

**IEEE Standard for
Information technology—
Telecommunications and information
exchange between systems—
Local and metropolitan area networks—
Specific requirements—**

**Part 15.4: Wireless Medium Access Control
(MAC) and Physical Layer (PHY)
Specifications for Low-Rate Wireless
Personal Area Networks (WPANs)**

1. Overview

1.1 General

Wireless personal area networks (WPANs) are used to convey information over relatively short distances. Unlike wireless local area networks (WLANs), connections effected via WPANs involve little or no infrastructure. This feature allows small, power-efficient, inexpensive solutions to be implemented for a wide range of devices.

This document defines a standard for a low-rate WPAN (LR-WPAN).

1.2 Scope

The scope of this revision is to produce specific enhancements and corrections to IEEE Std 802.15.4, all of which will be backwards compatible with IEEE Std 802.15.4-2003. These enhancements and corrections include resolving ambiguities, reducing unnecessary complexity, increasing flexibility in security key usage, considerations for newly available frequency allocations, and others.

IEEE Std 802.15.4 defines the physical layer (PHY) and medium access control (MAC) sublayer specifications for low-data-rate wireless connectivity with fixed, portable, and moving devices with no battery or very limited battery consumption requirements typically operating in the personal operating space (POS) of 10 m. It is foreseen that, depending on the application, a longer range at a lower data rate may be an acceptable tradeoff.

It is the intent of this revision to work toward a level of coexistence with other wireless devices in conjunction with Coexistence Task Groups, such as IEEE 802.15.2 and IEEE 802.11/ETSI-BRAN/MMAC 5GSG.

1.3 Purpose

The purpose of this revision is to extend the market applicability of IEEE Std 802.15.4 and to remove ambiguities in the standard. Implementations of the 2003 edition of this standard have revealed potential areas of improvements. Additional frequency bands are being made available in various countries that are attractive for this application space.

2. Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

FIPS Pub 197, Advanced Encryption Standard (AES).¹

IEEE Std 802[®], IEEE Standards for Local and Metropolitan Area Networks: Overview and Architecture.²

ISO/IEC 8802-2 (IEEE Std 802.2[™]), Information technology — Telecommunications and information exchange between systems — Local and metropolitan area networks — Specific requirements — Part 2: Logical link control.³

ISO/IEC 9646-7 (ITU-T Rec. X.296), Information technology — Open systems interconnection — Conformance testing methodology and framework — Part 7: Implementation conformance statements.

¹FIPS publications are available from the National Technical Information Service (NTIS), U. S. Dept. of Commerce, 5285 Port Royal Rd., Springfield, VA 22161 (<http://www.ntis.org/>).

²IEEE Publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, Piscataway, NJ 08854, USA (<http://standards.ieee.org>).

³ISO/IEC publications are available from the ISO Central Secretariat, Case Postale 56, 1 rue de Varembé, CH-1211, Genève 20, Switzerland/Suisse (<http://www.iso.ch/>). ISO/IEC publications are also available in the United States from Global Engineering Documents, 15 Inverness Way East, Englewood, Colorado 80112, USA (<http://global.ihs.com/>). Electronic copies are available in the United States from the American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (<http://www.ansi.org/>).

3. Definitions

For the purposes of this standard, the following terms and definitions apply. Terms not defined in this clause can be found in the *The Authoritative Dictionary of IEEE Standards Terms*, Seventh Edition [B3].⁴

3.1 alternate personal area network (PAN) coordinator: A coordinator that is capable of replacing the PAN coordinator, if the PAN coordinator leaves the network for any reason. A PAN can have zero or more alternate PAN coordinators.

3.2 association: The service used to establish membership for a device in a wireless personal area network (WPAN).

3.3 authentication tag: Information that allows the verification of authenticity of a message.

3.4 beacon-enabled personal area network (PAN): A PAN in which all coordinators emit regular beacons, i.e., have beacon order $< 0x0F$.

3.5 block cipher: A cryptographic function that operates on strings of fixed size.

3.6 block size: The size, in bits, of the strings on which a block cipher operates.

3.7 contention access period (CAP): The period of time immediately following a beacon frame during which devices wishing to transmit will compete for channel access using a slotted carrier sense multiple access with collision avoidance (CSMA-CA) mechanism.

3.8 contention access period (CAP) symbol: A symbol period occurring during the CAP.

3.9 coordinator: A full-function device (FFD) capable of relaying messages. If a coordinator is the principal controller of a personal area network (PAN), it is called the PAN coordinator.

3.10 data authentication: The process whereby an entity receiving a message corroborates evidence about the true source of the information in the message and, thereby, evidence that the message has not been modified in transit.

3.11 data authenticity: Assurance about the source of information.

3.12 device: Any entity containing an implementation of the IEEE 802.15.4 medium access control (MAC) and physical interface to the wireless medium. A device may be a reduced-function device (RFD) or a full-function device (FFD).

3.13 encryption: The transformation of a message into a new representation so that privileged information is required to recover the original representation.

3.14 frame: The format of aggregated bits from a medium access control (MAC) sublayer entity that are transmitted together in time.

3.15 full-function device (FFD): A device capable of operating as a coordinator.

3.16 group key: A key that is known only to a set of devices.

3.17 idle period: A duration of time where no transceiver activity is scheduled to take place.

⁴The numbers in brackets correspond to the numbers of the bibliography in Annex G.

