Entity (SME). The SME is independent from the physical and MAC layers and is not specified in the standard. In general, the station management entity is responsible for tasks such as the gathering of information from the PHY and MAC layer management entities, and similarly setting the value of certain parameters.

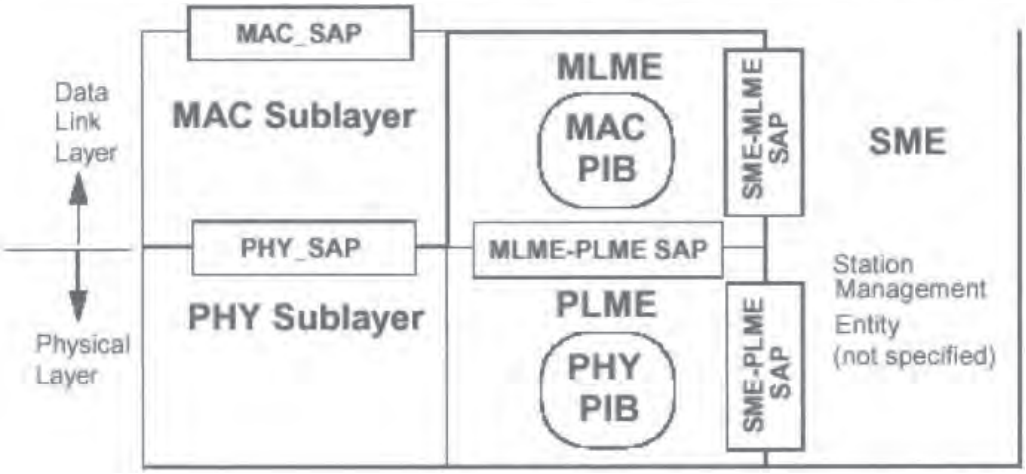

Figure 3–12: IEEE 802.15.3 Reference model

Figure 3–12 also depicts the relationship among management entities. The physical layer can be considered as providing services to the MAC layer. The five different entities interact via service access points (SAP). The format of these interactions is defined explicitly within 802.15.3. Defined primitives are exchanged over the service access points. The primitives fall into two basic categories: (1) primitives that support peer-to-peer interactions; and (2) primitives that have local significance and support sublayer-to-sublayer interactions. While the format of the primitives is defined, most interactions are not defined explicitly within the standard. For example, the interfaces between the MAC and its management layer, and the PHY and its management layer are not defined. Figure 3–12 is a conceptual diagram. The physical layer can be controlled by the MAC and the MAC management layer. In this case, there is no physical layer management. Vendors may opt to keep the service access points nonobservable. In this case, these vendors can implement these interfaces in any way they wish.

INTEL-1020
U.S. Patent 8,219,129

Like the other wireless standards within IEEE 802, IEEE 802.15.3 appears to higher layers starting with the logical link control (LLC), as an IEEE 802-style LAN. However, above the MAC layer, there may not be a 802.2 layer, but a USB or IEEE 1394 convergence layer.

The next section of this chapter is devoted to the MAC layer of IEEE 802.15.3. It is followed by a section devoted to the physical layer.

MAC layer

The IEEE 802.15.3 MAC provides for the following functionality:

1) Allow stations to form and terminate personal area networks.
2) Authenticate stations with each other.
3) Transport data between stations with multiple QoS levels and optional privacy.
4) Allow stations (including the coordinator) to minimize power requirements and still maintain the network.
5) A procedure for the coordinator to pass coordination to another station in the WPAN.

The piconet of IEEE 802.15.3 is shown in Figure 3–13. There are two types of devices: basic devices and piconet coordinators (PNC). One piconet can have at most 255 devices. The PNC (1) provides the basic timing for the piconet with the beacon; (2) manages the QoS; (3) manages power-save modes; and (4) handles the authentication. All devices are not required to be able to act as piconet coordinators. In this way, to achieve a lower cost, some devices do not have full functionality.

Similar to IEEE 802.15.1, the roles of a device and the piconet coordinator can be exchanged on the fly, i.e., a device can take over as a coordinator. This flexibility will be necessary when a coordinating station is removed from the WPAN—for example, when it powers down or moves out of range. Exchanging the roles of the coordinator and a device can take place if the coordinator determines that another station is better suited to perform this role. Note that, similar to IEEE 802.15.1, the coordinator handover requires support at the application layer, beyond the features of the standard. For example, changing the coordinator will usually mean changing the QoS parameters.

INTEL-1020
U.S. Patent 8,219,129

A device can start its own piconet or associate with an existing piconet. If a device wants to start its own piconet, it must make sure there is not an IEEE 802.15.3 piconet already started in the same channel. Before starting a piconet, the device will use passive scanning to detect an existing piconet. During passive scanning, the device will look for beacon frames from a PNC. Devices search for the piconet by starting from the first channel and traversing through all the indexed channels available in the PHY. If there are available channels, the device can take the responsibilities of a PNC, choose a channel, and start sending beacons to start its own piconet. If no channels are available, the device can start a dependent piconet.

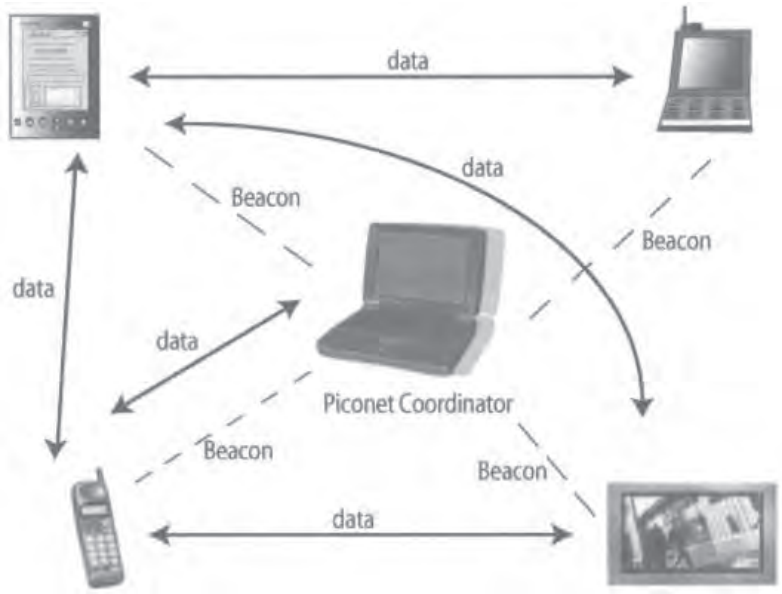

Figure 3–13: The piconet structure of IEEE 802.15.3

Figure 3–13 shows a piconet with no dependent piconets. More complex ad-hoc networks can also be formed, similar to the scatternet architecture of IEEE 802.15.1. A device that participates in a piconet can form a dependent piconet. The dependent piconet depends on time allocation from the parent piconet and is synchronized to the parent piconet. The two types of dependent piconets are child and neighbor piconets. A parent piconet can have more than one child piconet, and a child piconet, in turn, can have a child or neighbor piconet of its own. One piconet can have simultaneously more than one child and neighbor piconets.

INTEL-1020
U.S. Patent 8,219,129

The child and neighbor piconets are autonomous piconets and have distinct piconet identifiers (PiconetIDs). They depend on a dedicated time slot called channel time assignment (CTA) from the parent piconet. The coordinator of the child piconet is a device in the parent piconet. Therefore, the child piconet extends the area of coverage of the main piconet. A parent/child piconet combination is architecturally just like a IEEE 802.15.1 scatternet. A parent/child piconet combination poses the question of efficient routing. (This question has not been studied and is a very suitable topic for future research.) The coordinator of the neighbor piconet is not participating in the parent piconet. Therefore, the dependent piconet does not extend the area of coverage of the parent piconet. Dependent piconets are normally used to create another piconet even when no channels are available. Following a time-division multiplexing principle, dependent piconets work in the same frequency spectrum as the independent piconet. The superframe of the dependent piconet fits in one guaranteed time slot of the parent piconet. A device that wants to form a child piconet of its own will send a request for a CTA to its PNC.

Piconets are discovered by passive scanning, in which devices listen for a beacon. A station desiring to join an IEEE 802.15.3 WPAN will set its receiver to periodically listen on the various channels for a beacon. If the device finds a beacon, and the beacon indicates a network of interest, it will attempt to authenticate with the coordinator. Security in all standards belonging to the IEEE 802.15 family is provided on the assumption that there is no connection to an external network. This is an important difference between IEEE 802.15 and IEEE 802.11. In addition, generally IEEE 802.15 devices have less computational power and less memory. It is also assumed that an attacker can have more computational power and more memory. There are two security modes—mode 0 and mode 1. In mode 0, authentication and encryption are not required. The only security service available in mode 0 is an access control list (ACL). The security services provided in mode 1 are access control list, mutual authentication, key establishment, key transport, and verifying the authenticity of keys, data encryption, data integrity protection, beacon integrity protection, freshness protection, and command integrity protection.

Stations wishing to join the WPAN authenticate with the coordinator and then with any other station with which communication is required. This authentication can be "open," i.e., only the public key information of the station to which

INTEL-1020
U.S. Patent 8,219,129

authentication is attempted is used. It may also require that a secret PIN is known and transferred during the exchange, thus allowing controlled authentication. A secret key is generated to allow privacy for any subsequent data exchange. Upon successful authentication, the station is considered to be in the WPAN. It will also exchange capability information with the coordinator. This capability information includes all the PHY data rates supported by the station, its power management status, whether the station can be a coordinator, buffer space, etc. As a result of this exchange, a coordinator handover may occur. All the frames between the PNC and the device before the completion of the association of the device are exchanged in the contention access period only. To participate in a piconet, a device must join by association. The unregistered device initiates the association process by sending an association request. This association request is sent during the contention access period. When the PNC receives a valid association request, it sends an acknowledgment first, and after some time sends the association response, indicating acceptance or rejection of the request. If the association is successful, the coordinator informs the new device about the services offered by the piconet, and the new device informs the coordinator of the services that it offers. As a result the beacon is changed to include information about the new device, so that the rest of the devices are also informed about the services that the new device offers. In a very similar way, when a device leaves the piconet, the information about this device is removed from the beacons.

Security mode 0 is mandatory to implement. Security mode 1 is optional. Mode 1 offers three security suites: elliptic curve Menezes-Qu-Vanstone (ECMQV) Koblitz-283, NTRUEncrypt 251-1, and RSA-OAEP 1024-1. The ECMQV security suite has three subsuites: ECMQV manual that does not accept certificates, ECMQV implicit, and ECMQV X.509 that accept certificates. RSA-OAEP 1024-1 has two subsuites, one that does not accept certificates and another that does. Devices that support mode 1 must implement at least one security subsuite. The ECMQV Koblitz-283 security suite provides 128-bit security, the NTRUEncrypt 251-1 security suite provides 80-bit security, and the RSA-OAEP 1024-1 security suite provides 80-bit security.

All security suites support CTR + CBC – MAC (CCM) combined encryption and authentication. The CCM mechanism consists of the generation of an integrity code followed by the encryption of plaintext data and the integrity code. The output consists of the encrypted data and the encrypted integrity code. The

INTEL-1020
U.S. Patent 8,219,129

symmetric authentication operation used in this security suite consists of the generation of an integrity code, using a block cipher in CBC mode, computed on a nonce, followed by (optional) padded authentication data, followed by (optional) padded plaintext data. The verification operation consists of the computation of this integrity code and a comparison to the received integrity code. The symmetric encryption operation used in this security suite consists of the generation of a key stream (using a block cipher in counter mode with a given key) and nonce, and performing an XOR of the key stream with the plaintext and integrity code. The decryption operation consists of the generation of the key stream and the XOR of the key stream with the cipher-text to obtain the plaintext and integrity code. CCM is parameterized by the use of the Advanced Encryption Standard (AES) with the length field L set to 2 bytes, the length of the authentication field M set to 8 bytes, and the length of the nonce set to 13 bytes. AES is 128 bit, performed as specified in NIST FIPS Pub 197 [B113].

