16/64**	**0/8/16/64**
Source PAN ID	Destination PAN ID	Source Address	Destination Address

Figure 48ud—Addressing field

The content of this field shall be set according to the addresses contained in the MHR of the MPDU being fragmented.

5.2.4.23.7 PHY-dependent Parameters field

The value of the PHY-dependent Parameters field depends upon the PHY being used. The possible values are implementation dependent. Table 4j shows the format for the LECIM DSSS PHY defined in 19.1; Table 4k shows the format for the LECIM FSK PHY defined in 19.2.

Table 4j—PHY dependent fragment context parameters for LECIM DSSS PHY

Parameter	Bit position	Valid range	Parameter description
TBD	TBD	TBD	TBD

Table 4k—PHY dependent fragment context parameters for LECIM FSK PHY

Parameter	Bit position	Valid range	Parameter description
Slot # /channel #	TBD	S0/C3:S8/C1:S16/C7:S24/C2:S32/C5	Where network sets slot durations at 50 ms and channel page to 9
Sync info	TBD	TBD	TBD
Time out period	TBD	0–0x3c	>= 60 slots or 3 seconds (0xff = no period defined)

5.3 MAC command frames

Change Table 5 (the entire table is not shown) as indicated:

Table 5—MAC command frames

Command frame identifier	Command name	RFD		Subclause
		Tx	Rx	
TBD	HWSL wakeup			5.3.12.2
TBD	HWSL data request	X		5.3.12.3
TBD	Channel switching notification			5.3.14
TBD–0xff ~~0x21–0x3f~~	Reserved			—
~~0x44–0x5f~~	~~Reserved~~			—
~~0x61–0x62~~	~~Reserved~~			—
~~0x64–0xff~~	~~Reserved~~			—

5.3.12 LE commands

Insert the following new subclauses (5.3.12.2, 5.3.12.3) after 5.3.12.1:

5.3.12.2 HWSL wakeup command

TBD

5.3.12.3 HWSL data request command

TBD

Insert the following new subclauses (5.3.14–5.4.2.3) after 5.3.13.3.2:

5.3.14 Channel switching notification command

The channel switching notification command is sent to a device to request that the device switch channels in order to receive a fragmented frame on the new channel; both the originating and recipient devices switch to the new channel prior to transmission of the fragmented frame. All devices that support fragmentation shall be capable of transmitting this command. The channel switching notification command shall be formatted as illustrated in Figure 59de.

Octets: variable	**1**	**1**	**1**
MHR fields	Command Frame Identifier	Sequence ID	Channel Offset

Figure 59de—Channel switching notification command format

Add text explaining how all the fields in the command frame are set.

5.4 MPDU fragmentation

When MPDU fragmentation is enabled, the completed MPDU is processed into a sequence of fragment cells. The context of the fragment sequence is established between the initiating device and the recipient device prior to transmission. Each fragment containing a fragment check sequence, fragment descriptor, and fragment content is packaged into a PPDU. Certain MHR fields may be transformed or elided in order to reduce the size of the fragment.

5.4.1 MPDU PHY adaptation, fragmentation and reassembly

5.4.1.1 Fragment sequence context

The fragment sequence context is established by transmitting a fragment context frame containing an MPDU Fragment Sequence Context Description IE, as described in 5.2.4.23. A fragment context frame is any directed MAC command or data frame which contains an MPDU Fragment Sequence Context Description IE, and a frame shall contain exactly one such IE.

The fragment context frame initiates the transaction and establishes the initial state for the MPDU sequence transaction. The fragment context frame shall be transmitted with the Acknowledge Request field set to one. If an acknowledgment is not received, the fragment context frame shall be retransmitted up to *macMaxFrameRetries* times as needed. If an acknowledgment is received, the initiating device transmits the fragments until either the transaction is complete or the transmission is aborted.

Upon reception of the fragment context frame, the information contained within the frame is associated with the value of the Transaction ID field in the MPDU Fragment Context Description IE, and that ID value is used to identify subsequent fragments in the sequence. If a fragment cell is not received within *aMPDUFragTimeout*, the fragmentation transaction shall be terminated.

When the LECIM DSSS PHY is in use, a unique spreading code or codes may be used between the coordinator and end point, in which case the context of the transaction is established uniquely by code separation and the Transaction ID field may be elided from each fragment cell. The Transaction ID field value of zero shall be used for this purpose. A Transaction ID field with a value of zero indicates that fragment cells do not contain a transaction ID and shall be used only with PHYs that support other means to establish point-to-point unique context.

5.4.1.2 Fragment cell formats

The fragment cell is depicted in Figure 59df.

Bits: 0/10	**5**	**1**	**variable**	**16/24/32**
Transaction ID	Fragment Number	Extension	Fragment Data	Fragment Validation

Figure 59df—Fragment cell general form

The Transaction ID (TID) field, when present, shall contain the value assigned to the transaction context, as indicated in the fragment context frame. When context is unambiguously known via other means provided by the PHY in use, the TID field may be suppressed. Upon reception, if the TID field contains a value other than the TID of a currently active transaction, the cell is ignored (i.e., not acknowledged and not counted to reset the transaction timeout).

The Fragment Number field identifies which fragment in the sequence the data part contains. A Fragment Number field value of zero shall be used to indicate a terminated transaction. Upon reception of a cell with the TID field equal to zero, the receiving devices will invalidate the transaction context; if subsequent cells are received with the same TID field value prior to a new fragment context frame, they may be ignored. Upon MPDU reassembly, the fragmented data shall be placed in order according to fragment number.

The Extension field is used to indicate an extended cell descriptor and is reserved for future versions of this standard.

The Fragment Data field contains the part of the fragmented MPDU indicated by the Fragment Number field. The size of the data field depends on the configuration of the PHY in use. For the LECIM DSSS PHY, the data field may be 15 to 23 octets in length. For the LECIM FSK PHY, the data field may be from 19 to TBD octets in length.

The Fragment Validation field is used to validate the received fragment cell. It shall be calculated as defined for the TBD length CRC according to 5.2.1.9.

5.4.1.3 Fragmentation

The MPDU is prepared for fragment transmission according to the following steps:

a) Determine the fragment context using the MHR fields (i.e., source addressing, destination addressing, and data request parameters).
b) Construct the fragment context frame, as described in 5.4.1.1.
c) Elide/compress the MHR fields that are effectively transmitted in the fragment context frame.

d) Divide the remaining MPDU into fragment cells of the size supported by the current PHY configuration. All fragments, with the exception of the final fragment, contain the maximum number of data octets. For PHY configurations that use a fixed PPDU size (i.e., no PPDU length field transmitted), the final fragment data is padded with *macMPDUFragPadValue,* which may be a PHY-dependent value. The Fragment Validation field for the final fragment is calculated including the pad octets.

e) Transmit the fragment context frame (retransmit as necessary).

f) Upon acknowledgment of the fragment context frame, transmit the fragment cells. After I-ACK interval fragment cells have been transmitted, wait for the I-ACK. Retransmit the cell preceding the I-ACK if the acknowledge is not received with the I-ACK timeout.

g) Upon transmission of the final fragment cell and/or reception of the final I-ACK as appropriate, the MPDU level acknowledgment is performed as described in 5.1.6.

Fragments are transmitted in the order shown in Figure 59dg. The I-ACK is described in 5.4.2.1. If the I-ACK retransmission count is exceeded during the transaction, the transaction is terminated and a fragment cell with the Fragment Number field set to zero is transmitted to signal the receiving device.

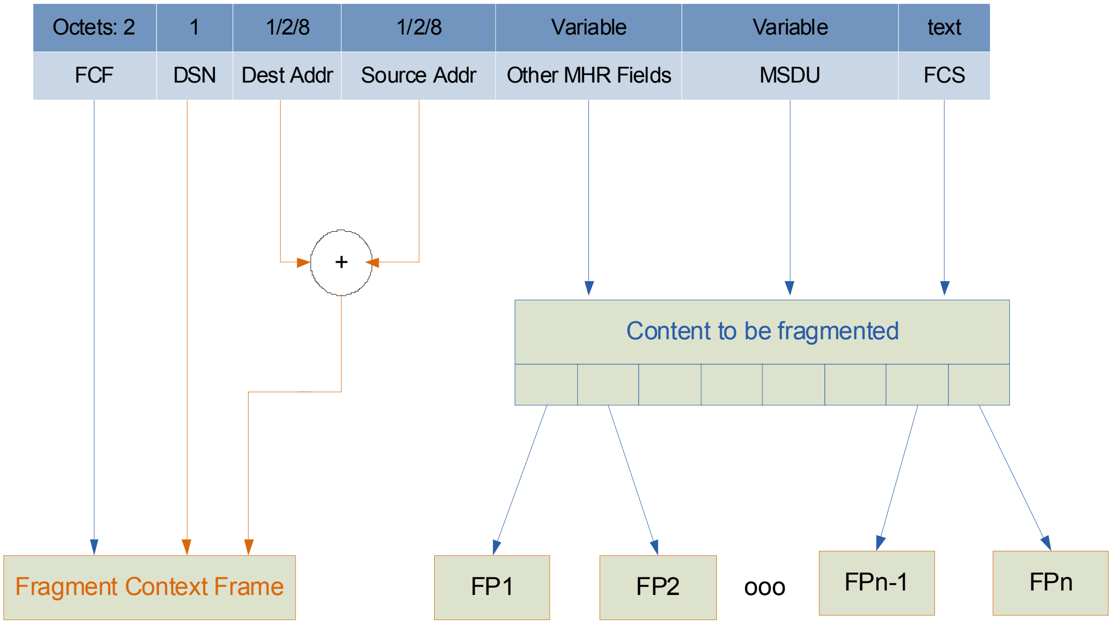

Figure 59dg—Fragmentation process overview

5.4.1.4 Reassembly

Upon reception of the fragment context frame, the transaction state is initialized for a new MPDU fragment sequence transaction, and the fragment context frame is acknowledged. Each received fragment cell is placed into the reassembled MPDU based on the value of the corresponding Fragment Number field. I-ACKs are generated according to 5.4.2.1. When the final fragment is received and validated, MPDU validation proceeds according to 5.1.6.