3.18 key: Privileged information that may be used, for example, to protect information from disclosure to, and/or undetectable modification by, parties that do not have access to this privileged information.

3.19 key establishment: A process whereby two or more parties establish a key.

3.20 keying material: The combination of a key and associated security information (e.g., a nonce value).

3.21 key management: The collection of processes for the establishment and maintenance of keying relationships over a system's lifetime.

3.22 key sharing group: A set of devices that share a key.

3.23 local clock: The symbol clock internal to a device.

3.24 link key: A secret key that is shared between precisely two devices.

3.25 minimum security level: Indication of minimum protection required on information in transit.

3.26 mobile device: A device whose logical location in the network may change during use.

3.27 m-sequence: Maximal length linear feedback shift register sequence.

3.28 nonbeacon-enabled personal area network (PAN): A PAN in which coordinators do not emit regular beacons, i.e., have beacon order = 0x0F.

3.29 nonce: A nonrepeating value, such as an increasing counter, a sufficiently long random string, or a timestamp.

3.30 orphaned device: A device that has lost contact with its associated coordinator.

3.31 packet: The formatted, aggregated bits that are transmitted together in time across the physical medium.

3.32 payload data: The contents of a data message that is being transmitted.

3.33 personal area network (PAN) coordinator: A coordinator that is the principal controller of a PAN. An IEEE 802.15.4 network has exactly one PAN coordinator.

3.34 personal operating space (POS): The space about a person or object that is typically about 10 m in all directions and envelops the person or object whether stationary or in motion.

3.35 plain text: A string of unscrambled information.

3.36 protection: The combination of security services provided for information in transit, such as confidentiality, data authenticity, and/or replay protection.

3.37 radio sphere of influence: The region of space throughout which a radio may successfully communicate with other like radios.

3.38 reduced-function device (RFD): A device that is not capable of operating as a coordinator.

3.39 security level: Indication of purported protection applied to information in transit.

3.40 self-healing: The ability of the network to detect, and recover from, faults appearing in either network nodes or communication links, without human intervention.

3.41 self-organizing: The ability of network nodes to detect the presence of other nodes and to organize into a structured, functioning network without human intervention.

3.42 symmetric key: A secret key that is shared between two or more parties that may be used for encryption/decryption or integrity protection/integrity verification, depending on its intended use.

3.43 transaction: The exchange of related, consecutive frames between two peer medium access control (MAC) entities, required for a successful transmission of a MAC command or data frame.

3.44 transaction queue: A list of the pending transactions, which are to be sent using indirect transmission, that are initiated by the medium access control (MAC) sublayer of a given coordinator. The transaction queue is maintained by that coordinator while the transactions are in progress, and its length is implementation-dependent but must be at least one.

4. Acronyms and abbreviations

AES	advanced encryption standard
ASK	amplitude shift keying
AWGN	additive white Gaussian noise
AWN	affected wireless network
BE	backoff exponent
BER	bit error rate
BI	beacon interval
BLE	battery life extension
BO	beacon order
BPSK	binary phase-shift keying
BSN	beacon sequence number
CAP	contention access period
CBC-MAC	cipher block chaining message authentication code
CCA	clear channel assessment
CCM	counter with CBC-MAC (mode of operation)
CCM*	extension of CCM
CFP	contention-free period
CRC	cyclic redundancy check
CSMA-CA	carrier sense multiple access with collision avoidance
CTR	counter mode
CW	contention window (length)
DSN	data sequence number
DSSS	direct sequence spread spectrum
ED	energy detection
EIRP	effective isotropic radiated power
EMC	electromagnetic compatibility
ERP	effective radiated power
EVM	error-vector magnitude
FCS	frame check sequence
FFD	full-function device
FH	frequency hopping
FHSS	frequency hopping spread spectrum
GTS	guaranteed time slot
IFS	interframe space or spacing
ISM	industrial, scientific, and medical
IUT	implementation under test
IWN	interfering wireless network
LIFS	long interframe spacing
LLC	logical link control
LQI	link quality indication
LPDU	LLC protocol data unit
LR-WPAN	low-rate wireless personal area network
LSB	least significant bit
MAC	medium access control

MCPS	MAC common part sublayer
MCPS-SAP	MAC common part sublayer service access point
MFR	MAC footer
MHR	MAC header
MIC	message integrity code
MLME	MAC sublayer management entity
MLME-SAP	MAC sublayer management entity service access point
MSB	most significant bit
MPDU	MAC protocol data unit
MSDU	MAC service data unit
NB	number of backoff (periods)
OCDM	orthogonal code division multiplexing
O-QPSK	offset quadrature phase-shift keying
OSI	open systems interconnection
PAN	personal area network
PC	personal computer
PD	PHY data
PD-SAP	PHY data service access point
PER	packet error rate
PHR	PHY header
PHY	physical layer
PIB	PAN information base
PICS	protocol implementation conformance statement
PLME	physical layer management entity
PLME-SAP	physical layer management entity service access point
PN	pseudo-random noise
POS	personal operating space
PPDU	PHY protocol data unit
PSD	power spectral density
PSDU	PHY service data unit
PSSS	parallel sequence spread spectrum
RF	radio frequency
RFD	reduced-function device
RX	receive or receiver
SD	superframe duration
SER	symbol error rate
SFD	start-of-frame delimiter
SHR	synchronization header
SIFS	short interframe spacing
SIR	signal-to-interference ratio
SNR	signal-to-noise ratio
SO	superframe order
SPDU	SSCS protocol data units
SRD	short-range device
SSCS	service-specific convergence sublayer
SUT	system under test