The ECMQV Koblitz-283 security suite uses a cryptographic hash function. All operations are on elliptic curves. An elliptic curve key pair consists of an integer q and a point Q on the curve determined by multiplying the generating point G of the curve by this integer, i.e., $Q = qG$. Here, Q is called the public key, and q is called the private key. The ECMQV algorithm can use different types of certificates. The manual certificate is a concatenation of the public key and the device address. The manual certificate is not a digital certificate, because the binding between a public key and its purported owner is to be established and verified by noncryptographic means. The implicit certificate scheme used with the ECMQV security suite is the implicit certificate scheme with associated implicit certificate generation protocol and implicit certificate processing protocol ANSI X9.63-2001. The explicit certificate scheme used with the ECMQV security suite is the ECC X509 certificate scheme, with associated ECC X509 certificate generation protocol and ECC X509[10] certificate processing protocol, as specified in IETF documents RFC 2459 [B85], RFC 3279 [B86], and RFC 3280 [B87]. The certificate is generated and verified using the digital

> An implicit, manual, or digital ECC certificate is verified by first extracting the public key and then extracting the identifying string of its purported owner. Then, the identifying string U of the purported owner is compared with the device address.

INTEL-1020
U.S. Patent 8,219,129

signature algorithm and verification protocol ECDSA, which is fully specified in ANSI X9.62-1998.

The NTRUEncrypt 251-1 security suite uses a cryptographic hash algorithm as specified in the FIPS 180-2 draft standard [B44]. This security suite has a certificate-based security subsuite. The digital certificate format and verification protocol are as in [B87].

The RSA-OAEP 1024-1 security suite is based on modular exponentiation. An RSA public key in this security suite consists of the modulus n and the public exponent e. The private key is information that allows its owner to find e-th roots mod n. The raw public key is the modulus concatenated with the exponent. The encryption scheme in this security suite is as specified in [B118].

Authentication is mutual and based on public-key cryptography. However, to achieve mutual authentication with this system, the devices must trust that the key is indeed associated with the intended device. This trust can be achieved in several ways. One way is to use access control lists (ACLs). The ACL consists of device addresses and other device-specific information such as keys. Before the authentication, the ACL can have information about the key and its associated device ID. For example, some vendors can preload their products with appropriate ACLs to facilitate authentication with certain other devices. Another way to establish trust is by the verification of digital certificates at the time of authentication. Yet another approach is by user action, where, for example, the user pushes buttons to facilitate the authentication process. All of these processes are outside of the 802.15.3 standard.

As a result of the authentication, shared keys are established. Data encryption is performed by a symmetric cipher. Data integrity and sequential freshness are other security services. The beacon may be integrity-protected to assure the devices that indeed, the beacon comes from the coordinator. To maintain security, any device that has performed mutual authentication with another device, can request that the other device periodically transmit a secure frame to indicate its continued presence in the piconet. The standard also ensures that only current

[10] X.509 is part of the ITU X.500 series of recommendations. It specifies the form of the authentication information, and the way to use this information to perform authentication. The certificate can be freely communicated. It contains a user's public key that has been signed by a certificate authority.

INTEL-1020
U.S. Patent 8,219,129

members of the piconet can transmit and receive data. Therefore, when a device joins or leaves the piconet, the group keys are changed.

Medium access mechanism

To support QoS efficiently and to maintain the very successful 802-style CSMA protocol, after considerable deliberation and compromise, IEEE 802.15.3 decided to make the medium access mechanism a combination of CSMA and TDMA. In this way, the advantages of both access mechanisms are combined, and the MAC can efficiently transport asynchronous and isochronous data. The MAC superframe structure consists of three parts: a network beacon interval, a contention access period (CAP), and a contention-free period (CFP), as shown in Figure 3–14. The minimum superframe duration is 512 μs and the maximum superframe duration is 65,535 μs. The beacon contains information about the duration of the superframe, the duration of the CAP, and the contention-free period. The beacon also has information about the identity of the network so that devices may join.

The beacon is also used to synchronize the network coordinator and the devices. Because the medium access mechanism is a combination of CSMA and TDMA, synchronization is very important. The network coordinator contains the reference clock. The beacon sent at the beginning of every superframe contains as a time-stamp the value of the coordinator timer. This time stamp value is equal to the value of coordinator's local timer value at the time of transmission of the first bit of the time-stamp field. Because the frame checking is done after receiving the entire frame, all devices make a copy of the local timer and keep the copy until the end of the beacon frame and adjust it to make it equal to the clock of the coordinator. Unassociated devices that wish to join the piconet synchronize using the same mechanism.

A device must hear at least one beacon before sending its request for association. All devices resynchronize their clocks with the coordinator when they receive the beacon. In addition, the child and neighbor piconet coordinators synchronize their clocks with the coordinator of the parent piconet. The superframe duration can be changed at the discretion of the network coordinator. The contention-free period is reserved for channel time allocations (CTA), including management CTA (MCTA). The boundary between the CAP and CFP periods is dynamically

INTEL-1020
U.S. Patent 8,219,129

adjustable. The CAP period is reserved for transmitting non-QoS data frames. The medium-access mechanism during the CAP period is CSMA/CA. The remaining duration of the superframe is reserved for CTA to carry data frames with specific QoS provisions. Finally, in power-save mode, all QoS features are maintained.

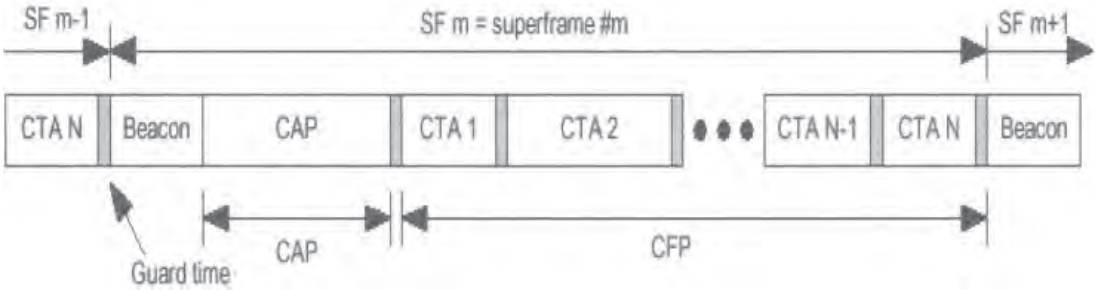

Figure 3–14: Superframe of a IEEE 802.15.3

Four interframe spaces (IFS) are defined: short (SIFS), minimum (MIFS), backoff (BIFS), and retransmission (RIFS). The actual values of these interframe spaces are PHY dependent. Both in the CAP and the CFP, the acknowledgment starts within SIFS after the end of the transmission of the previous frame. Similarly, SIFS duration is allowed between a frame that does not expect an immediate response and the next successive frame. During the CFP, all devices use RIFS for retransmissions. The beacon is transmitted RIFS after the previous transmission. When acknowledgment is not required, in an CTA during the CFP, the transmission of two successive frames will be separated by MIFS.

All data is exchanged in a peer-to-peer manner. Allocating a one-time CTA is similar to polling. The device requests from the coordinator the total amount of time necessary for communication. The PNC grants or rejects this one-time request. Unlike IEEE 802.15.1 all traffic during both the CAP and the CFP is exchanged between peer devices. This ultimately improves latency and the QoS.

It should be noted that the CSMA/CA mechanism in the CAP is slightly different from the one used in IEEE 802.11. In particular, during the CAP no

There are three methods of communication. First, a device can contend for the medium during the CAP using the CSMA/CA scheme. Second, the device can have allocated CTA on a periodic basis for periodic traffic, and third, the CTA may not be on a periodic basis.

INTEL-1020
U.S. Patent 8,219,129

transmission can extend past the duration of the CAP. A device is not allowed to extend its transmissions that started during the CAP in to CFP. Hence, once device decrements its backoff counter to zero, it will check whether there is enough time in the CAP for the transmission of current frame and SIFS. If there is not enough room for this entire frame exchange sequence, then the device will abort the transmission and choose a backoff value with minimum allowed backoff-window size.

A backoff value is randomly (uniformly distributed) chosen in a certain range. The retry count is zero for the first transmission attempt. A backoff counter is maintained, and it is decremented only when the medium is idle for the entire interval minus the RX-TX turnaround time. Whenever the channel is busy, the backoff counter is suspended. When the backoff counter reaches zero, the device will transmit its frame. The backoff counter is also suspended outside CAP duration. In this way, the backoff counter is maintained across multiple superframes and is not reset with every beacon. Immediate acknowledgment is sent without the backoff procedure. The retransmission and the timeout are the same as in IEEE 802.11. The backoff is exponential, where the size of the backoff window is doubled after every unsuccessful transmission attempt. Because the beacon is used for synchronization, the backoff procedure is not applied to the transmission of the beacon. During CAP, a device is allowed to transmit one frame at a time with backoff being applied to every frame. After transmission, if the device does not receive the expected immediate response (acknowledgment), the retry count is incremented. If the maximum number of retries has not been exceeded, the backoff procedure is again resumed. Instead of DIFS, as in IEEE 802.11, a device first waits for a duration equal to BIFS before beginning the backoff procedure.

The PNC divides the CFP into several guaranteed time slots (GTS). Each GTS is a single time slot with guaranteed start time and a guaranteed time duration reserved within the CFP. Hence, a device that is allocated with a GTS is guaranteed that no other devices will compete for the channel during the indicated time and duration of the GTS. During the CFP channel access is entirely controlled by the network coordinator. The PNC assigns channel time allocations to individual devices in a TDMA fashion. The PNC allocates these guaranteed time slots based on the currently pending requests by all the devices in the piconet and the available channel time within CFP. If a device needs channel time during the

INTEL-1020
U.S. Patent 8,219,129

contention-free period, it makes a request to the PNC. These requests take place during the contention access period. There are two types of requests—requests for isochronous or periodic channel time and requests for aperiodic channel time. Rather than requesting recurring channel time, an asynchronous channel time request asks for a total amount of time. The coordinator evaluates the current usage for each new request for QoS and determines if the service can be granted. If the resources are available, the PNC allocates time in a GTS for the device. Note that the allocation by the PNC may not exactly match the duration of time requested by the device. If the requirements for the data change, then the device can request a change to the allocation. The allocated CTA can be shortened or terminated by the transmitting device or the PNC. Upon grant from the coordinator, the device will begin using the CFP in the *next* superframe. The coordinator will assign a stream index to the requesting station. Because all communication is between peer devices, CTAs are specified by a source address and a destination address. Devices cannot extend their transmissions beyond their allocated CTA, including an immediate acknowledgment (if required).

Special CTAs for commands are called management CTAs. They are distinguished by the source address and destination address being the PNC. Adjacent CTAs are separated by guard times. In addition, after the last transmission in a CTA and the beginning of the guard period, there is time period equal to SIFS.

A device with a GTS may or may not make use of all the allocated time duration within the GTS. What to transmit during a GTS is determined locally by the device depending on priority. The devices associated with a PNC can send their changes in channel time requirement. The slot assignments within CFP are based on the current pending requests from all the devices and the currently available channel time within CFP. The slot assignments may change from superframe to superframe as felt required by the PNC.

In no case can a device extend its transmissions beyond the end of the time slot where transmission began. Therefore, the source device will always check whether there is enough time in the time slot for the transmission of the current frame plus an amount of time equal to SIFS. In addition, if an immediate acknowledgment is expected for that frame, the remaining time in the time slot must be large enough to accommodate the current frame, the immediate

INTEL-1020
U.S. Patent 8,219,129

acknowledgment, and an amount of time equal to twice SIFS. If there is not enough room for this entire frame exchange sequence, then this device must abort the transmission and leave the remaining duration of the time slot unused. Within the CFP, however, when an acknowledgment is expected but not received within RIFS, the source device will retransmit the frame as long as there is enough time during its CTA.

