

If DEV-2 is successful in joining the parent piconet, then the next problem that is unique to dependent PNC hand-over is handing over control of the CTA to the new dependent PNC. Because DEV-2 has passed its new DEVID to the current dependent PNC in the PNC Handover Response command, the current dependent PNC should not make a mistake in selecting the correct DEVID to send in the Channel Time Request command. However, the PNC may reject the handover of the CTA control if any of the following are true:

- The DEVID specified in the target ID list does not correspond to an associated DEV.
- The DEVID specified in the target ID list does not correspond either to a secure member of the piconet or a NbrID and secure membership is required in the parent piconet.
- The DEVID specified in the target ID list is in power-save mode.
- The PNC is terminating the stream.

The PNC should have time available to allocate the CTA because it is already allocating the same amount of time for the dependent piconet. If the parent PNC rejects the request to hand over control of the CTA to the new DEV-2 PNC, the dependent PNC will send a PNC Handover Request command to DEV-2 with the Handover Status field set to 1 to indicate that the hand-over process is being cancelled, as illustrated in Figure 5-9.

If the dependent PNC cancels the hand-over process, DEV-2 may choose to disassociate from the parent piconet because there is no longer a need for it to be a part of the parent piconet. If DEV-2 joined as a neighbor PNC, it should disassociate from the parent piconet when the hand-over process is cancelled to free up that resource for other DEVs that may want to form a neighbor piconet.




Figure 5-9: Hand-over of dependent PNC fails because of failure to hand over control of the dependent piconet CTA

Chapter 6 Security

Wireless networks are by their nature a broadcast medium. They provide no intrinsic PHY level security, at least for standardized solutions. Occasionally, people will mistake the use of a technology by the military, e.g., spread-spectrum or ultra-wideband, to indicate that the PHY has some sort of built-in security. However, the military does not publish interoperability guides to its secure systems, whereas that is exactly the purpose of a standard. Standardized PHYs allow anyone reasonably skilled in the art to successfully extract the data from the wireless medium. In fact, many wireless networking cards have a "promiscuous mode" wherein the radio demodulates all of the frames on the air and passes them up for further analysis.

Manufacturers and people in the press have erroneously claimed from time to time that standardized solutions for CDMA (IS-95) [B28], DSSS (802.11b) [B13], FHSS (802.11 and RangeLAN) [B12] or even ultra wideband (UWB) PHYs offer some sort of security inherent in the PHY layer signaling.²⁹ The reality is that none of these PHYs intrinsically offer even a modest level of security. Instead, the only strong security available to a wireless network is the tried and true method of using bit-level encryption techniques on the data. This is actually good news because the security provided by these techniques is well studied and understood by experts in the field. Very strong security can be achieved with a minimal impact in terms of logic gates and power in an IC implementation, even at data rates as high as 500 Mb/s.

INTRODUCTION AND HISTORY

The security in 802.15.3 is based on three capabilities:

- a) *Authentication and authorization*: The ability to determine the identity of a DEV and decide if it should be allowed in the piconet.

²⁹ Quantum cryptography that can actually detect when the PHY layer signaling is modified is an exception to this rule. It may be that PHYs based on these techniques will offer security that is inherent in the PHY signaling.

- b) *Data integrity*: The ability to determine the true source of the data and to verify that the data was not changed once it had been sent.
- c) *Data privacy*: The ability to prevent unauthorized DEVs from gaining access to data that is sent over the air.

The latter two capabilities are provided by the symmetric key encryption suite (see "Symmetric key security suite" on page 195), which provides MAC level encryption and/or integrity checking of the data.³⁰ The authentication and authorization portion of the draft caused the biggest problem with the development of the draft as well as some of the most contentious debates at meetings and on

the e-mail reflector. The initial draft released for working group ballot, D09 (i.e., the ninth version of the draft), did not include any security suites to enable authentication (or even encryption) but rather included variable length frames that could be used by a security process, e.g., some form of public-key cryptography, to implement authentication. The actual process for performing the authentication was determined to be out of scope.³¹

However, in the first working group ballot, a sizeable percentage of the voters (about 12%) requested that the draft include a specific authentication process to enable interoperability among consumer electronic devices. To resolve the comments of these voters, the task group began a process of a call for proposals and a downselection process to select a single public-key cryptography security suite and symmetric key suite for inclusion in subsequent versions of the draft. Two public-key suites were proposed, one based on elliptic curve cryptography (ECC) and digital certificates proposed by individuals from Certicom™ and Motorola® and another one based on the short lattice vector problem proposed by



The final draft of the standard was numbered D17 and so one might surmise that there were 17 released versions of the draft. In fact, there were 21 released versions of the draft. Following D09, the draft versions adopted hexadecimal notation. This provided some humor as the next draft was appropriately named D0A (i.e., Dead On Arrival). This was followed by D0B, D0C, D0D and then D10 and the rest followed numerically. (Versions D0E and D0F were skipped.)

³⁰ The Integrity Code field in 802.15.3 is normally referred to as a message authentication code or MAC in cryptographic circles. However, the acronym MAC is already being used by the standard in a few places, so IEEE Std 802.15.3 refers to this field as the IC instead.

³¹ This was in November 2001; this same conclusion was to be reached in January 2003 when the draft comes full circle on this issue with draft D16.

individuals from Ntru™. In addition, two symmetric key security suites were proposed, the advanced encryption standard (AES-128) [B22] and another based on triple data encryption standard (triple-DES).

The selection of the symmetric key suite was relatively straightforward and uncontentious, and AES-128 was selected to provide the symmetric key services. Likewise, the group felt strongly that digital certificates should be an option and not a requirement for the standard. The public-key security suites resulted in a much more difficult selection process. The task group was concerned because the encryption algorithm used by Ntru was patented as were certain protocols often used to implement ECC public-key cryptography. With the possibility of required royalty payments on the line for both sides, the task group selected ECC but failed to get the 75% approval³² required to go forward. The group worked out a compromise; the technical editor would put in an ECC security suite that had the least amount of essential IP into the next draft as the mandatory security suite, and Ntru would be included as an optional security suite. A due date was provided for the submission of the ECC suite, and if a complete solution was not provided to the technical editor by this date, the Ntru security suite would become the mandatory security suite.

Two proposals for an ECC security suite were received by the deadline, and a working group letter ballot had to be held to determine the one that would be used in the draft, delaying the release of the draft and the start of another working group letter ballot for at least two months. Based on the result of the working group ballot, one of the two ECC suites was put into the draft as mandatory for the next recirculation ballot. This, predictably, resulted in more no votes that had to be resolved. Finally, a compromise was struck. Three public-key security suites would be provided, ECMQV Koblitz-283[B27], NTRUEncrypt 251 [B3], and RSA-OAEP 1024 [B4], all optional for a DEV to implement. The only restriction was that if a DEV implemented security, it had to support at least one of the three security suites.

Although it seemed that this would finally settle the question of security suites in the standard, it came to light during sponsor ballot that authentication and authorization were indeed outside of the scope of a MAC and a PHY. In fact, the

³² The motion failed by one vote.

security suites made mention of the fact that authorization and even some parts of the authentication protocol, such as digital certificates, were outside of the scope of the standard. Rather than risk the standard being sent back from either the IEEE 802 Executive Committee or the IEEE-SASB Standards Review Committee (RevCom), the task group decided to remove the authentication portion of the security suites and retain only the frame formats with variable length fields that would support any type of authentication protocol.

This exercise was not entirely futile. The process of selecting a security suite provided a focus on the effects of security on the piconet. The process also helped to solidify the types of frames required for authentication as well as defining the methods used to update keys in a secure relationship. IEEE Std 802.15.3 also added peer-to-peer security to the group level security that was first proposed.

Finally, IEEE Std 802.15.3 led the way by adopting counter plus chaining block cipher message authentication code (CTR + CBC-MAC or CCM) mode for AES-128 using a single key. This was later adopted by both 802.15.4 and 802.11i as the mandatory encryption algorithm for those standards. Overall, the process of selecting, reviewing, and evaluating the security suites helped to improve the overall quality of the security services in the 802.15.3 draft standard. It also provided an opportunity for the members of the 802.15.3 task group to become better educated on the implementation considerations for security.



One of the consequences of purging the security suites from the the standard is that the security terminology was changed. For example, anywhere that the standard says "secure membership" it really means "authenticated and authorized." Likewise, the "key originator" is really the "security manager." The Security Message command is intended to be used for 3 and 4 pass authentication protocols and the Security Information command is intended to carry the ACL and the associated authentication information.

SECURITY MODES AND POLICIES

The 802.15.3 standard operates in one of two security modes, mode 0 and mode 1. When the piconet is operating in mode 0, no security is provided and the piconet is wide open. In this mode, individual DEVs may form secure relationships or the higher layers may form secure links, e.g., a virtual private network (VPN), to protect the privacy and integrity of the data. Piconet-wide information such as the

beacon or commands sent to or by the PNC are not protected in any fashion in mode 0. In mode 1, all commands and the beacon have their integrity protected and all data are encrypted to provide privacy and have their integrity protected. All of the DEVs in the piconet have access to same group key, so each DEV in the piconet will have to evaluate if this level of security is acceptable for its requirements.

The security model for 802.15.3 has a DEV that acts as the key originator (KO) for other DEVs in that security relationship. The key originator keeps track of the keys and forces an update to them whenever required by a security event. When the piconet is operating in mode 1, the PNC acts as the key originator for that group security relationship.

Security services provided in mode 1

The following security services are provided in mode 1:

- a) *Key transport*: The key originator in a relationship can update the symmetric keys that are used, and the DEVs in the relationship can also request new keys if they detect that the encryption keys have changed.
- b) *Data encryption*: This service provides privacy for the data that are transmitted in the piconet. If the piconet group key is used, then every DEV in the piconet will be able to read and understand the data. The data can also be protected using a key that is only shared between two DEVs.
- c) *Data integrity*: Frames that require data integrity use an Integrity Code field that is part of the frame format to protect the data from modification and to verify the source of the data.
- d) *Beacon integrity protection*: The information in the beacon is not encrypted because it is used by DEVs that are searching for piconets to join. The information in the beacon is protected with an Integrity Code field to protect the information from modification and to verify that it was sent by the PNC.
- e) *Freshness protection*: Every beacon sent by a PNC includes a time token that is used by the encryption algorithm as one of the inputs used to create the nonce. The time token is included in the beacon regardless of the security mode. This enables peer-to-peer security in a piconet that is operating in mode 0.

- f) *Command integrity protection*: Like the beacon, commands sent in a piconet operating in mode 1 include an Integrity Code field to verify that the command has not been modified and to verify the identity of the source of the command. Most of the commands sent in the piconet are protected. Exceptions to this are for commands that are sent before a DEV can obtain secure membership, e.g., Association Request command, Association Response command, and Disassociation command.

The following security services are outside of the scope of IEEE Std 802.15.3 and so are not provided with mode 1:

- a) *Access control list*: The access control list provides a method for the key originator to determine which DEVs are authorized to operate in the piconet.
- b) *Mutual authentication*: This process verifies the identity of the parties in a security transaction. This process will provide the management keys required in the security relationship to be able to update the symmetric keys.
- c) *Key establishment*: Generally, this will be a part of the authentication process. In the process of authenticating with another DEV, the DEVs will create a shared secret that can then be used to generate the management keys for the relationship. Because this is tied to the authentication process, it is outside of the scope of the standard.
- d) *Authorization*: This is one of the most important security services and it is normally provided by the end user at the application level. In certain applications, e.g., a home theater system, it may be possible for the manufacturer to pre-program the authorizations required to enable secure operation of the DEVs.

The security service listed in item d), authorization, is one that has never really been within the scope of the standard. Although it is possible to provide automatic methods to authenticate a DEV using digital certificates, this still does not provide any information regarding authorization. For example, a person purchases an extra pair of wireless speakers for their home audio system. When the user turns on the speakers, they provide a digital certificate to the audio system that proves that they were produced by "Manufacturer A" and provides some sort of identification, i.e., the MAC address. However, it is not possible for the audio

system to know that these speakers are indeed the speakers that the person just picked up at their local electronics store and are not, for example, the speakers in the next-door neighbor's house that happen to be made by the same manufacturer. At some point, a person will need to intervene in some fashion to declare that the new DEVs are either allowed in the piconet or that they will be refused membership.

There are a variety of methods that can be used to authorize a DEV for operation in a given piconet. In a cordless phone, the handset often uses the physical contacts of the battery charger to send information that authorizes the handset to work with that base station. This is a sufficient level of security because if the malicious user could gain physical control of both the handset and the base station, the physical security of the home would have already been severely compromised. Another method for authorization is for the audio system to display a unique number sent by the new DEV that is also provided in the documentation for the product. The user just needs to verify that the number matches what is printed on the product in order to authorize the DEV.

Security policies

An important part of the 802.15.3 security architecture is the security policies that are required to protect the information in the piconet. These policies dictate specific security-related actions in response to a security event in the piconet. One of the key security policies is to make sure that the DEVs that are supposed to be members of the security relationship are still participating in the piconet. This is similar to the ATP being used to determine the DEVs that are still participating in the piconet. If the key originator has not received a frame of some sort from the DEV within that time period, it will deauthenticate the DEV. If the key originator is the PNC, the time period is the ATP and the result of a time out will be a disassociation, which also includes a logical deauthentication.

If the membership of the group changes, i.e., a DEV joins or leaves the group, the key originator is required to change the symmetric keys. The KO signals this change by changing the security ID (SECID) used in its frames and then uses the Distribute Key Request command to send the new keys to the DEVs in this security relationship. The SECID is used in the frame to indicate the key that is to be used to either check the data integrity or to decrypt the frame. If the KO is the

PNC, then the events that change the group membership are the association and disassociation procedures. Every time that the PNC either associates a new DEV or disassociates an older one, it must change the keys and the SECID.

PNC handover also affects the makeup of the security relationship because the PNC is also the key originator for the piconet. After the PNC passes over control to the new PNC, all of the DEVs in the piconet will be required to establish a secure relationship with the new PNC. The DEVs need to do this so that they can get a copy of the management key that will be used for subsequent key exchanges. If the DEVs that are already members of the piconet have entries in the ACL, then the new PNC can use the ACL to determine if a DEV should be added to the piconet without necessarily needing input from the user.

As a part of the 802.15.3 security policies, there is a list of all of the possible frames and the keys required to send these frames. The two-and-a-half page table is not replicated here, but the rules in the table need to be followed in order to maintain the security of either the piconet or the security relationship. As mentioned earlier, some of the command frames are always sent with security turned off. In addition, a DEV may send a data in a non-secure data frame while the piconet is operating in Mode 1. One of the potential applications for this is streaming video. The content providers will likely provide end-to-end encryption of the data stream that starts and ends at the application layers. Encrypting this data, which are already encrypted, is a waste of clock cycles and octets in the frame. Instead, the DEV can send these data frames with security turned off to enhance the efficiency of the network.

When a DEV receives a secure beacon, it will first check to see that the time token in the beacon is strictly greater than the previous valid time token and is less than that value plus `mMaxTimeTokenChange`, a parameter with value of 65,535. The DEV will accept time tokens in that range because it may have missed some of the beacon frames due to interference or because the DEV was sleeping. If the SECID matches the DEV's current SECID for the piconet group key and the beacon did not generate an integrity check failure, then the DEV will update its value for the last valid time token. Because the time token is used to preserve the freshness of commands in the piconet, it is very important that the DEV not get too far out of sync with the PNC in its expectation for the time token number.

A DEV performs similar checks for the command frames that it receives, except that it uses the time token that is correct for the current superframe, even if the DEV did not correctly receive the beacon. The time token is sent only in the beacon and is not included in the secure frame fields.

SYMMETRIC KEY SECURITY SUITE

The data privacy and data integrity capabilities in this standard are provided by a symmetric key encryption suite. The encryption algorithm used is AES with 128-bit keys and 128-bit blocks. The AES algorithm is the result of a worldwide competition and downselection for a new encryption standard for the U.S. that could be used across a variety of applications. Accordingly, the algorithm was analyzed not just for its security properties, but also for the efficiency of its implementation. The AES algorithm has been subjected to intense scrutiny by the cryptographic community and is considered to be a very reliable algorithm.

However, as important as the algorithm is, equally important is the mode in which the algorithm is used. AES is often proposed to be used in open code book (OCB) mode. AES OCB has been proven to be secure, but it is also a patented mode of operation. The 802.15.3 Task Group was concerned with the potential cost of a patented solution, so they pursued a mode of operation that is thought to be IP free.³³ This led to the selection of the counter plus chaining block cipher message authentication code (CTR + CBC-MAC or CCM) mode.

Overview of AES CCM

The AES CCM mode is described in detail in the standard, but a few key features will be discussed here. The CCM mode uses the same encryption engine to calculate the Integrity Code field, which can save on design costs. However, to handle high-speed data, it may be necessary to have parallel encryption engines to quickly calculate the cipher text and the Integrity Code field.

The authentication operation begins with the generation of an Integrity Code from a nonce using the block cipher in CBC mode. This is followed by the optional

³³ It is very difficult, if not impossible, for anyone to guarantee that something is IP-free. The author is not providing any assurances as to the IP status of AES CCM, and implementers need to perform their own analysis.

padding authentication data followed by an optional padded plaintext data. At the receiving side, the same Integrity Code is generated and is then compared to the Integrity Code that is sent with the frame.

Note that the padding only occurs in the generation of the Integrity Code; the actual plaintext is not padded before it is encoded into ciphertext. The key here is that the ciphertext is always the same length as the plain text. If this was not the case, then an additional length field would have been needed in the secure frame formats to indicate the length of the actual plain text.

The selection of the nonce is important to the security characteristics of the symmetric algorithm. The key is to find a nonce that will not have a chance of being repeated before the keys can be changed. The 802.15.3 nonce consists of the SrcID, DestID, Time Token, Secure Frame Counter, and the Fragmentation Control field. The Time Token is 6 octets long and is

incremented with every beacon. Even with a 1-ms superframe duration, the time token will roll over only once in every 8,925 years, which should be long enough. However, to ensure that the nonce is unique for every frame sent in a superframe, each secure frame includes a secure frame counter that is incremented for every frame that is sent in the superframe. This allows a DEV to send up to 65,536 frames in a superframe for each DestID before it will need to stop to avoid reusing the nonce. The use of the SrcID allows multiple DEVs in the same piconet to use the same key without running the risk of two frames using the same nonce. The secure frame counter is incremented every time a secure frame is transmitted, even if it is a retransmission.

When the Integrity Code for a frame is calculated, the entire 10 octets of the MAC frame header is included in the calculation as is the SECID, secure frame counter, and the payload of the frame. If the frame that is going to be protected is an



There are two types of cryptographic keys defined in the standard: the management key and the data key. While both keys are symmetric keys, the management key is used only to encrypt new data keys for distribution to the DEVs in the security relationship. Each pair of DEVs in a security relationship has a unique set of management keys. On the other hand, within a security relationship, many DEVs might use the same data key. When a DEV leaves a security relationship, the security manager changes the data key using the management key. This prevents the DEV that left from gaining access to the new data key or any data protected by it.

encrypted data frame, the Integrity Code is computed over the encrypted data, not over the plaintext. If the Integrity Code was calculated over the plain text, then an attacker could instantly determine that a crack had succeeded when the Integrity Code check of the unencrypted data passed. Instead, the Integrity Code only protects information that is already available on the air and so does not give any information regarding the plain text or other secrets.

Key distribution

There are three types of key distribution supported by the standard:

- a) The PNC distributing keys to DEVs in the piconet
- b) A peer key originator distributing keys to DEVs in the security relationship
- c) A DEV requesting the latest key from the key originator

The simplest key distribution protocol is when the PNC needs to deliver new keys to the DEVs in the piconet. This process is illustrated in Figure 6-1.

When a peer key originator distributes a new key to another DEV in the security relationship, it requires a response from the DEV. This response has the effect of being a secure ACK, which is not otherwise provided in the standard. The frame exchange sequence for a peer key originator distributing a key is illustrated in Figure 6-2.

Finally, if a DEV determines that the SECID in the security relationship has changed, the DEV will need to request a new key from the key originator, as illustrated in Figure 6-3. A DEV can determine that the SECID changed for the piconet-wide key by an integrity check failure on the beacon along with a new SECID in the beacon. However, the DEV cannot be sure that the SECID really changed until it requests the new key from the key originator. Because the Integrity Code check failed, the DEV cannot be sure that the PNC actually sent the frame. In a peer-to-peer security relationship, the DEV determines that the SECID has changed when it receives a frame that uses a different SECID and also creates an integrity check failure.

Whenever the SECID changes, the keys change and so the DEV should get Integrity Code check failures whenever there is a new SECID in any of its security relationships.




Figure 6-1: Frame exchange for the PNC distributing a new key to a DEV in the piconet




Figure 6-2: Frame exchange for peer key originators distributing a new key to a DEV in the piconet




Figure 6-3: DEV requesting a new key in response to a change in the SECID of the relationship

SECURITY INFORMATION

The key originator will, in general, maintain an access control list to determine the DEVs that will be allowed to join a secure piconet. When the PNC is the key originator, this also implies determining which DEVs will be able to participate in the piconet as a DEV. Because the ACL includes not only DEVs that are currently in the piconet, but also DEVs that are allowed to join the piconet but are not currently members of the piconet; the list can be quite large. Also, the list is based on the DEV addresses rather than the DEVID because the DEVID can change as DEVs enter and leave the piconet.

When the standard was changed to remove all references to authentication, all references to an ACL were removed as well. Instead, the Security Information Request command and the Security Information Response command were created to facilitate the exchange of security information, in particular to enable DEVs to exchange the ACL. The security information can be retrieved from the key originator either for one DEV at a time or by requesting all of the entries that are maintained by the key originator. To receive a single security information entry from the key originator, a DEV sends the Security Information Request command

to the key originator with the Queried DEVID set to the ID of the DEV whose security information is being requested. The key originator will then respond with the security information entry that corresponds to the Queried DEVID. Note that in this case, the DEV is only allowed to request a security information entry for a DEV that is currently part of the piconet. The frame exchange sequence for requesting a single security information entry is illustrated in Figure 6-4.




Figure 6-4: Frame exchange for a DEV requesting a single security information entry from the key originator

When the current PNC hands over control of the piconet to another DEV, it may also want to hand over the ACL. One reason is that the authorization step will generally require the intervention of the user. Requiring this intervention goes against the idea that the hand-over process is seamless and does not interrupt activity in the piconet. The actual authorization protocol may involve bringing the DEVs into close proximity or even touching. If the user had to perform this procedure every time the PNC handed over, no one would want a secure 802.15.3 solution.

The new PNC requests the security information from the current PNC by sending an Security Information Request command to the PNC with the Queried DEVID field set to be the BcstID. The PNC will then respond with a Security Information commands that contain all of the security information records that the PNC is currently managing. The frame exchange sequence for requesting and receiving security information records from the PNC is illustrated in Figure 6-5.



The information in the Security Information command can be split across multiple commands. However, the command itself cannot be fragmented. Each Security Information command contains a sequence number as well as a field that indicates the total number of frames that will be sent.




Figure 6-5: Frame exchange for the new PNC requesting all of the security information entries from the PNC

INTEL

Cha

OVE

The 2
achie
result

The
for
ope
in o
thei
allo

An
pro
5 C
pro
res
poi
ho
pa

Chapter 7 2.4 GHz PHY

OVERVIEW

The 2.4 GHz PHY uses five modulation formats with an 11 Mbaud symbol rate to achieve scalable data rates from 11 to 55 Mb/s. The modulation types, coding, and resulting data rates are summarized in Table 7-1.