TRX	transceiver
TX	transmit or transmitter
WLAN	wireless local area network
WPAN	wireless personal area network

5. General description

5.1 Introduction

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

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

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

Two different device types can participate in an IEEE 802.15.4 network; a full-function device (FFD) and a reduced-function device (RFD). The FFD can operate in three modes serving as a personal area network (PAN) coordinator, a coordinator, or a device. An FFD can talk to RFDs or other FFDs, while an RFD can talk only to an FFD. An RFD is intended for applications that are extremely simple, such as a light switch or a passive infrared sensor; they do not have the need to send large amounts of data and may only associate with a single FFD at a time. Consequently, the RFD can be implemented using minimal resources and memory capacity.

This standard is backward-compatible to the 2003 edition; in other words, devices conforming to this standard are capable of joining and functioning in a PAN composed of devices conforming to IEEE Std 802.15.4-2003.

5.2 Components of the IEEE 802.15.4 WPAN

A system conforming to this standard consists of several components. The most basic is the device. A device may be an RFD or an FFD. Two or more devices within a POS communicating on the same physical channel constitute a WPAN. However, this WPAN shall include at least one FFD, operating as the PAN coordinator.

An IEEE 802.15.4 network is part of the WPAN family of standards although the coverage of the network may extend beyond the POS, which typically defines the WPAN.

A well-defined coverage area does not exist for wireless media because propagation characteristics are dynamic and uncertain. Small changes in position or direction may result in drastic differences in the signal strength or quality of the communication link. These effects occur whether a device is stationary or mobile, as moving objects may impact station-to-station propagation.

5.3 Network topologies

Depending on the application requirements, an IEEE 802.15.4 LR-WPAN may operate in either of two topologies: the star topology or the peer-to-peer topology. Both are shown in Figure 1. In the star topology the communication is established between devices and a single central controller, called the PAN coordinator. A device typically has some associated application and is either the initiation point or the termination point for network communications. A PAN coordinator may also have a specific application, but it can be used to initiate, terminate, or route communication around the network. The PAN coordinator is the primary controller of the PAN. All devices operating on a network of either topology shall have unique 64-bit addresses. This address may be used for direct communication within the PAN, or a short address may be allocated by the PAN coordinator when the device associates and used instead. The PAN coordinator might often be mains powered, while the devices will most likely be battery powered. Applications that benefit from a star topology include home automation, personal computer (PC) peripherals, toys and games, and personal health care.

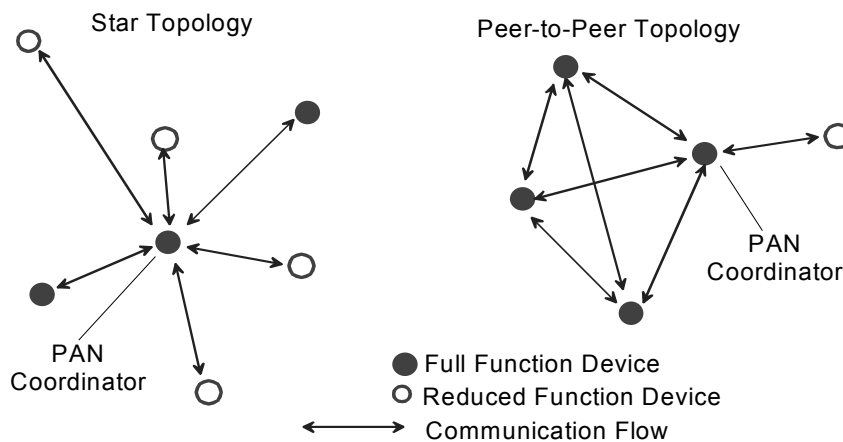


Figure 1—Star and peer-to-peer topology examples

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

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

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

5.3.1 Star network formation

The basic structure of a star network is illustrated in Figure 1. After an FFD is activated, it can establish its own network and become the PAN coordinator. All star networks operate independently from all other star networks currently in operation. This is achieved by choosing a PAN identifier that is not currently used by any other network within the radio sphere of influence. Once the PAN identifier is chosen, the PAN

coordinator allows other devices, potentially both FFDs and RFDs, to join its network. The higher layer can use the procedures described in 7.5.2 and 7.5.3 to form a star network.

5.3.2 Peer-to-peer network formation

In a peer-to-peer topology, each device is capable of communicating with any other device within its radio sphere of influence. One device is nominated as the PAN coordinator, for instance, by virtue of being the first device to communicate on the channel. Further network structures are constructed out of the peer-to-peer topology and it is possible to impose topological restrictions on the formation of the network.

An example of the use of the peer-to-peer communications topology is the cluster tree. The cluster tree network is a special case of a peer-to-peer network in which most devices are FFDs. An RFD connects to a cluster tree network as a leaf device at the end of a branch because RFDs do not allow other devices to associate. Any of the FFDs may act as a coordinator and provide synchronization services to other devices or other coordinators. Only one of these coordinators can be the overall PAN coordinator, which may have greater computational resources than any other device(s) in the PAN. The PAN coordinator forms the first cluster by choosing an unused PAN identifier and broadcasting beacon frames to neighboring devices. A contention resolution mechanism is required if two or more FFDs simultaneously attempt to establish themselves as PAN coordinators; however, such a mechanism is outside the scope of this standard. A candidate device receiving a beacon frame may request to join the network at the PAN coordinator. If the PAN coordinator permits the device to join, it adds the new device as a child device in its neighbor list. Then the newly joined device adds the PAN coordinator as its parent in its neighbor list and begins transmitting periodic beacons; other candidate devices may then join the network at that device. If the original candidate device is not able to join the network at the PAN coordinator, it will search for another parent device. The detailed procedures describing how a PAN is started and how devices join a PAN are found in 7.5.2 and 7.5.3.