5.4.2 Fragment acknowledgment and retransmission

Two levels of fragment acknowledgment are provided: acknowledgment of fragments during the transfer process (i.e., incremental acknowledgment), which provide “progress reports”; and acknowledgment of the reassembled MPDU. In each acknowledgment level, the status of individual fragments is indicated and the initiating device can retransmit only those fragments that were not received and validated.

5.4.2.1 Incremental fragment acknowledgment (I-ACK)

The I-ACK reports status indicating which fragments have been successfully received up to that point, and it is generated incrementally during the fragment sequence transfer.

5.4.2.1.1 I-ACK format

The I-ACK includes the Fragment Status field, constructed as shown in Figure 59dh.

Bits: 4	1/2/3/4	16
IACK Content	Fragment Status Flags (Set 0–Set 3)	Validation

Figure 59dh—I-ACK Fragment Status field

The IACK Content field is shown in Table 7c. This field indicates which fragment status flags are included. A value of one in a bit position indicates that the corresponding set of eight status flags is present; a value of zero in a bit position indicates that the corresponding set of eight status flags is absent. Setting all bit positions to zero indicates an aborted transaction. Bit b_0 is transmitted first in time.

Table 7c—IACK Content field

Bit position	Description
b_0	Indicates whether fragment status flags 0–7 are present
b_1	Indicates whether fragment status flags 8–15 are present
b_2	Indicates whether fragment status flags 16–23 are present
b_3	Indicates whether fragment status flags 24–31 are present

The Fragment Status Flags field indicates the status of received fragments up to the current point in the transaction. The status flags are grouped into four sets of eight 1-bit flags. Flags for fragment numbers 0–7 are contained in Set 0, flags for fragment numbers 8–15 are contained in Set 1, flags for fragment numbers 16–23 are contained in Set 2, and flags for fragment numbers 23–31 are contained in Set 3. Within each set, the individual flags are ordered such that s_0, the first bit transmitted/received in time, corresponds to the lowest numbered fragment number in the set. When more than one set is included in the I-ACK, the lowest numbered set is transmitted first in time, so that the correspond fragment numbers go from low to high as transmitted.

The Validation field is used to validate the received I-ACK. It shall be calculated as defined for the TBD length CRC according to 5.2.1.9.

5.4.2.1.2 I-ACK overview

The interval of the I-ACK is determined by the IACKinterval parameter of the MCPS-DATA.request primitive, which is transmitted to the receiving device with the fragmentation sequence set-up message (xref). Upon completion of transmission of each IACKinterval fragment cell, the initiating device will suspend transfer and wait *macIACKtimeout* for the expected I-ACK. Upon reception of the I-ACK, fragments indicated as not received correctly shall be retransmitted. The number of retransmissions is limited by *macMaxFrameRetries*. If an I-ACK has not been received following *macIACKtimeout*, the initiator will retransmit the last fragment sent and wait for the I-ACK again, repeating this process up to *macMaxFrameRetries* times.

Upon receipt of the I-ACK, the initiator of the fragment sequence will examine the Fragment Status field and shall retransmit the fragments that are not indicated as successfully received (i.e., retry fragments) following the I-ACK. The fragments to be retransmitted shall be transmitted in the order of initial transmission, followed by the next k fragments in sequence, where $k = (\text{IACKInterval} - \text{ number of retry fragments})$.

5.4.2.2 Aggregated MPDU transfer acknowledgment

If the received MPDU has its Acknowledgment Request field set to one in the MHR, the generated acknowledgment (using the enhanced acknowledgment) will include a Fragment Status IE, constructed and transmitted as described here.

If *macFragmentSequExtAck* is FALSE, an MPDU acknowledgment is generated once the reassembly of the MPDU is completed and address filtering, if enabled, is completed. The Fragment Status IE is populated with the status of each fragment in the sequence and the final FCS.

If *macFragmentSequExtAck* is TRUE, the MPDU higher layer may become involved in the acknowledgment processing. The recipient device will, upon receiving the final fragment, generate an acknowledgment to the originator with the Frame Pending field set and also generate an MCPS-DATA indication containing the reassembled MPDU with the fragment sequence status information, as described in 6.3.3. Upon completion of higher layer processing, which is out of scope of this standard, the higher layer may use the MCPS-EXT-ACK.request, which initiates generation of the MPDU acknowledgment frame containing the status and feedback information provided with the service parameters.

5.4.2.3 Channel switching for fragment sequence exchange

Given the potentially long duration of the MPDU transaction in time, there is a possibility that channel conditions may change significantly. The higher layer may decide that the channel is becoming unusable and desire to change to another channel for subsequent transactions. The channel switch notification process provides this capability.

The channel switch notification (CSN) command is a directed command frame that facilitates changing channel or PHY parameters between the sending and receiving nodes. The CSN command is sent by the recipient of a fragment context frame to the originator. To initiate the switch channel, the information of the new channel will be included in the CSN command. The switch is initiated by the higher layer via the MLME-CHAN-NOTIFY service. When initiated, the CSN command will be transmitted following the aggregated MPDU acknowledgment. The AR field of CSN command is set to one. The originator shall acknowledge reception of the CSN command, and the originator and recipient shall switch to the new channel indicated in CSN command. [add MSCs and references]

The CSN command will affect only the device sending it and the device receiving it. The CSN command shall not be transmitted with the broadcast PANID and/or broadcast destination address. The higher layer network management entity controls which channel and/or PHY configurations are used to communicate with which neighbors; the process by which this is done is outside the scope of this standard.

In the event that a CSN command is not acknowledged, the channel switch shall not be performed. In the event that communication is not re-established after either a channel switch or aborted channel switch, the device shall revert to the prior channel and PHY configuration after the *macCSNeffectTimeout*. The originator should perform handshake with the recipient prior to transmission by sending the first fragment to the recipient and receiving acknowledgment. This confirmation process is outside the scope of this standard.

6. MAC services

6.2 MAC management service

Insert the following new row at the end of Table 8, and insert the following new subclauses (6.2.22–6.2.22.3) after 6.2.21.3.4:

Table 8—Summary of the primitives accessed through the MLME-SAP

Name	Request	Indication	Response	Confirm
MLME-CHAN-NOTIFY	6.2.22.1	6.2.22.3		6.2.22.2

6.2.22 Primitives for support of fragmentation

These primitives are used to manage fragment transmissions between two devices.

6.2.22.1 MLME-CHAN-NOTIFY.request

The MLME-CHAN-NOTIFY.request primitive is used by one device to notify a second device to switch channels in order to receive a fragmented frame from the first device on the new channel.

The semantics of this primitive are:

```
MLME-CHAN-NOTIFY.request        (
                                DstAddrMode,
                                DstAddr,
                                ChannelOffset,
                                SequenceID,
                                TransactionID
                                )
```

The primitive parameters are defined in Table 44aa.

On receipt of the MLME-CHAN-NOTIFY.request primitive, the channel switching process is initiated, as defined in [xref].

6.2.22.2 MLME-CHAN-NOTIFY.confirm

The MLME-CHAN-NOTIFY.confirm primitive is used to inform the next higher layer of the initiating device whether the channel switching notification command was transmitted successfully.

The semantics of this primitive are:

```
MLME-CHAN-NOTIFY.confirm        (
                                status
                                )
```

Table 44aa—MLME-CHAN-NOTIFY.request parameters

Name	Type	Valid range	Description
DstAddrMode	Enumeration	NO_ADDRESS, SHORT_ADDRESS, EXTENDED_ADDRESS	The destination addressing mode for this primitive. This parameter is only used when a fragment transaction is not in progress.
DstAddr	Device address	As specified by the DstAddrMode parameter	The individual device address of the device for which the frame was intended. This parameter is only used when a fragment transaction is not in progress.
ChannelOffset	Integer	0x00–0xff	The offset value of the channel sequence of the PAN.
SequenceID	Integer	0x00–0xff	The ID of the fragment sequence.
TransactionID	Integer	0x0000–0xffff	The transaction ID of the fragment sequence. A value of zero indicates that the parameter will not be included in the fragment cells. This parameter is only used when a fragment transaction is in progress.

The primitive parameters are defined in Table 44ab.

Table 44ab—MLME-CHAN-NOTIFY.confirm parameters

Name	Type	Valid range	Description
status	Enumeration	SUCCESS, CHANNEL_ACCESS_FAILURE, NO_ACK, UNSUPPORTED_LEGACY, INVALID_PARAMETER	The status of the attempt to transmit the channel switching notification command.

6.2.22.3 MLME-CHAN-NOTIFY.indication

The MLME-CHAN-NOTIFY.indication primitive is used to indicate the reception of a channel switching notification command.

The semantics of this primitive are:

```
MLME-CHAN-NOTIFY.indication     (
                                ChannelOffset,
                                SequenceID,
                                TransactionID
                                )
```

The primitive parameters are defined in Table 44ac.

Table 44ac—MLME-CHAN-NOTIFY.indication parameters

Name	Type	Valid range	Description
ChannelOffset	Integer	0x00–0xff	The offset value of the channel sequence of the PAN.
SequenceID	Integer	0x00–0xff	The ID of the fragment sequence.
TransactionID	Integer	0x0000–0xffff	The transaction ID of the received fragment sequence. A value of zero indicates that the parameter will not be included in the fragment cells. This parameter is only used when a fragment transaction is in progress.

6.3 MAC data service

6.3.1 MCPS-DATA.request

Insert the following new parameter at the end of the list in 6.3.1 (before the closing parenthesis):

IACKspan

Insert the following new rows at the end of Table 46:

Table 46—MCPS-DATA.request parameters

Name	Type	Valid range	Description
IACKspan	Integer	0x0000–TBD	If the value of the parameter is non-zero, specifies the number of fragments to send prior to waiting for an I-ACK. If the value of the parameter is zero, no I-ACK is requested and only an MPDU level acknowledgment is requested. This parameter is only valid when fragmentation of the MPDU is enabled.

Insert the following paragraph at the end of 6.3.2:

When fragmentation of the MPDU is enabled and the IACKspan parameter is set to a non-zero value, the I-ACK feature is enabled, as described in 5.4.2.1.

6.4 MAC constants and PIB attributes

6.4.1 MAC constants

Insert the following new row at the end of Table 51:

6.4.2 MAC PIB attributes

The first paragraph of 6.4.2 is reproduced here to assist the reader in understanding the notation used in Table 52. No changes are made to this paragraph.