Table 7-1: Modulation, coding, and data rates for 2.4 GHz PHY

Modulation type	Coding	Data rate
QPSK	8-state TCM	11 Mb/s
DQPSK	None	22 Mb/s
16-QAM	8-state TCM	33 Mb/s
32-QAM	8-state TCM	44 Mb/s
64-QAM	8-state TCM	55 Mb/s

The 2.4 GHz PHY was chosen over 5 GHz and ultra-wideband (UWB) proposals for a variety of reasons. The 2.4–2.4835 GHz band is available for unlicensed operation in most of the countries in the world. Although progress is being made in opening up the 5 GHz bands worldwide, there were still many limitations on their use in Europe and Asia. The allocations for UWB have only recently been allowed in the US and are still under consideration in other parts of the world.

Another inherent advantage of using the 2.4 GHz unlicensed band is that it provides a greater range for the same transmit power and receiver design than for 5 GHz radios. Although it is possible to add higher gain antennas to 5 GHz products to extend their range, as is sometimes done with 802.11a products, this results in a decrease in coverage in other areas. For example, some 802.11a access points (APs) use high gain dipole antennas that extend the range of the AP in the horizontal direction. However, this reduces the coverage in the vertical direction, particularly between floors in typical multi-story houses. In the case of UWB, the

lower transmit power available as well as the higher frequency range used for UWB (the FCC has allowed intentional emissions for UWB radiators in the 3.1–10.1 GHz range) greatly limits the range that is possible. In particular, at those frequencies and power levels, UWB systems will be limited to use only in a single, small room because of their short range and inability to penetrate walls.

The 2.4 GHz PHY, on the other hand, will easily be able to cover a typical single family home in the US. Using the 802.11b channel model for the 2.4 GHz band, the rate vs. range for an 802.15.3 radio is graphed in Figure 7–1 for both a low-power transmitter (+0 dBm) and a higher power (+14 dBm) one.




Figure 7–1: Rate vs. range for high and low power 802.15.3 2.4 GHz PHY

The calculations in Figure 7–1 use the 802.11 channel model (see “Assumptions for coexistence simulations” on page 278), a noise figure of 7 dB, an implementation loss of 4 dB, and 2 dB gain antennas at both the transmitter and the receiver. The noise figure is typical for 802.11b single-chip radios currently on

the market. The sensitivity measured at the antenna connector (i.e., excluding the receiver antenna gain) for this configuration is summarized in Table 7-2. A +14 dBm transmit power is typical for 802.11b radios. A power amplifier capable of this type of transmit power uses less than 150 mA at 3 V (for example, the Maxim® 2242 or the SiGe semiconductor PA2423MB). For the 0 dBm case, no power amplifier is required.

Table 7-2: Receiver sensitivity for typical 802.15.3 implementation^a

Modulation	Data rate	Implementation sensitivity
QPSK-TCM	11 Mb/s	-88 dBm
DQPSK	22 Mb/s	-81 dBm
16-QAM	33 Mb/s	-80 dBm
32-QAM	44 Mb/s	-77 dBm
64-QAM	55 Mb/s	-74 dBm

^a The standard has relaxed sensitivity requirements, these number represent values for a typical implementation.

Another advantage of a 2.4 GHz radio over 5 GHz or UWB radios is that RF-CMOS and SiGe designs for low-cost, single-chip 2.4 GHz RFICs are available from many sources. The combination of hype for Bluetooth and the brisk sales of 802.11b and 802.11g have encouraged a number of companies to invest in creating these RFICs. In the near future, 5 GHz RFICs will be widely available in low-cost processes from many different sources. Eventually, this will likely be true for UWB ICs, although it will likely take much longer than the ramp-up of 5 GHz technology because high-volume, low-price applications are a brand new market for UWB techniques.

GENERAL PHY REQUIREMENTS

Channel plan

The 2.4 GHz PHY has two channel plans depending on the environment in which it is operating. The center frequencies and channel IDs (CHNL_ID) are listed in

Table 7-3. The high-density channel plan provides four nonoverlapping channels for use when 802.15.3 is the main network in use. If, on the other hand, an 802.11 network is present, particularly if it is occupying 802.11b channel 6 (2.437 GHz center frequency), then the 802.15.3 network will work best with the 802.11b coexistence channel plan. The reason for this is that if the 802.15.3 network used channel 2, it would overlap with an 802.11b network in both channel 1 or channel 6 (802.11b channel numbers). Likewise, if the 802.15.3 network used channel 4, it would interfere with 802.11b networks and be subject to interference from 802.11b networks in both channel 6 and channel 11 (802.11b channel numbers).

Table 7-3: 2.4 GHz PHY channel plan

CHNL_ID	Center frequency	High density	802.11b coexistence
1	2.412 GHz	X	X
2	2.428 GHz	X	
3	2.437 GHz		X
4	2.445 GHz	X	
5	2.461 GHz	X	X

On the other hand, if the 802.15.3 network adopts a channel plan that most closely matches the 802.11b channel plan, it would overlap with at most one channel. If only two of the three 802.11b channels are being used, then the 802.15.3 PNC would search for the best channel, from an interference point of view, and move the piconet to that channel. IEEE Std 802.15.3 requires that if a DEV is capable of detecting an 802.11b network, it needs to adopt the 802.11b coexistence channel plan if it detects an 802.11b network and move the piconet away from channels that contain 802.11b WLANs.

Timing issues

In any radio systems, there are constraints on the timing that arise from the physical limitations of the actual design; e.g., the length of time it takes to change the radio from transmit mode to receive mode. The key timing parameters for the 2.4 GHz PHY are described in Table 7-4. The interframe spacings (IFS's) are used by the MAC in synchronizing with the piconet. The acronyms stand for

minimum IFS (MIFS), short IFS (SIFS), backoff IFS (BIFS), and the retransmission IFS (RIFS).

Table 7-4: Timing parameters for the 802.15.3 2.4 GHz PHY

PHY parameter	MAC parameter	Value	Description
Channel switch time	N/A	500 μ s	The time required to change physical layer channels.
pPHYMIFSTime	MIFS	2 μ s	The time required between successive transmissions.
pPHYSIFSTime	SIFS	10 μ s	The time required to switch from transmit to receive and vice versa.
pCCADetectTime	pBackoffSlot	7.3 μ s	The time required to determine that the channel is clear.
pPHYSIFSTime + pCCADetectTime	BIFS	17.3 μ s	The minimum time required to determine that channel is clear and turn around to begin transmitting.
2*pPHYSIFSTime + pCCADetectTime	RIFS	27.3 μ s	The amount of time that needs to expire before the source can send a retransmission in the CTAP.

The channel change timing should be fairly easy to achieve. For example, a Bluetooth radio is required to change frequencies within 224 μ s, whereas the channel change requirement for 802.15.3 is twice as large. The only time this number comes into play is when the PNC changes the operating channel of the piconet. The PNC knows that every DEV in the piconet can change to the new channel within 500 μ s. Thus, in the last superframe in the old channel, the PNC should not put in any CTAs in the last 500 μ s of the superframe so that all of the DEVs have time to switch to the new channel. If the PNC did provide CTAs during this time, the source and destination DEVs would not use that time for communications, but rather would spend the time switching to the new channel. Because channel change is done very infrequently, this time does not affect the overall throughput of the system.

On the other hand, the IFSs occur in essentially every superframe and so they have a direct impact on the efficiency and throughput of the system. The most

important parameter for this is the SIFS. This time is used between every CTA, for every Imm-ACK and Dly-ACK frame, in the backoff algorithm, and in the retransmission timing. Therefore, the PHY subcommittee of TG3 wanted to keep this time as short as possible without creating an unreasonable burden on the implementer. Since the 802.15.3 PHY has many characteristics in common with 802.11b, the subcommittee decided to use the same SIFS as 802.11b: 10 μ s. The receiver has to be ready before the transmitter begins sending the frame so the TX to RX turnaround is strictly less than the SIFS while the RX to TX turnaround is no less than the SIFS but less than the SIFS plus 1 μ s. In the 2.4 GHz PHY description, Clause 11 of the standard, the SIFS time is also referred to as the pPHYSIFSTime.

The MIFS, like the SIFS, has a direct impact on the throughput of the system. Ideally, the radio is able to start transmitting the next frame as soon as it is done with the transmission of the previous frame. In a practical sense, the radio needs a finite amount of time to complete the transmission/reception of one frame before it begins the transmission/reception of the next one. Figure 7-2 shows the percent difference between using a finite time for the MIFS vs. using a zero time for the MIFS as a function of the data rate. The frame has a 2044-octet payload, and the calculation includes the overhead for the PHY preamble, PHY header, MAC header, HCS, FCS, as well as MIFS spacing. Because the MIFS is only used for frames with no-ACK policy, the time for an Imm-ACK or Dly-ACK frame is not included.

In Figure 7-2, the graphs show the amount of throughput lost due to having a non-zero MIFS. Therefore, lower numbers in the graph are better because the goal is to have the minimum amount of overhead. When the MIFS is set at 2 μ s, the overhead for the MIFS is less than 2.5% for all data rates. However, for the 10- μ s and 20- μ s cases, the throughput lost to the MIFS can be as high as 16% at higher data rates. Although the current PHY does not support data rates this high, it is anticipated that a new PHY will be available in 3 to 4 years that could reach rates as high as 400 Mb/s. These reductions in the throughput are in addition to the overhead already in IEEE Std 802.15.3, and so any increase in the overhead is undesirable.




Figure 7-2: Percent throughput lost due to a non-zero MIFS

The effect of the MIFS value on the throughput of the data rates of the 2.4 GHz PHY is shown in Figure 7-3. The data payload is assumed to be the maximum allowed for this PHY, 2044 octets. For the lowest data rate, 11 Mb/s, the effect of the MIFS on the overall throughput is very low, less than 1% for a 10-μs MIFS and only slightly greater even for a 20-μs MIFS. The throughput decrease increases as the data rate increases. At 55 Mb/s, the difference between using a MIFS of 10 μs vs. 2 μs is about 2.0%, as shown in Figure 7-3. This is an additional 1.1 Mb/s of throughput. Although this improvement is not spectacular, it is important not to discard even small improvements in throughput if they can be implemented in a relatively simple fashion. From the PHY point of view, supporting a 2-μs MIFS is not any more difficult than supporting the TX ramp, which is of the same order. In fact, for a 2-μs MIFS, the transmitter does not get a chance to completely shut off, rather the ramp-down and ramp-up, each of which is 2 μs in duration, overlap in the interframe space. In this case, the input to the

transmitter would only be idled between frames and the rest of the TX chain would be left on. The receiver would simply be left on so that it could receive the next frame. As the environment changes slowly relative to the frame duration, keeping the AGC setting from the previous frame would enable fast acquisition of the next frame.




Figure 7-3: Percent throughput lost for each of the supported data rates at various MIFS durations

The MIFS affects the throughput less for larger data frames than for smaller data frames. Figure 7-4 shows the percent difference in throughput at 55 Mb/s due to a change from a 10-μs MIFS, i.e., having the MIFS equal to the SIFS, to a 2-μs MIFS. For larger data payload sizes, the advantage of decreasing the MIFS is diminished. Although the most efficient data transfer in an error-free environment is to use the longest possible frames, this is not true in the error-prone wireless environment. To keep the FER low, the application may use smaller frame sizes because a smaller frame size reduces the FER for a given BER. The improved FER can actually improve the overall throughput even though the overhead is

higher for the smaller frames. The application may use smaller data fragments as well, and so it would not be possible for the MAC to use the maximum frame size. For smaller data payloads, e.g., 512 octets, the improvement in throughput is 5.5% or about 3 Mb/s.




Figure 7-4: Percent throughput increase from changing the MIFS duration from 10 μ s to 2 μ s for various data payload sizes

Using the CAP for random access requires that the DEVs are able to determine if a frame is being sent over the air. Because the first part of a frame that is being transmitted is the CAZAC sequence, the standard requires that the DEVs can recognize the PHY preamble within five repetitions of the CAZAC sequence. This time, referred to as the pCCADetectTime, is therefore $5 \times 16/11$ or about 7.3 μ s.

The standard also defines two other interframe spacings, the backoff IFS (BIFS) and the retransmission IFS (RIFS). The BIFS is used in the backoff algorithm to determine if another DEV is transmitting a frame. Once a DEV has determined that it is time for its turn to access the medium in the CAP, it needs to change from

receive to transmit and begin transmitting its frame. The other DEVs also need time to hear the start of the frame that is being transmitted. Thus, the BIFS is the sum of the SIFS time and the CCA detection time, i.e., $BIFS = pPHYSIFSTime + pCCADetectTime$, or about 17.3 μs . The BIFS is also used at the start of the CCA algorithm to ensure that a DEV that successfully sent a frame in the contention period is allowed time to also receive an Imm-ACK, if one was requested.

The RIFS is used for retransmissions in the CTAP. Retransmissions in the CAP are required to use the backoff algorithm. In fact, they are required to use a larger contention window, i.e., the exponential backoff (see "Contention access period (CAP)" on page 55). However, in a CTA (other than a contention-based MCTA), the source DEV is the only one that is allowed to transmit and it can begin its retransmission as soon as it is practical. If the source DEV does not hear an ACK to the frame, it would want to begin to retransmit as soon as possible. In the case of the 2.4 GHz PHY, the source DEV changes from transmit to receive (SIFS) and listens to hear the beginning of the ACK ($pCCADetectTime$). If it does not hear the beginning of the ACK, then it changes from receive to transmit (SIFS) and begins to send the retransmission. Thus, the RIFS is equal to $2 * SIFS + pCCADetectTime$.

There may be two reasons why the source did not get an ACK:

- a) The destination did not successfully receive the frame that was sent.
- b) The ACK from the destination is being sent, but the source DEV is unable to hear it due to bad channel conditions or because of interference.

In the first case, the most efficient procedure is to send the retransmission as soon as the source DEV determines that the ACK is not going to arrive. In the second case, if the source DEV uses a short RIFS, as defined for the 2.4 GHz PHY, then it will step on the ACK and the receiving DEV will not hear the retransmission either. For this case, the most efficient RIFS would be to choose it to be $2 * SIFS$ plus the duration of the ACK, either Imm-ACK or Dly-ACK. However, the CAZAC sequence has good cross and auto-correlation characteristics, and so it is highly unlikely that the source DEV would be unable to hear any of the CAZACs in the preamble of the ACK and still be able to hear the response. Because of these reasons, the TG decided to use CCA detection instead of waiting for an ACK duration.

Miscellaneous PHY requirements

Throughout the standard, all power measurements are referenced to the antenna connector and do not include the effects of the antenna. It is assumed that the test equipment is matched to the impedance at the antenna connector. If the antenna connector is not available, then the implementer is to estimate the antenna gain and subtract that from the measurement to get a value equivalent to that if the measurement had been made at the antenna-to-radio interface.

The operating temperature range for the measurements specified in this standard is 0 to -40°C . While most consumer devices will exceed this temperature range, the task group did not want to attempt to determine a single temperature range that would be applicable for all applications. For a device that is used primarily indoors, e.g., a television set, 0 to -40°C may be sufficient. On the other hand, automotive applications typically require a -40 to $+85^{\circ}\text{C}$ temperature range.

The MAC protocol allows the PNC to use either the CAP or contention-based MCTAs for DEVs to join the network and to communicate with the PNC. However, this flexibility would require that all DEVs implement both methods in order to communicate with the PNC. The alternative would be that a DEV would be unable to join a piconet because it supported a contention method that was different from the one supported by the PNC. Rather than allow this, the task group decided to require that all PNC-capable DEVs that implement the 2.4 GHz PHY allow the use of the CAP for DEVs to associate and send commands and/or data. The CAP is a very efficient method for channel access (see "Comparing the contention access methods" on page 60) and it benefits low-cost DEVs to implement only one of the two methods. The PNC can also allow DEVs to use MCTAs for contention-based access, but this is in addition to the CAP, not as a replacement for it. As a regular MCTA is exactly the same as a CTA, the PNC is always allowed to use these to communicate with a DEV.

PHY FRAME FORMAT

One of the key requirements for interoperability is to provide a specification for how the data are formatted. Although some of the other requirements involve performance expectations, it is not possible to build interoperable implementations if the devices do not use the same data formats. Thus, the

specification of the PHY frame format is one of the key parts of IEEE Std 802.15.3.

Stuff bits and tail symbols

Although the MAC always passes an integer number of octets to the PHY, these bits will not necessarily map into an integer number of symbols. Consider a 3-octet frame payload with a 4-octet FCS. The number of symbols required for this 56-bit MAC frame body are listed in Table 7-5.

Table 7-5: Symbols required to encode a 56-bit MAC frame body

Modulation	Data rate	Bits/symbol	Number of symbols
QPSK-TCM	11	1	56
DQPSK	22	2	28
16 QAM-TCM	33	3	$18\frac{2}{3}$
32 QAM-TCM	44	4	14
64 QAM-TCM	55	5	$11\frac{1}{5}$

Table 7-5 shows that sometimes the MAC frame will require a fractional number of symbols. The MAC frame is always passed as an integer number of octets, and so the 11-, 22-, and 44-Mb/s modes will always be able to encode this as an integer number of symbols. However, the 33- and 55-Mb/s modes are not guaranteed to give an integer number of symbols. For these modes, stuff bits (SBs) are added to the end of the MAC frame body to round out the modulated frame into an integer number of symbols. The transmitting PHY can add either ones or zeros for the stuff bits, and these are ignored by the receiving PHY. The stuff bits are not used as part of the FCS calculation (which technically is done by the MAC). The transmitter always adds just enough bits to complete the symbol. Thus, for the 33-Mb/s mode, it adds less than 3 bits (i.e., 0, 1, or 2), and for the 55-Mb/s mode, it adds less than 5 bits (i.e., 0, 1, 2, 3, or 4).

A typical modem also needs some time to terminate the demodulation process, so extra symbols, called the tail symbols (TS), are added to the end of the frame. For the TCM modes, these tail symbols are specified so that the trellis code always

terminates in the same state. Two tail symbols are required for the 16-, 32-, and 64- QAM-TCM modes while the QPSK-TCM mode requires three tail symbols. For the DQPSK mode, two symbols are provided to assist in the implementation of a differentially coherent demodulator. The overhead for the extra symbols is extremely low (<200 ns per frame) while simplifying the modem design by ending the frame in a known set of symbols.

Frame format

The 802.15.3 PHY is responsible for including any PHY-related information in the PHY header and for creating some type of an error check sequence for the combination of the MAC and PHY headers. In addition, for radio synchronization, the PHY prepends a preamble to the frame. The complete PHY frame format for the 22-, 33-, 44-, and 55-Mb/s modes is shown in Figure 7-5.




Figure 7-5: PHY frame formatting for 22-, 33-, 44-, and 55-Mb/s modes

The frame begins with the preamble, which allows the receiver to set its gain control setting, detect and correct for the frequency offset, lock onto the phase of the received signal, and perform any necessary channel equalization to enhance the radio's multipath performance. The modulation for the preamble is 11 Mbaud

QPSK (22 Mb/s). The preamble is followed by the two octet PHY header, which sets the seed for the scrambler, indicates the data rate of the frame payload, and sets the length of the frame payload. The MAC header is sent next, and the PHY calculates a CRC over the PHY and MAC headers and sends this as the two octet header check sequence (HCS). The PHY header, MAC header, and HCS are modulated with the base rate, 22-Mb/s DQPSK. The rest of the frame (if it does not have a zero length frame payload) is modulated at the rate indicated in the PHY header, either 22-, 33-, 44-, or 55-Mb/s. The frame payload and FCS are passed to the PHY from the MAC via the PHY SAP.

The PHY frame for the 11-Mb/s mode is somewhat different. The reason for this is that the 11-Mb/s mode has a much lower required SNR as compared with the 22-Mb/s mode and the error rate for the PHY and MAC headers would determine the overall FER for the 11-Mb/s mode frame. This would completely wipe out the range advantage that is given by using the 11-Mb/s mode. For example, when the signal to noise ratio (SNR) is 5.5 dB, the BER of the 11-Mb/s mode is 10^{-5} while the BER of the 22-Mb/s mode is 2.8×10^{-2} . The FER for a group of bits based on its BER is calculated from

$$\text{FER} = 1 - (1 - \text{BER})^{(\text{number of bits})} \quad (1)$$

Thus, the FER for the header and a 2048-byte frame body are

$$\text{FER}_{(\text{header})} = 1 - (1 - 2.8 \times 10^{-2})^{(12 \times 8)} = 94.5\% \quad (2)$$

$$\text{FER}_{(\text{MAC frame})} = 1 - (1 - 10^{-5})^{(1024 \times 8)} = 7.9\% \quad (3)$$

In this case, the overall FER would be 95%. The frame would almost never get through because the header would almost always be in error.

To fix this issue, the PHY header, MAC header, and HCS are sent a second time for 11 Mb/s mode frames, modulated with QPSK-TCM. If the receiver correctly receives the first header, it can ignore the second header. However, if the receiver gets an error in the first header, it will listen for the second header to see if the frame was sent at 11 Mb/s. Because the data rate is in the PHY header, it is possible that the error in the first header sequence is in the bit that indicates the PHY data rate. Thus, if there is an error in the first header, the receiver cannot be

sure of the data rate with which the frame is being sent and it must listen for the second header in case the frame is an 11-Mb/s frame.

The PHY frame format for the 11-Mb/s mode is illustrated in Figure 7-6.




Figure 7-6: PHY frame formatting for 11-Mb/s mode

PHY preamble

In order to synchronize the radio, a known sequence is sent as the first portion of the data frame. For the 2.4 GHz PHY, this sequence is a constant amplitude, zero-autocorrelation (CAZAC) sequence. For QPSK modulation, the sequence consists of 16 symbols with the phase indicated in Table 7-6. In the table, the phase shift from the previous symbol is also listed, which shows an interesting pattern. The symbols in the sequence consist of four repetitions of each of the four possible phase shifts. As opposed to a PN sequence, the CAZAC sequence has zero autocorrelation when the sequences are not aligned. In a PN sequence, the autocorrelation is generally $1/N$ when the sequences are offset where N is the length of the PN sequence.

Table 7-6: CAZAC sequence

CAZAC sequence element	Value	Degrees phase shift from previous symbol
C_0	$1 + j$	0
C_1	$1 + j$	0
C_2	$1 + j$	0
C_3	$1 + j$	0
C_4	$-1 + j$	90
C_5	$-1 - j$	90
C_6	$1 - j$	90
C_7	$1 + j$	90
C_8	$-1 - j$	180
C_9	$1 + j$	180
C_{10}	$-1 - j$	180
C_{11}	$1 + j$	180
C_{12}	$1 - j$	270
C_{13}	$-1 - j$	270
C_{14}	$-1 + j$	270
C_{15}	$1 + j$	270

The PHY preamble is made up of 12 repetitions of the CAZAC sequences, denoted as $P_0, P_1, P_2, \dots, P_{10}, E$. However, the twelfth repetition of the sequence, E , has each element of the CAZAC sequence negated, or equivalently, rotated by 180 degrees. This last sequence acts as a marker to tell the PHY when the first bit of the PHY header will be arriving. The complete physical layer preamble is shown in Figure 7-7.




Figure 7-7: Physical layer preamble format

Data size restrictions

The maximum frame size that is supported by the 2.4 GHz PHY is 2048 octets, including the FCS. Without encryption, this limits the amount of data in a single data frame to be less than or equal to 2044 octets. In addition, the maximum transfer unit size is limited to 2044 octets. This parameter places a restriction on the largest frame that can be given to the MAC for transfer over the PHY. The MAC may fragment the frame into smaller pieces to improve the probability of reception. There are a couple of reasons for this restriction. The first reason is that most protocols actually use smaller frame sizes, e.g., TCP, MPEG, etc. The second reason is that a DEV that is assembling frames as the destination needs to provide sufficient buffering to hold all of the fragments of the frame before it passes them up. In an earlier draft of the standard, the maximum frame size was 65,335 octets, which would require approximately 62 kbytes or storage for each stream that the DEV is receiving. For six streams, this requires a buffer size of over 360 kbytes, which is not consistent with a low-cost implementation.