The simplest form of a cluster tree network is a single cluster network, but larger networks are possible by forming a mesh of multiple neighboring clusters. Once predetermined application or network requirements are met, the first PAN coordinator may instruct a device to become the PAN coordinator of a new cluster adjacent to the first one. Other devices gradually connect and form a multicluster network structure, such as the one seen in Figure 2. The lines in Figure 2 represent the parent-child relationships of the devices and not the communication flow. The advantage of a multicluster structure is increased coverage area, while the disadvantage is an increase in message latency.

5.4 Architecture

The IEEE 802.15.4 architecture is defined in terms of a number of blocks in order to simplify the standard. These blocks are called layers. Each layer is responsible for one part of the standard and offers services to the higher layers. The layout of the blocks is based on the open systems interconnection (OSI) seven-layer model (see ISO/IEC 7498-1:1994 [B12]).

The interfaces between the layers serve to define the logical links that are described in this standard.

An LR-WPAN device comprises a PHY, which contains the radio frequency (RF) transceiver along with its low-level control mechanism, and a MAC sublayer that provides access to the physical channel for all types of transfer. Figure 3 shows these blocks in a graphical representation, which are described in more detail in 5.4.1 and 5.4.2.


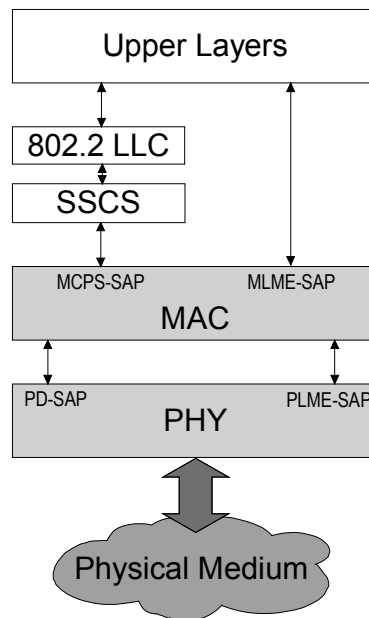


Figure 2—Cluster tree network



NOTE—For MCPS-SAP, see 7.1; for MLME-SAP, see 5.4.2; for PD-SAP, see 6.2; and for PLME-SAP, see 5.4.1.

Figure 3—LR-WPAN device architecture

The upper layers, shown in Figure 3, consist of a network layer, which provides network configuration, manipulation, and message routing, and an application layer, which provides the intended function of the device. The definition of these upper layers is outside the scope of this standard. An IEEE 802.2 Type 1 logical link control (LLC) can access the MAC sublayer through the service-specific convergence sublayer (SSCS), defined in Annex A. The LR-WPAN architecture can be implemented either as embedded devices or as devices requiring the support of an external device such as a PC.

5.4.1 Physical layer (PHY)

The PHY provides two services: the PHY data service and the PHY management service interfacing to the physical layer management entity (PLME) service access point (SAP) (known as the PLME-SAP). The PHY data service enables the transmission and reception of PHY protocol data units (PPDUs) across the physical radio channel.

Clause 6 contains the specifications for the PHY.

The features of the PHY are activation and deactivation of the radio transceiver, ED, LQI, channel selection, clear channel assessment (CCA), and transmitting as well as receiving packets across the physical medium. The radio operates at one or more of the following unlicensed bands:

- 868–868.6 MHz (e.g., Europe)
- 902–928 MHz (e.g., North America)
- 2400–2483.5 MHz (worldwide)

Refer to Annex F for an informative summary of regulatory requirements.

5.4.2 MAC sublayer

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

The features of the MAC sublayer are beacon management, channel access, GTS management, frame validation, acknowledged frame delivery, association, and disassociation. In addition, the MAC sublayer provides hooks for implementing application-appropriate security mechanisms.

Clause 7 contains the specifications for the MAC sublayer.

5.5 Functional overview

A brief overview of the general functions of a LR-WPAN is given in 5.5.1 through 5.5.6 and includes information on the superframe structure, the data transfer model, the frame structure, improving probability of successful delivery, power consumption considerations, and security.

5.5.1 Superframe structure

This standard allows the optional use of a superframe structure. The format of the superframe is defined by the coordinator. The superframe is bounded by network beacons sent by the coordinator (see Figure 4a) and is divided into 16 equally sized slots. Optionally, the superframe can have an active and an inactive portion (see Figure 4b). During the inactive portion, the coordinator may enter a low-power mode. The beacon frame is transmitted in the first slot of each superframe. If a coordinator does not wish to use a superframe structure, it will turn off the beacon transmissions. The beacons are used to synchronize the attached

devices, to identify the PAN, and to describe the structure of the superframes. Any device wishing to communicate during the contention access period (CAP) between two beacons competes with other devices using a slotted CSMA-CA mechanism. All transactions are completed by the time of the next network beacon.