Table 51—MAC sublayer constants

Constant	Description	Value
aMPDUFragTimeout	TBD	TBD

The MAC PIB comprises the attributes required to manage the MAC sublayer of a device. The attributes contained in the MAC PIB are presented in Table 52. Attributes marked with a dagger (†) are read-only attributes (i.e., attribute can only be set by the MAC sublayer), which can be read by the next higher layer using the MLME-GET.request primitive. All other attributes can be read or written by the next higher layer using the MLME-GET.request or MLME-SET.request primitives, respectively. Attributes marked with a diamond (◆) are optional for an RFD; attributes marked with an asterisk (*) are optional for both device types (i.e., RFD and FFD).

Change Table 52 (the entire table is not shown) as indicated. The description of macMaxFrameRetries is reproduced here to assist the reader. No change is made to this description.

Table 52—MAC PIB attributes

Attribute	Type	Range	Description	Default
macMaxFrameRetries	Integer	0–7	The maximum number of retries allowed after a transmission failure.	3
macLECIMAlohaBackoffSlot	TBD	TBD	Backoff period when priority access backoff mechanism is in use, as defined in 5.1.1.4.6.	TBD
macMPDUFragPadValue†	TBD	TBD	The value used to pad out the last fragment when MPDU fragmentation is enabled. See [TBD] for PHY specific values.	Dependent on currently selected PHY
macFragmentSequExtAck	Boolean	TBD	Controls the behavior of the aggregated MPDU transfer acknowledgment described in 5.4.2.2.	TBD
macCSNeffectTimeout	Integer	TBD	Timeout for the completion of the channel switching notification and handshake.	TBD
macIACKtimeout	Integer	TBD	The amount of time, in PHY symbol periods, to wait for an I-ACK after the transmission of the fragment cell for which the acknowledgement is expected.	Dependent on currently selected PHY

6.4.3 Calculating PHY dependent MAC PIB values

6.4.3.7 LE-specific MAC PIB attributes

Insert the following new rows at the end of Table 52j:

Table 52j—MAC PIB attributes

Attribute	Type	Range	Description	Default
macHWSLEnabled	Boolean	TRUE, FALSE	A value of TRUE indicates that HWSL mode is enabled. A value of FALSE indicates that it is disabled.	FALSE
macHWSLPeriod	Integer	0–65535	The HWSL sampled listening period measured in units of 10 symbols.	0
macHWSLMaxPeriod	Integer	0–65535	Maximum length of HWSL wakeup sequence measured in units of 10 symbols.	*macHWSL Period*
macHWSLFramePendingWaitTime	Integer	(*macMinLIFS Period* + maximum number of symbols per PPDU) – 65535	Specifies the length of time, in symbols, to keep the receiver on after receiving a data frame with the Frame Pending field of the Frame Control field set to one.	TBD
macHWSLWakeupInterval	TBD	TBD	Specifies the interval between two successive HWSL wakeup frames in an HWSL wakeup sequence.	TBD
macIRITPeriod	Integer	0x0000–0xffff	A value of zero indicates that I-RIT is disabled. A non-zero value specifies the interval, in symbol periods, from the end of the transmitted frame to the beginning of the I-RIT listening period.	0x00
macIRITListenDuration	Integer	0x00–0xff	The duration of listening time, in symbol periods, for which the receiver is listening for the beginning of a frame to receive.	0x64
macIRITEnabled	TBD	TBD	TBD	TBD

8. General PHY requirements

8.1 General requirements and definitions

Insert the following items at the end of the second list in 8.1:

LECIM DSSS PHY: TBD

LECIM FSK PHY: a multi-regional, frequency shift keying (FSK) PHY operating at over-the-air data rates in support of low energy, critical infrastructure monitoring (LECIM) applications.

8.1.1 Operating frequency range

Insert the following new rows at the end of table 66:

Table 66—Frequency bands and data rates

PHY (MHz)	Frequency band (MHz)	Spreading parameters		Data parameters		
		Chip rate (kchip/s)	Modulation	Bit rate (kb/s)	Symbol rate (ksymbol/s)	Symbols
470	470–510	—	GFSK/FSK	37.5	37.5	Binary
		—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary
780	779–787	—	GFSK/FSK	37.5	37.5	Binary
		—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary
863	863–870	—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary
915	902–928	—	GFSK/FSK	37.5	37.5	Binary
		—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary
917	917–923.5	—	GFSK/FSK	37.5	37.5	Binary
		—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary
920	920–928	—	GFSK/FSK	37.5	37.5	Binary
		—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary
2450	2400–2483.5	—	GFSK/FSK	37.5	37.5	Binary
		—	GFSK/FSK	25	25	Binary
		—	FSK	12.5*	12.5*	Binary

[*] Coordinator to end device only.

8.1.2 Channel assignments

<REVISIT> LECIM channel assignments match those used for the SUN PHY MR-FSK mode channel assignments. (Can this reference section 16? or do we want a separate channel page for LECIM?)

Table 1—Total number of channels and first channel center frequencies for LECIM FSK PHYs

Frequency band (MHz)	**Modulation (uplink/ downlink)**	***ChanSpacing* (MHz)**	***TotalNumChan***	***ChanCenterFreq*$_0$ (MHz)**
470–510	GFSK/FSK	0.2	199	470.2
779–787			39	779.2
863–870		0.1	69	863.075
902–928		0.2	129	902.2
917–923.5			32	917.1
920.5–923.5			15	920.6
2400–2483.5			416	2400.2

9. PHY services

9.2 PHY constants

9.3 PHY PIB attributes

Insert the following new rows at the end of Table 71:

Table 71—PHY PIB attributes

Attribute	Type	Range	Description
phyLECIMFSKPreambleLength	Integer	0–100	The number of 1-octet patterns, as defined in 19.2.1.1, in the preamble. This attribute is only valid for the LECIM FSK PHY.
phyLECIMFSKPSDUMod	Boolean	TRUE or FALSE	Indication of the type of modulation used. A value of TRUE indicates that P-GFSK/P-FSK is enabled for the PSDU. A value of FALSE indicates that GFSK/FSK modulation is enabled for the PSDU.
phyLECIMFSKSpreading	Boolean	TRUE or FALSE	A value of TRUE indicates that spreading is enabled. A value of FALSE indicates that spreading is disabled.
phyLECIMFSKSpreadingFactor	Integer	1, 2, 4, 8, 16	The spreading factor (SF) to be used when *phyLECIMFSKSpreading* is TRUE.
phyLECIMFSKScramblePSDU	Boolean	TRUE or FALSE	A value of FALSE indicates that data whitening of the PSDU is disabled. A value of TRUE indicates that data whitening of the PSDU is enabled. This attribute is only valid for the LECIM FSK PHY.
phyLECIMFECEnabled	Boolean	TRUE or FALSE	A value of TRUE indicates that FEC is turned on. A value of FALSE indicates that FEC is turned off. This attribute is only valid for the LECIM FSK PHY.
phyLECIMFSKInterleavingEnabled	Boolean	TRUE or FALSE	A value of TRUE indicates that interleaving is turned on. A value of FALSE indicates that interleaving is turned off. This attribute is only valid for the LECIM FSK PHY.

Insert after Clause 18 the following new clause (Clause 19):

19. LECSIM PHYs

Two PHYs are specified in order to support LECIM applications: direct sequence spread spectrum (DSSS; see 19.1) and frequency shift keying (FSK; see 19.2).

19.1 DSSS PHY specification

The direct sequence spread spectrum (DSSS) PHY is described in the following subclauses.

19.1.1 PPDU format for DSSS

For convenience, the PPDU structure is presented so that the leftmost field as written in this standard shall be transmitted or received first. All multiple octet fields shall be transmitted or received least significant octet first, and each octet shall be transmitted or received least significant bit (LSB) first.

The PPDU shall be formatted as illustrated in Figure 59.

Octets			
0/2/4	**0/1**	**0/1/2**	**Fixed/1–128/129–2048**
Preamble	SFD	See Figure 60	PSDU
SHR		PHR	PHY payload

Figure 59—Format of the LECIM DSSS PPDU

The relationship among the PPDU fields is given in Table 72.

Table 72—Relationship among the LECIM DSSS PPDU fields

Configuration	Preamble length (octets)	SFD length (octets)	PHR length (octets)	PSDU length (octets)
1	0	0	0	Fixed length
2	2/4	1	1	1–128
3	2/4	1	2	129–2048

Figure 60 shows the configuration of the PHR as a function of PHR length.

19.1.1.1 SHR

The synchronization header (SHR), if present, is used for obtaining frequency, symbol, and frame synchronization. It consists of the preamble and the start-of-frame delimiter (SFD). It is possible to recover a fixed length frame without the use of an SFD or SHR.

PHR length (octets)	**PHR contents**		
0	—		
1	0	Frame length (7 bits)	
2	1	Reserved (4 bits)	Frame length (11 bits)

Figure 60—PHR configuration for LECIM DSSS PHY

19.1.1.1.1 Preamble field

The Preamble field, if present, is used to obtain symbol timing and frequency offset. A preamble length of 0, 2, or 4 octets may be commissioned.

Preamble16 = [1 1 0 0 0 0 0 1 0 1 0 0 1 1 0]

Preamble32 = [T.B.D.]

19.1.1.1.2 SFD field

The SFD field, if present, indicates the beginning of the frame.

SFD = [T.B.D.]

19.1.1.2 PHR

The PHY header (PHR) is used to indicate the length of a variable length PHY payload. When the PHY payload is commissioned to a fixed size, the PHR is elided. For variable length PHY payloads of up to 128 octets, the PHR is one octet and represents a payload of $n + 1$ octets where $n = 0 \ldots 127$. For variable length PHY payloads of 129–2048 octets, the PHR is two octets. The bit definitions of the one and two octet PHRs are illustrated in Figure 60.

19.1.2 Modulation and spreading

19.1.2.1 Data rate

The data rate is band and/or region specific. Table 73 gives the frequency bands and data rates for the DSSS PHY.

The channelization for the 868 MHz band is as follows:

- Channel 0: 868.300 MHz
- Channel 1: 868.950 MHz
- Channel 2: 869.525 MHz

The channel numbering and spacing for the 902 MHz band are as follows:

$$904 + [(n - 1) \times 1.99\ \text{MHz}]$$

where $n = 1 \ldots 15$.