The MAC does not place a restriction on the smallest fragment. In theory, a DEV could send 1-octet fragments. To prevent this, the smallest fragment allowed for 2.4 GHz PHY implementations is 64 octets. The last fragment may be smaller than this value, and an unfragmented data frame may also be smaller, but the PHY is not allowed to fragment a frame into fragments smaller than this value.

MODULATION

In IEEE Std 802.15.3, five different modulation methods are used to get five data rates with one symbol rate. Figure 7-8 illustrates the signal constellations used in encoding bit streams into discrete signal levels that are sent over the air interface. The base modulation rate, which all DEVs are required to support, is the 22-Mb/s DQPSK uncoded modulation. The base rate is used for all broadcast frames, such as the beacon, because all DEVs are capable of understanding this modulation. The uncoded DQPSK modulation was chosen as the base rate because it is simple to implement and allows the development of compliant radios with a minimum cost. It is even possible to non-coherently demodulate the DQPSK signaling, which would allow extremely simple baseband processing.

However, multimedia applications require high data rates and connectivity over long distances. To achieve these goals, other modulations, up to 64-QAM, with coding in the form of an eight-state trellis code are used. The eight-state trellis encoder is illustrated in Figure 7-9. Note that only the lower two bits for the QAM modulations are used in the actual encoder. The other bits are used to select the symbol from the subsets that are selected by the trellis coder. The set partitioning is described in more detail in the standard. The set partitioning was chosen such that the decimal representation of the bit mapping goes from low to high as the constellation points are traced from center outward. This rule ensures that the decimal representations of the bit mappings from 0 to 15 belong to 16-QAM constellation, and 0 to 31 belong to 32-QAM constellation, and 0 to 63 belong to 64-QAM constellation.




Figure 7-8: DQPSK, QPSK, 16/32/64-QAM signal constellations




Figure 7-9: QPSK, 16/32/64-QAM eight-state trellis encoder

The final mapping of bits to symbols for the TCM encodings is shown in Figure 7-10. Also shown in the figure are the eight subsets, D0-D7, that are used in creating the mapping from the input bits to the trellis-coded symbols.




Figure 7-10: 16/32/64-QAM constellation bit mappings

The drawing in Figure 7-8 shows the outer points of the 64-QAM constellation at seven times the distance of the QPSK constellation points. If this was implemented in the radio, the average power in the 64-QAM portion of the frame would be much greater than the power in the preamble, which would cause problems with the receiver's automatic gain control (AGC) algorithms. It would also make the ranges for the 22-Mb/s links closer to the 55-Mb/s links by decreasing the range of the 22-Mb/s mode. To handle the differences in the modulation, the transmitter is required to scale the constellation points such that the average power in the preamble and headers is the same as the average power

d during the transmission of the data in the frame. This scaling is similar to that used in 802.11a, which also had DQPSK headers followed by up to 16-, 32-, and 64-QAM modulation. The constellation points in Figure 7-8, $I + jQ$, are multiplied by a scaling factor, K_{MOD} , to give the output point, d . The values for K_{MOD} for each of the modulations are listed in Table 7-7.

Table 7-7: Normalization factor for PHY modulation formats

Modulation	K_{MOD}
DQPSK	1
QPSK-TCM	1
16 QAM-TCM	$1/(\sqrt{5})$
32 QAM-TCM	$1/(\sqrt{10})$
64 QAM-TCM	$1/(\sqrt{21})$

The K_{MOD} factor can be calculated by computing the average power in the constellation and setting the average power for DQPSK to be unity. The power of an individual point is the square of the Euclidean distance because the constellation is plotted in voltage. The average power, $P_{\text{avg(mod)}}$, then is this number divided by the number of points in the constellation, as detailed below:

$$P_{\text{avg(mod)}} = \frac{\sum_{n=0}^{N_{\text{points}}-1} (I_n^2 + Q_n^2)}{N_{\text{points}}} \quad (4)$$

The scaling factor is the square root of the ratio of the average power in the DQPSK modulation versus the average power in the other modulation, i.e.,

$$K_{\text{MOD}} = \sqrt{\frac{P_{\text{avg(DQPSK)}}}{P_{\text{avg(mod)}}}} \quad (5)$$

The values for the total power and average power calculated using this formula for the defined modulations are listed in Table 7-8.

Table 7-8: Total and average power in the constellation

Modulation	Total power	Average power
DQPSK	8	2
16 QAM	160	10
32 QAM	640	20
64 QAM	2688	42

RECEIVER PERFORMANCE

During the development of IEEE Std 802.15.3, there was a consistent focus on keeping the requirements relaxed so that implementers could create low-cost, low-power solutions to best serve the portable, consumer electronics market. This desire to allow for innovative designs from implementers is reflected in the receiver specifications. For example, the standard does not include any intermodulation specifications. DEVs compliant to the standard operate in a highly dynamic environment that varies greatly from user to user. The large disparity between the different use cases made it impossible to generate an interference model that would be applicable to a large number of users. Instead, the intermodulation performance requirements are left to the implementers and users to determine through market forces.

Thus, the receiver and transmitter specifications are focused on guaranteeing a minimum level of performance along with any requirements necessary to guarantee interoperability. The PHY subcommittee defined an error-rate criterion that is used for many of the receiver performance tests. This condition requires the DEV to achieve an FER of 8% with 1024 octet frames. Using Equation (1), this works out to a BER of approximately 10^{-5} . The PHY subcommittee also chose a maximum receiver noise figure of 12 dB and a typical implementation loss of 4 dB. The implementation loss is intended to include imperfections in the analog and digital portions of the receiver. Implementers are free to split this margin in any way they choose. For example, they may use a 14 dB system noise figure and 2 dB of implementation loss or 8 dB of system noise figure and 8 dB of implementation loss. Using these numbers, the minimum required sensitivity for a compliant 802.15.3 receiver is given in Table 7-9. Also listed in the table are numbers for the sensitivity of a typical implementation. It is likely that most

implementations will achieve these numbers, although some very low-cost devices may choose values closer to the reference sensitivity.

Table 7-9: Reference sensitivity levels for modulation formats

Modulation	Reference sensitivity	Typical implementation sensitivity
QPSK-TCM	-82 dBm	-88 dBm
DQPSK	-75 dBm	-81 dBm
16-QAM-TCM	-74 dBm	-80 dBm
32-QAM-TCM	-71 dBm	-77 dBm
64-QAM-TCM	-68 dBm	-74 dBm

The SNR required to achieve the error rate criterion is listed for each of the modulations in Table 7-10. Using TCM provides a significant reduction in the SNR number while modestly reducing the data rate. In particular, the 33 Mb/s mode achieves almost the same sensitivity as the DQPSK mode with 50% higher data rate. The TCM encoder is quite simple to realize in the transmitter (see Figure 7-9); however, the decoder is more complex, and a good demodulator requires good design and a non-negligible amount of die area. Because of this, the DQPSK modes are uncoded to allow very simple radios that are able to maintain connectivity with other high-speed 802.15.3 DEVS.

Table 7-10: Required SNR to achieve the error-rate criteria for each of the modulation formats

Modulation	SNR required
QPSK-TCM	5.5 dB
DQPSK	12.6 dB
16-QAM-TCM	13.5 dB
32-QAM-TCM	16.6 dB
64-QAM-TCM	19.8 dB

Figure 7-11 compares the reference sensitivity of 802.15.3 with 802.11b and 802.11a for the data rates supported by those standards. The 802.11g draft currently adopts the reference sensitivities for 802.11a for the mandatory enhanced data rates, so they are the same as the values given in the figure for 802.11a. In IEEE Std 802.11a, the authors included their assumptions—10 dB of noise figure and 5 dB of implementation loss, or 1 dB less than that assumed for the 802.15.3 PHY. The assumptions are not listed in 802.11b or for the original 802.11 DSSS PHY. Taking into account the 1-dB difference in the assumptions of the 802.11a and 802.15.3, the reference sensitivity for 802.15.3 is about 4 dB better than that for 802.11a due mainly to the improved SNR performance of the trellis-coded modulation vs. the combination of the OFDM modulation and convolutional code. The only 802.15.3 modulation that comes close to the 802.11a reference sensitivity is the 22 Mb/s mode, which is essentially the same in the graph but about 1 dB better when the difference in the assumptions are taken into account.

The 802.11b results are somewhat misleading in that the writers of the standard used relatively loose requirements in specifying the sensitivity. For example, the SNR difference between 1 Mb/s Barker coding (1 dB for 10^{-5} BER) and 11 Mb/s CCK coding (7.6 dB for 10^{-5} BER) should have degraded the reference sensitivity by 6.6 dB rather than the 4 dB in the standard. One of the reasons for this may have been that when the initial 802.11 standard was written (1990–1997), low-cost radios were not as sophisticated and so it was more difficult to meet the sensitivity number. By the time that 802.11b was published, 1999, it was much easier to create low-cost radios that also had pretty good performance. In fact, most 802.11b implementations achieve around -90 dBm for the 1 Mb/s mode and about -82 dBm for the 11 Mb/s mode, which brings the values more in line with 802.15.3 reference sensitivity numbers. However, typical 802.15.3 implementations will exceed the reference sensitivity number. Based on the required SNR, an 11 Mb/s 802.15.3 radio should have about 2 dB better performance than an 802.11b radio at the same data rate.

One of the other important parameters for a radio is its ability to reject jamming signals. In IEEE Std 802.15.3, as in other standards, this performance is simulated by introducing another instance of the signal, uncorrelated in time with the desired signal, in one of the other channels and determining the error rate performance in the desired channel. Because the jammers in the 2.4 GHz band can be quite




Figure 7-11: Comparison of receiver sensitivities for 802.11a, 802.11b, and 802.15.3

varied, e.g., 802.11 FHSS, 802.11b, 802.15.1, cordless phones, microwave ovens, and so on, it would be much too difficult to both specify performance criteria for each of these jammers as well as for implementers to test to all of these. Instead, an 802.15.3-compliant transmitter is introduced in an adjacent channel with random DQPSK data uncorrelated in time with desired signal. The desired signal is set at 6 dB above the reference sensitivity to remove the effects of thermal noise. The jamming signal is then put into one of the other channels, either adjacent or alternate, with a power level equal to the reference sensitivity plus the value indicated in Table 7-11.

For example, for 16 QAM-TCM, the desired signal is set at -68 dBm (-74 dBm + 6 dBm) while the interfering signal is set to either -49 dBm in the adjacent channel or to -34 dBm in the alternate channel. A summary of the values for the desired signal and the jamming signal power is given in Table 7-12.

For the purposes of the test, the four-channel plan is used (i.e., channels 1, 2, 4, and 5) and an adjacent channel is defined as a channel that is one removed from the desired channel while the alternate channel is at least one channel further out. A list of adjacent and alternate channels for the jamming resistance tests is given in Table 7-13.

Table 7-11: Receiver jamming resistance requirements

Modulation format	Adjacent channel rejection	Alternate channel rejection
QPSK-TCM	33 dB	48 dB
DQPSK	26 dB	41 dB
16 QAM-TCM	25 dB	40 dB
32 QAM-TCM	22 dB	37 dB
64 QAM-TCM	19 dB	34 dB

Table 7-12: Power levels for desired and jamming signals

Modulation format	Desired signal	Adjacent channel power	Alternate channel power
QPSK-TCM	-76 dBm	-49 dBm	-34 dBm
DQPSK	-69 dBm	-49 dBm	-34 dBm
16 QAM-TCM	-68 dBm	-49 dBm	-34 dBm
32 QAM-TCM	-65 dBm	-49 dBm	-34 dBm
64 QAM-TCM	-62 dBm	-49 dBm	-34 dBm

TRANSMITTER PERFORMANCE

The transmitted signal is band limited to allow up to four channels in the 83.5 MHz allocated for unlicensed operation in the 2.4 GHz band. Although the FCC has allocated spectrum from 2.4 GHz–2.4835 GHz, there are “restricted bands” at 2.31–2.39 GHz and at 2.4835–2.5 GHz. The FCC restricts the field strength emitted by an intentional radiator in these bands. The allowed signal level in the restricted band is approximately -42 dBm equivalent isotropically radiated power (EIRP) in a 1-MHz bandwidth. EIRP means the power input to an ideal

Table 7-13: Adjacent and alternate channels for receiver jamming resistance test

Desired channel number	Adjacent channel number	Alternate channel number
1	2	4, 5
2	1, 4	5
4	2, 5	1
5	4	1, 2

isotropic radiator that results in the same field strength. As 802.15.3 devices will typically transmit about +14 dBm in a 15-MHz bandwidth, the in-band power is about +3 dBm/MHz. To achieve the FCC requirements, the transmitter will need to ensure that the transmit signal power in the restricted band is down by approximately 45 dB. The closest restricted band is at the lower band edge, separated by 22 MHz from the center frequency of channel 1. The higher frequency restricted band is 21.5 MHz from the center frequency of channel 5.

A compliant 802.15.3 2.4 GHz PHY transmitter is required to meet the power spectral density limits specified in Table 7-14. These limits were developed to balance the desire for piconets that overlapped in space on separate channels while keeping the transmitter complexity low. Decreased limits (i.e., lower allowed power) improve the performance of 802.15.3 networks that are located near each other in space but are on different channels. However, as this limit is decreased, the linearity requirements for the PA are increased. If the limits are too low, special filtering and even different transmitter architectures would be required to achieve the required performance.

Table 7-14: Transmit PSD limits

Frequency	Relative limit
$7.5 \text{ MHz} < f - f_c < 15 \text{ MHz}$	-30 dBr
$15 \text{ MHz} < f - f_c < 22 \text{ MHz}$	$-10/7[f - f_c \text{ (MHz)} + 13] \text{ dBr}$
$22 \text{ MHz} < f - f_c $	-50 dBr

A graphical representation of the transmit PSD is shown in Figure 7-12. The transmit spectral requirements for 802.15.3 are more stringent than either 802.11b or 802.11a. There are two reasons for this: the first is to allow up to four channels in the 73 MHz of usable bandwidth and the second is to improve the coexistence properties of 802.15.3 with other wireless standards.




Figure 7-12: Transmit power spectral density mask

The maximum transmitter power is not set by IEEE Std 802.15.3. Rather, implementers are free to use any power up to the limits allowed in the geographical region where they want their products to operate. A summary of the requirements in four of the geographic regions are listed in Table 7-15. As indicated in the table, there are two rules that are applicable to 2.4 GHz PHY in the U.S.: 15.247 and 15.249 [B1]. The 15.247 rules originally required that the device use a form of spread spectrum, either frequency hopping or direct sequence. Direct sequence spread spectrum (DSSS) systems had to pass a jamming resistance test to be certified under this ruling. However, in early 2002, the rules were modified to change from requiring a DSSS system to only requiring digital modulation. In addition, the jamming margin test was replaced by a power spectral density limit to ensure that the power was spread over a reasonable bandwidth. The higher power allowed under the new ruling allows 802.15.3 systems to achieve data rates and ranges in excess of those offered by either 802.11b or 802.11a.

Table 7-15: Maximum transmit power levels

Geographical Region	Power limit	Regulatory document
Japan	10 mW	ARIB STD-T66 [B2]
Europe (except Spain and France)	100 mW EIRP 10 mW/MHz peak power density	ETSI EN 300 328 [B6], ETSI ETS 300 826 [B7]
USA	50 mV/m at 3 m in at least a 1 MHz resolution bandwidth ^a	47 CFR 15.249 [B1]
USA	1 W total power and less than 8 dBm in any 3 kHz bandwidth, reductions for antenna gains > 6 dB	47 CFR 15.247 [B1]

^a Electric field strength measurement rather than conducted power measurement.

IEEE Std 802.15.3 also places a requirement on the signal fidelity of the transmitted signal. This requirement is measured by comparing the actual transmitted constellation to the ideal constellation. This measurement, called the error vector magnitude (EVM), is illustrated in Figure 7-13. The EVM performance is very important because a high EVM will make it much more difficult for the receiver to successfully demodulate the transmitted signal.

The calculation for the EVM is defined as:

$$EVM = \sqrt{\frac{\frac{1}{N} \sum_{j=1}^N (\delta I_j^2 + \delta Q_j^2)}{S_{max}^2}} \times 100\% \quad (6)$$

where S_{max} is the magnitude of the vector to the outer-most constellation point and $(\delta I_j, \delta Q_j)$ is the error vector as shown in Figure 7-13. The allowed EVM values for 802.15.3 DEVs are listed in Table 7-16.




Figure 7-13: Error vector calculation

Table 7-16: EVM values for various modulations

Modulation	r (dB)	SNR (dB)	EVM (%)
64 QAM	3.6796	26	3.3
32 QAM	3.3606	23	4.8
16 QAM	2.5527	20	7.5
DQPSK/QPSK-TCM	1.7700	19	9.2
QPSK-TCM	1.7700	12	20.0

The EVM values shown in Table 7-16 have been calculated using the indicated SNR values in the table. These SNR values incorporate the TCM coding gain (except for QPSK). The values for the SNR in the table are at least 6 dB greater than the SNR required to achieve a 10^{-5} BER. The equation used to calculate the EVM comes directly from vector considerations of the signal plus noise in relationship to the ideal constellation points of any M-QAM signal. This equation for calculating the EVM is:

$$\text{SNR}_{\text{M-QAM}} = -(r + 20 \log [\text{EVM}_{\text{M}} / (100\%)]) \quad (7)$$

where r is the peak-to-average energy ratio. The value for 32-QAM is an approximation because the constellation is a "cross;" thus, some penalty is incorporated due to an increased encoding and decoding complexity among the constellation points; i.e., not all constellation points are separated by the same Hamming distance. The minus sign in Equation (7) is necessary because SNR is the ratio of signal to noise, whereas EVM is the ratio of noise to signal.

REGULATORY AND REQUIREMENTS

The 802.15.3 2.4 GHz PHY was designed such that it could be used essentially worldwide. However, IEEE Std 802.15.3 does not specify the performance required to achieve regulatory compliance in the various geographic regions where this band is allocated for unlicensed use. For example, the requirements for spurious emissions are not listed in the standard; instead, the implementer is required to adhere to the regulations appropriate for the regions in which they wish to sell their products.

The list of the geographical regions and the appropriate regulatory agencies are listed in Table 7-17.

Table 7-17: Geographic regions and regulatory bodies

Region	Approval standard	Document	Approval authority
Europe ^a	European Telecommunications Standards Institute (ETSI)	ETS 300-328, ETS 300-826	National type approval authorities
Japan	Association of Radio Industries and Businesses (ARIB)	ARIB STD-T66	Ministry of Post and Telecommunications (MPT)
United States	Federal Communications Commission (FCC)	47 CFR, Part 15, Sections 15.205, 15.209, 15.247, 15.249	FCC
Canada	Industry Canada (IC)	RSS-210 [B25]	IC

^a Except France and Spain.

Many of the regulatory documents are available online as well. Table 7-18 lists the location of the documents for the various approval authorities.

Table 7-18: Online document locations for regulatory agencies

Organization	Documents available from
ETSI	http://www.etsi.org
ARIB	http://www.arib.or.jp ^a
FCC	http://www.fcc.gov/oet/info/rules/Welcome.html
IC	http://strategis.ic.gc.ca/engdoc/main.html

^a ARIB documents can be difficult to obtain.

DELAY SPREAD PERFORMANCE

One of the PHY characteristics that was evaluated in the downselection process was the ability of the proposed PHY to withstand delay spread. Delay spread, or multipath, is the result of reflections of the transmitted signal from objects in the environment that reach the receiver. As an example, consider only two paths, the direct path to the receiver and a reflected path, as shown in Figure 7-14. In the figure, the straight line represents the direct ray or the line of sight (LOS) signal. The other two paths, d_1 and d_2 , combine to provide the indirect ray or the non-line of sight (NLOS) signal. Because the NLOS signal travels farther and will generally lose energy in the reflection, it is usually a lower amplitude than the LOS signal.




Figure 7-14: Geometry for two-ray multipath

If the path difference between the two rays is an odd multiple of one-half the wavelength, then the amplitude of the signal from the direct path will be reduced by the amplitude of the reflected wave. If the direct and reflected signals are nearly the same in amplitude, the cancellation will be nearly complete. The fading from this effect can be as much as 30 or 40 dB in a narrow bandwidth. Generally, the deeper the fade, the more narrow it will be in both frequency and position. For a 2.45 GHz frequency, one-half of the wavelength in free space is 61 mm or about 2.4 inches. The result of this is that typically, there are many deep nulls (greater than 20 dB) in the 2.4–2.4835 GHz frequency range for typical environments. Because both the radios and the objects around them may be moving, the frequency of these nulls will change as a function of time.

If, on the other hand, the time difference between the paths of the two rays is on the order of a symbol period, then the received signal will be distorted because symbol n will overlap with symbol $(n + 1)$, increasing the BER of the channel through intersymbol interference (ISI). A rule of thumb for the limit on this time difference is that it should be less than 15% to 25% of the symbol period. The limits on the delay spread, using this rough criteria, for 1 Mbaud (Msymbol/s) and 11-Mbaud systems, are summarized in Table 7–19.

Table 7–19: Delay spread limits for two symbol rates

Symbol rate	Symbol duration	15% time	15% distance	25% time	25% distance
1 Mbaud	1000 ns	150 ns	45 m	250 ns	75 m
11 Mbaud	91 ns	14 ns	4.1 m	23 ns	6.8 m

The actual delay spread that can be tolerated also depends on the packet length, type of modulation, protocol, desired packet error rate, type of receiver, and the actual distribution, in amplitude and time, of the delayed signals. In addition, because the LOS signal is usually much stronger than the NLOS signal, the effect of delay spread will be worse if there is not a direct ray from the transmitter to the receiver.

Mitigating the effects of delay spread

Unlike most noise sources in radio systems, increasing the transmitted signal power does not improve the BER once delay spread has become the dominant error term. The reason for this is that the interfering signal is generated by the transmitter, and so increasing the transmit power also increases the interfering power. Decreasing the noise figure of the system also has no effect because the thermal noise power is dwarfed by the power of the interferer.

There are techniques for minimizing the effects of delay spread. These include the following:

- Equalization, both fixed and adaptive (e.g., Viterbi, decision feedback equalizer, linear equalizer, etc.)
- Spreading the signal bandwidth, e.g., DSSS and OFDM
- Coding techniques, such as FEC and interleaving
- Directivity by using higher gain antennas
- Diversity, either space, polarization, frequency, or a combination of these

Coding modifies the signal to give it properties (such as good auto-correlation) that are resistant to fading. This approach consumes more bandwidth in order to reduce fading. Equalization, on the other hand, sums the received signal with copies of the signal that are delayed in time and are adjusted in amplitude to undo the effects of multipath. Equalization methods must be adaptive to handle the dynamic nature of the channel. Equalization techniques can be very effective, but they add complexity and cost to the system. Directivity is a simple solution, but it requires that the transmitter and receiver are "pointed" at each other and it results in larger antennas. Diversity uses multiple sources for the received signal, choosing the "best" source to make the decision. Diversity increases the complexity, cost, and size of the system. Spatial diversity with antenna patterns that are nearly orthogonal in azimuth and elevation allow higher directivity without giving up near omnidirectional coverage.