Figure 4—Superframe structure without GTSs

For low-latency applications or applications requiring specific data bandwidth, the PAN coordinator may dedicate portions of the active superframe to that application. These portions are called guaranteed time slots (GTSs). The GTSs form the contention-free period (CFP), which always appears at the end of the active superframe starting at a slot boundary immediately following the CAP, as shown in Figure 5. The PAN coordinator may allocate up to seven of these GTSs, and a GTS may occupy more than one slot period. However, a sufficient portion of the CAP remains for contention-based access of other networked devices or new devices wishing to join the network. All contention-based transactions is completed before the CFP begins. Also each device transmitting in a GTS ensures that its transaction is complete before the time of the next GTS or the end of the CFP. More information on the superframe structure can be found in 7.5.1.1.

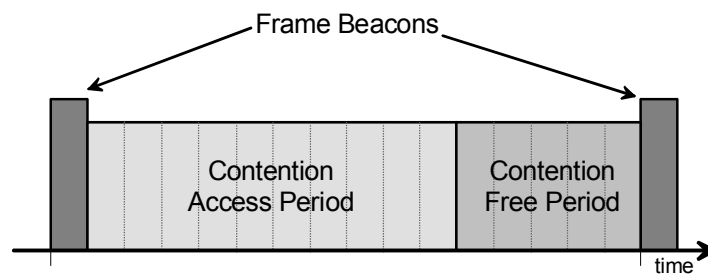


Figure 5—Superframe structure with GTSs

5.5.2 Data transfer model

Three types of data transfer transactions exist. The first one is the data transfer to a coordinator in which a device transmits the data. The second transaction is the data transfer from a coordinator in which the device receives the data. The third transaction is the data transfer between two peer devices. In star topology, only two of these transactions are used because data may be exchanged only between the coordinator and a device. In a peer-to-peer topology, data may be exchanged between any two devices on the network; consequently all three transactions may be used in this topology.

The mechanisms for each transfer type depend on whether the network supports the transmission of beacons. A beacon-enabled PAN is used in networks that either require synchronization or support for low-latency devices, such as PC peripherals. If the network does not need synchronization or support for low-latency devices, it can elect not to use the beacon for normal transfers. However, the beacon is still required for network discovery. The structure of the frames used for the data transfer is specified in 7.2.

5.5.2.1 Data transfer to a coordinator

When a device wishes to transfer data to a coordinator in a beacon-enabled PAN, it first listens for the network beacon. When the beacon is found, the device synchronizes to the superframe structure. At the appropriate time, the device transmits its data frame, using slotted CSMA-CA, to the coordinator. The coordinator may acknowledge the successful reception of the data by transmitting an optional acknowledgment frame. This sequence is summarized in Figure 6.




Figure 6—Communication to a coordinator in a beacon-enabled PAN

When a device wishes to transfer data in a nonbeacon-enabled PAN, it simply transmits its data frame, using unslotted CSMA-CA, to the coordinator. The coordinator acknowledges the successful reception of the data by transmitting an optional acknowledgment frame. The transaction is now complete. This sequence is summarized in Figure 7.




Figure 7—Communication to a coordinator in a nonbeacon-enabled PAN

5.5.2.2 Data transfer from a coordinator

When the coordinator wishes to transfer data to a device in a beacon-enabled PAN, it indicates in the network beacon that the data message is pending. The device periodically listens to the network beacon and, if a message is pending, transmits a MAC command requesting the data, using slotted CSMA-CA. The coordinator acknowledges the successful reception of the data request by transmitting an acknowledgment frame. The pending data frame is then sent using slotted CSMA-CA or, if possible, immediately after the acknowledgment (see 7.5.6.3). The device may acknowledge the successful reception of the data by transmitting an optional acknowledgment frame. The transaction is now complete. Upon successful completion of the data transaction, the message is removed from the list of pending messages in the beacon. This sequence is summarized in Figure 8.




Figure 8—Communication from a coordinator in a beacon-enabled PAN

When a coordinator wishes to transfer data to a device in a nonbeacon-enabled PAN, it stores the data for the appropriate device to make contact and request the data. A device may make contact by transmitting a MAC command requesting the data, using unslotted CSMA-CA, to its coordinator at an application-defined rate. The coordinator acknowledges the successful reception of the data request by transmitting an acknowledgment frame. If a data frame is pending, the coordinator transmits the data frame, using unslotted CSMA-CA, to the device. If a data frame is not pending, the coordinator indicates this fact either in the acknowledgment frame following the data request or in a data frame with a zero-length payload (see 7.5.6.3). If requested, the device acknowledges the successful reception of the data frame by transmitting an acknowledgment frame. This sequence is summarized in Figure 9.




Figure 9—Communication from a coordinator in a nonbeacon-enabled PAN

5.5.2.3 Peer-to-peer data transfers

In a peer-to-peer PAN, every device may communicate with every other device in its radio sphere of influence. In order to do this effectively, the devices wishing to communicate will need to either receive constantly or synchronize with each other. In the former case, the device can simply transmit its data using unslotted CSMA-CA. In the latter case, other measures need to be taken in order to achieve synchronization. Such measures are beyond the scope of this standard.