Information about the CTAs is included in the beacon. However, due to characteristics of the wireless environment, in any superframe there may be one or more devices in the piconet that receives the beacon in error, depending on their location and interference. If a device misses the beacon, it will not transmit in its CTA. To maintain QoS, even if a device misses the beacon for several superframes, it may still transmit in its CTA, if the CTA is pseudostatic. If a device misses the beacon for longer periods, it will assume that the piconet no longer exists.

Acknowledgment and retransmission

There are three types of acknowledgments: immediate, delayed, implied, and in addition, there is simply no acknowledgment. A transmitted frame with an acknowledgment policy set to indicate no acknowledgment is not acknowledged. The transmitting device assumes that the transmission has been successful. Broadcast and multicast frames are some of the frames that do not require acknowledgments. If the policy is immediate acknowledgment, then following a correct receipt of a frame, the recipient starts the transmission of the acknowledgment within a time period equal to the SIFS. Delayed acknowledgment allows the receiver to group the acknowledgment into retransmission request command. Delayed acknowledgments can be sent anytime before the expiry of the retransmission window. The transmitter can solicit a delayed acknowledgment. Delayed acknowledgment should be used only for periodic traffic. The third type of acknowledgment, implied acknowledgment, is applicable only during the contention-free period. Implied acknowledgment is not used during the CAP. Implied acknowledgment allows the receiver to send another data packet in response. This response data frame is sent after a SIFS period, just like immediate acknowledgment. In turn, the response data frame can set its own acknowledgment policy, which can be of any type except implied.

INTEL-1020
U.S. Patent 8,219,129

During the CAP, retransmissions (up to a certain maximum number) are performed according to the backoff rules. During the time slots within the contention-free period when either immediate acknowledgment or implied acknowledgment is expected, the transmitting device performs CCA between the end of SIFS and the end of RIFS to detect the start of the acknowledgment frame. If an acknowledgment is not detected, the source device can start the retransmission of the frame (or new transmission) at the end of RIFS as long as there is enough channel time remaining in the time slot for the entire frame exchange. When an immediate acknowledgment is expected and received, the transmitting device waits for the duration of immediate acknowledgment frame plus SIFS before attempting another transmission. When an implied acknowledgment is expected and received, the source device waits for at least the duration that was indicated in its immediately previous transmission or the CCA indicating that the channel is idle.

Network Robustness

IEEE 802.15.3 includes measures for network robustness. These measures are transmitter power control, dynamic channel selection, and network coordinator handover.

Dynamic channel selection and transmitter power control are parts of 802.15.3. Two types of transmitter power control are available. During the CAP, the power level is fixed. This prevents some devices from unfairly gaining access to the wireless medium, due to their higher transmit power. During the CFP, however, the power levels are adjustable. The two devices that communicate during a CTA can determine the power level to be used.

This technology also supports dynamic channel selection. To overcome the problems due to overlapping piconets and interference in a given channel, the PNC may choose to move the operations of the piconet to a new channel. Once the coordinator has started a piconet, it will periodically allocate channel time in the CFP to scan the channel for other piconets. This considerably improves the coexistence performance. If the coordinator finds another piconet, it can become the child or neighbor piconet of the other piconet, or it can change to a different channel.

INTEL-1020
U.S. Patent 8,219,129

The network coordinator monitors channel conditions by performing passive scanning or gathering information from the other devices. If it decides that channel conditions are not appropriate, it can dynamically change the channel the piconet is operating on without any disruption in services. The PNC collects the channel status from the devices to arrive at this decision. There is a channel status request command to request that the devices inform the PNC of their channel status. The algorithm to use the channel status information and decide how to perform the change of channel is beyond the scope of 802.15.3.

If the decision is made by the PNC to change the channel, the PNC keeps the piconet quiet by not transmitting any beacon for one or more beacon intervals. Within that quiet time, the PNC may change to one or more other channels to check if one of the other channels is better than the current channel, and then return back to the current channel. If no beacons are received for a certain duration, the devices assume that they are disconnected and restart the association process. If the PNC returns to the current channel within the indicated timeout, the PNC sends a beacon to cancel the quiet state of the piconet. That beacon can include the channel change element indicating the new channel and the timeout for changing channels. The devices that received the beacon with the channel change element will change the channel to the indicated new channel within the indicated time period. They will start waiting for beacons in the new channel. If another channel is not available, the PNC may decide to stay in the same channel. In this case, it will not send a channel change element in its beacon following the cancellation of quiet state of the piconet.

The devices in the piconet can use power-saving techniques to reduce their power consumption. Each device in the piconet can choose transmission power based on the current channel conditions. Furthermore, stations not using a data stream can reduce their power requirements by turning off their physical layer until the next superframe begins. Stations need only be awake during the beacon and the subsequent RTS window of the CAP. If devices are known to have their PHY inactive, the RTS window is used to notify the station of a message. If no message for a station is indicated in the RTS window, and the station does not have activity for the CFP, it can power down until the next beacon. The power state of a device is communicated to other devices by streamless data exchanges. These streamless data exchanges take place during the contention access period. Data can be exchanged in an allocated stream in the CFP, or for small amounts of data without

INTEL-1020
U.S. Patent 8,219,129

QoS requirements, in the data window of the CAP. In the second case, if the destination station is not known to be powered up, the RTS window is also used, and the data will be exchanged only if the destination responds.

Every device can be in one of two states—awake or sleep. Devices that transmit or receive are in the active state. There are four modes: active, hibernate, piconet synchronized power save (PSPS) and synchronous power save (SPS). Devices that are in the hibernate, PSPS, or SPS modes can be awakened by a awake beacon. The awake beacon for devices in the PSPS and SPS modes is periodic. In SPS mode, several devices form a group and can be awakened simultaneously. In hibernate mode, devices conserve power and do not listen for beacons. These modes alone are not enough for battery-powered devices. In addition, these devices shut down during parts of the superframe when they neither transmit nor receive. For example, devices will shut down during the GTS of other devices.

Note that there is a device information request command that each device can use to obtain the information about other devices in the piconet. In this way, peer discovery and peer capability discovery are performed. Thus, the PNC knows the capabilities of the devices in the network.

If the current piconet coordinator is about to be powered down or is moving out of range, it must attempt to choose a device that is capable of being a PNC as its successor. The PNC will send a coordination-handover command to its chosen device with an indication of the handover timeout. The device must accept the nomination and obtain the device information from the current PNC within the indicated timeout period. The new PNC must announce its new responsibility as PNC by sending its first beacon at the first expected beacon transmission time. When the coordination handover is successful, the association of the remaining devices with the piconet is unaffected, and hence they are not required to reassociate with the new PNC. Note that the coordinator handover need not always necessarily stop all the stream transmission.

Each MAC frame consists of a header, a variable-length frame body, and a frame-check sequence (FCS). The header includes a header check sequence, which is verified by the PHY, i.e., the PHY passes on to the MAC frames, which pass the HCS test. The FCS is a 32-bit cyclic redundancy code (CRC). The CRC is calculated over the entire frame body. The maximum size of a MAC frame is 2048

INTEL-1020
U.S. Patent 8,219,129

bytes and the minimum frame body is zero bytes. Fragmentation is performed at the transmitting device on all frames whose size is greater than a certain fragmentation threshold. All fragments except the last must be of the same size. There is no theoretical limit for the size of frames with this fragmentation mechanism.

Physical layer for the 2.4 GHz ISM Band

At present the IEEE 802.15.3 physical layer operates in the 2.4 GHz band. Other physical layers are currently under development. The symbol rate is 11 M symbols per second. Six distinct modulation formats are supported at this symbol rate: uncoded DQPSK, QPSK, trellis-coded QPSK, and 16/32/64-QAM. To achieve low-cost and low-power consumption, OFDM was rejected and IEEE 802.15.3 settled on a single-carrier physical layer. The IEEE 802.15.3 signals occupy a bandwidth of 15 MHz. The total number of channels is five, but they can be used in two different ways, or modes. The first is the high-density mode, which allocates four channels centered at 2.412 GHz, 2.428 GHz, 2.445 GHz, and 2.461 GHz. The second is an IEEE 802.11b coexistence mode, which allocates three channels centered at 2.412 GHz, 2.437 GHz, and 2.461 GHz. A device that is scanning the spectrum for beacons will scan all channels. Furthermore, a device that wants to start a piconet can scan the IEEE 802.11b channels for IEEE 802.11b networks, if it is capable of detecting IEEE 802.11b. If a coordinator detects an IEEE 802.11b network, it will use the IEEE 802.11b coexistence mode.

If a device supports a given QAM modulation format, it also must support all of the lower modulation formats. The base rate of the IEEE 802.15.3 2.4 GHz PHY is 22 Mb/s, uncoded DQPSK mode. DQPSK is used instead of the BPSK to reduce the overhead due to the duration of the PHY and MAC headers. Differential encoding also allows noncoherent receivers to be used. A device can send a frame with one of the supported data rates to a destination device only when the destination device is known to support that rate. All group-addressed frames, regardless of their type, are sent at the base rate that is supported by the piconet because all devices must be able to receive these frames.

Figure 3–15 illustrates the signal constellations used in encoding bit streams into discrete signal levels sent through the air medium. The average power within a packet including the PHY header is required to be a constant, regardless of the

INTEL-1020
U.S. Patent 8,219,129

modulation format. Thus, the constellation must be scaled by a normalization factor such that the PHY header and the frame body have the same average power.

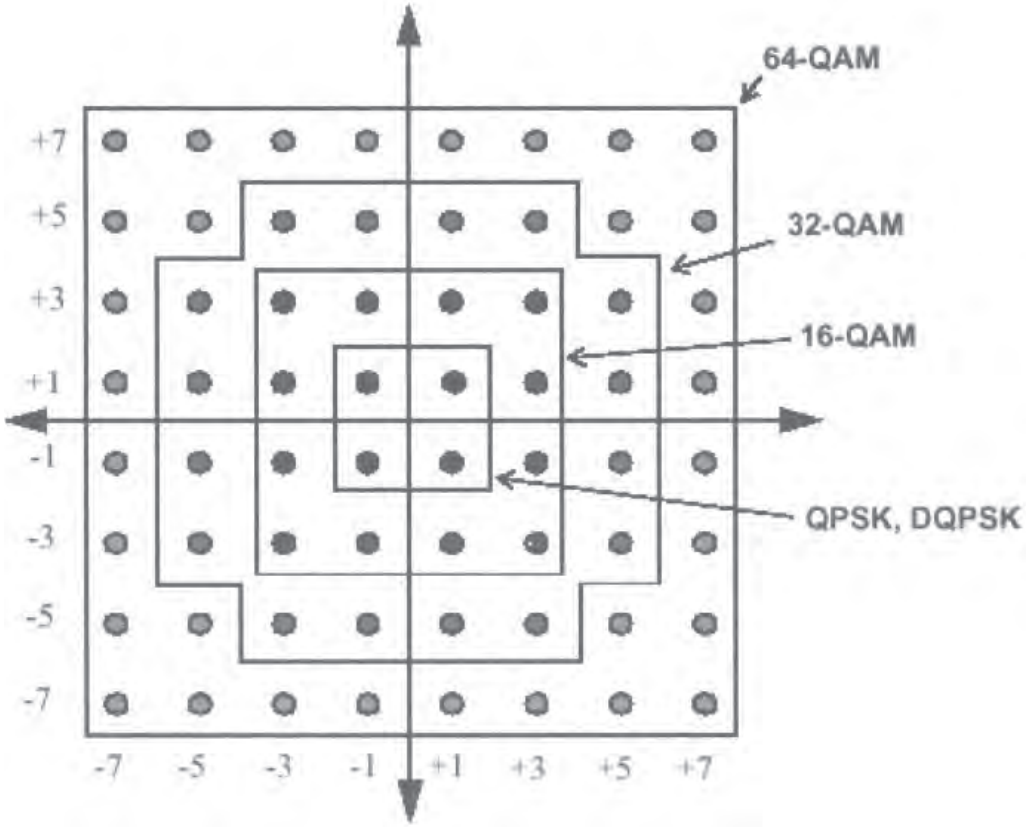

Figure 3–15: DQPSK, QPSK, and QAM constellation diagram

The higher data rates of 33–55 Mb/s are achieved by using 16, 32, or 64-QAM schemes with trellis coding. Coded modulation schemes such as trellis coding are known to provide better coding performance [B134]. In trellis coding, not all information bits are encoded, which results in an overall coding gain. The trellis code that is applied to QPSK and 16/32/64 QAM is a 2-D eight-state (2D-8S) trellis code, as shown in Figure 3–16. The constellation points of the modulation are partitioned into eight subsets: D0, D1, D2, … D7. Each subset contains two, four, and eight symbols for 16/32/64-QAM, respectively. According to this trellis code for 16/32/64 QAM, the lower-order three bits select one of these eight subsets. The remaining higher-order bits select one of the symbols from each

INTEL-1020
U.S. Patent 8,219,129

subset. In the trellis-coded QPSK mode, the lower-order two bits select one of the four subsets D0, … D3, each containing a single symbol.