Fading channel model used for 802.15.3

To model the effects of multipath, an exponentially decaying Rayleigh fading channel was chosen as the environmental model to compare the proposals. The model was originally proposed in IEEE P802.11-97/96³⁴. The channel is assumed to be static throughout the frame and is generated independently for each frame. An example of the channel impulse response using this model is illustrated in Figure 7-15. In the figure, the black arrows illustrate the average magnitude of the signal over time. The gray arrows illustrate the magnitude of a specific random realization of the channel. The time positions of black and gray samples are staggered for clarity only. In the model, the channel has values only at the sampling instants, i.e., at $t = T_s$.




Figure 7-15: Channel impulse response used for 802.15.3 evaluation

³⁴ This is an 802.11 Working Group document number which indicates that this was submitted in 1997 and was the 96th document that year. The document can be located via the 802.11 Working Group web page, http://grouper.ieee.org/groups/802/11/Documents/DocumentArchives/1997_docs in the file 97sep.zip with the file name 70967.doc.

The impulse response of the channel, h_i , is composed of complex samples with random uniformly distributed phase and Rayleigh distributed magnitude with average power decaying exponentially, as shown in Figure 7-15. The impulse response is then given by

$$h_i = N(0, \frac{1}{2}\sigma_k^2) + jN(0, \frac{1}{2}\sigma_k^2) \quad (8)$$

$$\sigma_k^2 = \sigma_0^2 e^{-k(T_s/T_{\text{RMS}})} \quad (9)$$

$$\sigma_0^2 = e^{-T_s/T_{\text{RMS}}} \quad (10)$$

where

$$N(0, \frac{1}{2}\sigma_k^2) \quad (11)$$

is a zero mean Gaussian random variable with variance

$$\frac{1}{2}\sigma_k^2, \quad (12)$$

and σ_0^2 is chosen so that the condition

$$\sum \sigma_k^2 = 1 \quad (13)$$

is satisfied to ensure the same average received power.

The sampling time, T_s , in the simulation is assumed to be shorter than a symbol time (or chip time) by at least a factor of four. In simulations T_s is typically a submultiple of the symbol duration. The number of samples to be taken in the impulse response should ensure sufficient decay of the impulse response tail, e.g., $k_{\text{max}} = 10T_{\text{RMS}}/T_s$.

Defining delay spread performance

Although it is common to quote delay spread performance in terms of time, e.g., "Multipath delay spread tolerance > 250 ns,"³⁵ that single number does not even begin to describe the delay spread performance of a radio system. The channel model is one of the key items that needs to be fully specified in order to determine the delay spread performance of a system. Different architectures and modulations can have markedly different performance characteristics depending on the channel model that is chosen. The channel model will either have to specify the exact number and type of rays (distance and amplitude) in the model or it will have to assume a distribution (typically either Rician or Rayleigh) and specify the rms delay spread as well as the standard deviation of the delay spread. In addition to that, the introduction of delay spread will increase the FER of the system so another key aspect of the specification is the acceptable level of degradation of the system FER due to delay spread. Finally, the received signal level needs to be specified.

The 802.15.3 PHY subcommittee formed a subgroup to draft a concise definition of delay spread so that all of the proposals could be evaluated against the same criteria. This definition and the results of the evaluations are contained in the IEEE TG3 document 00/110r14P802-15_TG3-Criteria-Definitions.doc. The test required the following steps:

- a) Generate 1000 random channels according to environment model, for the desired delay spread T_{RMS} .
- b) Apply a signal gain, such that the average RF signal level of the 1000 random channels is 14 dB higher than the RF signal level required to achieve a 1% FER in an AWGN channel. Because the channel model is Rayleigh faded, each instance of the channel may be either above or below the average power.
- c) Simulate the FER for each of the 1000 random channels. The simulation should include AWGN arising from the RF circuit (thermal noise + noise figure). Perfect symbol and carrier synchronization may be assumed. Equalizers must perform frame-based coefficient adaptation. The

³⁵ From Texas Instruments' product bulletin on the ACX100 802.11b baseband processor, http://focus.ti.com/pdfs/vf/bband/ti_acx100.pdf.

simulation should not include advanced channel selection and/or diversity techniques that alter the distribution of the channel model. Use one of the two simulation methods defined below:

- 1) *Direct measurement of FER*: Simulate at least 1000 frames per random channel. Each frame should consist of the proposed preamble, header, tail, and 512 bytes of data. Directly calculate the FER for each random channel.
- 2) *Measurement of FER by BER*: Simulate the BER for each random channel. The measurement may be performed on continuously transmitted data; however, equalizer adaptation must be performed on the proposed preamble/header. Convert the BER to an FER, assuming 512 bytes of data, preamble, header, and tail bits.
- d) Discard the results of the 50 channels (5%) with the highest FER. Find the maximum FER of the remaining channels.
- e) Repeat steps a) through d) for different values of T_{RMS} . T_{RMS} values of 10, 25, and 40 ns must be simulated. Additional values of T_{RMS} may also be simulated to demonstrate the robustness of the system: 50, 75, 100 ns, and so on. The delay spread tolerance is the maximum value of T_{RMS} in which the maximum FER over 95% of the channels is at most 1%. All lower values of T_{RMS} must also achieve 1% FER.

Delay spread measurements

With the channel model and simulation procedure, the only element left was to determine an appropriate value of T_{RMS} that reflected typical values in the home environment. A review of the literature indicated that a value of about 25 ns was appropriate. Huang and Khayata [B9] performed delay spread measurements at 915 MHz, 2.44 GHz, and 5.8 GHz with dipole antennas. The areas that they measured included office space (with cubicles), library, auditorium, classrooms, and a basement with large motors and electrical equipment. Their observation of the delay spread was as follows (note, F1 is 915 MHz):

It can be observed from Fig. 1 that delay spreads are under 100 ns in most of the cases. The median level of delay spreads is under 50 ns in most of the cases. Delay spreads do not change much as a function of

frequency except in the Auditorium area. We found, by examining individual multipath delay profiles, that radio wave tends to be reflected more by building walls and other internal structures at F1 than at the two higher frequencies. ... At all three frequencies (more observed at higher frequencies), the regression curves show that delay spreads increase slowly with respect to distance.

They also noted that, "The linear regression graphs also shows that the delay spread at 5.8 GHz lower than that at 2.44 GHz for short distances. It increases more rapidly, however, and appears to overtake between 128 and 265 feet." These measurements give a value of less than 50 ns for environments, e.g., office areas, that should have much worse delay spread than a typical home. John C. Stein in an article on the Harris (now Conexant®) web page [B26] indicated that in an office environment, "Typical values for indoor spreading are less than 100 nanoseconds," which is similar to the results obtained by Huang and Khataya.

Nobles et al. [B23] measured delay spread in an office building at three frequencies. They used swept frequency measurements with an IFFT to calculate the time domain data. They found the following (note that [2–4] below are references in the paper):

The values of the mean rms delay spread (10–50 ns) are typical of those found in similar indoor environments [2–4]. For location T1 the LOS values are very similar to those for location T2 despite the quite different placing of the transmitter within the Communications Lab. In both cases mean values of rms delay spread reduce with increasing frequency.

The authors provided both the mean and the standard deviation of the measured delay spread at three different locations for both LOS and NLOS configurations. The results from the measurements are summarized in Table 7–20 (in the table, the columns labeled s.d. are for the standard deviation). These measurements support a more benign view of the delay spread environment. The measurements are somewhat unique in two ways. The first is that the authors measured three frequencies that are dramatically different. The second difference is that the authors used a swept frequency measurement rather than a pulse response type measurement.

One of the best known researchers in the area of channel measurements is Dr. Theodore S. Rappaport, now a professor at the University of Texas at Austin. In a 1992 paper, Rappaport and Hawbaker [B24] reported on measurements taken at five sites: a typical sports arena (Site A), two (dissimilar) factory buildings (Sites B and C), a six-story closed-plan office (Site D), and a large single-story grocery store (Site E). The measurements were taken at 1.3 and 4 GHz, and they provided cumulative distribution functions (CDF's) of the delay spread. The authors summarize the results for the delay spread measurements by stating the following (note: σ_{tau} is the rms delay spread and [1] is a reference to the thesis of Hawbaker [B8]):

Most σ_{tau} values above 85 ns were measured in Site B, the large open-plan factory that contains metal machinery and inventory. Most values of σ_{tau} below 50 ns are for the closed-plan office building (Site D) which contained partitioned walls and standard dropped-ceilings. Median values of σ_{tau} range from as low as 22 ns for Site D to 67 ns for Site B. Extensive delay spread results are given in [1], and show very little difference between 1.3 GHz and 4.0 GHz, or between high and low base station antenna heights.

Table 7-20: Measured mean and standard deviation of delay spread in three locations from Nobles, et al. [B23]

Location		Frequency					
		2-2.5 GHz		5-6.5 GHz		17-17.5 GHz	
		mean	s.d.	mean	s. d.	mean	s. d.
T1	LOS	34.5	13.2	14.4	3.0	11.0	2.2
	NLOS	39.6	10.5	22.0	4.5	15.3	8.4
T2	LOS	37.5	11.8	14.8	3.0	11.6	1.6
	NLOS	37.7	12.0	21.2	9.4	19.2	10.5
T3	LOS	49.0	5.4	15.7	11.7	26.9	17.7
	NLOS	37.6	17.0	22.1	5.5	27.6	10.6

The rms delay spread measured in the paper is summarized in Table 7-21. Four antenna configurations were used, and measurements were made at different heights. These measurements characterize the channel in an office and industrial environment. These types of environments tend to have larger rms delay spread due to the high reflectivity of equipment and cubicles (which have metal mesh inside).

Table 7-21: Delay spread measurements from Rappaport and Hawbaker [B24]

Configuration and elevation	σ_{tau} = rms delay spread (ns)			
	Directional LP ^a	Directional CP ^b	Omni LP Vertical	Omni LP Horizontal
LOS - 9.40 m	5.1	1.9	14.0	12.9
LOS - 15.2 m	13.0	2.9	34.4	28.7
LOS - 31.1 m	4.3	2.1	33.6	29.2
NLOS - 34.1 m	33.9	46.1	57.8	61.4
NLOS - 31.7 m	25.6	44.2	79.9	66.9
NLOS - 13.6 m	34.2	26.3	55.1	62.7

^a LP - A linearly polarized (LP) antenna

^b CP - A circularly polarized (CP) antenna

In a paper by MacLellan et al. [B20], the delay spread and path loss in four woodframe residential homes were measured in the 915 MHz and 1900 MHz bands. For the path loss measurements, the authors determined an effective path loss exponent and the standard deviation for the homes and for the average of all of homes. In terms of delay spread, the authors state:

The multipath measurements suggest that the time dispersion in both LOS and NLOS topographies may be slightly less at 1900 MHz than at 915 MHz. Typical rms delay spreads for LOS topographies were 9.76 ns for 915 MHz and 8.26 ns for 1900 MHz. Similarly, average rms delays for all NLOS measurements were 14.39 ns for 915 MHz and 12.47 ns for 1900 MHz.

The average standard deviation of rms delay spread from the paper is listed in Table 7-22.

Table 7-22: Standard deviation of delay spread measurements from MacLellan et al. [B20]

Condition	915 MHz	1900 MHz
LOS: Indoor-to-Indoor	2.89	2.42
LOS: Outdoor-to-Indoor	1.71	3.77
NLOS: Indoor-to-Indoor	6.51	4.61
NLOS: Outdoor-to-Indoor	2.62	2.87

Very recently, the Intel® Corporation contracted with the University of Southern California to perform extensive channel measurements in a residential environment. The database of the measurements is available on line at the Intel UWB Database [B19] and contains nearly 900 residential channel measurements that span a frequency band of 2–8 GHz with a resolution of 3.75 MHz. The website allows the user to request delay spread and path loss data for a given frequency range. Using a frequency range of 2.4 GHz to 2.4835 GHz, the results for LOS and NLOS measurements are presented in Table 7-23. The database returns three types of calculations, frequency domain, time domain and windowed time domain with RMS values for all three and mean values for the latter two. The largest value of either the RMS or mean delay spread is less than 20 ns for both LOS and NLOS measurements.

Based on the measured data presented in these papers, a 25 ns rms delay spread for the home environment is a reasonable number.

RADIO ARCHITECTURES

The 802.15.3 PHY was developed to enable low-cost implementations. However, in order to realize these cost savings, the radio designer needs to choose an architecture that meets the requirements of 802.15.3 and is cheap to build. The choice of a radio architecture will change the impact of these requirements on the design of the radio.

Table 7-23: Average delay spread measurements from Intel UWB Database [B19]

Calculation method	LOS	NLOS
Mean τ_{RMS} (frequency domain approach)	18.03 ns	18.48 ns
Mean τ_{RMS} (time domain approach)	10.21 ns	13.53 ns
Mean τ_{RMS} (windowed time domain approach)	7.51 ns	11.93 ns
Mean τ_{mean} (time domain approach)	10.45 ns	16.30 ns
Mean τ_{mean} (windowed time domain approach)	6.97 ns	12.92 ns

The discussion in this section centers on the architectures for the receiver. In most cases, the transmitter architecture will be similar. The transmit-to-receive turn-around time is very short, which requires that the TX and RX frequency sources be locked simultaneously. Other than introducing an extra set of synthesized sources, and thereby increasing the cost of the transceiver, the receiver VCOs will likely need to be reused for the transmit chain. Thus, the frequency plan for the transmitter would likely be the same as for the receiver.

Superheterodyne

The 802.15.3 standard's 2.4 GHz PHY was based on many of the characteristics of 802.11b. Current (as of 2002) 802.11b radios on the market use a superheterodyne architecture with an IF in the 350–450 MHz range, as shown in Figure 7-16. This architecture allows for excellent channel filtering and eases the problem of I/Q phase and amplitude imbalance. One reason for this is that the I/Q split is accomplished at a lower frequency. One method for generating good quadrature out of the second LO is to run the VCO at four times the frequency of the second LO and use a divide by four circuit. A divide by four allows near-perfect quadrature because the phase difference of the LO is independent of the duty cycle of the VCO. In the case of a divide by two, one of the phases is derived from the rising edge and the other from the falling edge. If the VCO has a perfect 50% duty cycle, then this will give accurate quadrature phase. In the case of the divide by four, both phases can be derived from either the rising or falling edges of the VCO, and so the result is independent of the duty cycle of the VCO. With a 350–450 MHz IF, the VCO frequency would be in the range of 1.4–1.8 GHz. This will make the

VCO smaller and potentially easier to realize on-chip due to the reduced inductor and varactor size required for this frequency.




Figure 7-16: Superheterodyne radio architecture

However, the superheterodyne architecture does require a channel select filter with a fairly narrow bandwidth ($<5\%$). A narrow bandpass filter requires high-Q components, which leads to the requirement for a surface acoustic wave (SAW) filter. In the past, SAW filters have been large and expensive. The SAW industry has been able to reduce the size somewhat and has made great advances in reducing the price. In fact, a SAW filter for an 802.11b application is now under \$1.00³⁶ in large volumes. Although the size of the filter has decreased, it is still almost as large as the package for a direct conversion radio IC. In some applications, e.g., an 802.15.3 radio in a TV, the extra size may not matter, whereas in other applications, e.g., a digital camera, the space may be very important.

Because the channel select filter is very effective at removing jamming signals, the third-order intercept is not greatly affected by parts of the receiver after the filter. This allows the designer to use lower current in these circuits that could lead to some power savings. Another advantage of the superheterodyne architecture is that the channel select filter will normally remove essentially all of the adjacent and alternate channel interferers. This reduces the dynamic range requirement of the analog baseband and the ADC in the demodulator.

³⁶ Indicated in U.S. dollars.

Another advantage of the SAW filter is that the requirements for the on-chip analog baseband filters and the digital filters in the modem can be greatly relaxed. In most single-chip RFIC designs, the baseband filters are one of the largest parts of the RFIC in terms of die area. Thus, the extra cost of the SAW filter will be slightly offset by the reduced size of the baseband filters. The only baseband filtering required is to keep the noise bandwidth of the baseband below a reasonable value and to enhance the stability of the variable gain amplifiers (VGAs). Thus, the actual corner frequency of the analog baseband filters does not affect the chip performance and so would tend to help increase the yield of the chip and therefore reduce its cost. To keep high yields, other architectures like direct conversion, use calibration circuits to accurately set the corner frequency of the baseband filters using the crystal oscillator as the reference. In a superheterodyne architecture, the area for this circuit would not be needed.

If, on the other hand, the implementer desires better jamming performance, the superheterodyne architecture also allows better filtering because it is done at different frequencies. Even at low frequencies on the chip, in the 10's of MHz, the isolation between structures is finite and so the ultimate attenuation of the filter is finite due to the limited isolation in the package and in the die itself. One of the traditional reasons to use a superheterodyne receiver is to provide very high attenuation of close-in jammers, sometimes as much as 80 dB.

One of the disadvantages of the superheterodyne receiver architecture is that it often requires an image filter between the LNA and the first downconverter. If the LNA gain at the image frequency is approximately the same as it is in-band, then the noise-figure of the mixer will be reduced by 3 dB due to the additional noise at the image frequency that is mixed down to the IF. Normally, this image filter is a ceramic filter that provides 20 to 30 dB attenuation of the image frequency.

However, typically, the LNA will have less gain at the image frequency than it does in-band and so the image will not noticeably affect the mixer noise figure. IEEE Std 802.15.3 does not have any requirements for rejecting a jammer outside of the allocated 2.4 GHz band. The requirements for this performance criteria will need to be determined by the implementer with input from the customer. If the requirements are relaxed, then it will be possible to achieve the image filtering requirements through the design of the LNA's bias and output matching networks rather than using a ceramic image filter.

Direct conversion

Although the direct conversion or zero IF (ZIF) architecture has received a great deal of attention in recent years as it is applied to new low-cost radios, it is really one of the oldest radio architectures. In older papers, this architecture is often referred to as a homodyne system. The appeal of the ZIF architecture is that it requires a minimum of external components, as shown in Figure 7-17. There is no image frequency for the ZIF architecture because the LO frequency is the same as the RF frequency. Because of this, there is no requirement for an image filter at the RF. The channel selectivity for the receiver is accomplished on the chip with analog baseband low-pass filters. The matched filter in a digital demodulator will also assist in filtering out the adjacent channel signals due to its low-pass nature. The I/Q down-conversion process only needs to preserve the signal quality and is not required to remove the image frequency. Thus, the requirements on I and Q phase and amplitude balance are usually more relaxed than are those for a receiver that implements image rejection with the mixers, e.g., the low IF architecture.




Figure 7-17: Direct conversion radio architecture

However, the ZIF architecture creates new challenges for the radio architect that are quite different from those of the superheterodyne architecture. For example, although the two-tone, third-order intermodulation test is very important for superheterodyne architectures, a single-tone, second-order intermodulation creates problems for the ZIF architecture. In the superheterodyne architecture, two tones (or one tone and an interfering signal) separated from the desired center frequency by δf and $2*\delta f$ will create a third-order product at the output of the of the amplifier at a frequency equal to the RF or at the output of the mixer at a

frequency equal to the IF. The amplitude of this signal depends on the amplitude of the input signals and the cascaded third-order input intercept point (IIP3) of the receiver up to the channel select filter.

Although IIP3 performance is also an issue in the ZIF architecture, the IIP2 performance is generally a more important problem. In the ZIF receiver, signals created by mismatches in the I and Q mixers and by energy from the LO that leaks out through the RF chain and is

reflected back into the input of the mixer and create a DC voltage at the output of the mixer. Although this may be a very small offset in voltage, perhaps only 1 mV, the baseband amplifiers often have greater than 60 dB of power gain or a factor of 1000 in voltage. Thus, a 1 mV offset at the output of the mixer will turn into a 1 V offset at the output of the baseband filter/VGA combination, which would destroy the dynamic range of the ADC. The IIP2 performance of a mixer generally depends on the matching of the components in the mixer, which can be very good if the designer pays close attention to the layout.



The second-order products of the LNA will also create a DC voltage at the output of the LNA, but because essentially all LNA designs are AC coupled (i.e., they have a series capacitor on the output), the DC voltage is completely blocked. Thus, the IIP2 performance of a typical LNA is essentially infinite.

In addition to requiring a high IP2 mixer, the ZIF architecture can deal with the DC offset problem in one of two ways. The first technique is to require high-pass characteristics from the baseband circuitry, usually by including a series capacitor. However, if the high-pass corner is too high, then the receiver performance will be degraded due to the missing signal energy that was filtered by the high-pass corner. If the high-pass corner is too low, then the filter response time will be too slow and none of the frames will be successfully received. The high-pass corner normally occurs in the 10's of kHz, which requires at least two large value capacitors, one for I and one for Q³⁷. Depending on the fab process that is used for the RFIC design, these capacitors may need to be off-chip. The second technique for handling the DC offset problem is to use a DC compensation loop in the baseband filters and VGAs. The design of these loops is well-known, and the

³⁷ For a differential baseband implementation, which is typical on single chip radios, it would require twice as many components or four capacitors plus at least four pins in the package to access them.

implementation takes up a relatively small portion of the overall analog baseband die area.

However, removing the DC offset will also desensitize the receiver to signal energy around DC as well. For signals without DC content, such as orthogonal FSK or the Barker code in 802.11b, there is essentially no information at DC and this energy can be removed without affecting the performance of the demodulator. In the case of 802.15.1, the modulation is nonorthogonal FSK, and so the down-converted signal contains significant DC components. This modulation choice combined with a very short preamble made it difficult for implementers to use ZIF as an architecture for 802.15.1. Although some implementers were successful with the ZIF architecture, many opted for other methods such as low IF to work around this issue.

The direct conversion architecture also requires that the implementer create good quadrature at the RF frequency. For other standards, e.g., 802.15.1 or 802.11b, the phase and amplitude balance requirements are not very strict. The higher-order modulations of the 2.4 GHz PHY of 802.15.3, e.g., 32 and 64 QAM, on the other hand, have a relatively challenging quadrature phase and amplitude requirements to provide good signal quality on transmit (i.e., low EVM) as well as enabling low BERs for the receiver. In general, creating the quadrature signals from the VCO becomes more difficult as the frequency increases, and so trying to create this at 2.4 GHz is more difficult than creating it at a lower IF.

Walking IF

The walking IF architecture is a dual-conversion architecture that uses only one synthesizer and does not have a fixed IF frequency. Instead, because both LOs change when the channel changes, the IF will change as well. Although the IF changes with the RF frequency, the total IF bandwidth is not as large as the total RF bandwidth. The walking IF architecture, illustrated in Figure 7-18, has the advantage of using only one VCO and synthesizer while gaining many of the advantages of a superheterodyne system.

One of the key design decisions in a walking IF architecture is the choice of the divider ratio used to generate the second LO. If a divide by two is selected, then the first LO is two-thirds of the RF frequency and the second LO is one-third of the RF. This would be referred to as a two-thirds walking IF architecture. As with




Figure 7-18: Walking IF radio architecture

a superheterodyne architecture, the designer needs to consider the frequencies where there may be interference when selecting the IF frequency as well as considering the location of possible spurs. A four-fifths walking IF would use a first LO of 1.9296 to 1.9696 GHz and have an IF center frequency that varied from 482.4 to 492.4 MHz. With the 15 MHz signal bandwidth, the image filter would need to span approximately 25 MHz at a center frequency of 487 MHz. An advantage of the four-fifths walking IF architecture is that the divide by four circuit for the second LO would provide nearly perfect quadrature signals for the final down-conversion to the I and Q baseband signals.