5.5.3 Frame structure

The frame structures have been designed to keep the complexity to a minimum while at the same time making them sufficiently robust for transmission on a noisy channel. Each successive protocol layer adds to the structure with layer-specific headers and footers. This standard defines four frame structures:

- A beacon frame, used by a coordinator to transmit beacons
- A data frame, used for all transfers of data
- An acknowledgment frame, used for confirming successful frame reception
- A MAC command frame, used for handling all MAC peer entity control transfers

The structure of each of the four frame types is described in 5.5.3.1 through 5.5.3.4. The diagrams in these subclauses illustrate the fields that are added by each layer of the protocol.

5.5.3.1 Beacon frame

Figure 10 shows the structure of the beacon frame, which originates from within the MAC sublayer. A coordinator can transmit network beacons in a beacon-enabled PAN. The MAC payload contains the superframe specification, GTS fields, pending address fields, and beacon payload (see 7.2.2.1). The MAC payload is prefixed with a MAC header (MHR) and appended with a MAC footer (MFR). The MHR contains the MAC Frame Control field, beacon sequence number (BSN), addressing fields, and optionally the auxiliary security header. The MFR contains a 16-bit frame check sequence (FCS). The MHR, MAC payload, and MFR together form the MAC beacon frame (i.e., MPDU).

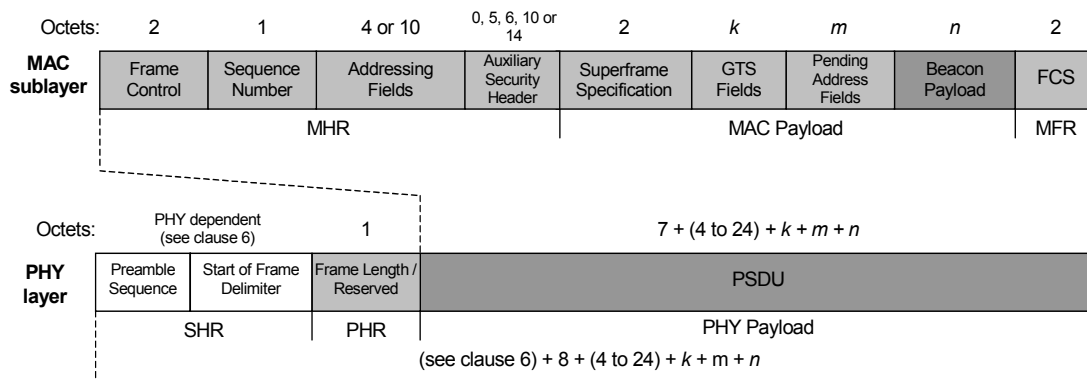


Figure 10—Schematic view of the beacon frame and the PHY packet

The MAC beacon frame is then passed to the PHY as the PHY service data unit (PSDU), which becomes the PHY payload. The PHY payload is prefixed with a synchronization header (SHR), containing the Preamble Sequence and Start-of-Frame Delimiter (SFD) fields, and a PHY header (PHR) containing the length of the PHY payload in octets. The SHR, PHR, and PHY payload together form the PHY packet (i.e., PPDU).

5.5.3.2 Data frame

Figure 11 shows the structure of the data frame, which originates from the upper layers.

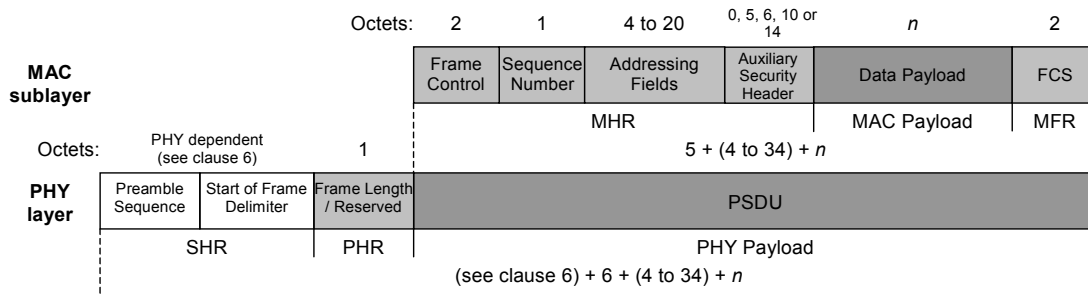


Figure 11—Schematic view of the data frame and the PHY packet

The data payload is passed to the MAC sublayer and is referred to as the MAC service data unit (MSDU). The MAC payload is prefixed with an MHR and appended with an MFR. The MHR contains the Frame Control field, data sequence number (DSN), addressing fields, and optionally the auxiliary security header. The MFR is composed of a 16-bit FCS. The MHR, MAC payload, and MFR together form the MAC data frame, (i.e., MPDU).

The MPDU is passed to the PHY as the PSDU, which becomes the PHY payload. The PHY payload is prefixed with an SHR, containing the Preamble Sequence and SFD fields, and a PHR containing the length of the PHY payload in octets. The preamble sequence and the data SFD enable the receiver to achieve symbol synchronization. The SHR, PHR, and PHY payload together form the PHY packet, (i.e., PPDU).

5.5.3.3 Acknowledgment frame

Figure 12 shows the structure of the acknowledgment frame, which originates from within the MAC sublayer. The MAC acknowledgment frame is constructed from an MHR and an MFR; it has no MAC payload. The MHR contains the MAC Frame Control field and DSN. The MFR is composed of a 16-bit FCS. The MHR and MFR together form the MAC acknowledgment frame (i.e., MPDU).