For 64-QAM, five input bits per symbol are encoded into six output bits. For 32-QAM, in a similar fashion, four input bits per symbol are encoded into five output bits, and for the case of 16-QAM symbol mapping, three input bits per symbol interval are encoded into four output bits. Finally, as shown in Figure 3–16, a 16/32/64 QAM constellation point is selected.

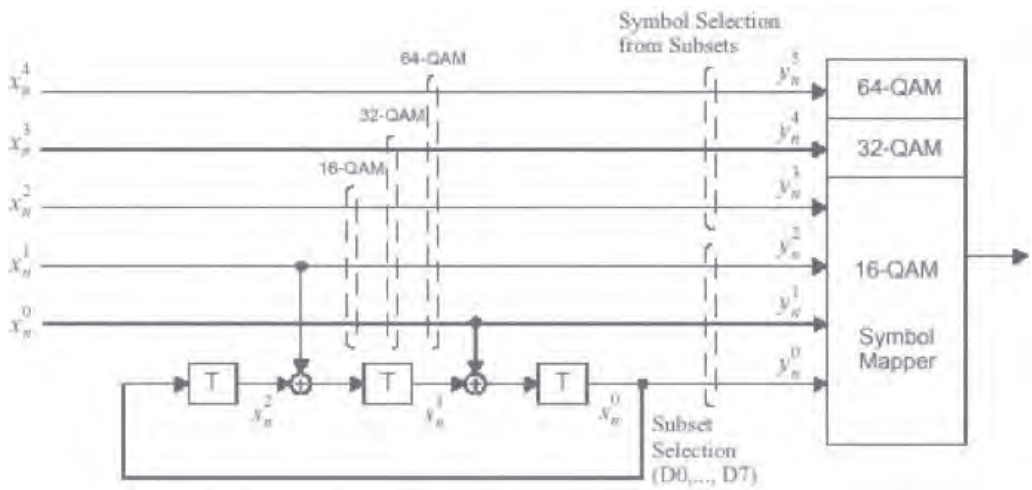

Figure 3–16: Trellis coding diagram for 16/32/64 QAM

For a given output bit vector of the trellis encoder, a 16/32/64-QAM constellation point is selected based on the set partitioning rule illustrated in Figure 3–17 and for QPSK as illustrated in Figure 3–18. These figures also show the bit-to-symbol mappings. The lower (rightmost on the figures) three bits correspond to the subset numbers. The higher 3 bits are assigned within each signal subset (D0, ...D7) such that decimal representation of the bit mapping goes from low to high as the constellation points are traced from bottom to top and from left to right. This rule ensures that the decimal representations of the bit mappings from 0–15 belong to 16-QAM, and 0–31 belong to 32-QAM, and 0–63 belong to 64-QAM.

INTEL-1020
U.S. Patent 8,219,129

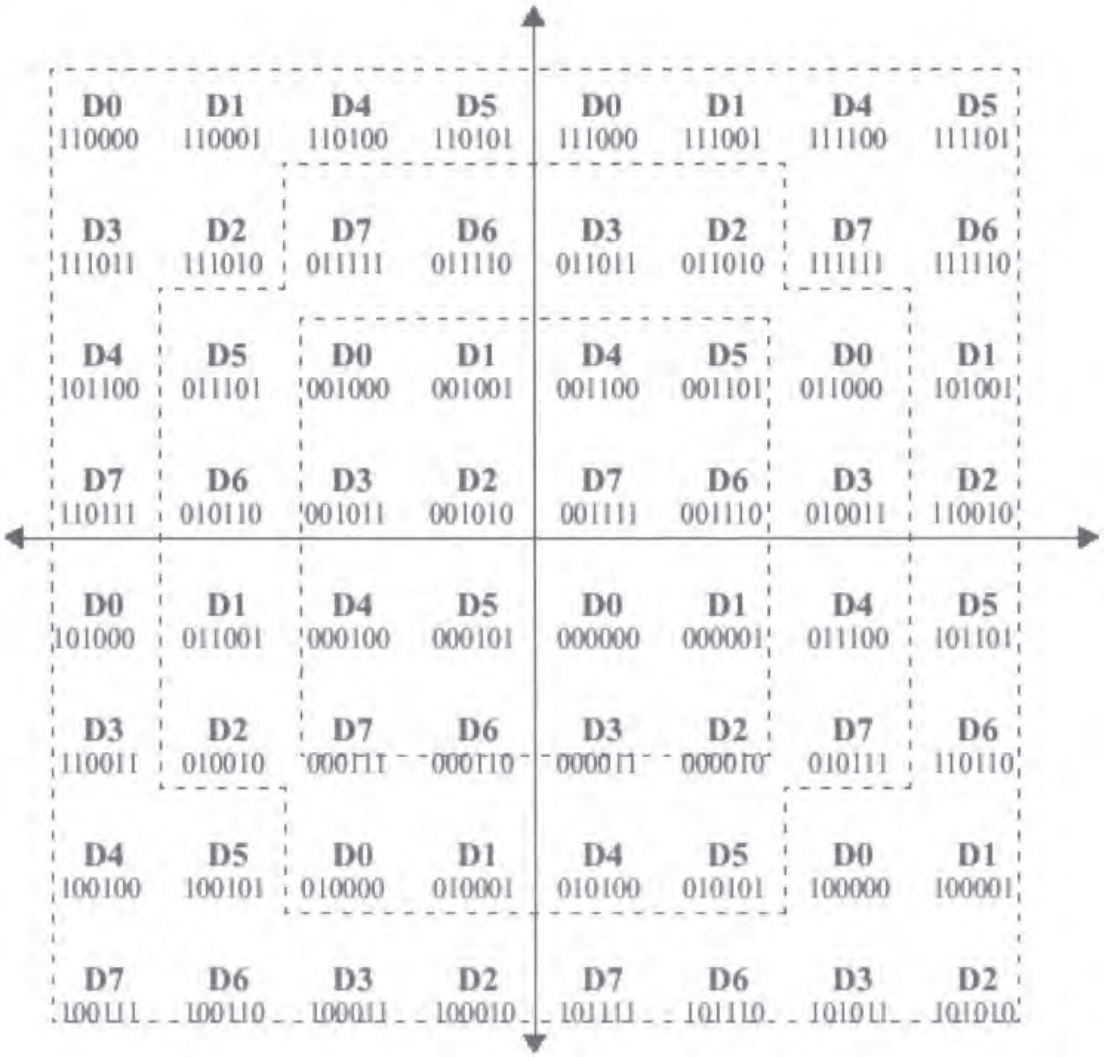

Figure 3–17: Partitioning of the constellation points into eight subsets and bit-to-symbol mapping for 16/32/64 QAM

INTEL-1020
U.S. Patent 8,219,129

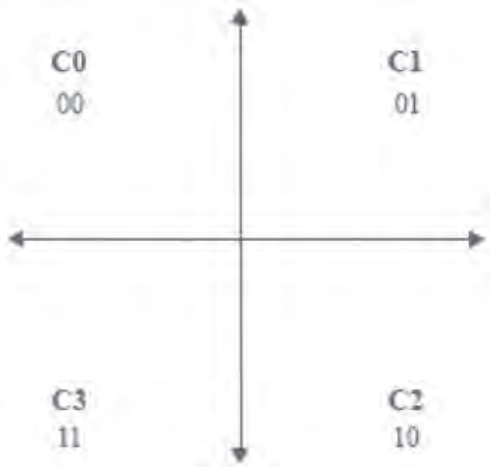

Figure 3–18: Bit-to-symbol mapping and set partitioning for QPSK

The differential encoding applies only to the QPSK mode. In this mode, the entire frame, with the exception of the PHY preamble, is encoded differentially. The differential encoding is provided to allow for noncoherent and, therefore, lower cost receiver implementations.

The IEEE 802.15.3 physical layer adds a PHY header to the MAC header, and calculates the header check sequence (HCS) over the combined MAC and PHY headers. The PHY header has information about the length and the data rate of the payload. The HCS is appended to the end of the MAC header. While the maximum MAC frame is 2048 bytes, this includes the frame body and FCS but not the PHY preamble, PHY header, MAC header, or HCS. The header check sequence contains the 16-bit CRC for the combined PHY and MAC headers. This 16-bit CRC is the 1's complement of the remainder generated by the modulo-2 division of the combined PHY and MAC headers by the polynomial $x^{16} + x^{12} + x^5 + 1$, as in the CCITT CRC-16, and the same as in IEEE 802.11b.

A frame sent over the air is illustrated in Figure 3–19 for data rates of 22, 33, 44, and 55 Mb/s, and in Figure 3–20 for the data rate of 11 Mb/s.

At the rate of 11 Mb/s, two repetitions of the PHY header—MAC header and HCS—are transmitted, as shown in Figure 3–20. The first repetition is sent using 22 Mb/s DQPSK. The second repetition is sent using 11 Mb/s QPSK + TCM. This

INTEL-1020
U.S. Patent 8,219,129

repetition makes the header more robust, which is desirable because the header is more important.