Table 73—Frequency bands and data rates for LECIM DSSS PHY

PHY (MHz)	Frequency band (MHz)	Region/availability	Chip rate (kchip/s)	Modulation
400	400–470	South Korea	100 (12.5 kHz channel bonded)	BPSK/?
470	470–510	China		BPSK/O-QPSK
				BPSK/O-QPSK
				BPSK/O-QPSK
780	779–787	China		BPSK/O-QPSK
				BPSK/O-QPSK
				BPSK/O-QPSK
868	863–870	EU/CEPT	100	BPSK/O-QPSK
902	902–928	Americas, Austrailia	1000/?	BPSK/O-QPSK
917	917–923.5	South Korea		BPSK/O-QPSK
				BPSK/O-QPSK
				BPSK/O-QPSK
920	920–928	Japan	200	BPSK/O-QPSK
			600	BPSK/O-QPSK
			1000	BPSK/O-QPSK
2450	2400–2483.5	Worldwide	1000, 2000?	BPSK/O-QPSK

The channel numbering and spacing for the 2400 MHz band are as follows:

$$2402 + [(n - 1) \times 1.99 \text{ MHz}]$$

where $n = 1 \ldots 41$.

The 1.99 MHz spacing is used to minimize false lock and interference from spurious.

19.1.2.2 Reference modulator diagram

The functional block diagram in Figure 61 is provided as a reference for specifying the LECIM DSSS PHY modulation. All binary data contained in the SHR, PHR, and PSDU shall be encoded using the modulation shown in Figure 61.

19.1.2.3 Convolutional forward error correction (FEC) encoding

The convolutional encoder is the same as specified in the IEEE Std 802.11^{TM}-2007.

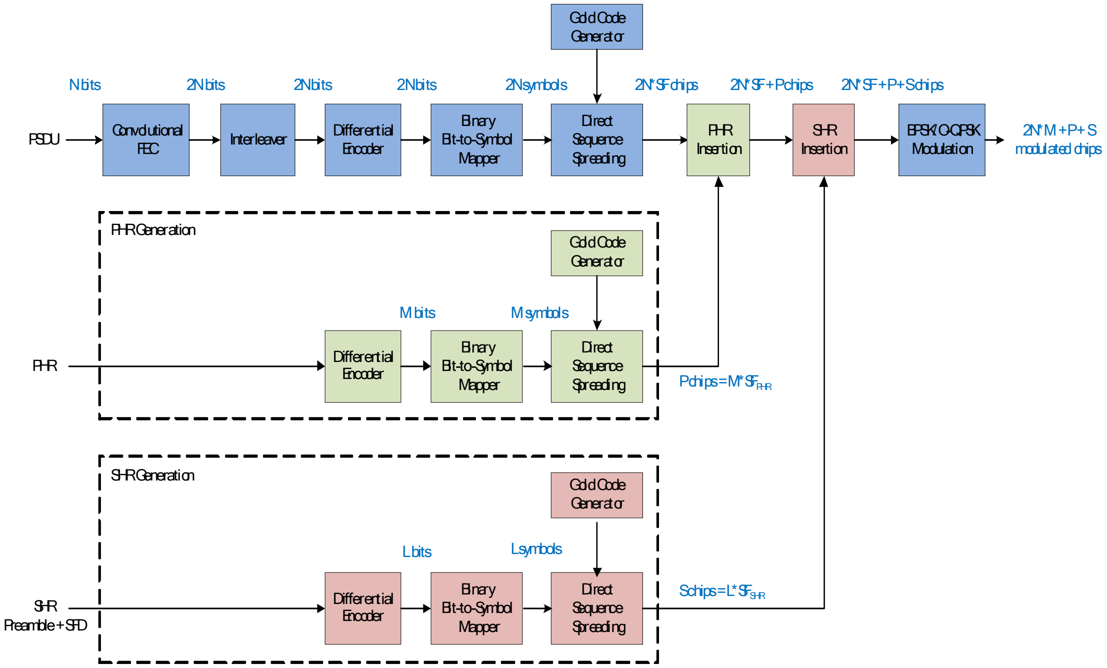

Figure 61—LECIM DSSS reference modulator diagram

The PSDU and tail/pad parts shall be coded with a convolutional encoder of coding rate R = 1/2. The convolutional encoder shall use the industry standard generator polynomials $g_0 = 133_8$ and $g_1 = 171_8$ of rate R = 1/2, as shown in Figure 62.

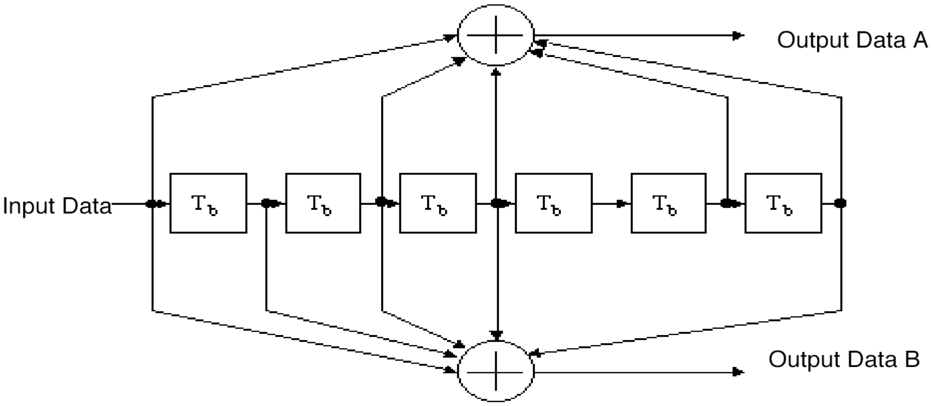

Figure 62—Convolutional encoder (k=7) for LECIM DSSS PHY

19.1.2.4 Interleaver

The output of the convolutional coder is interleaved using a pruned bit reversal interleaving (PBRI) algorithm.

The text that follows contains examples of bit reverse interleavers for three fragment sizes (256, 384, 512 symbols). Fragment sizes that are not powers of two (e.g., 384) employ pruning.

19.1.2.4.1 256 symbol fragment size

If the input sequence into the interleaver is represented by

$$[S_0\ S_1 \ldots S_{255}]$$

Then the output sequence of the interleaver can be described as

$$[S_0\ S_{16} \ldots S_N\ S_{255}]$$

The value N for the M^{th} output is determined as the bit-reversal of the value M.

Representing the value M as a binary representation

$$M = [m_7\ m_6 \ldots m_0]$$

where m_i are the binary digits, then

$$N = [m_0\ m_1 \ldots m_7]$$

where M is incremented sequentially from 0 to 255.

The sequence of N is shown in Table 74.

Table 74—Sequence of *N* for 256 symbol fragment size

Octet: 0	000	128	064	192	032	160	096	224
1	016	144	080	208	048	176	112	240
2	008	136	072	200	040	168	104	232
3	024	152	088	216	056	184	120	248
4	004	132	068	196	036	164	100	228
5	020	148	084	212	052	180	116	244
6	012	140	076	204	044	172	108	236
7	028	156	092	220	060	188	124	252
8	002	130	066	194	034	162	098	226

Table 74—Sequence of *N* for 256 symbol fragment size

9	018	146	082	210	050	178	114	242
10	010	138	074	202	042	170	106	234
11	026	154	090	218	058	186	122	250
12	006	134	070	198	038	166	102	230
13	022	150	086	214	054	182	118	246
14	014	142	078	206	046	174	110	238
15	030	158	094	222	062	190	126	254
16	001	129	065	193	033	161	097	225
17	017	145	081	209	049	177	113	241
18	009	137	073	201	041	169	105	233
19	025	153	089	217	057	185	121	249
20	005	133	069	197	037	165	101	229
21	021	149	085	213	053	181	117	245
22	013	141	077	205	045	173	109	237
23	029	157	093	221	061	189	125	253
24	003	131	067	195	035	163	099	227
25	019	147	083	211	051	179	115	243
26	011	139	075	203	043	171	107	235
27	027	155	091	219	059	187	123	251
28	007	135	071	199	039	167	103	231
29	023	151	087	215	055	183	119	247
30	015	143	079	207	047	175	111	239
31	031	159	095	223	063	191	127	255

19.1.2.4.2 384 symbol fragment size

If the input sequence into the interleaver is represented by

$$[S_0\ S_{1\ldots}S_{383}] \qquad (1)$$

Then the output sequence of the interleaver can be described as

$$[S_0\ S_{16}\ldots S_N\ S_{383}]$$

Representing the value M as a binary representation

$$M' = [m'_8\ m'_6\ldots m'_0]$$

where m_i are the binary digits, then

$$N = [m'_0\ m'_1\ldots m'_8]$$

where M is incremented sequentially from 0 to 512 and M' are the ordered set of M whose corresponding N is less than 384 (this is the pruning process).

The sequence of N is shown in Table 75.

Table 75—Sequence of *N* for 384 symbol fragment size (pruned)

Octet: 0	000	256	128	064	320	192	032	288
1	160	096	352	224	016	272	144	080
2	336	208	048	304	176	112	368	240
3	008	264	136	072	328	200	040	296
4	168	104	360	232	024	280	152	088
5	344	216	056	312	184	120	376	248
6	004	260	132	068	324	196	036	292
7	164	100	356	228	020	276	148	084
8	340	212	052	308	180	116	372	244
9	012	268	140	076	332	204	044	300
10	172	108	364	236	028	284	156	092
11	348	220	060	316	188	124	380	252
12	002	258	130	066	322	194	034	290
13	162	098	354	226	018	274	146	082
14	338	210	050	306	178	114	370	242
15	010	266	138	074	330	202	042	298
16	170	106	362	234	026	282	154	090
17	346	218	058	314	186	122	378	250
18	006	262	134	070	326	198	038	294

Table 75—Sequence of *N* for 384 symbol fragment size (pruned)

19	166	102	358	230	022	278	150	086
20	342	214	054	310	182	118	374	246
21	014	270	142	078	334	206	046	302
22	174	110	366	238	030	286	158	094
23	350	222	062	318	190	126	382	254
24	001	257	129	065	321	193	033	289
25	161	097	353	225	017	273	145	081
26	337	209	049	305	177	113	369	241
27	009	265	137	073	329	201	041	297
28	169	105	361	233	025	281	153	089
29	345	217	057	313	185	121	377	249
30	005	261	133	069	325	197	037	293
31	165	101	357	229	021	277	149	085
32	341	213	053	309	181	117	373	245
33	013	269	141	077	333	205	045	301
34	173	109	365	237	029	285	157	093
35	349	221	061	317	189	125	381	253
36	003	259	131	067	323	195	035	291
37	163	099	355	227	019	275	147	083
38	339	211	051	307	179	115	371	243
39	011	267	139	075	331	203	043	299
40	171	107	363	235	027	283	155	091
41	347	219	059	315	187	123	379	251
42	007	263	135	071	327	199	039	295
43	167	103	359	231	023	279	151	087
44	343	215	055	311	183	119	375	247

Table 75—Sequence of *N* for 384 symbol fragment size (pruned)

45	015	271	143	079	335	207	047	303
46	175	111	367	239	031	287	159	095
47	351	223	063	319	191	127	383	255

19.1.2.4.3 512 symbol fragment size

If the input sequence into the interleaver is represented by

$$[S_0\ S_1 \ldots S_{255}]$$

Then the output sequence of the interleaver can be described as

$$[S_0\ S_{16} \ldots S_N\ S_{255}]$$

The value N for the M^{th} output is determined as the bit-reversal of the value M.