The frequencies required for a walking IF architecture using low-side injection for the mixer can be determined from the following equations:

$$F_{LO1} = F_{RF} \frac{(N-1)}{N} \quad (14)$$

$$F_{LO2} = F_{IF} = \frac{F_{RF}}{N} \quad (15)$$

$$BW_{IF} = \frac{BW_{RF}}{N} \quad (16)$$

A walking IF radio can also be implemented using high-side injection for the first down-mixer. The IF frequencies and the IF bandwidth are the same in this case,

only the LO is different. The LO frequency required for the walking IF using high-side injection for the mixer is:

$$F_{LO1} = F_{RF} \frac{(N+1)}{N} \quad (17)$$

Because the IF is moving in frequency, it is not possible to completely filter the adjacent channel jammers with the channel select filter. Instead, this filter only attenuates jammers in the alternate channel and those at larger frequency separation. If the second mixer has sufficient dynamic range, it is possible to use only a simple filter at the IF and perform the jammer rejection with the analog baseband and digital filters in the modem. For 802.15.3, the IF bandwidth required for the walking IF architecture, 25 MHz, is not much larger than the bandwidth required for a superheterodyne architecture, 15 MHz. On the other hand, for 802.15.1, the bandwidth for the walking IF, 17 MHz, is much larger than for the superheterodyne channel select filter, 1 MHz. So, from the point of view of the complexity of the channel select filters, the walking IF architecture is a better choice for 802.15.3 than it would be for 802.15.1.

Although the walking IF architecture does eliminate one VCO and one synthesizer as compared with a superheterodyne architecture, it has some drawbacks that are similar to the superheterodyne architecture. For example, it requires some form of image rejection, either in the form of an image filter or with an image reject mixer. The walking IF architecture also requires an IF channel select filter that is either off-chip or requires relatively large on-chip components. Because the walking IF uses two LOs, it will generate more spurious responses, and so the implementer will need to pay careful attention to the frequency plan.

Another disadvantage of the walking IF architecture is that it requires a PLL step size that is smaller than the channel spacing. For example, in 802.11b and 802.15.3, the PLL is required to support frequency step sizes of 1 MHz in order to be able to tune to the channel frequencies of 2.412, 2.437, and 2.462 GHz. On the other hand, a four-fifths walking IF PLL would have to support a frequency step size of 800 kHz in order to create the LO frequencies of 1.9296, 1.9496, and 1.9696 GHz. If the PLL uses integer division, the comparison frequency with the crystal could then be no higher than 800 kHz. This lower-comparison frequency will result in a slower lock time for the synthesizer, which can be important in

frequency-hopping systems like 802.11 FHSS and 802.15.1. In the case of 802.15.3, this is not an issue as the frequency switch time is relatively long (500 μ s) because the WPAN changes frequencies fairly infrequently. The smaller PLL step size can also be alleviated by using a fractional-N synthesizer instead of an integer-N synthesizer.

Low IF

An alternative to either ZIF or the relatively high IF of either the dual-conversion or walking IF architectures is to use a low IF. If the IF is equal to one-half of the RF bandwidth, then the architecture is referred to as a very low IF (VLIF) architecture. For the 2.4 GHz PHY, the IF for a VLIF architecture would be 7.5 MHz. In terms of the components required, the VLIF, as illustrated in Figure 7-19, is similar to the ZIF architecture in Figure 7-17. The key difference between the two is that in the case of the low IF or VLIF, the adjacent or alternate channel jammers will be at the image frequency and the analog baseband filters will not be able to filter out this interference. Thus, the LIF or VLIF architecture needs to provide a much higher level of image rejection.

In Figure 7-19, the image rejection is obtained using complex filters. A filter that has a real time-domain response will have an even amplitude response in frequency; i.e., negative frequencies are attenuated the same amount as the positive frequencies. A complex filter, on the other hand, can be constructed to have an odd response in the frequency domain; i.e., the positive frequencies will be passed, whereas the negative frequencies are attenuated or vice versa. The complex filter is realized by feeding part of the I channel back in to the Q channel and vice versa. The filters are more difficult to implement than is the standard analog low-pass filter due to the higher complexity, and most designers are not familiar with the techniques required to implement them. However, some companies have implemented this approach in 802.15.1 receivers.

As an alternative to complex filters, the receiver can also reject the image frequencies by performing a complex multiplication on the I and Q channels, as shown in Figure 7-20. In this implementation, the complex IF signal is multiplied by a complex LO that converts the IF signal to baseband. This multiplication shifts the baseband spectrum from the IF to DC without folding the image back into the pass band. Normally, this multiplication is performed digitally in the

modem, e.g., with a CORDIC³⁸ algorithm, after the baseband I and Q channels have been sampled with an ADC. Because the final down-conversion is performed in the digital domain, the yields from this part of the system are very high due to the repeatability of modern digital IC processes.




Figure 7-19: Low IF radio architecture




Figure 7-20: Alternate low IF radio architecture

³⁸ CORDIC is an acronym for COordinate Rotation Digital Compute, which is a class of shift-add algorithms for rotating vectors in a plane.

The complex multiplication requires a total of four real multiplications and two additions which can be seen by multiplying two complex numbers, e.g.:

$$\gamma_1 \times \gamma_2 = (\alpha_1 + j\beta_1) \times (\alpha_2 + j\beta_2) = \alpha_1\alpha_2 - \beta_1\beta_2 + j(\alpha_1\beta_2 - \alpha_2\beta_1) \quad (18)$$

Because of this property, the complex multiplication is sometimes represented by four digital mixers and could be implemented with four analog mixers as well.

Another method to obtain the image rejection required for a low-IF receiver is to use the Weaver architecture [B30], as illustrated in Figure 7-21. In this architecture, two sets of analog mixers are used to down-convert the signal to a low IF without the image. The demodulator will need to either subsample the IF or mix it down to baseband. As the image has been removed, the demodulator does not need to implement a full complex multiplication but instead only needs to perform an I/Q down-conversion. One of the advantages of this architecture is that only one ADC is required to digitize the signal for the demodulator.

The low-IF receiver is more challenging to design than is the ZIF receiver due to the increased image rejection requirements. However, the low IF does not suffer from the DC offset problems that complicate the ZIF design. In addition, it does not require external filters, as in the case of the superheterodyne and walking IF architectures.




Figure 7-21: Low IF using the Weaver architecture

Summary of radio architectures

The choice of a radio architecture depends on many different factors. Part of the considerations are the characteristics of the modulation type and frequency. Perhaps equally important, however, are the skills of the RFIC design team that is going to do the actual design. No radio architecture will simultaneously solve all of the problems; instead, each one has a different area in which the designers will find difficulty.



It is the author's opinion that selecting a particular radio architecture will not remove the pain in a design; it only allows you to determine where you will get it. TANSTAAFL is especially true in RF design. (TANSTAAFL = There ain't no such thing as a free lunch.)

Table 7-24 summarizes some of the characteristics of each of the radio architectures discussed in this section.

Any one of the architectures discussed in this section can be used to design a low-cost 802.15.3 standard 2.4 GHz PHY, although the architectures that require a minimum of external filters, i.e., the ZIF and LIF/VLIF, will probably be the ones most likely to be implemented for the low-cost, relaxed performance required for 802.15.3.

Table 7-24: Comparison of some receiver architectures

Characteristic	Super-het	Walking IF	ZIF	LIF/VLIF
Off-chip filters?	Yes ^a	Yes ^a	No	No
Off-chip image filter?	Yes ^b	Yes ^b	No	No
Strong image reject performance?	No	No	No	Yes
DC offset correction required?	No	No	Yes	No
Complexity	Highest	High	Lowest	Low
Jamming resistance	Best	Good	Fair	Fair

^a In the case of the superheterodyne and walking IF architectures, the channel select filter can be replaced with on-chip filtering and baseband low-pass filters. However, this approach requires substantially more dynamic range out of the second set of mixers and discards some of the performance advantages of these architectures.

^b Depending on the actual design and performance requirements, these filters could be implemented on the chip.

Chapter 8 Interfacing to 802.15.3

IEEE Std 802.15.3 specifies MAC and PHY layers to provide interoperability between compliant DEVs that implement the standard. In addition to the requirements for the MAC and PHY, a typical implementation also includes, at least conceptually, a variety of management entities that control the MAC and PHY and provide interfaces to allow higher layers to control the operation of the MAC and PHY. Each of these management entities provides a layer management service interface, referred to as a service access point (SAP) through which the management entity is controlled. Figure 8-1 provides an illustration of the various management entities and the SAPs that are used for communication and control.




Figure 8-1: Reference model used in IEEE Std 802.15.3

The reference model in Figure 8-1 represents only one possible arrangement of the layers and management entities. An implementer is only required to provide a minimum interface at the frame convergence sublayer (FCSL) SAP and at the PHY antenna. None of the other interfaces are mandatory, and so they can be implemented in a manner that is the most efficient for the implementer. This model is relevant to the text and descriptions used in the standard as it provides a reference for the terms and divisions that were used while developing the standard.

The device management entity (DME) is intended to provide "glue" between the various layers and to implement all of the undefined control functions for the various layers. The exact function of the DME is not specified in the standard. Instead, the DME provides a location where all of the undefined procedures of the standard exist. In general, if the actions by the DEV or the PNC are fully-defined in the standard, then they will exist in either the PHY or the MAC. Any implementation-dependent functionality in the standard resides, at least conceptually, in the DME. The only exception to this rule is the channel time allocation (CTA) algorithm in the PNC. Instead of passing the request for channel time to the DME, the standard indicates that PNC's MAC/MLME makes the decision on allocating channel time, but it does not indicate the method that will be used to make the decision. The Task Group wanted to keep the CTA algorithm as an implementation dependent function for a couple of reasons. First, leaving this outside of the standard allows companies to innovate in this area to provide differentiation to products that would otherwise be functionally the same. A PNC with a good CTA algorithm will be able to better maintain the QoS in piconet and to provide greater overall throughput. Both of these are qualities of the piconet that are visible to the end user in the form of the performance of their applications.

The second reason that the CTA algorithm is not specified in the standard is that no member of the group proposed one for inclusion in the standard. Most companies consider their CTA algorithm to be a proprietary advantage in developing solutions, and so they do not want to publish their methods. In addition, even if CTA algorithms had been proposed, it is unlikely that the task group would have been able to come to agreement on one algorithm to require in the standard. The "correct" algorithm depends on the type of traffic that the PNC expects to support, the level of complexity that can be supported by the price of the implementation, and so on.

Although some interfaces between the management entities are defined in the standard, other key interfaces are only implied and are not specified in the standard. For example, the FCSL to DME interface is not defined, and yet this interface is required for the creation of streams. Isochronous data come into the FCSL, which maps them into stream indices, but it is the DME in the standard that requests, modifies, and terminates all isochronous stream allocations.

Another missing interface is the interface between either the DME or the FCSL and the higher layers to enable them to request an isochronous stream connection. The current 802.2 FCSL only defines an interface to pass individual frames that do not have a relationship to any of the other frames that might be passed to the FCSL. In the case of a stream, the frames that enter the FCSL really belong to one of the allocated streams in the superframe, but there is not an interface that allows the FCSL to be able to map the incoming data frames into the existing streams in the piconet.

The standard also does not specify the interface between the MLME and the MAC or between the PLME and the PHY. In modern implementations, the layer management entities are tightly bound to their respective layers, and there is no need to specify a separate interface that will never be exposed in practical implementations.

The SAPs that are defined in the standard are the following:

- FCSL SAP
- MAC SAP
- PHY SAP
- MLME SAP
- PLME SAP
- MLME-PLME SAP

The PLME SAP and the MLME-PLME SAP use identical primitives, and so they are viewed instead as a single SAP, referred to as the PLME SAP, that can be used either by the MLME or the DME to control the behavior of the PHY. Of course, the actual implementation does not have to implement separate interfaces or to even support the PLME SAP as an external interface.

The MLME SAP in the standard is a relatively detailed interface (it has 22 subclauses, most with more than one subclause describing a different primitive). In fact, it is much more detailed than is the equivalent interface defined for 802.11. The Task Group spent an enormous amount of time on specifying this interface, which is optional anyway, to enable the formal verification of the standard using methods like the structured definition language (SDL). SDL simulators allow a protocol to be modeled and tested to find deadlocks, live locks, and other problems in the protocol in a formal manner. SDL modeling tools also allow the generation of message sequence charts (MSCs) from the SDL model as well as generating object code that can be used to instantiate the protocol. The MLME SAP can be used with an SDL modeling tool to model and test the 802.15.3 protocol as it is written in the standard (with the exception of the CTA algorithm, which is in the MLME but is not specified).

Quite a bit of effort was spent in trying to keep the MLME SAP up to date with the protocol description contained in the clauses for the frame formats (Clause 7) and the MAC functional description (Clause 8) in the standard. The MLME SAP can be viewed as being derived from the other two clauses rather than introducing its own functionality. Because it is a derived clause, many of the technical comments received in the letter ballots had to do with changes in the protocol that were specified in Clause 7 and Clause 8 but were not reflected by changes in Clause 6. In order to create an interoperable, 802.15.3-compliant DEV, the implementer only needs to refer to Clause 7, Clause 8, Clause 9 (if security is implemented), and Clause 11. The information in Clause 6 for the MLME SAP is only a guide.

The standard explicitly states that the interfaces described in the layer management clause (Clause 6) do not have to be implemented as detailed in the standard if they are not exposed as external interfaces. Even if the interfaces are exposed, the implementer "should" support the primitives defined in the clause, but they are not required to support those primitives. If the implementer does support the primitives defined in the layer management clause, they are not restricted from creating additional primitives that may be required to complete the implementation or to add vendor specific functionality to the MAC.

THE PIBS AND THEIR INTERFACE

The PAN information base (PIB) contains information that is similar to that found in the management information base (MIB) used in other standards like 802.11.

The standard uses a different name to differentiate between the data that might be managed over a network using SNMP, i.e., the MIB, versus information that is controlled only by the local DEV and its applications, i.e., the PIB. The PIB is not intended to be used for centralized management because the assumption of IEEE Std 802.15.3 is that the DEVs often operate in piconets that do not have connections with the larger networks, e.g., LANs, MANs, or WANs.

The information in the PIB can be written using either the MLME-SET.request primitive or the PLME-SET.request primitive. Likewise, the MLME-GET.request primitive or the PLME-GET.request primitive are used to retrieve information from the PIB. These primitives use the same syntax, and so they are generically referred to in the standard as XX-GET and XX-SET. The MLME primitives retrieve or set information in the MAC PIB, whereas the PLME versions of the primitive are used to retrieve or set information in the PHY PIB. The MAC and PHY PIBs can be thought of as being contained in MLME and PLME, respectively. Thus, the PIBs contain information that will typically be used by either the MAC or the PHY in its normal operation.

The PIB entries may be either read only, write only, or read-write, depending on how they are defined in the PIB tables in the standard. If the XX-SET.request primitive attempts to write a read-only PIB entry, it will generate an error. Likewise, if the XX-GET.request primitive tries to retrieve the value of a write-only PIB entry, it will also generate an error. A PIB entry that has the characteristic of write-only restricts the operation of the XX-GET primitive. The PLME or MLME can read or write this PIB entry; it is only the DME that is forbidden to read it.³⁹

³⁹ Some engineers at a semiconductor company were almost successful in releasing an official data sheet for "write-only memory" as a joke, but were caught by a unusually vigilant marketing staff.

MLME SAP

The MLME SAP, as described earlier, provides a relatively complete list of primitives that can be used to model the defined behavior of the standard. Because these optional primitives are linked to the normative behavior of the frame exchanges in the wireless medium, they are also referenced in the MAC functional description (Clause 8 in the standard) in both the text and the MSCs. However, the presence of these primitives in the text of Clause 8 or in the MSCs of Clause 8 should not be viewed as requiring a particular implementation of this interface. As it is stated in Clause 6, this interface is optional and implementers are free to create any division of functionality within the MAC and to create an interface that is appropriate to that division.

Therefore, the MLME SAP should not be viewed as an indication of the correct architecture for the MAC or as the point at which the MAC is split into hardware and software. An implementer may create a product in which essentially all functions of the MAC and DME are in hardware, whereas another implementer may realize essentially all of the functions in software. Either implementation is correct as long as it adheres to the frame formats in Clause 7 and the sequence of frames sent in the wireless medium as defined in Clause 8 of the standard. The primitives provided by the MLME SAP are summarized in Table 8-1. Because the MLME SAP is simply a derived interface, the functionality provided by these has already been described in Chapter 4, Chapter 5, and Chapter 6 of this handbook.

The description of the MLME SAP primitives in Clause 6 in the standard makes the mistake of trying to also describe some of the functionality of the DME, which is outside of the scope of the standard. All of the MLME SAP primitives in Clause 6 include subclauses for "when generated" and "effect of receipt." However, these subclauses are only relevant when they refer to the operation of the MAC and cannot apply to the operation of the DME, which is not supposed to be specified. Accordingly, the XX.request and XX-response primitives should have only the "effect of receipt" subclause because this describes what the MAC will do when it receives these primitives. In a similar fashion, the XX.indication and XX.confirm primitives should only specify the "when generated" subclause and not the "effect of receipt" subclause because these indicate an action that the MAC has taken. The DME's response to the XX.indication and XX.confirm primitives is not

Table 8-1: Summary of MLME primitives

Name	Request	Indication	Response	Confirm
MLME-RESET	X	-	-	-
MLME-SCAN	X	-	-	X
MLME-START	X	-	-	X
MLME-START-DEPENDENT	X	-	-	X
MLME-SYNCH	X	-	-	X
MLME-ATP-EXPIRED	-	X	-	-
MLME-ASSOCIATE	X	X	X	X
MLME-DEV-ASSOCIATION-INFO	-	X	-	-
MLME-DISASSOCIATE	X	X	-	X
MLME-REQUEST-KEY	X	X	X	X
MLME-DISTRIBUTE-KEY	X	X	X	X
MLME-DE-AUTHENTICATE	X	X	-	X
MLME-MEMBERSHIP-UPDATE	X	-	-	-
MLME-SECURITY-ERROR	-	X	-	-
MLME-SECURITY-MESSAGE	X	X	-	X
MLME-PNC-HANDOVER	X	X	X	X
MLME-NEW-PNC	-	X	-	-
MLME-PNC-INFO	X	X	X	X
MLME-SECURITY-INFO	X	X	X	X
MLME-CREATE-ASIE	X	-	-	X
MLME-RECEIVE-ASIE	-	X	-	-
MLME-PROBE	X	X	X	X
MLME-ANNOUNCE	X	X	-	X
MLME-PICONET-SERVICES	-	X	X	X
MLME-CREATE-STREAM	X	-	-	X
MLME-MODIFY-STREAM	X	-	-	X
MLME-TERMINATE-STREAM	X	X	-	X
MLME-MULTICAST-RX-SETUP	X	-	-	-
MLME-CHANNEL-STATUS	X	X	X	X
MLME-REMOTE-SCAN	X	X	X	X
MLME-PICONET-PARM-CHANGE	X	-	-	X
MLME-TX-POWER-CHANGE	X	X	-	X
MLME-PS-SET-INFORMATION	X	-	-	X
MLME-SPS-CONFIGURE	X	-	-	X
MLME-PM-MODE-CHANGE	X	-	-	X
MLME-PM-MODE-ACTIVE	-	X	-	-

defined in the standard. When this mistake was realized, it was too late in the process to clean it up without potentially introducing more technical errors in the standard. Instead, this clean up of the layer management clause will have to wait until an amendment project authorization request (PAR⁴⁰) is approved for IEEE Std 802.15.3.

With those caveats in mind, the MLME SAP does provide some additional insight into the functions that are fully defined in the standard versus functions that are either implementation specific or are outside of the scope of the standard. For example, consider the primitives associated with power management, MLME-PM-XX, MLME-SPS-XX, and MLME-PS-XX. The standard includes the descriptions of the primitives that are generated by the power save DEV, the XX.request and XX.confirm, because the DEV determines its power-save mode characteristics in a manner that is outside of the scope of the standard. The DEV will want to consider the characteristics of its PHY in terms of AWAKE state versus SLEEP state power drain as well as the potential traffic that it could be sending or receiving.

On the other hand, the PNC's participation in the power-save mode process does not have any MLME SAP primitives associated with it. A quick examination of the MSCs in the power management subclause of Clause 8 will show that there is absolutely no interaction with the PNC DME. In fact, the PNC DME appears in only one of the MSCs in the subclause, and in that instance, it could have been deleted as well. The behavior of the PNC for supporting the power-save modes is very straightforward and is well defined in the standard. Most of its activity is spent setting bits in IEs, placing these in the beacon, and updating the next wake beacon from time to time. The PNC will generally accept all requests by the DEVs (unless they make a mistake or get mixed up in the order that the frames are supposed to be sent in the wireless medium). Thus, no MLME SAP primitives are required by the PNC to support the power-save mode DEVs.

⁴⁰ A PAR is a project authorization request, the means by which standards projects are started in the IEEE Standards Association. PARs define the scope, purpose, and contact points for the new project.

PLME SAP

The PLME SAP represents another conceptual interface in the standard. This interface provides higher-level capabilities to control the PHY. A total of five primitives in three classes are defined for this SAP. The list of supported primitives is given in Table 8-2.

Table 8-2: Summary of PLME SAP primitives

Name	Request	Indication	Response	Confirm
PLME-RESET	X	—	—	X
PLME-TESTMODE	X	—	—	X
PLME-TESTOUTPUT	X	—	—	—

The PLME primitives are pretty self-explanatory. The PLME-RESET causes the PHY to clear any settings and set the PHY in the off state (i.e., neither transmitting nor receiving). The other two primitives are provided to enable testing of the PHY without enabling the MAC functionality. The PLME-TESTMODE primitive is used to set up the PHY for transmitting one of three test patterns, and the PLME-TESTOUTPUT primitive instructs the PHY to begin the transmission. These primitives only provide a small portion of the test characteristics that would normally be implemented in a design. An actual implementation would also have receiver test modes and possibly a set of loop-back mode for self-testing of the PHY.

MAC SAP

The MAC SAP provides an interface between one or more FCSL layers and the MAC sublayer. The FCSL takes frames passed from higher layers and identified with source and destination addresses and maps this data into the SrcID, DestID, and stream index to be used in transmitting the data. The MAC SAP then provides six primitives, three each for isochronous and asynchronous data, which are summarized in Table 8-3.

Table 8-3: Summary of MAC SAP primitives

Name	Request	Indication	Response	Confirm
MAC-ASYNC-DATA	X	X	—	X
MAC-ISOCH-DATA	X	X	—	X

The asynchronous primitives specify the target and destination IDs of the data, whereas the isochronous primitives use only the stream index. The stream index is sufficient for isochronous streams because each assigned stream index applies to only one SrcID/DestID pair. Asynchronous data, on the other hand, shares a single stream index, and so it is not possible to determine the SrcID/DestID pair from the stream index alone. Each of the XX.request primitives also includes a TransmissionTimeout parameter that limits the length of time that a DEV will attempt to send the frame before it gives up on it and discards the frame.

PHY SAP

In IEEE Std 802.11, the PHY is split into two sublayers, the physical layer convergence protocol (PLCP) and the physical medium-dependent (PMD) sublayers. The initial IEEE Std 802.11 supported three different, non-interoperable PHYs,⁴¹ and so the working group used the PLCP to provide an interface between the three incompatible layers and the single MAC.

Instead of this approach, IEEE Std 802.15.3 defined only the PHY layer and did not explicitly split it up into separate sublayers. The functionality equivalent to the PLCP and PMD sublayers of 802.11 is included in the 802.15.3 PHY layer; it just is not split into two different sublayers. However, this does not mean that IEEE Std 802.15.3 does not support multiple PHYs. The PHY SAP interface provides sufficient abstraction so that other future PHYs can be implemented with the 802.15.3 MAC. This process is already underway with the development of 802.15.3a, a higher-rate PHY for the 802.15.3 MAC.

⁴¹ IEEE Std 802.11 now defines a total of nine different PHYs, none of which are mandatory and six of which, 802.11 DSSS, 802.11b (2 PHYs), and 802.11g (3 PHYs) are interoperable at some level.