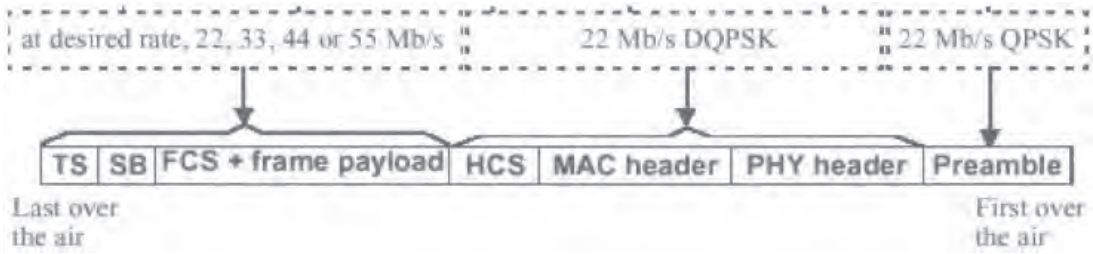

Figure 3–19: Transmitted frame for data rates of 22, 33, 44, and 55 Mb/s

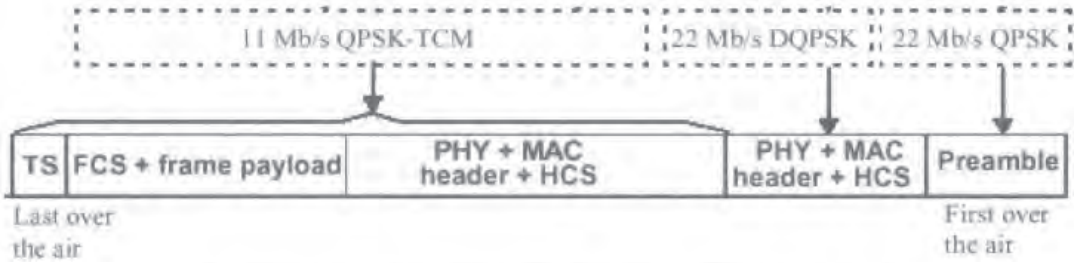

Figure 3–20: Transmitted frame at 11 Mb/s

A PHY preamble is added to aid receiver algorithms related to synchronization, carrier offset estimation, and channel estimation and equalization. The preamble consists of successively appending 10 periods of a special sequence. The sequence is called constant-amplitude zero autocorrelation (CAZAC) sequence. Each period of the CAZAC sequence contains 16 complex symbols, C_0, $C_1,\ldots C_{15}$, where

$$C_i = 1 + j,\ i = 0,\ldots,3, \text{ and } i = 7,9,11,15$$

$$C_i = -1 + j,\ i = 4,14$$

$$C_i = -1 - j,\ i = 5,8,10,13$$

$$C_i = 1 - j,\ i = 6,12$$

Clearly these complex symbols are of QPSK type. The first nine periods of the preamble are a concatenation of the CAZAC sequence. In the tenth period, each

INTEL-1020
U.S. Patent 8,219,129

element of the CAZAC sequence is negated, or equivalently, rotated by 180 degrees to denote end of preamble.

The MAC header, HCS, and payload are scrambled, while the PHY preamble and PHY header are not scrambled. The scrambler is described by the polynomial $1 + x^{14} + x^{15}$. There are four initialization vectors for the scrambler. For the first frame, the scrambler is initialized with 0011111111111111, for the second, 0111111111111111, for the third, 1011111111111111, for the fourth, 1111111111111111. After the fourth frame the initialization pattern periodically repeats. The first two bits of the initialization vector are sent to help the receiver estimate the initialization vector.

If the number of payload bits does not correspond to integer number of symbols, then stuff bits (SB) are added. The data rate of 11 Mb/s corresponds to one bit per symbol, and therefore stuff bits will not be needed. The use of TCM requires tail symbols. These tail symbols are used to terminate the encoded trellis sequence in a known state to aid the decoding process. At the end of each frame, the trellis should be terminated in state 0. For 11 Mb/s QPSK-TCM format 3 symbols, each containing 1 bit, are appended to the end of the MAC frame body. For 16/32/64 QAM modulation two tail symbols are appended to the end of the payload to terminate the encoded trellis sequence in state 0.

The receiver sensitivity is defined as the minimum received signal level in dBm that guarantees a packet error rate of less than 8 percent. For the different modulation schemes, receiver sensitivities range from –68 dBm to –82 dBm.

For the maximum transmit power, IEEE 802.15.3 devices must satisfy FCC regulations in 15.249.

Low-Rate WPAN

Overview and architecture

After the development of Bluetooth, it became clear that the one technology-fits-all approach is not appropriate. Many industrial, agricultural, medical, and vehicular applications, such as sensors, meter reading, smart tags/badges, and home automation, require short-range wireless connectivity that is different from Bluetooth. Two parameters are very important for these applications—ultra low power consumption and ultra low cost. Power

INTEL-1020
U.S. Patent 8,219,129

consumption must be low so that batteries can last several months or longer. Data rate is not that important for these applications, and they are willing to trade it for very long battery life. Bluetooth or the high-rate WPAN technology cannot address these applications. IEEE 802 addressed these market needs by developing low-rate WPAN technology. The standard is called IEEE 802.15.4 and (similarly to the other 802 standards) defines the physical and medium access control layers. In addition to defining these two layers, the standard aims to achieve a certain level of coexistence with other wireless devices.

> The data rate may be low, but the QoS requirements ensure that wireless devices built using this technology can carry signals such as voice.

In general, low-rate WPAN devices can be fixed, portable, and/or moving. The main objective of the 802.15.4 standard is a simple, but flexible and efficient, protocol. It was decided that a IEEE 802.15.4 network is composed of two different device types—a full-function device and a reduced-function device. The full-function device can operate in three modes—as a network coordinator (controlling the network), a coordinator, or a network device. It can dynamically switch its mode of operation. Reduced-function devices cannot control the network and can be extremely simple, which is appropriate for devices such as a light switch or a passive infrared sensor. Reduced-function devices cannot communicate with other reduced-function devices; they can only communicate with full-function devices. As a result, the protocol for reduced function devices can be implemented using minimal resources and memory capacity.

Low-rate WPAN devices achieve data rates of 250 Kb/s, 40 Kb/s, and 20 Kb/s, use CSMA/CA channel access, have power management, and use 16 channels in the 2.4 GHz ISM band. The remainder of this section is devoted to the network topologies for low-rate WPAN, the protocol architecture, and an overview of the constituent protocols.

There are two basic network architectures for low-rate WPAN—star and peer-to-peer, as shown in Figure 3–21. In the star topology, the network nodes can only communicate with the single central controller, i.e., the network coordinator. A network device has some associated function and is either the initiation point or the termination point for network communication. A network coordinator can be used to initiate, terminate, or route communication around the network. The

INTEL-1020
U.S. Patent 8,219,129

network coordinator does not need to be a host PC. One difference compared with other technologies is that while the nodes are generally bidirectional, unidirectional nodes are permitted with extremely limited network access. Messages between two network devices can be exchanged through virtual peer-to-peer links, where the message frames are routed though the network coordinator. The star topology is appropriate for PC peripherals and home automation products. In addition to the two basic network types, more complex network architectures are possible, such as cluster-tree networks. Therefore, 802.15.4 supports the architectural requirements of wireless sensor networks.

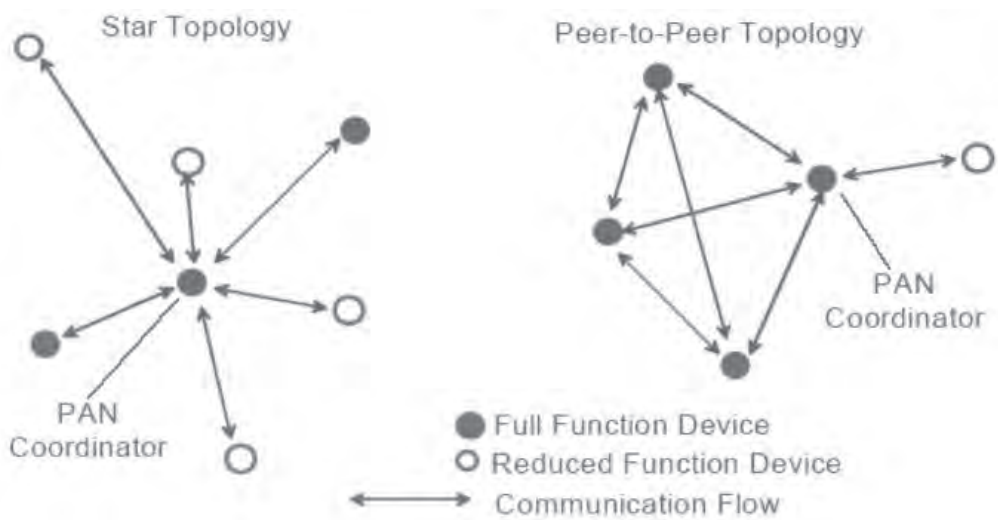

Figure 3–21: Topology of a low-rate WPAN

In the peer-to-peer topology, any node can communicate with any other nodes as long as they are in range of one another. Therefore, the peer-to-peer topology requires only full-function devices. In this mode, more complex network topologies, such as the cluster tree topology, are possible. Applications such as intelligent agriculture, industrial control and monitoring, wireless sensor networks, asset and inventory tracking, and security would benefit from such network topology. Its advantage is that it is ad hoc, self-organizing, and self-healing. It allows multihops to route messages from any node on the network to any other node on the network. While 802.15.4 allows this type of network, the formation and management of a cluster-tree network are performed by a layer above the MAC layer of 802 standards. Furthermore, both the star and

INTEL-1020
U.S. Patent 8,219,129

peer-to-peer architectures can be connected to other networks. For example, if the PAN coordinator is a computer, the computer can be connected to other wireless or wired networks.

The LR-WPAN standard consists of MAC and a PHY. A LR-WPAN device is comprised of a physical layer, which contains the RF transceiver and any low-level control of the transceiver, and a medium access control sublayer. The layers interface by one or more service access points.

Power consumption is paramount in this technology. In many applications that use IEEE 802.15.4, the devices will be battery powered. Reduced power consumption is achieved in two ways. First, the IEEE 802.15.4 protocol is developed so that devices require little power to operate. Second, in the physical implementation of 802.15.4, additional power management techniques can be used.

The additional power management techniques are one area of vendor differentiation in the marketplace. Note that some devices could potentially be mains powered. Battery-powered nodes will typically use duty-cycling to reduce power consumption. In fact, these devices will spend most of their existence in a sleep state, not participating in the network. However, each network device must periodically listen to the RF channel for a network beacon in order to determine whether a message is pending from a network coordinator or another device, and to stay synchronized to the network. This mechanism allows the application designer to decide on the balance between battery consumption and message delay latency. Mains-powered network devices have the option to listen to the RF channel continuously.

PHY Layer

The features of the PHY layer are activation and deactivation of the radio transceiver, channel adjustment, and clear channel assessment. To achieve low cost, the radio of the on-air protocol stack operates in the 2.4 GHz ISM band, available worldwide, and in lower ISM bands such as the 868 MHz band, available in Europe, and the 915 MHz band available in the U.S. (Table 3–4). A total of 27 channels, numbered 0–26, are available in these three frequency bands. Sixteen channels are available in the 2.4 GHz band, 10 in the 915 MHz band, and one in the 868 MHz band.

INTEL-1020
U.S. Patent 8,219,129

The center frequencies of these channels are:

f_c = 868.3 MHz in the 868–868.6 MHz band, available in Europe,

$f_c = 906 + 2(k - 1)$ MHz, $k = 1,\ldots,10$, in the 900 MHz ISM,

$f_c = 2405 + 5(k - 11)$ MHz, $k = 11,\ldots,26$, in the 2400 MHz band.