Representing the value M as a binary representation

$$M = [m_8\ m_6 \ldots m_0]$$

where m_i are the binary digits, then

$$N = [m_0\ m_1 \ldots m_8]$$

where M is incremented sequentially from 0 to 511.

The sequence of N is shown in Table 76.

Table 76—Sequence of *N* for 512 symbol fragment size

Octet: 0	000	256	128	384	064	320	192	448
1	032	288	160	416	096	352	224	480
2	016	272	144	400	080	336	208	464
3	048	304	176	432	112	368	240	496
4	008	264	136	392	072	328	200	456
5	040	296	168	424	104	360	232	488
6	024	280	152	408	088	344	216	472
7	056	312	184	440	120	376	248	504

Table 76—Sequence of *N* for 512 symbol fragment size

8	004	260	132	388	068	324	196	452
9	036	292	164	420	100	356	228	484
10	020	276	148	404	084	340	212	468
11	052	308	180	436	116	372	244	500
12	012	268	140	396	076	332	204	460
13	044	300	172	428	108	364	236	492
14	028	284	156	412	092	348	220	476
15	060	316	188	444	124	380	252	508
16	002	258	130	386	066	322	194	450
17	034	290	162	418	098	354	226	482
18	018	274	146	402	082	338	210	466
19	050	306	178	434	114	370	242	498
20	010	266	138	394	074	330	202	458
21	042	298	170	426	106	362	234	490
22	026	282	154	410	090	346	218	474
23	058	314	186	442	122	378	250	506
24	006	262	134	390	070	326	198	454
25	038	294	166	422	102	358	230	486
26	022	278	150	406	086	342	214	470
27	054	310	182	438	118	374	246	502
28	014	270	142	398	078	334	206	462
29	046	302	174	430	110	366	238	494
30	030	286	158	414	094	350	222	478
31	062	318	190	446	126	382	254	510
32	001	257	129	385	065	321	193	449
33	033	289	161	417	097	353	225	481
34	017	273	145	401	081	337	209	465
35	049	305	177	433	113	369	241	497

Table 76—Sequence of *N* for 512 symbol fragment size

36	009	265	137	393	073	329	201	457
37	041	297	169	425	105	361	233	489
38	025	281	153	409	089	345	217	473
39	057	313	185	441	121	377	249	505
40	005	261	133	389	069	325	197	453
41	037	293	165	421	101	357	229	485
42	021	277	149	405	085	341	213	469
43	053	309	181	437	117	373	245	501
44	013	269	141	397	077	333	205	461
45	045	301	173	429	109	365	237	493
46	029	285	157	413	093	349	221	477
47	061	317	189	445	125	381	253	509
48	003	259	131	387	067	323	195	451
49	035	291	163	419	099	355	227	483
50	019	275	147	403	083	339	211	467
51	051	307	179	435	115	371	243	499
52	011	267	139	395	075	331	203	459
53	043	299	171	427	107	363	235	491
54	027	283	155	411	091	347	219	475
55	059	315	187	443	123	379	251	507
56	007	263	135	391	071	327	199	455
57	039	295	167	423	103	359	231	487
58	023	279	151	407	087	343	215	471
59	055	311	183	439	119	375	247	503
60	015	271	143	399	079	335	207	463
61	047	303	175	431	111	367	239	495
62	031	287	159	415	095	351	223	479
63	063	319	191	447	127	383	255	511

19.1.2.5 Differential encoding

The differential encoding of the DSSS PHY is described in 11.2.3.

19.1.2.6 Bit-to-symbol and symbol-to-chip encoding

The bit-to-symbol mapper converts bits into binary symbols through the mapping:

$$x[n] = \begin{cases} 1, \text{ if } b[n] = 0 \\ -1, \text{ if } b[n] = 1 \end{cases}$$

These binary symbols are then spread to chip-rate with spreading factor SF. This process is illustrated explicitly in Figure 63 where SF = 8. The symbols are first up-sampled SF times and interpolated using a scaled boxcar filter, as shown in Figure 64, i.e., the symbol is repeated SF times at chip-rate. Note that this is a mathematical representation of the direct sequence spreading operation. This process can be implemented in an alternative manner that is mathematically equivalent. The up-sampled symbols are multiplied by a specified Gold code to create the spread signal.

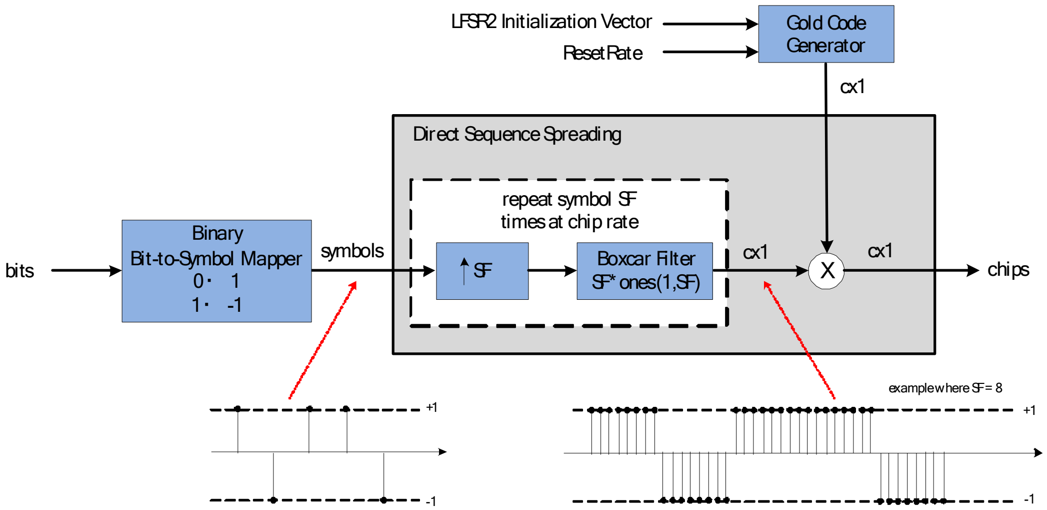

Figure 63—Bit-to-chip diagram for LECIM DSSS PHY

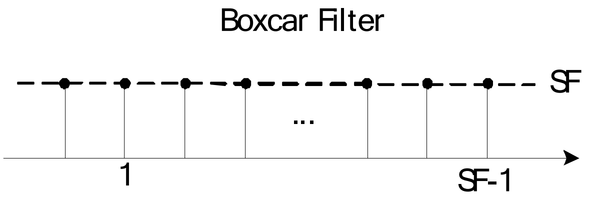

Figure 64—Boxcar filter

19.1.2.6.1 Gold code generator

Gold code sequences are a large family of easily parameterized PN sequences with good periodic cross-correlation and off-peak auto-correlation properties. A Gold code sequence is derived from the binary addition (XOR) of two maximum length sequences (*m*-sequences, or MLS), as illustrated in Figure 65. The *m*-sequences are generated using Fibonacci linear feedback shift registers (LFSR). Each LFSR is constructed from primitive (or prime) polynomials over Galois field 2 (GF[2]). The resulting sequences thus constitute segments of a set of Gold sequences. The specific *m*-sequences that follow are the preferred pair as described in the 3rd Generation Partnership Project (3GPP) Technical Specification 25.213. The Gold sequence can be parameterized by setting the initialization vector of LFSR2 to different values (LFSR1 is always initialized to 0x1).

- $m = 25$ (length of LSFR)
- $n = 2m - 1 = 33,554,431$ (length of Gold code)

- $n + 2 = 33{,}554{,}433$ (total Gold sequences) = a, b, $a \times b$, $a \times Tb$, $a \times T2b$, …

LFSR (MLS) generator polynomials:

- $p1(x) = x^{25} + x^3 + 1$
- $p2(x) = x^{25} + x^3 + x^2 + x + 1$

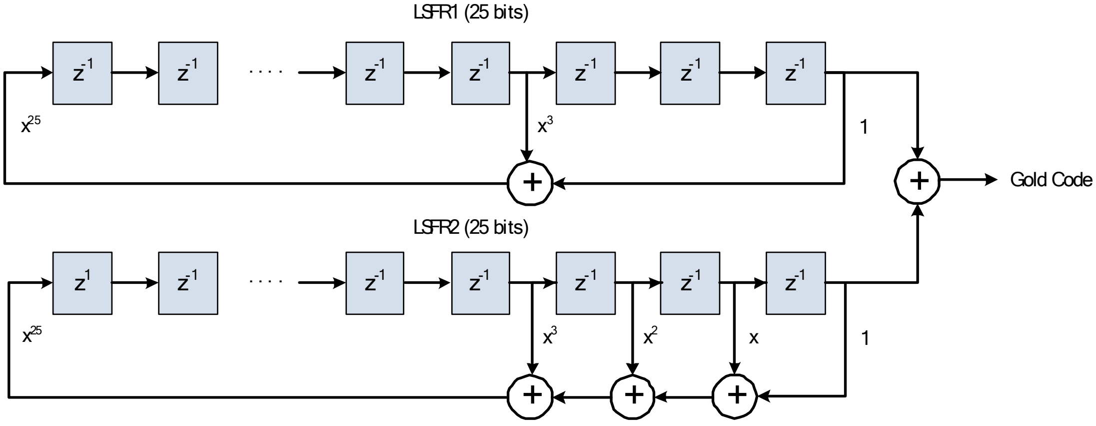

Figure 65—Gold code generator

19.1.2.6.2 OVSF code generator

The orthogonal variable spreading factor (OVSF) code is the same as the Walsh code, except that each sequence has a different index number in the code set, which is a result of their different generation algorithms.