The development process of this new PHY uses the PHY SAP to define the demarcation between the MAC and the PHY. For example, all of the proposals will be evaluated based on the throughput that is delivered at the PHY SAP interface. This includes the overhead for the PHY frame, on-channel timing, and any forward error correction (FEC) if it is included in the proposals.

The PHY SAP provides primitives in nine categories that relate to the operation of the PHY. The primitives that are provided by the PHY SAP are summarized in Table 8-4.

Table 8-4: Summary of PHY SAP primitives

Name	Request	Indication	Response	Confirm
PHY-DATA	X	X	—	X
PHY-TX-START	X	—	—	X
PHY-TX-END	X	—	—	X
PHY-CCA-START	X	—	—	X
PHY-CCA-END	X	—	—	X
PHY-CCA	—	X	—	—
PHY-RX-START	X	X	—	X
PHY-RX-END	X	X	—	X
PHY-PS	X	—	—	X

The PHY-TX-START, PHY-TX-END, PHY-RX-START, and PHY-RX-END are self-explanatory; they provide methods to begin and end transmissions and receptions. The PHY-RX-START and PHY-RX-END primitives also provide indications because the DEV will not necessarily know when a frame that is intended for it will arrive at the receiver. In addition, the PHY-RX-END will return additional information regarding the attempt to receive the frame. For example, it will indicate if the signal was lost in the middle of a frame, that the format of the frame was wrong, or that the frame was sent with an unsupported data rate.

The data rate and MAC header are passed to the PHY with the PHY-TX-START.request and are received from the PHY with the PHY-RX-START.indication. The PHY-RX-START.indication also returns the received signal strength indication (RSSI) measured during the MAC and PHY headers, whereas the PHY-RX-END.indication returns the link quality indication (LQI).

Once the PHY is in either transmit or receive mode, the PHY-DATA primitive is used to either pass data to the PHY to be sent or to pass data from the PHY to the MAC as it is received. The data are passed in this primitive one octet at a time, but an implementation may choose to use a different method to transfer the data between the PHY and the MAC.

The CCA primitives are used in the CAP as well as in the retransmission timing. The MAC begins the CCA process by sending the PHY-CCA-START.request and ends the process by sending the PHY-CCA-END.request. While the process is running, the MAC will receive PHY-CCA.indications periodically that indicate the current status of the channel, either busy or idle.

THE FCSL

The frame convergence sublayer (FCSL⁴²) was developed by the 802.15.3 Task Group when it became apparent that the IEEE 802.2 LLC would not provide an interface that could support all of the types of multimedia data that IEEE Std 802.15.3 was intended to provide. The IEEE 802.2 LLC is a data-centric layer that works well for asynchronous data traffic in a network. However, it was not designed with streams of isochronous data in mind. There are other higher layer standards, e.g., IEEE Std 1394 and USB, that do provide support for the creation, modification, and termination streams of isochronous data. Because an interface for data and legacy equipment is a requirement for the standard, IEEE Std 802.15.3 needed to provide the standard IEEE 802.2 LLC interface. However, to support alternative methods of negotiating stream connections, IEEE Std 802.15.3 created a thin layer, the FCSL, that provides a method for multiple LLC layers to interface with the 802.15.3 MAC. This interface is illustrated in Figure 8-2.

⁴² For the benefit of the curious reader, the acronym is FCSL instead of FCS because the frame-check sequence preceded the naming of the FCSL by at least a year. The FCSL was originally named the service-specific convergence sublayer (SSCS), but it was renamed in one of the Working Group letter ballots.

The FCSL performs the following functions when it receives a protocol data unit (PDU) from an upper layer:

- Maps the source addresses and destination address into a SrcID and Dest-ID.
- Uses the priority and other information to determine the correct stream index.
- If the frame needs a stream allocation, it sends a request to the DME to create a new stream.
- Delivers the PDU to the MAC SAP either as an asynchronous MSDU with a SrcID and DestID or as an isochronous MSDU with an associated stream index.

When the FCSL receives an MSDU from the MAC SAP, it reforms the following actions:

- Maps the received DestID and SrcID or the stream index into a source address and a destination address.
- Delivers the PDU to the appropriate higher layer.




Figure 8-2: Frame convergence sublayer model

IEEE Std 802.15.3 defines only one FCSL, the IEEE 802.2 FCSL. Other FCSLs may be defined as a part of a later MAC enhancement PAR, via an industry SIG or by another standards group, e.g., IEEE Std 1394.⁴³

⁴³ The IEEE Std 1394 is often referred to as FireWire™.

Trade Association (TA) has recently started a project (see [B21]) to write a protocol adaptation layer (PAL) for 802.15.3. A 1394 PAL is roughly equivalent to the 802.15.3 FCSL. When this project is completed, it should provide a support for the upper layer parts of 1394 on the 802.15.3 MAC and PHY.

The IEEE 802.2™ [B11] FCSL provides QoS support by either using a best effort scheme or by using the hierarchical IEEE 802.1D™ [B10] QoS priority scheme. The best effort scheme treats all frames the same; i.e., there is no attempt to deliver some frames with less delay than other frames. The 802.1D QoS priority scheme defines eight possible priority levels with seven defined traffic types. The 802.1D™ traffic types are listed in Table 8-5.

Table 8-5: IEEE Std 802.1D traffic types

User priority	Traffic type	Used for	Comments
0 (default)	Best effort (BE)	Asynchronous	Default level
1	Background (BK)	Asynchronous	—
2	—	Spare	Currently not assigned
3	Excellent effort (EE)	Isochronous	For valued customers
4	Controlled load (CL)	Isochronous	Traffic will have to conform to some higher protocol layer admission control
5	Video (VI)	Isochronous	< 100 ms delay and jitter
6	Voice (VO)	Isochronous	< 10 ms delay and jitter
7	Network control (NC)	—	—

The IEEE 802.2 FCSL provides only three primitives to send and receive data, MA-UNITDATA.request, MA-UNITDATA.indication, and MA-UNITDATA-STATUS.indication. The MA-UNITDATA.request primitive is used to send a PDU, whereas the MA-UNITDATA.indication is used to notify the higher layers that a PDU has been received. The MA-UNITDATA-STATUS.indication is used to report the status of the transmission attempt for the PDU from the immediately-preceding MAC-UNITDATA.request.

Chapter 9 Coexistence mechanisms

INTRODUCTION

The wireless medium is, by its nature, a shared medium. Because of this, the development of the 802.15.3 standard considered various methods of coexistence and adopted various methods and recommendations that enable compliant devices (DEVs) to achieve a level of coexistence with other IEEE standards. There are many definitions that have been proposed for coexistence, but for the purpose of IEEE Std 802.15.3, coexistence is defined as "The ability of one system to perform a task in a given shared environment where other systems have an ability to perform their tasks and may or may not be using the same set of rules." This definition recognizes that the impact of two systems on each other is not a single number, but rather a group of performance curves that depend on the physical configuration of the networks, the amount and type of traffic on the networks, and the specific choices made by an implementer.

In general, as soon as any two wireless systems that are operating in the same general frequency are brought into proximity with each other, there is some degradation in the throughput of at least one of the systems. Thus, it is not realistic to expect that there would be no effect on either system by the presence of another. Even if the systems cooperate with each other, the methods used only minimize the reduction in throughput, they do not to eliminate it entirely. For example, if the networks share the medium agreeing to a time-sharing protocol, e.g., time division multiplexing (TDMA), the throughput of the systems is reduced by the amount of time that is allocated to the other network. However, this sharing can improve the aggregate throughput as compared with the case when the two networks have not synchronized their timing.

COEXISTENCE TECHNIQUES IN 802.15.3

The coexistence approach for 802.15.3 focused on a variety of areas, ranging from the selection of the PHY channels to the ability to coordinate timing with networks that are not 802.15.3 compliant. In all, 802.15.3 identified a total of

eight different techniques that can be used to improve the level of coexistence among 802.15.3 networks and with other types of wireless networks. These techniques include the following:

- Passive scanning
- Dynamic channel selection
- The ability to request channel quality information
- Link quality and RSSI
- A channel plan that minimizes channel overlap
- Transmit power control
- Lower impact transmit spectral mask
- Neighbor piconet capability

Although all of these techniques are in 802.15.3, not all of them are applicable to improve the level of coexistence with other wireless networks. Table 9-1 lists the coexistence techniques and indicates if they are applicable to a particular IEEE standard. In the table and throughout this chapter, 802.11b also implies 802.11 DSSS and 802.11g because of the similarities in the PHYs. Likewise, 802.15.1 also includes the Bluetooth specification because it is the IEEE standardization of that specification and includes 802.11 FHSS by virtue of the similarities of the PHYs.

Table 9-1: Applicability of coexistence techniques to IEEE standards

Technique	802.11b	802.15.1
Passive scanning	X	X
Dynamic channel selection	X	X
Requesting channel quality information	X	X
Link quality and RSSI	X	X
Channel plan to minimize overlap	X	—
Transmit power control	X	X
Lower impact transmit spectral mask	X	X
Neighbor piconet capability	X	—

Passive scanning

When an 802.15.3 DEV wants to either create or join a piconet, it is required to passively scan a potential channel before attempting to start a piconet. Passive scanning means that the process of starting or joining a piconet will not cause an interruption to existing networks. If the DEV is trying to start a new piconet, it will be looking for a channel that is relatively quiet. This increases the probability that it will not start up a piconet on a channel that is already in use by another wireless network. In addition, an 802.15.3 DEV may be able to determine that an 802.11b WLAN is present in the area. In this case, the 802.15.3 DEV will avoid using the channel in use by the 802.11b WLAN.

The ability to request channel quality information

In order to improve coexistence, it is useful to have information about the amount of interference in each of the channels as well as information about other wireless networks present in the channel. In the case of 802.15.3, not only does the PNC sense the channel in its area, but it is also capable of asking any other DEV to respond with its own estimate of the channel status. This command is useful for detecting coexistence problems in remote DEVs by the PNC or other DEVs that are unable to detect an interference environment (for example, during a passive scan).

Dynamic channel selection

The 802.15.3 piconet coordinator (PNC) periodically requests channel status information from the DEVs in the piconet. The DEVs respond with statistics about how many frames have been successfully transmitted or received. This information allows the PNC to determine that the channel is having problems, as it would when an 802.11b WLAN is present. If the PNC determines that the channel is getting too much interference, it searches the other channels to find one that has a lower level of interference. If the PNC finds a channel with less interference, then the PNC moves the piconet to a quieter channel. Thus, if an 802.11b network is present, the 802.15.3 piconet would change channels to avoid interfering with 802.11b. Likewise, if an 802.15.1 PAN is present and is using adaptive frequency hopping to avoid a portion of the spectrum, the PNC could move the piconet to the quiet portion of the spectrum.

Link quality and RSSI

The 2.4 GHz PHY specifies that a DEV returns the RSSI and, for the higher order modulations, an estimate of the link quality. The RSSI provides an estimate of the strength of the received signal, which can be used for transmit power control. The RSSI combined with the link quality indication (LQI) provides a method to differentiate between low signal power and interference causing the loss of frames. For example, if the RSSI is low and frames are being lost, then the cause is low receive power. On the other hand, if the RSSI is relatively high, but the LQI is low, that would indicate the possibility of interference in the channel.

Channel plan that minimizes channel overlap

The channel plan for the 2.4 GHz PHY balances the requirement of four simultaneous piconets with the desire to coexist with other wireless standards, such as 802.11b. To do this, two channel sets are available (see "Channel plan" on page 205). The high-density channel plan was selected to allow four piconets to operate in relatively close proximity. The reduced transmit spectral mask of 802.15.3 allows one more non-overlapping channel as compared with 802.11b. The 802.11b coexistence channel plan, on the other hand, only allocates three channels so that it matches up with the frequencies for 802.11b.

The PNC chooses the channel plan that it will use for the piconet. If the PNC does not detect an 802.11b WLAN, then one of the channels in the high-density channel plan would be used for its piconet. On the other hand, if the PNC detects an 802.11b WLAN, then it would choose the 802.11b coexistence channel plan. The reason for this is that each of the two center channels in the high-density channel plan would overlap two 802.11b channels. The 802.11b coexistence channel plan aligns the channels so that 802.15.3 operation in one of these channels would only affect a single 802.11b channel. Table 9-2 shows how the 802.11b coexistence channel plan for 802.15.3 compares with the channels for 802.11b in North America and Europe. By choosing the "quietest" one, the 802.15.3 piconet would minimize its impact on the 802.11b networks.

If the PNC has already selected the high-density channel plan and is operating on one of the two inner channels, i.e., either channel 2 or 4, it is able to switch channel plans if it detects an 802.11b WLAN in its operational area. In this case, the PNC would use the piconet parameter change process to move the piconet to a

new channel in the 802.11b coexistence channel plan. Note that channel 1 or channel 5 is the same in both plans, so piconets operating on either of these two channels would not necessarily have to change channels if an 802.11b WLAN was detected (unless it is on the same channel).

Table 9-2: Comparison of 802.15.3 and 802.11b channel center frequencies

802.15.3 channel center frequency	802.11b North American channel center frequency	802.11b European channel center frequency
2.412 GHz	2.412 GHz	2.412 GHz
2.437 GHz	2.437 GHz	2.442 GHz
2.462 GHz	2.462 GHz	2.472 GHz

Transmit power control

IEEE Std 802.15.3 provides three methods for controlling transmit power. The first method is that the PNC is able to set a maximum power level for the beacon, CAP, and directed MTSS. For the 2.4 GHz PHY, this can be set as low as 0 dBm. This allows 802.15.3 piconets to reduce the interference they create for other networks while maintaining the operation of the piconet.

The second transmit power control method allows DEVs in a CTA to request a change in the transmit power of the remote DEV for that link. The originating DEV sends a Transmit Power Change command to the target DEV that tells the target DEV the amount to change its power up or down, depending on the received signal power. Thus, two DEVs that are relatively close to each other are able to both save power and reduce the interference to other networks while maintaining a high-quality link. The third method for controlling the transmitter power is that DEVs are allowed to change their transmit power based on their own estimation of the channel.

Lower impact transmit spectral mask

The transmit spectral mask for 802.15.3 was chosen as a balance of complexity, performance, and coexistence. A wider transmit spectrum reduces the filter

requirements, thereby reducing cost. It also results in a slightly better BER performance for a given SNR. However, a wider transmit spectrum also causes more problems to networks operating in adjacent and alternate channels. It also makes the receiver more susceptible to other interference sources. Because of this, this standard chose a narrower transmit spectral mask than is used for either 802.11b or 802.11g, as shown in Figure 9-1.




Figure 9-1: 802.15.3 and 802.11b transmit spectral mask

As the figure shows, the 802.15.3 transmit spectral mask is significantly smaller than the one required for either 802.11b or 802.11g. The smaller mask reduces the interference that 802.15.3 DEVs cause other wireless networks. It also allows narrower receive filters for 802.15.3 DEVs, which reduces the interference caused by other wireless networks in an 802.15.3 piconet.

Neighbor piconet capability

The neighbor piconet capability allows a DEV, which may not be fully 802.15.3 compliant, to request time to operate a network that is collocated in frequency with the 802.15.3 network. Although current 802.11b radios do not implement this functionality, it is relatively easy to build 802.11b radios that could support enough of a subset of 802.15.3 to request the neighbor piconet capability. One reason that this can be done is that the 802.15.3 PHY has many similarities with 802.11b that make it easier to build dual-mode radios.

One of the similarities is that 802.11b and 802.15.3 share the same frequency band, which makes interoperability of radio modules much simpler. Also, the 802.15.3 PHY layer uses 11-Mbaud DQPSK modulation for the base rate, which is the same as the chip rate and modulation for 802.11b. However, 802.11b uses either a Barker code, CCK or PBCC as a spreading code, none of which are a part of the 802.15.3 standard. The 802.15.3 PHY was also chosen with the same frequency accuracy, allowing the reuse of reference frequency source and frequency synthesizers. In 802.11b coexistence mode, 802.15.3 uses the same frequency plan as 802.11b. In addition, the synthesizers that would normally be used in either 802.11b or 802.15.3 would be capable of 1-MHz frequency step size and so would be capable of supporting either frequency plan. The RX/TX turnaround time is also the same for both protocols. However, the TX/RX turnaround for 802.11b is 5 μ s vs. 10 μ s for 802.15.3, which could have an impact on the architecture of a dual-mode radio. To summarize, the similarities between 802.15.3 and 802.11b are as follows:

- DQPSK modulation
- 11-Mbaud symbol (chip) rate.
- Frequency and symbol timing accuracy of ± 25 ppm.
- RX/TX turnaround time
- Power ramp up/down
- Frequency plan

Some of the differences include:

- Barker, CCK or PBCC spreading code
- Power spectral density
- Performance criteria (e.g. sensitivity, jamming resistance, etc.)
- TX/RX turnaround time.
- PHY preamble, header, frame structure
- MAC

A dual-mode 802.11b/802.15.3 access point could request a CTA for a neighbor piconet to share time with an existing 802.15.3 piconet. The AP would use the

PCF and NAV to set aside time for the operation of the 802.15.3 piconet, while maintaining clear operation for part of the time for the 802.11b WLAN.

COEXISTENCE RESULTS

The mechanisms for coexistence in 802.15.3 will enable other WPANs and WLANs to operate in close proximity to 802.15.3 piconets. This section presents some of the calculated results for 802.15.3 piconets operating in proximity with other wireless networks.

Assumptions for coexistence simulations

Although it is possible to do an in-depth simulation and modeling of the 802.15.3 MAC and 2.4 GHz PHY combination, the results would depend heavily on the type and quantity of traffic assumed in each of the wireless networks. Instead, 802.15.3 used a simple PHY layer model, proposed in IEEE Std 802.15.2 [B15], to gauge the impact of 802.15.3 piconets sharing operational areas with 802.15.1 piconets and 802.11b networks.

The calculations rely on IEEE Std 802.15.1 [B14], IEEE Std 802.11 [B12], IEEE Std 802.11b [B13], and IEEE Std 802.15.2 [B15]. For the calculations used to determine the level of coexistence, the following assumptions have been made:

- a) The sensitivity for each of the receivers is the reference sensitivity given in each of the standards, i.e.,
 - 1) -70 dBm for 802.15.1,
 - 2) -76 dBm for 802.11b 11 Mb/s CCK, and
 - 3) -75 dBm for 802.15.3 22 Mb/s DQPSK.
- b) The received power at the desired receiver is 10 dB above the receiver sensitivity. The level of 10 dB was selected because a 10-dB margin results in 10% frame error ratio (FER) in a Raleigh fading channel. Reliable communications without interference would require at least this margin. The distance between the desired transmitter and receiver is not directly specified by this requirement; instead, the transmitter power and the channel model listed below would be used to determine the resulting distance.
- c) The transmitter power for each of the protocols is

- 1) 0 dBm for 802.15.1,
- 2) +14 dBm for 802.11b, and
- 3) +8 dBm for 802.15.3.

The path loss model is one that was proposed for 802.11 and used by 802.15.2:

$$d = 10^{\frac{(P_t - P_r - 40.2)}{20}} \quad \text{for } d < 8 \text{ m}$$

$$d = 8 \times 10^{\frac{(P_t - P_r - 58.5)}{33}} \quad \text{for } d > 8 \text{ m}$$

The path loss is plotted versus the range in Figure 9-2.

- d) The receiver bandwidths are based on the requirements in the standard:
 - 1) 1 MHz for 802.15.1,
 - 2) 22 MHz for 802.11b, and
 - 3) 15 MHz for 802.15.3.
- e) The transmitter spectral masks are the maximum allowed in the standards. This is a very pessimistic assumption because the transmitter sideband spectrum will generally be significantly lower than the spectral mask over most of the frequencies. There are usually only narrow peaks that come close to the required limits. The subclauses that define the transmitter spectral mask for the three standards are as follows:
 - 1) Subclause 7.2.3.1 for 802.15.1,
 - 2) Subclause 18.4.7.3 for 802.11b, and
 - 3) Subclause 11.5.3 for 802.15.3.
- f) The energy from the interfering signal affects the desired signal in a manner equivalent to additive white Gaussian noise (AWGN) in the same bandwidth.
- g) The 802.15.3 piconet operates with the 802.11b coexistence channel plan as described in subclause 11.2.3 of IEEE Std 802.15.3.

For all of the results in this section, the FER calculations assume that one of the two systems is constantly transmitting, while the other one is constantly receiving. If both systems are either both transmitting or both receiving, then the FER will be essentially unaffected by the presence of the other system. For random transmission and reception, this means that the actual FER will only be about one-half of the values presented in the figures. In addition, the MAC protocol and the traffic patterns will affect the actual throughput that is possible with the collocated networks.




Figure 9-2: Path loss versus distance for empirical model

BER calculations

The BER calculations for these simulations are based on the formulas suggested in IEEE Std 802.15.2 [B15]. For 802.11b and 802.15.3, the basic BER calculation depends on the theoretical formula for the BER of QPSK modulation, which in turn depends on the value of the Q function. The Q function is defined as the area

under the tail of the Gaussian probability density function with zero mean and unit variance.

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{t^2}{2}} dt$$

Although the Q function is not commonly found in software packages, it can be calculated from the complementary error function, erfc , which is more common and can be found in spreadsheets.

$$Q(x) = 0.5 \text{erfc}\left(\frac{x}{\sqrt{2}}\right)$$

If neither the Q function nor the erfc are available, the Q function may be approximated using the following fifth-order approximation for values of x greater than 1.

$$Q(x) = \left(\frac{1}{\sqrt{2\pi}}\right) e^{-\frac{x^2}{2}} \times \left(\frac{8 + 9x^2 + x^4}{15x + 10x^3 + x^5}\right)$$

The probability of error in the various modulations can be calculated from the following equations (see [B15] for more details on the derivations):

$$\text{BER}_{22 \text{ Mb/s } 802.15.3} = Q(\sqrt{\text{SNR}})$$

$$\text{BER}_{1 \text{ Mb/s } 802.11} = Q(\sqrt{11 \times \text{SNR}})$$

$$\text{BER}_{2 \text{ Mb/s } 802.11} = Q(\sqrt{5.5 \times \text{SNR}})$$

$$\text{BER}_{5.5 \text{ Mb/s } 802.11} = \frac{8}{15} [14 \times Q(\sqrt{8 \times \text{SNR}}) + Q(\sqrt{16 \times \text{SNR}})]$$



The equations here were derived assuming CCK could be modeled as a block code. This assumption is invalid for small values of the SNR and so these equations will give poor results for low SNR values.

$$\begin{aligned} \text{BER}_{11 \text{ Mb/s } 802.11} = & \frac{128}{255} [24 \times Q(\sqrt{4 \times \text{SNR}}) + 16 \times Q(\sqrt{6 \times \text{SNR}}) \\ & + 174 \times Q(\sqrt{8 \times \text{SNR}}) + 16 \times Q(\sqrt{10 \times \text{SNR}}) \\ & + 24 \times Q(\sqrt{12 \times \text{SNR}}) + Q(\sqrt{16 \times \text{SNR}})] \end{aligned}$$

Determining the BER for 802.15.1 is a little more difficult. The equation for the BER of orthogonal FSK is given by

$$\text{BER}_{802.15.1} = 0.5 \times e^{(0.5 \times \text{SNR})}$$

Although this is a very simple equation, 802.15.1 uses non-orthogonal FSK, which does not have a closed-form solution. However, this approximation is sufficient for the present analysis.