The physical layer in the 2.4 GHz band uses 16-ary quasi-orthogonal modulation. Four information bits select one of the 16 nearly orthogonal PN sequences to be transmitted. The complex chip sequence is modulated on the RF carrier using offset-QPSK modulation. Even-indexed chips are modulated onto the I carrier, and odd-indexed chips are modulated onto the Q carrier. According to offset QPSK, the Q-carrier will be delayed with respect to the I-carrier by half of the chip period. Each PN sequence has 32 chips. These 16 sequences are:

1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0 1 1 1 0
1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0
0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0
0 0 1 0 0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1
0 1 0 1 0 0 1 0 0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0 0 0 1 1
0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1 1 1 0 0
1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0 1 1 1 0 1 1 0 1 1 0 0 1
1 0 0 1 1 1 0 0 0 0 1 1 0 1 0 1 0 0 1 0 0 0 1 0 1 1 1 0 1 1 0 1
1 0 0 0 1 1 0 0 1 0 0 1 0 1 1 0 0 0 0 0 0 1 1 1 0 1 1 1 1 0 1 1
1 0 1 1 1 0 0 0 1 1 0 0 1 0 0 1 0 1 1 0 0 0 0 0 0 1 1 1 0 1 1 1
0 1 1 1 1 0 1 1 1 0 0 0 1 1 0 0 1 0 0 1 0 1 1 0 0 0 0 0 0 1 1 1
0 1 1 1 0 1 1 1 1 0 1 1 1 0 0 0 1 1 0 0 1 0 0 1 0 1 1 0 0 0 0 0
0 0 0 0 0 1 1 1 0 1 1 1 1 0 1 1 1 0 0 0 1 1 0 0 1 0 0 1 0 1 1 0
0 1 1 0 0 0 0 0 0 1 1 1 0 1 1 1 1 0 1 1 1 0 0 0 1 1 0 0 1 0 0 1
1 0 0 1 0 1 1 0 0 0 0 0 0 0 1 1 1 0 1 1 1 1 0 1 1 1 0 0 0 1 1 0 0
1 1 0 0 1 0 0 1 0 1 1 0 0 0 0 0 0 1 1 1 0 1 1 1 1 0 1 1 1 0 0 0

INTEL-1020
U.S. Patent 8,219,129

Table 3–4: Frequency bands and data rates

Band	Chip rate	Modulation	Bit rate	Symbol rate
868–868.6 MHz	300 kchip/s	BPSK	20 Kb/s	20 K
902–928 MHz	600 kchip/s	BPSK	40 Kb/s	40 K
2400–2483.5 MHz	2 Mchip/s	O-QPSK	250 Kb/s	62.5 K

The symbol rate is 62.5 Ksymbols/s, or alternatively, 250 Kb/s. Because each symbol is represented by a 32-chip sequence, the chip rate is 32 times the symbol rate, i.e., the chip rate is 2 Mchips/s. In offset QPSK, the Q-phase is delayed compared with the I-phase by one-half of a chip period with respect to the I-phase chip sequence. Half-sine pulse shaping is used, where

$$p(t) = \sin\frac{\pi t}{2T_c} \text{ when } 0 \le t \le 2T_c, \text{ and } p(t) = 0 \text{ otherwise.} \qquad \text{Eq. 3–4}$$

Figure 3-22 shows an example of chip sequences with half-sine pulse shaping. The synchronization header consists of the concatenation of 32 binary zeros and the 8-bit start-of-frame delimiter 11100101. The PHY header contains the number of bytes in the payload. The maximum packet size is 127 bytes. The preamble sequence is used by the transceiver to obtain chip and symbol synchronization with an incoming message; it also provides a settling period for automatic gain control (AGC) circuitry.

Figure 3–22: Example of chip sequences with half-sine pulse shaping

INTEL-1020
U.S. Patent 8,219,129

In the 868 MHz band, the data rate is 20 Kb/s, and in the 915 MHz band, 40 Kb/s. DSSS and BPSK are used in these bands. First, the bits are subject to differential encoding, which is the exclusive OR operation of a raw bit with the previous encoded bit. Then each bit is mapped to a 15-chip PN sequence. If the input bit is zero, it is mapped to 1 1 1 1 0 1 0 1 1 0 0 1 0 0 0. If the bit is one, it is mapped to 0 0 0 0 1 0 1 0 0 1 1 0 1 1 1.

The chip sequence is modulated using BPSK with raised cosine pulse shaping with a roll-off factor equal to 1. This pulse shape is described by the following equation:

$$p(t) = \frac{\sin(\pi t/T_c)}{(\pi t)/T_c}\frac{\cos(\pi t/T_c)}{1-4t^2/T_c^2} \qquad \text{Eq. 3–5}$$

Note that the 2 Mchip/s OQPSK waveform of the IEEE 802.15.4 2400 MHz air interface, with 6 dB bandwidth of about 1.5 MHz, and the 600 kchips/s BPSK waveform of the 902–928 MHz air interface, with bandwidth of about 600 kHz, both meet the definition of digital modulation.

The chip rate is 300 Kchips/s for the 868 MHz band and 600 Kchips/s for the 915 MHz band. This chip rate is 15 times higher than the symbol rate.

The PHY must be able to perform clear channel assessment (CCA). There are three ways to perform CCA—energy thresholding, IEEE 802.15.4 signal detection, and a combination of both (IEEE 802.15.4 signal detection with energy above a threshold).

IEEE 802.15.4 MAC

The main features of the MAC sublayer are beacon management, channel access mechanism, dynamic channel selection, frame reception and acknowledgment, association, disassociation, and security. Security functions are necessary for IEEE 802.15.4 to succeed in the market. Therefore, access control lists and symmetric cryptography are supported.

The LR-WPAN uses two types of channel access, depending on the network configuration. Networks without a beacon use a CSMA/CA channel access mechanism. This CSMA/CA channel access mechanism is much like the mechanism in IEEE 802.11.

INTEL-1020
U.S. Patent 8,219,129

Networks with a beacon use a slotted CSMA/CA channel access protocol. The slotted CSMA/CA mechanism operates as follows. A station that wants to transmit senses the channel. If the channel is idle it will wait for a random number of slots and begin transmitting on the next available slot boundary. Backoff slots are aligned with the start of the beacon. The beacon itself and acknowledgments are sent without CSMA/CA.

There are three types of data transfer. The first two are from a RFD to a FFD and from a FFD to a RFD. One is from the network device to the network coordinator. The second is from the network coordinator to the network device. The mechanisms for each transfer type depend on whether the network is beacon-enabled or not. When a network device needs to transfer data to the network coordinator in a beacon-enabled network, it first listens for the network beacon. The beacon is necessary so that the network device can determine when it can transmit. In a non-beacon-enabled network, the network device directly uses the CSMA protocol to transmit its data to the network coordinator. The network coordinator acknowledges a successful transmission. The acknowledgment policy can be no acknowledgment.

The network coordinator can send data to a network device in one of the following two ways. In a beacon-enabled network, the coordinator indicates pending data for a network device in its network beacon. The network device requests the data and then the network coordinator sends the appropriate data packet. The network device acknowledges successful transmission. In a non-beacon-enabled network, the coordinator must keep the data until the appropriate network device makes contact and requests the data.

A network device can contact the network coordinator in two ways. First, the network device can poll the network coordinator at a specific rate. If data is pending, the network coordinator transmits the data. The second way follows data transfer from the device to the coordinator. Then the network coordinator can indicate in its acknowledgment that it has data for the network device.

Peer-to-peer data transfers follow the CSMA scheme.

The packet structures have been designed to keep the complexity to a minimum while at the same time making them sufficiently robust for transmitting on a noisy channel. The LR-WPAN defines four frame structures: beacon frame, data frame,

INTEL-1020
U.S. Patent 8,219,129

acknowledgment frame, and a MAC command frame. The beacon frame contains a beacon payload.

Every MAC frame consists of a MAC header, variable-length payload, and a footer consisting of the frame check sequence (FCS). The MAC header comprises frame control and addressing information. The frame check sequence contains an industry-standard 16-bit cyclic redundancy code, specified by the generator polynomial $G(x) = x^{16} + x^{12} + x^5 + 1$.

Suppose that $M(x) = b_0x^{k-1} + b_1x^{k-2} + \ldots + b_0$ is the polynomial representation of the sequence of bits that we want to perform CRC on. First, $M(x)x^{16}$ is divided by $G(x)$ modulo 2 to obtain $R(x)$. The coefficients of $R(x)$ are the value of the CRC. The MAC header contains the destination and source addresses and the CRC bit field.

Before devices can exchange data, a network must be formed. A low-rate WPAN can be started by a full-function device. The device that wants to start a PAN must first perform an active or passive channel scan. The active channel scan consists of transmitting the beacon request command in every channel. If there are network coordinators that are active, they, in response to a beacon request command, will start transmitting beacons. A passive channel scan is when a device simply listens for beacons in every channel. After completing the scan, if there are no coordinators in the channel where it wants to operate, the device must select a suitable PAN identifier and will start to consider itself as a network coordinator.

Device discovery is always important in PAN, because it enables a variety of applications. To associate with a network, a device does a scan to find existing networks to join. If more than one network is operating, the device must select which PAN to join. The algorithm according to which a PAN is selected is not specified and is left to the application developers. After selecting a network, the device can initiate association by sending an association request command during the contention access period of an existing PAN. When the coordinator receives the association request it will decide whether there are enough resources available for the device and assign a 16-bit address to the device.

Two types of access mechanisms are used, depending on whether the network is beacon-enabled or not. The beacon contains information about the structure of the

INTEL-1020
U.S. Patent 8,219,129

following frame. The beacon is transmitted without the use of CSMA. In the beacon, the network coordinator device provides information about the permitted packet size, the current network coordinator operational state, and the length of the contention period. Because these parameters can change, the network coordinator must keep the information contained in the beacon up-to-date.

The contention period, in which the medium access mechanism is CSMA/CA, starts immediately after the beacon. In the contention period all frames, except acknowledgments, use CSMA/CA to access the channel. Considering the CSMA/CA mechanism for low-rate WPANs, the MAC sublayer keeps track of the number of retry attempts. Initially the number of retry attempts is set to 0. The MAC then issues a clear channel assessment request to the physical layer. The physical layer senses the channel to determine whether it is free or busy. If no activity is detected, the device may transmit. If activity is detected the device backs off for some random period before beginning the CSMA/CA procedure again. If the channel is idle then the MAC will check whether the backoff timer has expired before transmitting. At every time slot the backoff timer is decremented by one, and clear channel assessment is performed. If the channel is idle and the backoff timer has expired the device can transmit. The MAC must ensure that the entire transmission, including an acknowledgment, if one is requested, can be accomplished before the end of contention access period. If the channel is busy, the number of backoff attempts is incremented by one. If this value is greater than the maximum number of backoff attempts then the transmission is considered to have failed. Acknowledgements and retransmissions are optional. Every frame has a field, called the acknowledgment field. If this bit is set, acknowledgment is required. Otherwise it is not. If acknowledgment is not required the frame is never retransmitted. Otherwise retransmissions will be performed, up to a certain maximum.

The LR-WPAN standard allows the optional use of a superframe structure. The super frame format is defined by the network-coordinator. A superframe consists of beacons, contention access period, and contention-free period as shown in Figure 3–23.

While this is similar to the superframe used in 802.15.3, the contention-free period is optional. A superframe may have active and inactive portion, as illustrated in Figure 3–24. During the inactive portion of a superframe, devices

INTEL-1020
U.S. Patent 8,219,129

can enter low-power mode. The beacons are sent by the network coordinator and are used to synchronize the devices and to identify the network. The channel access mechanism during the contention period is slotted CSMA/CA. All transactions must be completed by the time of the next network beacon. For low latency applications or applications requiring a specific data bandwidth, the network-coordinator may allocate dedicated slots, called guaranteed time slots (GTSs). The GTSs comprise the contention-free period. The number of guaranteed time slots in a frame is up to the network. If there is a contention-free period, all contention-based transactions must be completed before the first allocated slot begins. Also each device transmitting in an allocated slot must ensure that its transaction is complete before the next allocated slot or the next beacon begins. The beacon contains information about the length of the contention and contention-free periods. A beacon-enabled network is used for supporting low-latency devices such as PC peripherals. If the network does not need to support such devices, it can elect not to use the beacon for normal transfers. However, the beacon is still required for network connection purposes.