The Gold code is to be used inside a co-located orthogonal network (CLON) as a primary code. OVSF codes are to be used to preserve orthogonality to identify the CLONs and clusters. It will provide double protection from outside interference.

The OVSF code is a linear code over a binary alphabet that maps messages of length n to codewords of length $2n$, and is generated from a Hadamard matrix but with the permutation matrix concept.

To reconstruct the OVSF code, recursively define a sequence of codes C_i as follows. Let C_0 be the root [1]. Assuming that C_i has been defined, for $i < r$, define C_{i+1} by

$$C_{i+1} = \begin{cases} C_i C_i, & \text{if } x_i = 0 \\ C_i(-C_i), & \text{if } x_i = 1 \end{cases}$$

The code C_N has the specified spreading factor and code index.

OVSF codes can also be defined recursively by a tree structure, as shown in Figure 66.

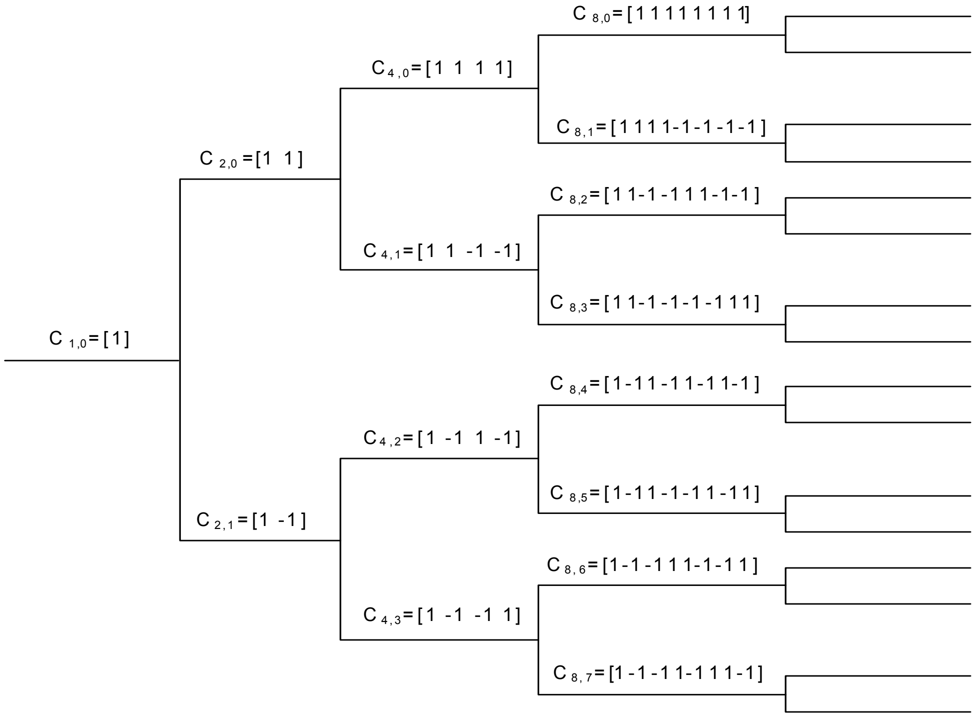

Figure 66—OVSF codes recursively defined by tree structure

19.1.2.7 BPSK/O-QPSK modulation

19.1.2.7.1 BPSK modulation

Binary phase-shift keying (BPSK) modulation for the DSSS PHY is described in 11.2.5.

The chip sequences are modulated onto the carrier using BPSK with raised cosine pulse shaping (roll-off factor = 1), where a chip value of one corresponds to a positive pulse and a chip value of zero corresponds to a negative pulse.

Chip rates/bands shown in Table TBD.

The pulse shape is described in 11.2.5.1.

During each symbol period, the least significant chip is transmitted first, and the most significant chip is transmitted last.

(further clarification on what is meant by least and most significant bit)

19.1.2.7.2 O-QPSK modulation

The chip sequences representing each data symbol are modulated onto the carrier using offset quadrature phase-shift keying (O-QPSK) with pulse shaping. For an even-indexed symbol, the even-indexed chips are modulated onto the in-phase (I) carrier, and odd-indexed chips are modulated onto the quadrature-phase (Q) carrier. For an odd-indexed symbol, even-indexed chips are modulated onto the quadrature-phase (Q) carrier, and odd-indexed chips are modulated onto the in-phase (I) carrier. To form the offset between I-phase and Q-phase chip modulation, the Q-phase chips shall be delayed by T_c with respect to the I-phase chips, as illustrated in Figure 67, where T_c is the inverse of the chip rate.

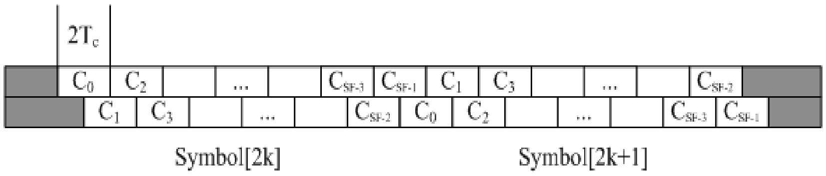

Figure 67—O-QPSK chip modulation

The pulse shape - 4g 16.3.2.13 Modulation parameters doe O-QPSK? Why multiple shapes?

The chip transmission order is the same as 10.2.7 - ?

During each symbol period, the least significant chip, C_0, is transmitted first, and the most significant chip, C_{SF-1}, is transmitted last.

(further clarification on what is meant by least and most significant bit or reference Figure 67 C_{SF-1})

19.2 FSK PHY specification

The frequency shift keying (FSK) PHY is described in the following subclauses.

19.2.1 PPDU format for FSK

The FSK PPDU shall support the format shown in Figure 68.

The synchronization header (SHR), PHY header (PHR), and PHY payload components are treated as bit strings of length n, numbered b_0 on the left and b_{n-1} on the right. When transmitted, they are processed b_0 first to b_{n-1} last, without regard to their content or structure.

All reserved fields shall be set to zero upon transmission and shall be ignored upon reception.

		Octets	
		N	**variable**
Preamble	SFD	As defined in 19.2.1.3	PSDU
SHR		PHR	PHY payload

Figure 68—Format of the LECIM FSK PPDU

19.2.1.1 Preamble field

The Preamble field shall contain *phyLECIMFSKPreambleLength* (as defined in 9.3) multiples of the 8-bit sequence "01010101."

Given the asymmetric nature of LECIM networks, greater capabilities of coordinators and low energy end devices, the range of preamble length is 0 to 100 octets. High functioning coordinators may need little or no preamble to synchronize, which reduces the transmit times of battery devices. A maximum preamble length of 100 is sufficient for the radios in end devices to synchronize for transmission.

19.2.1.2 SFD

The SFD shall be a 3-octet sequence, as shown in Table 77.

The SFD is transmitted starting from the leftmost bit (i.e., starting with b_0).

Table 77—SFD value for LECIM FSK PHY

Octets	1	2	3
Bit map	xxxxxxx	xxxxxxx	xxxxxxx

19.2.1.3 PHR

The formats of the PHR are shown for 127 and 2047 octet packets in Figure 69 and Figure 70, respectively. All multi-bit fields are unsigned integers and shall be processed MSB first.

The Frame Length field can be either 7 or 12 bits, for 127 and 2047 octet packets, respectively. The value of the Extension Bit field indicates which field length is used. The Frame Length field specifies the total number of octets contained in the PSDU (prior to FEC encoding, if enabled). The most significant bit (leftmost) shall be transmitted first.

It is important to note that LECIM networks are commissioned networks and strive to minimize energy consumption in battery-powered end devices. As such, not all parameters are signaled with bits in the PHR, but are instead assumed to be programmed into the network devices at commissioning. The parameters configuring the use of data whitening, FEC, interleaving, spreading, modulation type, and FCS length are considered commissioned parameters and are not signaled in the PHR.

Bit string index	0	1–7
Bit mapping	0	L_6–L_0
Field name	Extension Bit	Frame Length

Figure 69—PHR for 127 octet packet

19.2.1.4 PSDU field

The PSDU field carries the data of the PPDU.

Bit string index	0	1–3	1–12
Bit mapping	1	R_2–R_0	L_{11}–L_0
Field name	Extension Bit	Reserved	Frame Length

Figure 70—PHR for 2047 octet packet

19.2.2 Modulation and coding for FSK

The modulation for the FSK PHY shall be FSK/Gaussian FSK (GFSK) and position-based FSK (P-FSK)/position-based GFSK (P-GFSK).

Table 78 shows the modulation and channel parameters for the standard-defined PHY operating modes for the 863 MHz, 915 MHz, 917 MHz, 920 MHz, and 2450 MHz bands.

Although there are multiple data rates for each frequency band in Table 78, there is no over-the-air, dynamic data rate changing mechanism defined for this PHY. It is left to the system designer to select the appropriate data rates for the deployment during the design and commissioning of each specific network. The LECIM FSK PHY is not intended to be a multi-rate PHY with over-the-air signaling of changing data rates.