The MPDU is passed to the PHY as the PSDU, which becomes the PHY payload. The PHY payload is prefixed with the SHR, containing the Preamble Sequence and SFD fields, and the PHR containing the length of the PHY payload in octets. The SHR, PHR, and PHY payload together form the PHY packet, (i.e., PPDU).

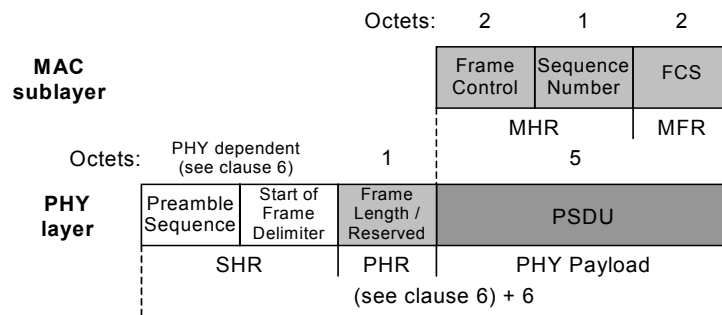


Figure 12—Schematic view of the acknowledgment frame and the PHY packet

5.5.3.4 MAC command frame

Figure 13 shows the structure of the MAC command frame, which originates from within the MAC sublayer. The MAC payload contains the Command Type field and the command payload (see 7.2.2.4). The MAC payload is prefixed with an MHR and appended with an MFR. The MHR contains the MAC Frame Control field, DSN, addressing fields, and optionally the auxiliary security header. The MFR contains a 16-bit FCS. The MHR, MAC payload, and MFR together form the MAC command frame, (i.e., MPDU).

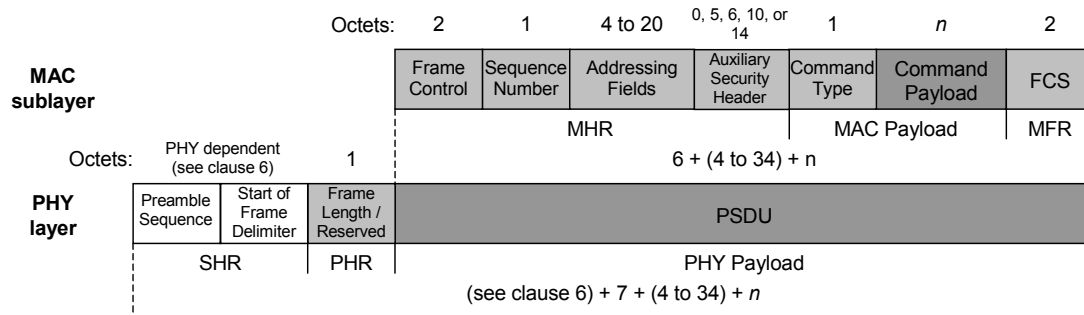


Figure 13—Schematic view of the MAC command frame and the PHY packet

The MPDU is then passed to the PHY as the PSDU, which becomes the PHY payload. The PHY payload is prefixed with an SHR, containing the Preamble Sequence and SFD fields, and a PHR containing the length of the PHY payload in octets. The preamble sequence enables the receiver to achieve symbol synchronization. The SHR, PHR, and PHY payload together form the PHY packet, (i.e., PPDU).

5.5.4 Improving probability of successful delivery

The IEEE 802.15.4 LR-WPAN employs various mechanisms to improve the probability of successful data transmission. These mechanisms are the CSMA-CA mechanism, frame acknowledgment, and data verification and are briefly discussed in 5.5.4.1 through 5.5.4.3.

5.5.4.1 CSMA-CA mechanism

The IEEE 802.15.4 LR-WPAN uses two types of channel access mechanism, depending on the network configuration. Nonbeacon-enabled PANs use an unslotted CSMA-CA channel access mechanism, as described in 7.5.1. Each time a device wishes to transmit data frames or MAC commands, it waits for a random period. If the channel is found to be idle, following the random backoff, the device transmits its data. If the channel is found to be busy following the random backoff, the device waits for another random period before trying to access the channel again. Acknowledgment frames are sent without using a CSMA-CA mechanism.

Beacon-enabled PANs use a slotted CSMA-CA channel access mechanism, where the backoff slots are aligned with the start of the beacon transmission. The backoff slots of all devices within one PAN are aligned to the PAN coordinator. Each time a device wishes to transmit data frames during the CAP, it locates the boundary of the next backoff slot and then waits for a random number of backoff slots. If the channel is busy, following this random backoff, the device waits for another random number of backoff slots before trying to access the channel again. If the channel is idle, the device begins transmitting on the next available backoff slot boundary. Acknowledgment and beacon frames are sent without using a CSMA-CA mechanism.

5.5.4.2 Frame acknowledgment

A successful reception and validation of a data or MAC command frame can be optionally confirmed with an acknowledgment, as described in 7.5.6.4. If the receiving device is unable to handle the received data frame for any reason, the message is not acknowledged.

If the originator does not receive an acknowledgment after some period, it assumes that the transmission was unsuccessful and retries the frame transmission. If an acknowledgment is still not received after several retries, the originator can choose either to terminate the transaction or to try again. When the acknowledgment is not required, the originator assumes the transmission was successful.

5.5.4.3 Data verification

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

The FCS mechanism is discussed in 7.2.1.9.