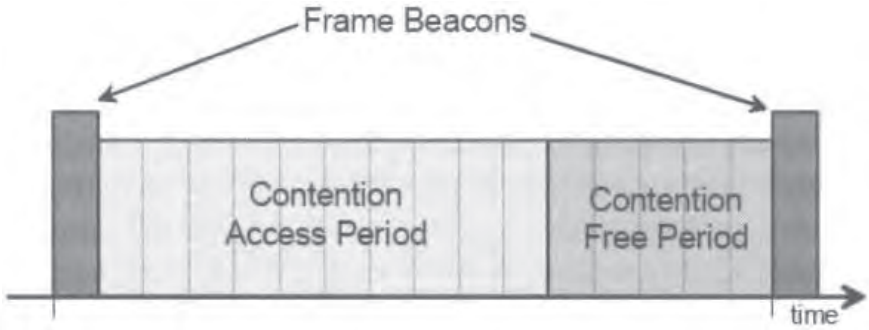

Figure 3–23: Superframe structure including an optional CFP

INTEL-1020
U.S. Patent 8,219,129

Figure 3–24: General superframe structure, with active and inactive periods

The GTSs follow immediately after the contention period. Within the GTS devices do not use CSMA/CA. GTSs are allocated exclusively for communication between the coordinator and a device. The device address and direction of communication (to or from the coordinator) uniquely identifies a GTS. GTS are allocated on a first-come, first-served principle. At any one time, there may be up to seven GTSs, because the device address field is three bits long. QoS can be changed on the fly—GTSs are allocated and deallocated dynamically at the discretion of the network coordinator. A device that has a GTS can also operate in the CAP. Only full-function devices can have GTSs, reduced-function devices cannot have GTSs. One full-function device can have one transmit GTS and/or one receive GTS. A device can request the allocation and deallocation of GTS. In response to a request, the network coordinator will grant a GTS if capacity is available. Note that after the deallocation of a GTS, the CFP can become fragmented. The network coordinator must ensure that there are no holes left by redundant slots. It may have to reallocate the GTS to make the contention-free period one contiguous block with no gaps. Reallocation is performed by appropriately informing the concerned devices. The set of slots must complete before the start of the next superframe. GTSs are not equal in length, and in this way devices with different bandwidth requirements can be accommodated.

In a network that does not use beacons, the entire period can be considered a contention access period. Devices operating in a non-beacon-enabled PAN will poll the coordinator for data on a periodic basis.

The access mechanism provides two types of QoS configurations. The first one is the priority in queuing. It is available in the contention-access period. The priority

INTEL-1020
U.S. Patent 8,219,129

flag in the frame control field of the MAC header indicates the priority of the higher layer assigned to the packet. The priority types are high and normal. The packet marked high should transmit before the normal packets if both types of packets are in the queue. Beacon-enabled networks can also have a contention-free period, which provides the second QoS mechanism.

Several features of the protocol enable very low power consumption—lower than any other IEEE 802 standard. A superframe has an active part and an inactive part. Devices can enter low-power mode during the inactive part. The boundary between the active and inactive part can be adjusted.

Because the network is wireless, it is possible for a device to stop receiving beacons from the coordinator for an extended period of time. Devices that become lost should not transmit to avoid conflicts on the wireless medium. A lost device can do an orphan scan to realign itself. The lost device sends an orphan notification command on each channel. If a coordinator receives the orphan notification command, it will check its device list to determine whether the orphaned device is a member of its PAN. If the device is a member of its PAN, the coordinator will send a coordinator realignment command back to the orphaned device. If the orphan scan is unsuccessful, the device will assume that it has permanently lost the connection.

Security

The MAC is responsible for providing security services. The three security modes in the order of increased security are: the unsecured mode, ACL mode, and secured mode. The unsecured mode is mandatory to implement. ACL mode and secured mode are optional, but devices that support secured mode also must support ACL mode. Security can be delegated to the higher layers. Unsecured mode does not provide any security.

Access control is a method where a device can select which devices it is willing to communicate with. In ACL mode, devices check for every frame even if it is received from a device in the access control list. The access control list is composed of devices with which a key is shared. There are no cryptographic operations in ACL mode. In ACL mode, there is only the access control security service. The secured mode has all four security services: access control, data encryption, frame integrity, and sequential freshness. However, key management

INTEL-1020
U.S. Patent 8,219,129

and device authentication are outside of the 802.15.4 standard and can be done only at higher layers.

> Each device that implements security must support one or more security suites. The security suite indicates the symmetric cryptographic algorithm and integrity code.

All security suites use AES. Secured mode also uses the access control list, but in addition, includes cryptographic protection. In secured mode, the access control list is checked when a frame is transmitted. Frames are not sent to a device not in the access control list. Every received frame is also checked to see whether it comes from a device in the access control list. Furthermore, when a frame is sent, the security suite is checked from the access control list. If the security suite defines encryption, encryption is performed on the MAC payload.

If the security suite defines integrity code, the integrity code is applied to the MAC header concatenated with the MAC payload. The result of the integrity code computation is placed in the payload in addition to any other data in the payload. Finally, the FCS is computed over the modified frame. At the receiver, first, the FCS is checked. Then, if the security suite defines encryption, the decryption operation is performed on the payload. Data encryption protects data from being read by parties without the cryptographic key. Data encryption may be performed on beacon payloads, command payloads, and data payloads. If the security suite defines an integrity code, the integrity code is checked. If any of the security operations fail, an error is generated. Similar to IEEE 802.11 and IEEE 802.15.3, frame integrity is a security service that uses a message integrity code (MIC) to protect data from being modified by parties without the cryptographic key. Integrity can be provided on data frames, beacon frames, and command frames. Sequential freshness is a security service that uses an ordered sequence of inputs to protect data with an old sequence value from being replayed. This provides evidence that the received data is newer than the last data received.

Security suites supported by 802.15.4 are described next. A security suite is a set of operations. The name of security suite indicates the symmetric cryptographic algorithm, mode, and integrity code bit length. The length of the integrity code is less than the block size of the cryptographic algorithm and it indicates the probability that a random guess of the integrity code would be correct. There are

INTEL-1020
U.S. Patent 8,219,129

seven security suites defined in the standard: AES-CCM-64 is mandatory to implement. Six more are optional.

The first security suite is AES-CTR. In the AES-CTR symmetric encryption algorithm, a key stream is generated first using a key and a nonce. Each block in the key stream is different from all other blocks. A block cipher is applied to the set of input blocks to produce a sequence of output blocks. These output blocks are XORed with the plaintext to produce the ciphertext. This encryption is symmetric and the decryption operation consists of the generation of the key stream and the XOR of the key stream with the ciphertext to recover the plaintext. Because all operations are performed on blocks of data this algorithm lends itself very well for a parallel implementation. AES-CTR provides access control, encryption, and sequential freshness, but it alone cannot provide frame integrity.

Frame integrity can be provided by the CBC-MAC symmetric authentication algorithm. CBC-MAC can also provide access control, but it alone cannot provide sequential freshness and encryption. The CBC-MAC security suite consists of performing authentication on the encoding of the length of the combined header and payload, concatenated with the header, concatenated with the payload. The CBC-MAC used in 802.15.4 consists of the generation of an integrity code using a block cipher in CBC mode. The CBC-MAC algorithm makes use of an underlying block cipher to provide data integrity on input data. The block cipher transforms (or encrypts) input vectors of the block size. The data to be authenticated is grouped into contiguous blocks, each with length equal to the block size. The verification operation consists of the computation of the integrity code and comparison to the received integrity code. Depending on the size of the integrity code, 128-bit, 64-bit, and 32-bit CBC-MAC algorithms are supported by IEEE 802.15.4.

Like IEEE 802.11, IEEE 802.15.4 is using the combined symmetric encryption and authentication algorithm called CTR + CBC – MAC (CCM) [B144] provides all four security services. Depending on the integrity code, 128-bit, 64-bit, and 32-bit CCM is supported. The block size in all cases is 128 bits. Extensions to other block sizes are possible, but are not straightforward. CCM is a generic authenticate-and-

> Note that in the CTR encryption algorithm, cipher text is produced from plain text, and the message integrity check is not encrypted.

INTEL-1020
U.S. Patent 8,219,129

encrypt block cipher mode. The CCM symmetric encryption and authentication consist of the generation of an integrity code followed by the encryption of plain-text data and the integrity code. The output is the encrypted data and the encrypted integrity code.

There are two parameters, M (the size of the authentication field) and L (the size of the length field). Larger values of M ensure greater protection against an attacker who can undetectably modify a message. At the same time, larger values of M lead to a larger message size. In 802.15.4 M is between 4 and 16 bytes and L is between 2 and 8 bytes. This accommodates a variety of applications. In CCM-128, L is 2 bytes and M is 16 bytes, in AES-CCM-64, M is 8 bytes and L is 2 bytes. In AES-CCM-32, M is 4 bytes, and L is 2.

To encrypt a message, the sender needs (1) an encryption key K and (2) a nonce of $N = 15 - L$ bytes. The nonce must be unique for the given encryption key. The message is divided into blocks of $l(m)$ bytes. There can be additional authenticated data that is not encrypted. The first step is to perform authentication. This is done using CBC-MAC. Then the message is encrypted with the CTR algorithm. The final result is the concatenation of the encrypted message, followed by the encrypted authentication value.

Coexistence issues involving IEEE 802.15.4

Clearly, IEEE 802.15.4 is yet another technology that operates in the 2400 MHz band. The channelization scheme of 802.15.4 in the 2400 MHz band is shown in Figure 3–25. As a result, a systematic investigation of coexistence is too complex. Because IEEE 802.15.2 does not consider IEEE 802.15.4, the coexistence performance of IEEE 802.15.4 with respect to IEEE 802.11b, IEEE 802.15.1 and IEEE 802.15.3 is analyzed in a separate annex to the standard. Generally IEEE 802.15.4 devices are expected to coexist well with other wireless systems.

The coexistence characteristics of 802.15.4 include the following:

- The medium access mechanism is CSMA-CA to avoid interference. In particular performing clear channel assessment by energy thresholding increases the coexistence performance. Note that in all coexistence scenarios combining energy level information and signal-to-noise ratio information

INTEL-1020
U.S. Patent 8,219,129

provides information whether the packet errors are a result of low signal levels or strong interference.

- To make the protocol frequency agile, 802.15.4 specifies dynamic channel selection (DCS). 802.15.4 devices will work in that portion of the ISM frequency bands that is free of interference.
- 802.15.4 uses direct-sequence spread-spectrum, which is robust, since the spreading factor is high. This will reduce substantially the impact of other wireless systems to 802.15.4. In terms of interference to others, 802.15.4 will be perceived by other wireless systems as wideband noise.
- 802.15.4 devices will have very low duty cycles (less than 1 percent) and output power less than 0 dBm. Although government regulations in the United States and around the world allow output power up to 1 W in the 2400 MHz band, 802.15.4 devices will transmit at much lower power levels. This will eliminate the need for a filter to suppress the out-of-band emissions. The power levels will be between –3 dBm and 10 dBm, 0 dBm being typical.

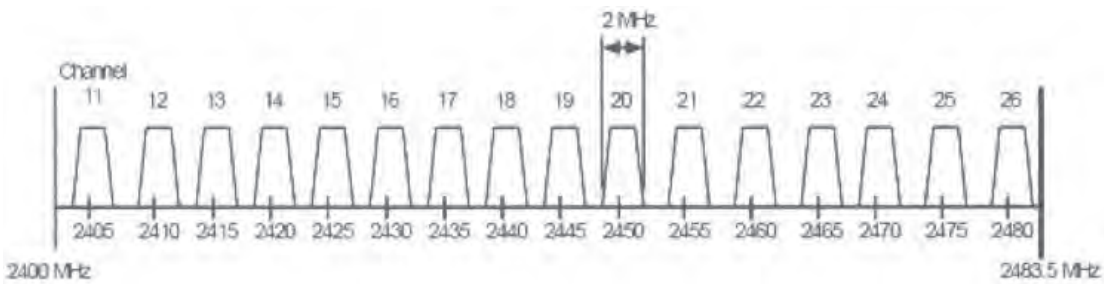

Figure 3–25: IEEE 802.15.4 channelization in the 2400 MHz band

The 1% Transmit Duty Cycle is the main limitation on operation within the 868.0–868.6 MHz band. However, it is the responsibility of the higher protocol layers to ensure that these parameters are satisfied.