Table 78—LECIM FSK modulation and channel parameters*

Frequency band (MHz)	Parameter	37.5 kbps	25 kbps	12.5 kbps
470–510 (China)	End device to coordinator	GFSK/P-GFSK	GFSK/P-GFSK	Not supported
	Coordinator to end device	FSK/P-FSK	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	200	200	200
779–787 (China)	End device to coordinator	GFSK/P-GFSK	GFSK/P-GFSK	Not supported
	Coordinator to end device	FSK/P-FSK	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	200	200	200
863–870 (Europe)	End device to coordinator	Not supportable due to regulations	GFSK/P-GFSK	Not supported
	Coordinator to end device	Not supportable due to regulations	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	100	100	100
902–928 (US ISM)	End device to coordinator	GFSK/P-GFSK	GFSK/P-GFSK	Not supported
	Coordinator to end device	FSK/P-FSK	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	200	200	200

Table 78—LECIM FSK modulation and channel parameters[*] *(continued)*

Frequency band (MHz)	**Parameter**	**37.5 kbps**	**25 kbps**	**12.5 kbps**
917–923.5 (Korea)	End device to coordinator	GFSK/P-GFSK	GFSK/P-GFSK	Not supported
	Coordinator to end device	FSK/P-FSK	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	200	200	200
920–928 (Japan)	End device to coordinator	GFSK/P-GFSK	GFSK/P-GFSK	Not supported
	Coordinator to end device	FSK/P-FSK	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	200	200	200
2400–2483.5 (Worldwide)	End device to coordinator	GFSK/P-GFSK	GFSK/P-GFSK	Not supported
	Coordinator to end device	FSK/P-FSK	FSK/P-FSK	FSK
	Modulation index	0.5	1.0	4.0
	Channel spacing (kHz)	200	200	200

[*] Data rates shown are over-the-air data rates (the data rate transmitted over the air regardless whether the FEC is enabled or not).

The symbol duration used for the MAC and PHY timing parameters are shown in Table 79.

Table 79—LECIM FSK symbol duration used for MAC and PHY timing parameters

Frequency band (MHz)	**FSK symbol timing used for MAC and PHY timing parameters (µs)**
470–510 (China)	26.67
779–787 (China)	40
863–870 (Europe)	40
902–928 (US ISM)	26.67
917–923.5 (Korea)	26.67
920–928 (Japan)	26.67
2400–2483.5 (Worldwide)	26.67

The use of P-FSK/P-GFSK modulation for PSDU data is controlled by the PIB attribute *phyLECIMFSKPSDUMod*, as defined in 9.3. The modulation for preamble, SFD, and PHR shall be FSK/GFSK regardless of the value of *phyLECIMFSKPSDUMod*.

FSK/GFSK encodes one bit by transmitting a frequency modulated signal $m(t)$ with duration T_S, i.e., $0 \le t < T_S$. P-FSK/P-GFSK encodes two bits by transmitting a FSK/GFSK modulated signal $m(t)$ with $\mathrm{T_s}$ duration in one of two possible positions (also known as time deviation), i.e., $0 \le t < T_S$ and $T_S \le t < 2T_S$.

19.2.2.1 Reference modulator diagram

The functional block diagram in Figure 71 is provided as a reference for specifying the FSK PHY data flow processing functions. The subclause number in each block refers to the subclause that describes that function. Each bit shall be processed using the bit order rules defined in 19.2.1.

When FEC is enabled, the PHR and PSDU shall be processed for coding as a single block of data, as described in 19.2.2.4. When data whitening is enabled, the scrambling shall be only applied over the PSDU, as described in 19.2.3. When spreading is enabled, the spreading shall be applied over the PHR and PSDU, as described in 19.2.2.6.

All fields in the PPDU shall use the same symbol rate and modulation order, unless otherwise specified elsewhere in this standard.

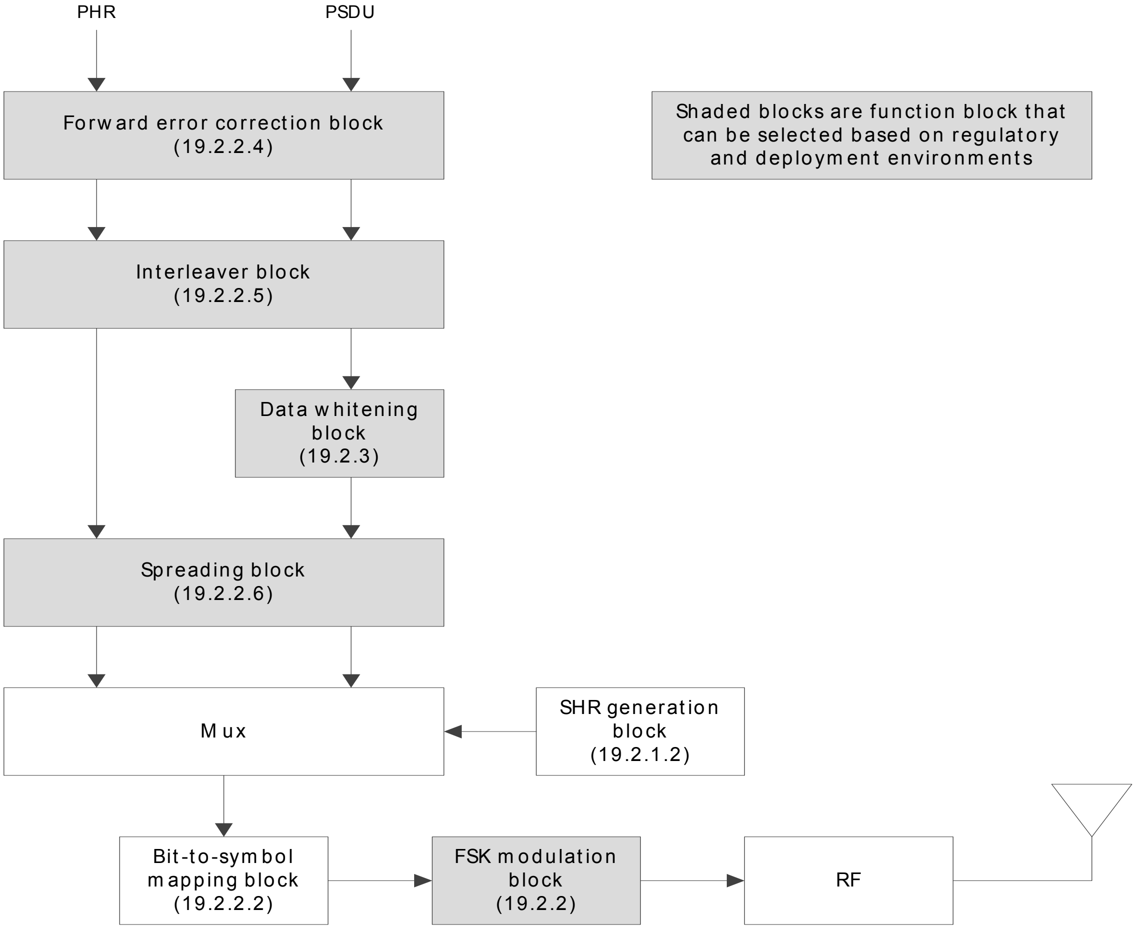

Figure 71—LECIM FSK reference modulator diagram

19.2.2.2 Bit-to-symbol mapping

The nominal frequency deviation, Δf, shall be

$$\frac{(\text{symbol rate} \times \text{modulation index})}{2}$$

The symbol encoding for FSK/GFSK and P-FSK/GFSK modulation is shown in Table 80 and Table 81, respectively, where the maximum frequency deviation, f_{dev}, is equal to Δf.

Table 80—FSK/GFSK symbol encoding

Symbol (b_0)	Frequency deviation	Time deviation
0	$-f_{dev}$	0
1	$+f_{dev}$	0

Table 81—P-FSK/P-GFSK symbol encoding

Symbol (b_0, b_1)	Frequency deviation	Time deviation
00	$-f_{dev}$	0
01	$-f_{dev}$	T_s
10	$+f_{dev}$	0
11	$+f_{dev}$	T_s

19.2.2.3 Modulation quality

Modulation quality shall be measured by observing the frequency deviation tolerance and the zero crossing tolerance of the eye diagram caused by a PN9 sequence of length 511 bits.

19.2.2.3.1 Frequency deviation tolerance

The GFSK modulation frequency tolerance is measured as a percentage of the frequency deviation dictated by the modulation index. The measured frequency deviation shall be ± 30% of the ideal frequency deviation, as shown in Figure 109 of 16.1.2.3.1. A binary one shall be represented by a positive frequency deviation, and a binary zero shall be represented by a negative frequency deviation.

The symbol timing shall be less than ± 20 ppm.

19.2.2.3.2 Zero crossing tolerance

The excursions for the zero crossings for all trajectories of the eye diagram shall be constrained as specified in 16.1.2.3.2.

19.2.2.4 Forward error correction

The FSK PHY shall perform FEC as defined in 16.3.2.6. The use of FEC is controlled by the PIB attribute *phyLECIMFECEnabled*, as defined in 9.3.

19.2.2.5 Code-symbol interleaving

The FSK PHY shall perform interleaving as defined in 16.1.2.5. The use of interleaving is controlled by the PIB attribute *phyLECIMFSKInterleavingEnabled*, as defined in 9.3.

19.2.2.6 Spreading

The use of spreading is controlled by the PIB attribute *phyLECIMFSKSpreading*, as defined in 9.3. The spreading factor (SF) can be 1, 2, 4, 8, or 16. The variable SF is indicated by the PIB attribute *phyLECIMFSKSpreadingFactor*, as defined in 9.3.

For spreading, a single input bit (b_0) is mapped into the spreading bits $(c_0, c_1, \ldots, c_{SF-1})$, as shown in Figure 72, and its mapping is represented in Table 82.

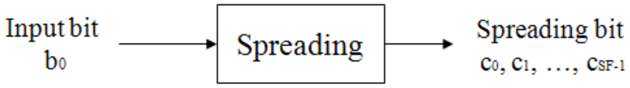

Figure 72—Spreading function

Table 82—Input bit to spreading bits mapping

Spreading factor (SF)	Input bit (b_0) = 0	Input bit (b_0) = 1
1	$(c_0) = 0$	$(c_0) = 1$
2	$(c_0, c_1) = 01$	$(c_0, c_1) = 10$
4	$(c_0, \ldots, c_3) = 0101$	$(c_0, \ldots, c_3) = 1010$
8	$(c_0, \ldots, c_7) = 0101\ 0101$	$(c_0, \ldots, c_7) = 1010\ 1010$
16	$(c_0, \ldots, c_{15}) = 0101\ 0101\ 0101\ 0101$	$(c_0, \ldots, c_{15}) = 1010\ 1010\ 1010\ 1010$

19.2.3 Data whitening for FSK

The FSK PHY may optionally perform data whitening as defined in 16.1.3. The use of data whitening is controlled by the PIB attribute *phyLECIMFSKScramblePSDU*, as defined in 9.3.

19.2.4 FSK PHY RF requirements

19.2.4.1 Operating frequency range

The FSK PHY operates in the bands given in Table 78.