5.5.5 Power consumption considerations

In many applications that use this standard, devices will be battery powered, and battery replacement or recharging in relatively short intervals is impractical. Therefore, the power consumption is of significant concern. This standard was developed with limited power supply availability in mind. However, the physical implementation of this standard will require additional power management considerations that are beyond the scope of this standard.

The protocol has been developed to favor battery-powered devices. However, in certain applications, some of these devices could potentially be mains powered. Battery-powered devices will require duty-cycling to reduce power consumption. These devices will spend most of their operational life in a sleep state; however, each device periodically listens to the RF channel in order to determine whether a message is pending. This mechanism allows the application designer to decide on the balance between battery consumption and message latency. Higher powered devices have the option of listening to the RF channel continuously.

5.5.6 Security

From a security perspective, wireless ad hoc networks are no different from any other wireless network. They are vulnerable to passive eavesdropping attacks and potentially even active tampering because physical access to the wire is not required to participate in communications. The very nature of ad hoc networks and their cost objectives impose additional security constraints, which perhaps make these networks the most difficult environments to secure. Devices are low-cost and have limited capabilities in terms of computing power, available storage, and power drain; and it cannot always be assumed they have a trusted computing base nor a high-quality random number generator aboard. Communications cannot rely on the online availability of a fixed infrastructure and might involve short-term relationships between devices that may never have communicated before. These constraints might severely limit the choice of cryptographic algorithms and protocols and would influence the design of the security architecture because the establishment and maintenance of trust relationships between devices need to be addressed with care. In addition, battery lifetime and cost constraints put severe limits on the security overhead these networks can tolerate, something that is of far less concern with higher bandwidth networks. Most of these security architectural elements can be implemented at higher layers and may, therefore, be considered to be outside the scope of this standard.

The cryptographic mechanism in this standard is based on symmetric-key cryptography and uses keys that are provided by higher layer processes. The establishment and maintenance of these keys are outside the

scope of this standard. The mechanism assumes a secure implementation of cryptographic operations and secure and authentic storage of keying material.

The cryptographic mechanism provides particular combinations of the following security services:

- *Data confidentiality*: Assurance that transmitted information is only disclosed to parties for which it is intended.
- *Data authenticity*: Assurance of the source of transmitted information (and, hereby, that information was not modified in transit).
- *Replay protection*: Assurance that duplicate information is detected.

The actual frame protection provided can be adapted on a frame-by-frame basis and allows for varying levels of data authenticity (to minimize security overhead in transmitted frames where required) and for optional data confidentiality. When nontrivial protection is required, replay protection is always provided.

Cryptographic frame protection may use a key shared between two peer devices (link key) or a key shared among a group of devices (group key), thus allowing some flexibility and application-specific tradeoffs between key storage and key maintenance costs versus the cryptographic protection provided. If a group key is used for peer-to-peer communication, protection is provided only against outsider devices and not against potential malicious devices in the key-sharing group.

For more detailed information on the cryptographic security mechanisms used for protected MAC frames following this standard, refer to Clause 7.

5.6 Concept of primitives

This subclause provides a brief overview of the concept of service primitives. Refer to ISO/IEC 8802.2⁵ for more detailed information.

The services of a layer are the capabilities it offers to the user in the next higher layer or sublayer by building its functions on the services of the next lower layer. This concept is illustrated in Figure 14, showing the service hierarchy and the relationship of the two correspondent N-users and their associated N-layer (or sublayer) peer protocol entities.

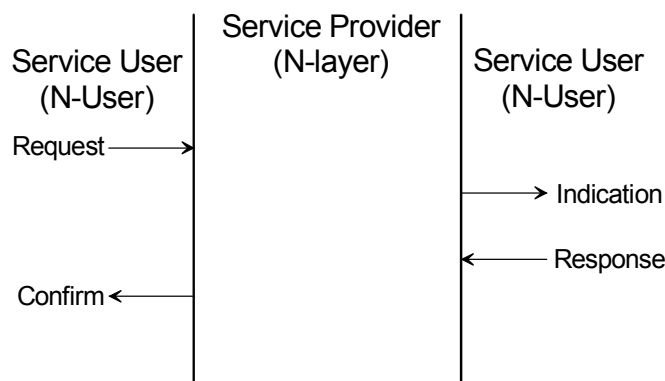


Figure 14—Service primitives

⁵For information on references, see Clause 2.

The services are specified by describing the information flow between the N-user and the N-layer. This information flow is modeled by discrete, instantaneous events, which characterize the provision of a service. Each event consists of passing a service primitive from one layer to the other through a layer SAP associated with an N-user. Service primitives convey the required information by providing a particular service. These service primitives are an abstraction because they specify only the provided service rather than the means by which it is provided. This definition is independent of any other interface implementation.

A service is specified by describing the service primitives and parameters that characterize it. A service may have one or more related primitives that constitute the activity that is related to that particular service. Each service primitive may have zero or more parameters that convey the information required to provide the service.

A primitive can be one of four generic types:

- *Request*: The request primitive is passed from the N-user to the N-layer to request that a service is initiated.
- *Indication*: The indication primitive is passed from the N-layer to the N-user to indicate an internal N-layer event that is significant to the N-user. This event may be logically related to a remote service request, or it may be caused by an N-layer internal event.
- *Response*: The response primitive is passed from the N-user to the N-layer to complete a procedure previously invoked by an indication primitive.
- *Confirm*: The confirm primitive is passed from the N-layer to the N-user to convey the results of one or more associated previous service requests.