Figure 9-3 illustrates the relationship between BER and SNR for 802.11b, 802.15.3 base rate, and 802.15.1. In the figure, the BER values for the 5.5 Mb/s and 11 Mb/s data rates of 802.11b are clipped at a probability of 0.5 because the BER equations for these data rates are inaccurate for low SNR values. As the SNR value decreases, the BER should approach a value of 0.5. However, the equations for the BER of 5.5 Mb/s and 11 Mb/s will approach values of 4 and 64 as the SNR approaches zero.⁴⁴ However, for high BER values, i.e., greater than 0.01, the FER is essentially 1 for frame sizes of larger than 256 octets, and so the inaccuracy in the BER calculation would not affect the FER values.

802.11b and 802.15.3

802.11b networks have become much more common in recent years, and so coexistence with these types of networks is important for 802.15.3 applications. Despite the relative success of 802.11b, it has not yet achieved widespread penetration into peoples' homes, which is the target market of 802.15.3 networks. In addition, although it is common for office deployments of 802.11b to use more than one of the available channels, in a typical home, there will be only one

⁴⁴ These values can be calculated by noting that the Q function is equal to 0.5 when its argument is zero.




Figure 9-3: BER versus SNR for 802.11b, 802.15.3, and 802.15.1

802.11b network, and so the other two channels will be available for the 802.15.3 WPAN to use. Thus, in many cases, interference from other sources, e.g., 2.4 GHz phones or microwave ovens, will be more important than will the interference from 802.11b WLANs.

If an 802.15.3 piconet operates in the same area as an 802.11b WLAN, the PNC will notice the reduced throughput and will search for the quietest channel. In addition, if it detects the presence of an 802.11b network, it will use the 802.11b coexistence channel plan to avoid overlapping with the WLAN. In addition, the PNC will rate that channel as the worst to avoid choosing the occupied channel.

The current IEEE Std 802.11b does not provide the ability for an AP to request neighbor piconet status, and it does not prohibit the AP from doing it either. As IEEE Std 802.15.3 is very new, there are not any products on the market that provide this capability. However, it is likely that vendors that want to serve both markets will develop a dual-mode AP that enables the two networks to share the channel capacity. One of the applications for this is a home where the consumer

electronic devices, e.g., televisions, CD players, camcorders, and so on, are networked with 802.15.3, but the network also needs to support an 802.11b laptop for Internet access. A dual-mode AP could allocate time for the 802.11 contention period and use contention-free period to set aside time for the 802.15.3 WPAN. One of the issues with this approach is that the end of the contention period in 802.11 is not fixed in time. Instead, it ends when the AP gets a chance to send a beacon after the end of a transmission by a STA. In this case, it would be possible for the 802.11 network to continue transmitting past the time allocated for it by the PNC.

Using the modeling assumptions stated in "Assumptions for coexistence simulations" on page 278, the degradation of the FER of an 802.15.3 piconet in the presence of an 802.11b WLAN is illustrated in Figure 9-4. In the graph, the adjacent channel is 25 MHz away, whereas the alternate channel is separated by 50 MHz.




Figure 9-4: FER results for 802.15.3 with 802.11b as the interferer

The adjacent and alternate channels for 802.15.3 as the receiver with 802.11b as the interferer are given in Table 9-3.

Table 9-3: Adjacent and alternate channels for 802.15.3 as receiver

802.15.3 channel	Adjacent 802.11b channel	Alternate 802.11b channel
1	6	11
3	1, 11	none
5	6	1

As Figure 9-4 shows, there is almost no effect from the 802.11b STA when it is in the alternate channel of the 802.15.3 DEV. However, when the 802.11b STA is in the adjacent channel, the 802.15.3 DEV shows a noticeable degradation in performance (i.e., an FER of >3%) when the separation is less than 8 m. The 802.11b STA causes a degradation in the 802.15.3 DEV's throughput due to its larger occupied bandwidth and higher transmit power.

The 802.15.3 DEV will degrade the FER of 802.11b STA as well. This is shown in in Figure 9-5 for the 802.15.3 DEV in the adjacent channel and in Figure 9-6 for the 802.15.3 DEV in the alternate channel. When the 802.15.3 DEV is in the adjacent channel, the 802.11b STA does not experience a noticeable degradation in the FER for any of the data rates until the range is reduced to less than 0.8 m. In the case when the 802.15.3 DEV is in the alternate channel, this is reduced to a 0.4 m range.

The 802.15.3 DEV has very low impact on the 802.11b STA for two reasons. The first reason is that the 802.15.3 DEV has a relatively narrow spectrum, and so it has very little transmit energy in the 802.11b channel. The second reason is that the 802.15.3 DEV uses a lower transmit power than does the 802.11b STA. The results for the adjacent and alternate channels are much closer in this case due to the more stringent spectral mask of 802.15.3. The adjacent and alternate channel numbers for 802.11b as the receiver with 802.15.3 as the interferer are given in Table 9-4.




Figure 9-5: FER results for 802.11b with 802.15.3 as the interferer in the adjacent channel

Table 9-4: Adjacent and alternate channels for 802.11b as receiver

802.11b channel	Adjacent 802.15.3 channel	Alternate 802.15.3 channel
1	3	5
6	1, 5	none
11	3	1

If the 802.15.3 DEV and the 802.11b STA are in the same channel, then the degradation is much more significant. The FER for this case is graphed in Figure 9-7. For both 802.11 and 802.15.3, the performance of the networks are dramatically reduced. In the case of the 802.11 STA, the performance begins to degrade when the distance is less than 60 m for 11 Mb/s and less than 50 m for the




Figure 9-6: FER results for 802.11b with 802.15.3 as the interferer in the alternate channel

1 Mb/s data rate. The 802.15.3 DEV, on the other hand, has its throughput affected up to a separation of almost 80 m.

The results for the co-channel interference illustrate the value provided by 802.15.3's ability to automatically change channels and to allocate time in the current channel to other networks. Although it is possible to change the channel for an 802.11b WLAN, it requires user intervention to change the setting on the AP. In 802.15.3, the PNC automatically changes the channel whenever the conditions require it. In addition, because all of the DEVs in the WPAN are notified of the upcoming change, the channel switch will have a minimal effect on the current connections.




Figure 9-7: FER results for 802.11b and 802.15.3 co-channel interference

802.15.1 and 802.11 FHSS overlapping with 802.15.3

802.15.1 and 802.11 FHSS systems have similar characteristics in terms of the coexistence performance with respect to 802.15.3. In practice, the protocol differences will give somewhat different results. In particular, the faster hopping time of 802.15.1 can cause more problems with long data frames, especially ones sent at the lower data rates. In addition, the TDMA method used in 802.15.1 causes more problems with the CSMA/CA-based access method used in 802.11 than with the TDMA method used in 802.15.3. The CSMA/CA STA will lose frames when the TDMA DEV overlaps transmissions. In addition, the CSMA/CA STA will back off on the TDMA DEV's transmissions, reducing the throughput. The backoff window is also increased when the frame is dropped due to the collision. For the present analysis, the effects of the protocol are ignored. As most of

the 2.4 GHz cordless phones are also FHSS systems, the results here are similar to those that would be obtained for interference with these systems.

For a frequency hopping system, the effect on the 802.15.3 FER will depend on the hop frequency. When the FHSS system is on a frequency that is far away from the 802.15.3 channel, the effect on either system will be low. On the other hand, when the hop frequency overlaps with the 802.15.3 channel, the effect is much worse because there is no filtering in either the transmitter or the receiver to reduce the interfering power. The overall FER can be calculated as the average of the FERs for each of the hop frequencies because the FHSS systems will visit each of the frequencies an equal number of times.⁴⁵

The average FER with 802.15.1 as the receiver and 802.15.3 as the interferer ("802.15.1 receiving") and with 802.15.3 as the receiver and 802.15.1 as the interferer ("802.15.3 receiving") is shown in Figure 9-8.

The FER curves are flat over a large separation distance. In this region, the systems do not affect each other when the hop frequency is outside of the 802.15.3 channel. On the other hand, when the hop frequency falls inside of the 802.15.3 channel bandwidth, both systems experience 100% lost frames for these distances. As the FHSS systems use 79 1 MHz wide channels, these overlap with the 15 MHz wide 802.15.3 channel on average 15 out of 79 times or about 19% of the time. This creates a 19% FER, as shown in the figure. Once the two systems have been separated by sufficient distance, the FER due to co-channel interference begins to drop and so the overall FER begins to fall as well.

The performance of the 802.15.3 system can be improved using either a notch filter or an equalizer that can remove some of the power in the interfering signal. If the FHSS system is collocated with the 802.15.3 system, then the frequency of the notch filter will be known. If the systems are not collocated, it is possible for the 802.15.3 system to learn the hopping pattern and predict the required location of the notch filter. Finally, if the interferer starts prior to the 802.15.3 frame, then the equalizer in the modem will tend to reduce the effect of the narrowband interferer.

⁴⁵ This is not true with 802.15.1 because the hop sequence is selected from a poorly-designed sequence generator that does not generate random sequences. However, the assumption that hop sequence for 802.15.1 is pseudo-random is sufficient for the present analysis.




Figure 9-8: Average FER results for 802.15.1 and 802.15.3



Summary

IEEE Std 802.15.3 is the first IEEE standard specifically designed to support wireless multimedia applications in low-cost, portable devices. Because of this, the standard is poised to become widely adopted in consumer electronic devices that need multimedia connectivity. The standard supports applications that require both high throughput and good QoS, e.g., wireless distribution of video and high-quality audio. The standard was also designed to provide high throughput with a latency that is low and predictable. These characteristics support applications such as telephony, video conferencing, gaming, remote control, and channel surfing. Although the standard has focused on support for wireless multimedia, asynchronous traffic is supported as well, providing high throughput for that traffic when the channel time is not required for higher priority streams.

During the development of the standard, the 802.15.3 Task Group kept its focus on enabling low-cost, portable implementations. Both the MAC and PHY were specified to allow compliant devices with minimal functionality for low-cost implementations, while providing the ability to create full-featured implementations that maintain compatibility. Significant effort went into the development and analysis of the power-saving modes that are required for portable applications. The power-save modes in the standard range from a very simple, unsynchronized deep-sleep mode to a very efficient mode that provides predictable latency and low-rate, low-power QoS communication.

The 802.15.3 MAC architecture is simple, but scalable, providing the structure necessary to develop complex applications. The most basic compliant entity is the DEV, which only needs to support CSMA/CA for data transfer. A DEV can also be made very complex, supporting the sourcing and sinking of multiple streams with efficient power savings, device discovery, and secure communications. Likewise, the simplest PNC only needs to support sending the beacon, associating a limited number of DEVs and maintaining a heartbeat counter for each of the associated DEVs. A PNC can also be a more complex implementation that is able to monitor the piconet conditions, instantly adapt the allocations to handle changes in the environment or in data, and control more than

one piconet. The standard has already planned for future growth by reserving fields and numbers for new amendments while also providing methods for implementors to differentiate their products without collisions that create interoperability problems.

REDDY

The combination of CSMA/CA and TDMA in the superframe supports the efficient transfer of commands and small amounts of data, while providing the guaranteed allocations necessary for QoS. The CAP is easy for the PNC to allocate and provides a simple contention method for DEVs to send commands or even small amounts of data. However, the bulk of the data in the superframe is carried in channel time allocations using TDMA as the access method. Using TDMA supports the requirements for QoS while providing throughput efficiencies that are 30% to 100% greater than using a CSMA/CA method. The combination of the 802.15.3 MAC and PHY allow piconet to achieve a throughput in excess of 80% of the actual data rate when all of the MAC and PHY overhead is included. Providing high throughput and guaranteed allocations are only two of the requirements for supporting QoS. IEEE Std 802.15.3 also allows a DEV to request allocations that control the latency of its data transmissions. The channel time allocation request is used by the DEV to request the spacing between allocations, either multiples of a superframe, in the case of sub-rate allocations, or multiple times in a superframe, in the case of super-rate allocations.

While the MAC sublayer provides the basis for delivering multimedia connections, it is the PHY that makes it wireless. The 2.4 GHz PHY is based on proven technology that has decades of experience in various implementations. The technology for 2.4 GHz unlicensed applications has matured to the point that single-chip radios are commonplace and complete radio solutions are both small in size as well as being very low cost. The supporting parts, e.g., filters, baluns, couplers, power amplifiers, etc., are also widely available at very low-cost points and in a variety of form factors. The 2.4 GHz band has the advantage of near world-wide availability for unlicensed applications. The 2.4 GHz band also has better range characteristics than either 5 GHz or UWB allocations because of the lower frequency and, in the case of UWB, much higher transmit power. Single carrier modulation was selected for the PHY because it can be easily implemented in very few gates. It also has very low required signal to noise ratios, extending the range for essentially no additional cost or power.

Because of the good propagation characteristics and wide availability of its frequency allocation, the 2.4 GHz band has attracted many different systems, including IEEE Std 802.11b, IEEE Std 802.11g, IEEE Std 802.15.1, and cordless telephones. The 802.15.3 standard uses both MAC and PHY techniques to enhance the coexistence of 802.15.3 piconets with these types of interferers.

MAC techniques for coexistence include passive scanning, dependent piconets, channel status request, remote channel scan, and dynamic channel selection. The PHY improves coexistence with the choice of operating frequencies, reduced transmit and receive spectral mask, link quality, RSSI and transmit power control. These techniques not only reduce the impact of 802.15.3 on existing wireless devices, they also improve the stability and throughput of 802.15.3 piconets.

IEEE Std 802.15.3 is uniquely capable of supporting high-throughput wireless multimedia connections between low-cost, low-power portable devices.

Applications that require good QoS are well supported by the standard. Although the MAC provides for simple, interoperable devices, it also supports interoperability for more full-featured implementations and applications. The flexibility of the standard makes it easy for a DEV to request and receive the exact type of channel time allocations needed to support the DEV's applications. The PHY portion of the standard builds on many years of radio design and leverages the popularity of current 2.4 GHz solutions in order to keep the cost of the PHY layer very low. The PHY allows simple low-cost implementations while also defining high-rate modes that allow for high-speed data connections. Coexistence with existing networks was taken into consideration in the development of the standard and so the performance of 802.15.3 piconets in the presence of interferers is very good. The many advantages of 802.15.3 will help to make it the default choice for low-power, low-cost, wireless multimedia connectivity in consumer electronic devices.

REDON

References

- [B1] 47 CFR Part 15, Radio Frequency Devices, Federal Communications Commission. <http://www.fcc.gov/oet/info/rules/Welcome.html>.
- [B2] ARIB STD-T66, "Second generation low-power data communication system/wireless LAN system," Association of Radio Industries and Businesses, <http://www.arib.or.jp>.
- [B3] Bailey, D., Singer, A. and Whyte, W., "NTRUEncrypt Security Suite," submission to IEEE Task Group 802.15.3, document number 02362r0P802-15_TG3-NTRUEncrypt-Security-Suite.pdf, <http://grouper.ieee.org/groups/802/15/pub/2003/Sep02/>.
- [B4] Bailey, D., Singer, A. and Whyte, W., "RSA Security Suite," submission to IEEE Task Group 802.15.3, document number 02363r0P802-15_TG3-RSA-Security-Suite.pdf, <http://grouper.ieee.org/groups/802/15/pub/2003/Sep02/>.
- [B5] Batra, Anuj, et. al., "Physical Layer Submission to 802.15 Task Group 3a: Time-Frequency Interleaved Orthogonal Frequency Division Multiplexing (TFI-OFDM)," submission to IEEE Task Group 802.15.3, document number 03141r3P802-15_TG3a-TI-CFP-Presentation.ppt, <http://grouper.ieee.org/groups/802/15/pub/2003/May03/>.
- [B6] ETSI EN 300 328 V1.4.1 (2003-04), "Electromagnetic compatibility and Radio spectrum Matters (ERM); Wideband Transmission systems; Data transmission equipment operating in the 2,4 GHz ISM band and using spread spectrum modulation techniques; Harmonized EN covering essential requirements under article 3.2 of the R&TTE Directive," European Telecommunications Standards Institute, <http://www.etsi.org>.
- [B7] ETSI ETS 300 826 ed.1 (1997-11), "Electromagnetic compatibility and Radio spectrum Matters (ERM); ElectroMagnetic Compatibility (EMC) standard for 2,4 GHz wideband transmission systems and High Performance Radio Local Area Network (HIPERLAN) equipment," European Telecommunications Standards Institute, <http://www.etsi.org>.

[B8] Hawbaker, D. A., "Indoor wide band radio wave propagation measurements and models at 1.3 and 4.0 GHz," Masters thesis, Electrical Engineering, Virginia Polytechnic Institute State University, May 1991.

[B9] Huang, C.-C., and Khayata, R., "Delay spreads and channel dynamics measurements at ISM bands," *SUPERCOMM/ICC '92 Digest*, Chicago, IL, pp. 1222-1226, June 1992.

[B10] IEEE Std 802.1D™, 1998 Edition, IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 3: Media Access Control (MAC) Bridges.⁴⁶

[B11] IEEE Std 802.2™, 1998 Edition (R2003), IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 2: Logical Link Control.

[B12] IEEE Std 802.11™, 1999 Edition (R2003), IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.

[B13] IEEE Std 802.11b™-1999 (R2003), Supplement to IEEE Standards for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher-Speed Physical Layer Extension in the 2.4 GHz Band.

[B14] IEEE Std 802.15.1™-2002, IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs).

⁴⁶ IEEE publications are available from the Institute of Electrical and Electronics Engineers, Inc., 445 Hoes Lane, P. O. Box 1331, Piscataway, NJ 08855-1331, USA. IEEE publications can be ordered on-line from the IEEE Standards website: <http://www.standards.ieee.org>.

[B15] IEEE Std 802.15.2™-2003, IEEE Recommended Practice for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 15.2: Coexistence of Wireless Personal Area Networks with Other Wireless Devices Operating in Unlicensed Frequency Bands.

[B16] IEEE 802.15.3™-2003, IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 15.3: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for High Rate Wireless Personal Area Networks (WPANs).

[B17] IEEE Std 1394a™-2000, IEEE Standard for a High-Performance Serial Bus—Amendment 1.

[B18] IEEE Std 1394b-2002, IEEE Standard for a High-Performance Serial Bus—Amendment 2.

[B19] Intel UWB Database, Ultra Lab, University of Southern California.
<http://impulse.usc.edu>.

[B20] MacLellan, J., Lam, S., and Lee, X., "Residential indoor RF channel characterization," *Proceedings 1993 IEEE Vehicular Technology Conference*, Secaucus, NJ, pp. 210–213, 1993.

[B21] Merrit, R., "Group hopes to leapfrog 802.11 for wireless video," *Electronic Engineering Times*, May 3, 2003, <http://www.eet.com/story/OEG20030502S0060>.

[B22] NIST FIPS Pub 197: Advanced Encryption Standard (AES), Federal Information Processing Standards Publication 197, U.S. Department of Commerce/NIST, November 26, 2001.

[B23] Nobles, P., Ashworth, D., and Halsall, F., "Propagation measurements in an indoor radio environment at 2, 5 and 17 GHz," *IEEE Colloquium on High Bit Rate UHS/SHF Channel Sounders—Technology and Measurement*, pp. 4/1–4/6, 1993.

- [B24] Rappaport, T. S., and Hawbaker, D. A., "Wide-band microwave propagation parameters using circular and linear polarized antennas for indoor wireless channels," *IEEE Transactions in Communications*, vol. 40, no. 2, pp. 240-245, Feb. 1992.
- [B25] RSS-210, "Specification Low-Power Licence-Exempt Radiocommunication Devices (All Frequency Bands)," Industry Canada, <http://www.strategis.ic.gc.ca/engdoc/main.html>.
- [B26] Stein, J. C., *Indoor Radio WLAN Performance Part II: Range Performance in a Dense Office Environment*, Harris (now Conexant) Semiconductor web page, <http://www.intersil.com/data/wp/WP0546.pdf>.
- [B27] Struik, R. and Rasor, G., "Mandatory ECC Security Algorithm Suite," submission to IEEE Task Group 802.15.3, document number 02200r3P802-15_TG3-Mandatory-ECC-Security-Algorithm-Suite.doc, <http://grouper.ieee.org/groups/802/15/pub/2003/Sep02/>.
- [B28] TIA/EIA-95-B: Mobile Station-Base Station Compatibility Standard for Wideband Spread Spectrum Cellular Systems (ANSI/TIA/EIA-95-B-99), Committee TR-45.5, February 1, 1999.
- [B29] "Title 47 of the Code of Federal Regulations, Part 15," <http://www.fcc.gov/oet/info/rules/>.
- [B30] Weaver, D. K., Jr., "A third method of generation and detection of single-sideband signals," *Proceedings of the IRE*, vol. 44, no. 12, pp. 1703-1705, June 1956.
- [B31] Yoshida, J., "Bluetooth gropes for identity," *Electronic Engineering Times*, page 1, June 23, 2003.
- [B32] Zyren, J., Enders, E. and Edmondson T., "802.11g Starts Answering WLAN Range Questions," *CommDesign.com*, January 14, 2003, <http://www.commsdesign.com/story/OEG20030114S0008>.

Glossary

ACTIVE mode: A power management mode in which the device (DEV) is in AWAKE state for every superframe to receive the beacon. The DEV is also in AWAKE state during every superframe for the contention access period (CAP) and any channel time allocations (CTAs) for which it is either the source or the destination.

ad hoc network: A network that is created without the intervention of the user. Ad-hoc networks are typically formed as they are needed and dissolved when the desired communication is complete.

association: In the context of IEEE Std 802.15.3, this is the process by which a device (DEV) joins an existing piconet. The association process allows the piconet controller (PNC) to refuse access to DEVs if its resources are limited. It is also used to assign the 8-bit device identifier (DEVID).

asynchronous power save (APS) mode: A power-save mode in which the device (DEV) is in SLEEP state and wakes only to send a frame to the piconet controller (PNC) to maintain its associated status. Unlike device synchronized power save (DPS) and piconet synchronized power save (PSPS) modes, a DEV in asynchronous power save (APS) mode is free to wake up in any superframe of its choosing to send a frame to maintain its association with the PNC and does not have to be in AWAKE state for any specific beacons.

AWAKE state: The condition of the device (DEV) when it is either transmitting or receiving.

channel time allocation (CTA): A specified duration of time assigned by the piconet controller (PNC) to a specific source device (DEV) during the superframe for its exclusive use. During this time, no other DEV in the piconet is allowed to begin a frame exchange. The destination can be either a single DEV, a group of DEVs (multicast), or all of the DEVs in the piconet (broadcast).

channel time allocation period (CTAP): The period of time in the 802.15.3 superframe set aside for channel time allocations. In general, devices (DEVs) do not contend for channel time during the CTAP, rather it is assigned to a DEV by

the piconet controller (PNC) in response to a request from that DEV. Channel time allocations use time division multiple access (TDMA), with the exception of association management channel time allocations (MCTAs) and open MCTAs, which use slotted aloha.

child piconet: A dependent piconet in which the piconet controller (PNC) is a full member of the parent piconet.

child piconet controller (PNC): A PNC that is a full member of another piconet, called the parent piconet, and synchronizes its piconet so that its communications occur only during a channel time allocation (CTA) granted to it by the parent PNC.

clear channel assessment (CCA): The process by which a radio determines that the wireless medium is clear. A radio performs CCA for 802.15.3 by enabling its receiver and searching either for the preamble or by determining that the power in the channel exceeds a specified level.

contention access period (CAP): The period of time in the 802.15.3 superframe set aside for devices (DEVs) to use collision sense multiple access with collision avoidance (CSMA/CA) to gain access to the medium. The CAP can be used to send data and commands.

dependent piconet: A piconet that operates during a channel time allocation in the superframe of the parent piconet. The parent piconet controller (PNC) assigns the channel time for the dependent piconet in response to a request from the dependent PNC. There are two types of dependent piconets defined in the IEEE Std 802.15.3, child and neighbor piconets.

dependent piconet controller (PNC): A PNC that has synchronized piconet timing to another piconet, called the parent piconet, so that the communications occur only during a channel time allocation (CTA) granted to it by the parent PNC. The dependent PNC can be either a child PNC or neighbor PNC depending on its membership status in the parent piconet.

device (DEV): Any of the participants in an IEEE Std 802.15.3 piconet. The piconet controller (PNC) maintains a DEV personality in the piconet even as it operates as the PNC.

device synchronized power save (DSPPS): A power save mode in which devices (DEVs) wake up periodically to listen to beacons. DSPPS mode differs

from piconet synchronized power save (PSPS) mode in that the DEV specifies the sleep interval for DSPS mode, while the piconet controller (PNC) determines it for PSPS mode.

disassociation: The process by which a device (DEV) either leaves the piconet or is removed from the piconet by the PNC.

extended beacon: A method for that allows the beacon transmissions to consist of relatively small frames for reliability. The extended beacon is composed of a beacon frame followed by one or more broadcast Announce commands and the appropriate interframe spacings.

fragmentation: The process by which the protocol splits up a larger data unit into smaller ones, adding information that will allow the receiving entity to reconstruct the original data unit.

frame: The basic on-air data unit for IEEE Std 802.15.3. It is composed of the PHY preamble, PHY header, MAC header, MAC header protection, data, security information (for secure frames), and frame check sequence.

integrity code (IC): A string of bits generated with a symmetric encrypting key that is typically appended to data in order to provide data integrity and source authentication similar to a digital signature.