19.2.4.2 Regulatory compliance

It is the responsibility of the implementer to verify and ensure that the device is in compliance with all regulatory requirements in the geographic region where the device is deployed or sold. Conformance with this standard does not guarantee compliance with the relevant regulatory requirements which may apply.

19.2.4.3 Radio frequency tolerance

The single-sided clock frequency tolerance T at the transmitter, in ppm, shall be as follows:

$$T = 20 \text{ ppm}$$

19.2.4.4 Channel switch time

Channel switch time shall be less than or equal to 500 µs. The channel switch time is defined as the time elapsed when changing to a new channel, including any required settling time.

19.2.4.5 Transmit spectral mask

Implementers are responsible to assure that the transmit spectral content conforms to all local regulations.

19.2.4.6 Receiver sensitivity

Under the conditions specified in 8.1.7, a compliant PHY device shall be capable of achieving a sensitivity of –95 dBm or better.

19.2.4.7 Receiver interference rejection

The minimum receiver interference rejection levels are given in Table 83. The adjacent channels are the ones on either side of the desired channel that are closest in frequency to the desired channel, and the alternate channels are one more removed from the adjacent channels. For example, when channel 15 is the desired channel, channel 14 and channel 16 are the adjacent channels, and channel 13 and channel 17 are the alternate channels.

Table 83—Minimum receiver interference rejection requirements

Adjacent channel rejection	Alternate channel rejection
10 dB	30 dB

The adjacent channel rejection shall be measured as follows. The desired signal shall be a compliant GFSK/ FSK PHY signal, as defined by 19.2.2, of pseudo-random data. The desired signal is input to the receiver at a level 3 dB greater than the maximum allowed receiver sensitivity given in 19.2.4.6.

In either the adjacent or the alternate channel, a compliant signal, as defined by 19.2.2, is input at the level specified in Table 83 relative to the desired signal. The test shall be performed for only one interfering signal at a time. The receiver shall meet the error rate criteria defined in 8.1.7 under these conditions.

19.2.4.8 Tx-to-Rx turnaround time

The FSK PHY shall meet the requirements for TX-to-RX turnaround time as defined in 8.2.1.

19.2.4.9 Rx-to-Tx turnaround time

The FSK PHY shall meet the requirements for RX-to-TX turnaround time as defined in 8.2.2.

19.2.4.10 Transmit power

A transmitter shall be capable of transmitting at least –3 dBm. The maximum transmit power is limited by local regulatory bodies.

Insert after Annex O the following new annex (Annex P):

Annex P

(informative)

Low Energy, Critical Infrastructure Monitoring Systems

P.1 Introduction

As per Wikipedia (http://en.wikipedia.org/wiki/Critical_infrastructure): Critical infrastructure is a term used by governments to describe assets that are essential for the functioning of a society and economy. Most commonly associated with the term are facilities for:

- Electricity generation, transmission, and distribution
- Gas production, transport, and distribution
- Oil and oil products production, transport, and distribution
- Telecommunication
- Water supply (e.g., drinking water, waste water/sewage, stemming of surface water [e.g., dikes and sluices])
- Agriculture, food production, and distribution
- Heating (e.g., natural gas, fuel oil, district heating)
- Public health (e.g., hospitals, ambulances)
- Transportation systems (e.g., fuel supply, railway network, airports, harbors, inland shipping)
- Financial services (e.g., banking, clearing)
- Security services (e.g., police, military)

P.1.1 LECIM characteristics

The LECIM portions of this standard form the MAC and PHY behaviors that implement a minimal network infrastructure, enables the collection of scheduled and event data from a large number of non-mains powered end points that are widely dispersed, or are in challenging propagation environments. To facilitate low energy operation necessary for multi-year battery life, MAC protocols minimize network maintenance traffic and device wake durations. In addition, LECIM addresses the changing propagation and interference environments encountered over many years.

The following is a list of LECIM characteristics and the underlying behaviors that form them:

a) Minimal infrastructure
    - Star topology, i.e., no repeaters are typically needed due long range.
    - Mains energy supply is only necessary for coordinator.
b) Commissioned network (not ad hoc)
    - Devices are configured specifically for the deployed network.
    - Devices are stateful, i.e., they are preconfigured with parameters that eliminate the need for wireless messages sending configuration information.
c) Long range
    - High receiver sensitivity, e.g., narrow bandwidth or high processing gain.

  - Interference robustness.
  - Challenging environments and widely dispersed devices.

d) Very limited energy supplied endpoints
  - Ten to twenty year life with no maintence, e.g., original battery must supply all energy for 20 years.
  - Energy harvesting with low power supplies, i.e., short and infrequent transmission and reception durations.

e) Significant difference between coordinator and endpoints
  - Does not preclude distributed systems.

f) Asymmetrical data flows
  - Sensor end point: up-link dominates data flow with limited down-link data needs.
  - Actuator end point: down-link dominates data flow with limited up-link data needs.

P.1.2 Use case examples

The following use cases exemplify LECIM applications.

P.1.2.1 Oil and gas pipeline monitoring

The key drivers of pipeline monitoring are as follows:
- Environmental protection
- Reliability (critical resources)
- Cost savings (increasing cost)
- Compliance (regulators)

P.1.2.2 Water leak detection

The key drivers of water leak detection are as follows:
- Permanent installation of large number of sensors underground
- Long range and ability to penetrate underground vaults
- Battery operated and long lifetime
  - Small data messages once per day and in case of alarm event (e.g., leak detected)
- Low installation cost (easy deployment) and low cost of maintenance

P.1.2.3 Soil monitoring

The key drivers of soil monitoring are as follows:
- Power consumption
  - Low-cost batteries that last over many years
- Networking
  - Long range links to cover large fields
  - Ability to use mesh or tree networking for complicated environment
  - Ability to connecting WPAN with mobile networks
- Reliability and cost
  - Very low maintence requirements

P.1.2.4 Inventory control - event driven with query

The application is for a warehouse floor with thousands of parts bins. Each bin has a battery operated RF link for communicating current quantity and changes in quantity to the central inventory control (CIC) system. Battery life is important.

Each bin contains only one part number. The RF link has an LCD display showing the quantity in the bin. It also has an “Increase Button” and a “Decrease Button.” When an operator adds units to the bin, he presses the Increase Button, and when parts are removed, he presses the Decrease Button. Each time a button is pressed, it generates an event to the RF module, which then transmits the change to the CIC. This would most likely use a contention access method for transmission, since events occur in an unscheduled manner.

The CIC receives events from all of the bins, as changes are made to the quantity contained in each bin. Both the local RF module and the CIC maintain the quantity in the bin.

For inventory auditing, it is necessary for the CIC to query each bin to check the quantity. This requires the CIC to initiate a transaction with each bin, either individually or as a broadcast/multicast message. The desire is to have all bins report within a reasonable time (minutes).

Also, since changes in quantity are event driven, the CIC need a means to query each bin to make sure that it is still operational and that no "change in quantity" events were missed.

To minimize battery drain, the LECIM device is only activated when necessary:

- A change in quantity as indicated by a button event
- Some type of synchronous sniff/query operation for receiving to queries from CIC
- Response to query messages

P.1.2.5 Building monitoring - time and event driven data with query

A building (or any structure) is being monitored by sensors that report measurement or state information over long periods, e.g., several minutes to several hours. There may also be sensors that report events or changes in state that are event driven and not time driven. Battery life is important.

Each measurement sensor is set to report its information at a certain interval, using either a GTS or the CAP. This gives very low duty cycle for normal operation, which is 99% of the usage. There may also be sensors that are event driven which report change in state, such as door open/closed, door locked/unlocked, switch on/off, etc. This is also low duty cycle.

Occasionally there is an event, maybe an emergency, where the central monitoring system must get readings from all sensors as soon as possible. The central controller must send a request to all sensors to report their current measurement or state. This requires a low latency response mechanism that can maintain long battery life.

P.1.3 LECIM behaviors

The following assumptions are essential to address the needs of LECIM applications:

- Commissioning
- Low energy
- Coverage extension

P.1.3.1 Commissioning

Commissioning by a professional installer allows the network to reduce the amount of data that must be sent by creating statefulness.

P.1.3.2 Low energy

LECIM applications require significantly low energy operation to be able to either last 20 years on original battery supply or energy harvesting mechanisms. Achieving low energy operation is made very difficult given the low data rates necessary for long range operation. Accordingly, LECIM networks must be able to elide any overhead octets not absolutely necessary to minimize transmit and receive durations, schedule link times to minimize device "on" durations, and maximize link reliability to minimize retransmissions.

P.1.3.3 Coverage extension

To keep infrastructure costs to a minimum, LECIM devices have large link margins to achieve long ranges without requiring mesh devices or repeater devices. Requiring mesh or repeater devices would increase the number of devices needed to sustain the network, increase costs by requiring renting or leasing space for those devices, and in most cases require mains power for these devices.

P.2 Functionality added: DSSS, FSK, fragmentation, frame priority, PIBs, IEs, attributes

P.2.1 DSSS

The DSSS devices used by LECIM networks differ from the other DSSS devices defined in this standard in that they have significant process gain to allow devices to receive messages with very low or negative carrier to noise ratios. High process gain also allows for code division multiple access (CDMA) operation to reduce the possibility of collisions.

P.2.2 FSK

The FSK devices for LECIM are typically narrow bandwidth (hence low data rate) devices that enable high sensitivities and many channels in order to reduce the possibility of collisions.

P.2.3 Fragmentation

The lower effective data rates resulting from high processing gain or narrow bandwidths dramatically decrease data rate, hence an increase in the bit times. The increase in bit times would make typical IEEE 802.15.4 frames so long in duration as to decrease their reliability and degrade coexistence. To keep each transmission duration sufficiently short, a method of fragmentation at the PHY/MAC level is required. The properties of this fragmentation method are to reduce all overhead to the minimum amount, by preconfiguring the link between two devices with the information necessary for reception and proper reassembly. This preconfiguration is done in order to eliminate the need to send this information in every packet.

P.2.4 Frame priority

Frame priority allows LECIM networks to exhibit low latencies for truly critical data messages versus those latencies for link maintenance or other lower priority messages.

P.2.5 PIB attributes

LECIM mechanisms and protocols require additional PIB attributes.