MAC protocol data unit (MPDU): A block of data exchanged between the MAC and the PHY either to be sent by the PHY or because it was received by the PHY.

MAC service data unit (MSDU): A block of data exchanged between the MAC and the frame convergence sublayer (FCSL) either to be sent over the wireless medium by the PHY or because it was received by the MAC from the PHY.

neighbor piconet: A dependent piconet in which the piconet controller (PNC) is not a full member of the parent piconet.

neighbor piconet controller (PNC): A PNC that has associated with another piconet, called the parent piconet, but which is not a full member of the parent piconet. The neighbor PNC synchronizes its wireless network with the parent piconet in such a way as to guarantee that transmissions from the neighbor piconet occur only during a channel time allocation (CTA) granted to it by the parent PNC.

parent piconet: A piconet that provides channel time allocations for another piconet, referred to as a dependent piconet, to allow it to operate in the same PHY channel.

piconet: The basic configuration of an IEEE Std 802.15.3 network. A piconet is formed when the piconet controller (PNC) begins sending beacons. All piconets have a PNC and at least one device (DEV) (which is just the PNC's DEV personality). The smallest non-trivial instantiation of an 802.15.3 piconet consists of the PNC and two DEVs, one of which is the PNC's DEV personality and one DEV that is distinct from the PNC.

piconet controller (PNC): A device (DEV) that is in charge of allocating channel time and allowing access to the piconet via association. The PNC can also act as the security coordinator for the piconet, performing authentication and key management for the members of the piconet.

piconet synchronized power save (PSPS): A power-save mode in which the device (DEV) is required to be in AWAKE state for beacons specified by the piconet controller (PNC). In PSPS mode, the PNC determines the periodicity (if any) of the wake beacons, as opposed to device synchronized power save (DSPS) mode where this is determined by the DEV.

power management (PM) mode: In IEEE Std 802.15.3, this is the current status of a device (DEV) with respect its desire to save power. 802.15.3 defines four PM modes: ACTIVE, device synchronized power save (DSPS), piconet synchronized power save (PSPS), and asynchronous power save (APS).

power-save (PS) mode: In IEEE Std 802.15.3, this refers to a mode in which the device (DEV) has reduced its activity in order to save power. 802.15.3 defines three power-save modes: device synchronized power save (DSPS), piconet synchronized power save (PSPS), and asynchronous power save (APS).

quality of service (QoS): A measure of the performance of a network with respect to certain basic performance criteria, including latency, reliability (low downtime), and error performance. This metric can be used to compare the ability of different types of networks with respect to delivering time-critical information.

secure frame: A command or data frame that has been protected with an integrity code and, in the case of data frames, encrypted to prevent eavesdropping.

secure membership: A device (DEV) achieves secure membership in the piconet when it has received a management key from the piconet controller (PNC). Once the DEV has the management key, it can also get the data key to encrypt and decrypt data as well as to authenticate commands.

secure piconet: A piconet in which the beacon and most command frames are protected with an integrity check and data frames can be encrypted with a common symmetric key pair.

SLEEP state: The condition where the device (DEV) is neither transmitting nor receiving.

stream: A unidirectional data connection between a source device (DEV) and one or more destination DEVs. Streams can be used for both isochronous and asynchronous data.

sub-rate allocation: A channel time allocation that occurs only once every n superframes ($n > 1$) where n is the CTA Rate Factor sent in the Channel Time Request command. For example, a CTA Rate Factor of 4 would result in an allocation that occurs in every fourth superframe. Sub-rate allocations are used for low-data rate streams or to maintain connectivity between sleeping devices (DEVs) in device synchronized power save (DPS) mode.

super-rate allocation: A channel time allocation that occurs at least once in every superframe.

superframe: The basic timing period for an IEEE Std 802.15.3 piconet. The superframe is composed of the beacon, the channel time allocation period, and optionally the contention access period.

synchronized power save (SPS): A power-save mode in which devices (DEVs) wake up regularly for the same beacon, called wake beacons, which occur at regular intervals, e.g., every eight superframes. There are two types of SPS modes defined in IEEE Std 802.15.3: device synchronized power save (DPS) and piconet synchronized power save (PSPS).

wake beacon: The beacon during which a synchronized power save (SPS) device (DEV), either piconet synchronized power save (PSPS) or device synchronized power save (DPS), will be in the AWAKE state. All of the members of a given SPS set are required to be in the AWAKE state for the wake beacon.

REDOY

Index

Numerics

802.11 3, 5, 6, 30, 31, 33, 45, 49, 104, 166, 203–208, 223, 230, 273, 274, 278, 279, 282–289

A

access control list (ACL) 190, 92, 194, 199, 200

acknowledgement (ACK) 60, 78, 81–86, 160, 170, 212

 delayed (Dly-ACK) 81–99, 119, 208, 212

 Dly-ACK negotiation 84–86

 immediate (Imm-ACK) 47–49, 61, 75, 80–83, 85, 94, 99, 119, 167, 208, 212

 no (no-ACK) 53, 55, 75, 81, 94–98, 149–150, 208

 policy 75, 80–86, 119

 request 81, 83, 85, 86

ACTIVE 40, 42, 43, 102, 103, 104, 106, 111–114, 118–129, 146, 150, 169, 179, 263

advanced encryption standard (AES) 189, 190, 195–197

advanced encryption systems (AES) 279

antenna gain 30, 205, 213, 231

association 26, 29, 32–35, 37, 56, 57, 58, 60, 61, 66, 73–78, 80, 81, 103, 145, 155, 159, 160, 168–173, 180, 183, 194

Association Request 168

association timeout period (ATP) 35, 41, 75, 76, 79, 80, 91, 103, 105, 127, 131, 146, 193

asynchronous power save (APS) 40, 41, 42, 102–107, 113, 114, 126–129

Authentication 279

authentication 6, 11, 12, 24, 70, 76, 187–195, 199

 mutual 192

authorization 70, 73, 187–189, 192, 193, 200

AWAKE 40, 43, 103, 147

B

backoff 54, 207

 CAP 56, 61, 158, 211, 212

 open MCTA 58

 slot 58, 61

 window 56, 288

base rate 10, 29, 38, 159, 160, 216, 220, 277, 282

beacon 26, 31, 32, 33, 43, 52, 64, 66, 71, 89, 103, 108, 110, 111, 113, 118, 123, 126, 134, 135, 136, 151, 153, 154, 157, 169, 170, 179, 197, 284

 announcement 125, 129, 130, 131

 extension 53, 161

- frame 53, 153, 160, 161, 194
 number 41, 44, 71, 103, 114, 124,
 126, 128, 132, 136, 137,
 151, 183
 wake 41, 42, 43, 104, 115, 124,
 129, 264
 BER 96, 210, 216, 224, 226, 232, 235,
 236, 240, 250, 276, 280-282
 bit error ratio
 See BER
 bit ordering 49, 50
 broadcast 29
 ID 34, 57, 81, 102, 103, 115, 141,
 148, 149, 181, 201
 traffic 42, 43, 124, 166
 BSID 44, 50, 53, 128, 134-136, 142,
 143, 145, 153
C
 channel
 adjacent 30
 alternate 30
 changing 273
 model 204, 237, 238, 239, 240, 278
 PHY 39, 40, 43
 plan 29, 205, 206, 274
 PNC channel scanning 150-152
 status 39, 148-150, 273
 switch time 207
 channel time 5, 6, 16, 26, 28, 32, 33,
 35, 39, 42
 asynchronous 89, 268
 requesting asynchronous 36
 terminating 36
 group allocation 88, 89, 90, 125
 individual allocation 88
 isochronous 268
 modifying 36
 requesting 35
 terminating 36
 request command (CTRq) 35, 36,
 37, 41, 88, 89, 90, 95, 167
 channel time allocation (CTA) 33-43,
 54, 55, 70, 82, 84, 86, 88, 95, 100, 102,
 103, 119-125, 133, 137, 142, 143, 144,
 145, 147, 151, 158, 170, 176, 177, 179,
 183, 184, 185, 212, 275, 277
 management (MCTA) 34, 57, 74,
 150, 157, 213
 open MCTA 57, 58, 61, 62
 private 58, 179
 pseudo-static 52, 59, 64, 65, 93,
 128, 138, 151, 167, 169,
 179
 rate factor 62, 120
 regular MCTA 34, 213
 sub-rate 42, 62, 93, 116-125
 super-rate 62, 179
 channel time allocation period
 (CTAP) 32, 33, 53, 55, 56-60, 64, 65,
 74, 207, 212
 child piconet
 See piconet, child

- child PNC
 - See piconet coordinator (PNC), child
- clear channel assessment (CCA) 34, 54, 55, 57, 60, 62, 212, 268
- coexistence 31, 165, 206, 230, 271–289
- commands
 - Announce 39, 53, 135, 138, 143–145, 156, 167
 - Association Request 58, 61, 74, 76, 167, 168, 192
 - Association Response 75, 76, 160, 169, 192
 - Channel Status Request 39, 40, 132, 139, 149
 - Channel Status Response 39, 139, 149
 - Channel Time Request 62, 65, 120, 122, 128, 167, 169, 179, 183, 185
 - Channel Time Response 170, 176, 179
 - Disassociation 35, 79, 160, 192
 - Disassociation Request 78, 79, 167, 176
 - Distribute Key Request 193
 - Piconet Services 139, 155
 - PM Mode Change 42, 110, 113, 123, 128
 - PNC Handover Information 70, 99, 181
 - PNC Handover Request 68, 71, 180, 183, 185
 - PNC Handover Response 71, 72, 178, 181, 184, 185
 - PNC Information 35, 39, 68, 71, 76, 77, 78, 99, 127, 139, 145–148, 159, 181
 - PNC Information Request 39, 131, 139, 147, 148, 159
 - Probe Request 39, 80, 129, 131, 138–143, 156, 159, 167
 - Probe Response 39, 138–141, 156, 167
 - PS Set Information Request 105
 - PS Set Information Response 70, 99, 105, 181
 - Remote Scan Request 39, 132, 139, 152
 - Remote Scan Response 135, 139, 152, 153, 154
 - Security Information 70, 190, 201
 - Security Information Request 70, 181, 199, 201
 - Security Information Response 199
 - Security Message 167, 190
 - SPS Configuration Request 42, 114, 116, 117
 - SPS Configuration Response 115, 117
 - Transmit Power Change 158, 275
 - Vendor Specific 161, 167

- constant amplitude zero auto
correlation (CAZAC) 211, 212, 217,
218
- constellations 220–223, 224, 231, 232
- contention access period (CAP) 9, 32,
33, 34, 35, 53–62, 63, 74, 81, 82, 84,
88, 99, 102, 103, 149, 157, 211, 213,
268, 275
- Continued Wake Beacon (CWB) 145
- cowboy mode 105
- CSMA/CA 33, 58, 99, 158, 288
- D**
- data key 196
- defragmentation 44, 96–99
- delay spread 234–244
 defining 239–240
 measurements 240–244
- dependent piconet
 See piconet, dependent
- dependent PNC
 See piconet coordinator (PNC),
 dependent 37
- DEV 16, 20, 25–28, 29, 33, 35, 36, 37
 life cycle 162–164
- device ID
 broadcast (BcstID) 34, 57, 81, 102,
 103, 131, 141, 148, 149,
 181, 201
 destination (DestID) 33, 57, 75, 82,
 89, 102, 103, 130, 131,
 135, 149, 151, 168, 170,
 183, 184, 196, 265, 269
 multicast (McastID) 81, 103, 131
 neighbor (NbrID) 81, 148, 166,
 169, 175, 181, 185
 piconet controller (PNCID) 33, 57,
 59, 66, 80, 148, 151, 160,
 166, 169, 170
 source (SrcID) 33, 34, 57, 76, 89,
 103, 135, 147, 151, 170,
 183, 184, 196, 265, 269
 unassociated (UnassocID) 34, 57,
 74, 75, 76, 81, 168, 184
- device synchronized power save
(DSPS) 11
 allocating channel time 119–123
 mode 40, 41, 102, 104, 112, 118–
 119, 123, 124, 126, 130
 set 104, 105, 106, 111, 113–117,
 125, 126
 creating 42, 116
 joining 41, 42, 116, 118
 leaving 117
- digital camera 14, 15, 155, 246
- disassociation 29, 35, 73, 78–80, 176,
193, 194
- Dly-ACK
 See acknowledgement (ACK),
 delayed (Dly-ACK)
- DME 258, 259, 261, 262, 264, 269
- duplicate detection 99
- E**
- ECC 188, 189
- encryption 28, 187–197, 219

EVM 231, 232, 233, 250

F

FCSL 96, 259, 265, 268–270

FER 39, 59, 93, 96, 148, 210, 216, 224, 239, 240, 278, 280, 282, 284, 285, 286, 289

fragmentation 44, 82, 86, 96–99, 100, 141, 146, 201, 211, 219, 281

frame

beacon 47, 53, 153, 160, 161, 194

check sequence (FCS) 61, 82, 84, 149, 208, 214, 216, 219

command 47, 58, 82, 160, 194, 195

data 47, 48, 86, 88, 99, 146, 160, 167, 194, 197, 210, 217, 219, 259, 288

delayed acknowledgement
(Dly-ACK) 47, 54, 57, 81–86, 98, 99, 119, 160, 208

immediate acknowledgement
(Imm-ACK) 47–49, 54, 55, 57, 61, 75, 76, 80–83, 160, 167, 208

maximum size 53, 211, 219

G

guard time 63–65, 84

H

hand-over

dependent 177–186

dependent PNC 68

PNC 66–72, 130

header 134

check sequence (HCS) 52, 61, 82, 84, 149, 208, 216

MAC 34, 52, 81, 84, 97, 134, 160, 196, 208, 213–216, 268

PHY 52, 84, 208, 213–219, 268, 277

I

IEEE

Registration Authority Committee
(RAC) 74, 134, 161

Standards Association (IEEE SA)
3, 11, 12, 190

Standards Review Committee
(RevCom) 12, 190, 281

Imm-ACK

See acknowledgement (ACK),
immediate (Imm-ACK)

information element (IE) 47, 129–136,
138–157

Application Specific 129, 131,
142, 143, 145

BISD 44, 50, 53, 128, 134–136,
142, 143, 145, 153

Capability 142, 143, 145, 159

Channel Time Allocation 141

Continued Wake Beacon (CWB)
121, 130, 142, 143

CTA Status 65, 131, 140, 141, 142,
143, 144, 145

CTA status 93, 130

- DEV Association 76, 77, 130, 131, 142, 143, 145, 169
- Overlapping PNID 134, 135, 142, 143, 144, 145
- Parent Piconet 53, 142, 143, 145, 153, 176
- Pending Channel Time Map (PCTM) 41, 123, 128, 130, 131, 142, 143, 145
- Piconet Parameter Change 44, 52, 103, 126–132, 136, 137, 138, 141–145, 151
- Piconet Services 139, 142, 143, 145, 155, 156, 157
- PNC Handover 71, 103, 130, 131, 134, 142, 143, 145
- PNC Shutdown 73, 130, 131, 142, 143, 145
- PS Status 43, 110, 113, 114, 116, 130, 131, 141, 142, 143, 144, 145
- Transmit Power Parameters 142, 143, 145
- Vendor Specific 76, 131, 140, 142, 143, 145, 161
- interframe spacing (IFS) 53–55
- BIFS 54, 56, 207, 211, 212
- MIFS 53, 54, 207–210
- RIFS 54, 55, 207, 211, 212
- SIFS 54, 56, 61–65, 82, 83, 94, 95, 207, 208, 210, 212
- K**
- key
- data 196
- distribution 197
- management 196
- symmetric 28, 188, 189, 195–197
- key originator 28, 66, 190, 191, 192, 193, 197, 199, 200
- L**
- line of sight (LOS) 234, 235, 241, 242, 243, 244
- local area network (LAN) 1, 23, 261
- M**
- MAC address 34, 35, 39, 70, 74, 134, 136, 146, 192
- MAC protocol data unit (MPDU) 82, 96, 97
- MAC service data unit (MSDU) 82, 86, 96, 97, 98, 100, 149, 269
- management key 196
- maximum drift 64, 65, 281
- medium access control (MAC) 47–164
- metropolitan area network (MAN) 1, 23, 261
- MLME 139, 258–264
- modulation 220–224
- DQPSK 10, 45, 203, 205, 215, 216, 220, 223–224, 225, 227, 228, 232, 277, 278
- QAM 10, 203, 205, 215, 220–224, 225, 228, 232, 250

- QPSK 10, 203, 205, 215, 216, 224–228, 232, 280
- trellis coded 10, 30, 203, 205, 214, 216, 222, 223, 225, 227, 228, 232
- MP3 player 4, 25
- multicast 29
 - ID 81, 103, 131
 - traffic 42
- multirate 159–160
- N
- neighbor piconet
 - See* piconet, neighbor
- neighbor PNC
 - See* piconet coordinator (PNC), neighbor 37
- Next Awake Beacon 41, 42, 113
- no-ACK 94
- noise figure 30, 204, 224, 226, 236, 239, 247
- non-line of sight (NLOS) 241, 242, 243, 244
- NTRUEncrypt 189
- O
- octet 47
 - ordering 50
- originating DEV 275
- OUI 134, 155, 156, 161, 162
- P
- path loss model 279
- photo kiosk 14, 15
- physical layer (PHY) 203–256
 - frame format 213–219, 277
 - header 52, 84, 208, 213–218, 268, 277
 - preamble 51, 55, 60, 84, 208, 211, 215, 217–219, 277
- PIB 261
 - MAC 261
 - PHY 261
- piconet 25–28, 154–157
 - child 28, 36–39, 48, 154, 170, 175, 176, 183
 - starting 170–172
 - dependent 27, 28, 36–39, 48, 73, 165–186
 - starting 168–173
 - stopping 176–177
 - identifier (PNID) 44, 128, 134–136, 144, 153, 170
 - independent 48
 - joining 26–29, 34, 66, 73–81, 163, 166, 168, 199
 - membership 26–28, 39
 - neighbor 28, 38–39, 48, 154, 173–176
 - starting 168, 173
 - parent 26, 37, 48, 171–183

- shutdown
 with dependents 38–39, 174–176
 starting 31–32, 66
 timing 51–65
- piconet coordinator (PNC) 25–28, 44, 47, 66–73, 145–148, 150–152
 child 37, 48, 72, 168, 171, 172, 173, 177, 280
 dependent 73
 hand-over 66–72, 130
 dependent 68, 177–186
 neighbor 48, 72, 173, 181, 281
 parent 48, 73, 171, 174–179, 181
 shutdown 130
- piconet parameters, changing 128–138
- piconet-synchronized power save (PSPS) 43
 mode 40, 43, 102–114, 124, 128–130
 set 43, 104, 107, 111, 113, 117
 join 124
- PLME 259–261
- PNC Des-Mode 67
- power management (PM) mode 40–43
 ACTIVE 40, 43, 102, 103, 106, 111–114, 118–129
 APS 40, 41, 42, 102–107, 113, 114, 126–129
 DSPS 40, 41, 42, 104, 112, 118–119, 123, 124, 126, 130
- PSPS 43, 102–114, 118, 124–126, 128–130
- SPS 102, 104–107, 111
- power save (PS)
 efficiency 107–110
 mode 41, 144, 170
 set 68, 70, 102, 104, 105, 106, 114
- R**
- radio architectures 244–256
 direct conversion 248–250
 low IF 253–255, 256
 superheterodyne 245–247
 walking IF 250–253
 Weaver architecture 255
 ZIF 248–250, 256
- range
 extending with dependent piconets 165
 performance of 2.4 GHz PHY 203, 204
- recommended practice 278
- retransmission 16, 54, 60, 63, 82, 92, 99, 149, 196, 207, 211, 212, 268
- S**
- secure membership 73
- security 187–201
 information 32, 68, 102, 181, 199–201
 inherent in the PHY layer signaling 187

- mode 0 190
 - mode 1 190–193
 - modes 190–193
 - policies 67, 166, 193–195
 - sensitivity 60, 126, 158, 205
 - reference 224, 226, 278
 - typical 205, 224
 - service
 - discovery protocol 155
 - piconet 154–157
 - service access point (SAP) 257
 - FCSL 259
 - MAC 259, 265
 - MLME 259, 262–264
 - MLME-PLME 259
 - PHY 216, 259, 266, 266–268
 - PLME 259, 265
 - SLEEP 40, 103
 - slotted aloha 34, 56, 58, 60, 62
 - SNR 216, 225, 226, 232, 233, 276, 281, 282, 283
 - stream index 35, 36, 44, 57, 88, 98, 100, 141, 147, 169, 265, 266, 269
 - stuff bits 214
 - sub-rate 35, 42, 43, 62, 82, 93, 110, 116–125
 - superframe 20, 26, 32–34, 35, 37, 51–65, 81, 84, 88, 89, 92, 103, 110, 116, 120, 122, 124, 127, 132, 151, 171, 173, 175, 177, 207
 - changing duration 128, 136–138, 151
 - duration 16, 43, 92, 94, 108, 196
 - moving 137, 151
 - super-rate 17, 35, 120, 125, 169, 179
 - symmetric key
 - See* key, symmetric
 - synchronized power save (SPS) 283
 - joining 106, 111, 117
 - mode 104–107, 111
 - set 102, 103, 105, 111, 114–117
- T**
- tail symbol 214
 - target DEV 49, 91, 122, 144, 149, 179, 275
 - telephone
 - cellular 16, 17
 - cordless 17, 31
 - television 4, 20, 21, 22, 155, 213
 - HDTV 15
 - wireless connectivity 21
 - time token 191, 194, 196
 - transmit power control 275
 - transmit PSD 230
 - transmitter power 30, 67, 157, 230, 275, 278
 - trellis coded modulation (TCM) 10, 30, 203, 205, 214, 216, 222, 223, 225, 227, 228, 232

U

UWB 10, 60, 187, 203, 204, 205

W

Wake Beacon

Interval 109, 110, 115, 122, 125,
126wake superframe 103, 119, 120, 121,
125

Weaver architecture 255

wide area network (WAN) 23, 261

wireless personal area network
(WPAN) 1, 13, 24, 31, 36, 38, 39, 158,
253, 278, 283, 284, 287