In terms of impact to other wireless systems, the bandwidth of a IEEE 802.15.4 signal is wide and typically only a fraction of it will be received by other wireless devices. Because IEEE 802.15.1 uses frequency-hopping, the design of IEEE 802.15.4 ensures that interference will affect only three out of the 79 hopping frequencies. IEEE 802.11b will combat the interference from IEEE 802.15.4 by the processing gain of the IEEE 802.11b spreading process.

INTEL-1020
U.S. Patent 8,219,129

Dynamic channel selection is discussed next. In a star network, the central network controller initiates dynamic channel selection if the current QoS requirements cannot be satisfied due to adverse conditions of the channel. The controller collects channel status information from the network devices to arrive at this decision. The algorithm required to use the channel status information and decide change of channel is not specified. If the central network controller decides to change the channel, first it will look for other available channels. If it finds an available channel, it will begin sending beacons with the new channel information and the timeout for changing the channel. The devices that received a beacon with new channel information will change to the new channel within the indicated timeout duration and wait for beacons in the new channel.

For a cluster tree network, the cluster head gathers the information to determine if the channel quality is degrading. The cluster head keeps a list of available clear channels and makes the decision whether or not to initiate a channel change. As in a star network, the actual decision-making algorithm to determine whether or not a channel change is necessary is beyond the scope of this standard. If the cluster head decides to change the channel, the cluster head may transmit a message on the original channel that contains the new channel information to all members of the cluster. Because the original channel is impaired, new channel information is transmitted three times by each receiving device to increase the probability of reception. It is not expected that a receiving device will receive all three transmissions. After the three transmissions, devices change to the new channel.

> In both star and cluster tree topologies, there may be devices that were unable to receive the new channel information and that will become lost after the channel switch.

WPAN TECHNOLOGY AND BUSINESS TRENDS

The first WPAN technology, Bluetooth, was developed because some cell phone manufacturers wanted to increase the value of their main product by making it capable of communicating with devices other than phones. At present high-end cell phones on the market have Bluetooth functionality. Within a couple of years almost all phones will have Bluetooth. The original intention of the Bluetooth SIG was for Bluetooth to achieve a price of $5 and become a truly ubiquitous

INTEL-1020
U.S. Patent 8,219,129

technology. According to U.S. antitrust laws, price discussions within IEEE 802 cannot take place, and therefore, the 802.15 Working Group has never discussed prices. While the price target of $5 seems too aggressive, perhaps as a consequence of the fundamental principles of communications, an increasing number of products are expected to have Bluetooth. In many of these products Bluetooth will be embedded, which will be a feature differentiating them in the marketplace. Higher-priced devices can absorb the cost of Bluetooth, and as a result some say that the cost of Bluetooth is zero. It seems reasonable that Bluetooth will succeed as a cable replacement technology. Bluetooth PC cards and various computer peripheral devices are among the first to include Bluetooth, in addition to cell phones. Similar to 802.11, some laptops are manufactured with an integrated Bluetooth unit. PDAs such as Palm and Handspring increasingly come with wireless communication capability. It must be noted that some powerful PDAs come with 802.11. Outside of computing and communication devices, in the near future a number of other devices and equipment will have a number of Bluetooth units.

Recently 802.15.3 began work on a new physical layer 802.15.3a, providing data rates in excess of 100 Mb/s, and up to 480 Mb/s. This physical layer will be based on ultra-wide band (UWB) technology and will attempt to enable true "cable replacement". Ultra-wide band is one of the most over-hyped wireless technologies in the last few years and the industry interest in 802.15.3a is very large. The standardization work within 802.15.3a involves most of the major semiconductor companies such as Intel, Texas Instruments, and Motorola. Start-up companies focused on UWB technology such as Time Domain®, XtremeSpectrum®, and Staccato Communications, Inc.® also participate.

The concept of UWB originated about 100 years ago, due to Marconi. The modern UWB technology was developed in the 1980s and 1990s. The main motivation initially was military communication systems, since FCC rules did not allow UWB operation for commercial use outside of the military. Since UWB operates over a very wide bandwidth there was a concern about interference to licensed users "and other important radio operations". In February 2002, the FCC amended the Part 15 regulations, allowing UWB to operate over a band that is 7500 MHz wide (3100–10600 MHz) on an unlicensed basis with average radiated emission limit of –41.3 dBm/MHz. The peak emission limit is 60 dB above the average. Note that UWB is defined by the FCC as any signal having a fractional

INTEL-1020
U.S. Patent 8,219,129

bandwidth greater than 20 percent and occupying a band of 500 MHz or more. Fractional bandwidth is the ratio of baseband bandwidth to RF carrier frequency. The band is bounded by the points that are 10 dB below the highest radiated emission. Note in particular that the modulation used is not relevant to the definition of UWB. The rules allow any modulation scheme to be used, as long as the fractional bandwidth is greater than 20 percent and the occupied bandwidth is greater than 500 MHz. At present other regulatory bodies around the world have not sanctioned the operation of UWB devices. While favorable rulings are expected, the allowed frequency band of operation may be different. (See Table 3–5 below.)

Table 3–5: FCC regulations for the mean EIRP of UWB systems

Frequency band (MHz)	Mean EIRP (dBm/MHz) (indoor/handheld)
960–1610	–75.3/–75.3
1610–1900	–53.3/–63.3
1900–3100	–51.3/–61.3
3100–10600	–41.3/–41.3
>10600	–51.3/–61.3

Traditionally UWB operation has been achieved with pulse-position modulation (PPM), where very short pulses are used. These pulses must have zero-mean, because an antenna cannot radiate a signal at zero frequency. If the pulse width is shorter compared to the delay spread then good performance in multipath channels can be achieved. In UWB, the coherence bandwidth is smaller than the signal bandwidth, and the coherence time is larger than the symbol duration [B117], [B139], [B140].

Initially there were 31 proposals submitted to 802.15.3a from various companies. The number of proposals was quickly reduced to two main technologies, as a result of dropouts and mergers. The first main proposal was based on impulse radio technology [B105]. This proposal uses BPSK and QPSK modulation with DSSS with 24 chips per symbol, leading to high spread-spectrum processing gain. It uses the entire band allowed by the FCC with the exception of the UNII

INTEL-1020
U.S. Patent 8,219,129

frequencies in the 5GHz region. Although not mandated by the FCC, this is chosen to enable coexistence with WLAN devices. To satisfy the FCC regulations the baseband signal according to the proposal can be 1.368 GHz or 2.736 GHz. The second main proposal was a merger of several proposals. According to this proposal the baseband signal is 528 MHz. The entire 7.5 GHz-wide band is divided into four smaller bands, with frequency hopping between these smaller bands used to facilitate multiple access. The carriers are modulated using QPSK, followed by 128-point FFT. Like the first proposal, the second proposal also avoids using the 5 GHz UNII bands. The performance comparisons between these two main proposals follow many of the comparisons between single-carrier and multi-carrier technologies.

IEEE Std 802.15.3 and IEEE Std 802.15.4 have not yet made it to the mass-market stage and represent interesting opportunities. IEEE 802.15.3 together with 802.11 will be a competing technology to be used in consumer electronic products for the home, such as TV sets and digital cameras. Digital cameras are typically connected via a cable to a computer for image and video communication. This cable can be conveniently replaced by a wireless technology, and Bluetooth is not fast enough.

IEEE 802.15.4 seems a strong competitor in a new and fast-growing market—wireless sensors. Wireless sensors have many applications. 802.15.4 products can be implemented in a way that requires two AA batteries for one year of operation. Examples of embedded devices suitable for low-rate WPAN are PIR sensors, cordless switches, intelligent remote controls (network-device) and security, lighting and heating, ventilating, and air-conditioning (HVAC) controllers (network-coordinator). Similar sensors can be used on buildings, bridges, and other structures to monitor their seismic movements. In the near future wireless sensors will find a variety of applications in industrial control and monitoring, intelligent agriculture, supply chain management and asset tracking, health monitoring, security and military sensing. Some companies believe that low data-rate applications actually represent a bigger market opportunity than high data-rate wireless personal area networks. Some early technological leadership has been established by companies like Motorola, and several start-up companies including Staccato Communications of San Diego, CA, and S5, Inc., of Draper, UT. To promote this technology and ensure interoperability among products from different vendors, interested companies have formed an industry consortium

INTEL-1020
U.S. Patent 8,219,129

called ZigBee Alliance™ of San Ramon, CA [B63]. The relationship between 802.15.4 and ZigBee is similar to the relationship between 802.11 and the Wi-Fi Alliance of Mountain View, CA.

Ultra-wide band technology is not only suitable for high data-rate applications. It is also suitable for low data-rate, low power applications. 802.15.4 is currently exploring another physical layer to be called 802.15.4a, which will likely be based on UWB technology [B3].

The WPAN devices can be implemented in two types of devices —embedded (non-host enabled) and host enabled. The embedded node architecture consists of a user interface, the actual device application, a network layer, and an on-air protocol stack. The API connects the device application with the network layer, providing a consistent application interface to all nodes. The network layer definition, the device application, and its user interface are outside the scope of the standard and will typically be provided by the end-application developer. The on-air protocol stack consists of a LLC sublayer, a MAC sublayer, and a PHY layer. A network-device may not need any of the upper layers.

For host-enabled devices to interface between the WPAN device and the external host a host stack is necessary. The interface can be via any suitable interconnect protocols such as RS232 or USB. Examples of host-enabled nodes are USB hubs and Internet portals. This architecture consists of a local user interface, the actual device application, an on-air protocol stack and a host stack providing inter-connectivity with the host. The interface between the application and the on-air protocol stack is via the API. The on-air protocol stack allows an application to access a WPAN. The user interface is dependent on the function of the application and the amount of information that needs to be communicated or extracted from the consumer. In network-devices, the user interface can be as simple as a single button, which can be used as a stimulus for certain networking functions, such as joining or leaving a network. Host enabled network-coordinator may require a local user interface. This will be used when the external host is not connected since the network-coordinator must still be able to function in the absence of a host. The user interface is not specified in this standard.

It is worth examining again the question of whether WPANs and WLANs are competing or complementing technologies. It would seem that WPAN devices used as cable replacement avoid competition with WLAN devices because of low

INTEL-1020
U.S. Patent 8,219,129

power consumption and cost. However, WPANs can cover a much larger range if they transmit at the maximum allowed output power. Therefore, WPANs can be used for other applications, such as infrastructure computing and home networking. In these applications WPANs can potentially compete with WLAN technologies. However, this competition will likely be won by WLAN. 802.11 has established itself as the most successful wireless standard. Many consumers use 802.11 in their offices and when they bring their laptops home expect to be able to connect to other 802.11 devices. Therefore, 802.11 has an advantage for home networking applications as well. There are three 802.11 physical layers in the marketplace—802.11b, 802.11a, and 802.11g. The cost of 802.11b is falling rapidly, and is approaching the cost of WPAN devices. The cost of 802.11a and 802.11g is also expected to decrease in the next couple of years. In the near future consumer electronics devices such as TV sets, DVD players, MP3 players, musical instruments, etc., will have the capability for wireless communication. In the long run, other consumer devices and kitchen appliances will have wireless communication capability. Since on one hand the majority of these devices are mains-powered, while on the other hand the range is relatively small, it remains to be seen whether vendors will choose WPAN over WLAN.

INTEL-1020
U.S. Patent 8,219,129

Copyright © 2004 IEEE

Chapter 4 Air Interface for Fixed Broadband Wireless Access Systems

OVERVIEW

Over the last several years, the area of broadband data communications experienced very high growth. The term "broadband" is usually taken to mean the capability to deliver significant bandwidth to each user, much higher than narrowband voiceband modems. Following ITU terminology, the term "broadband" means transmission rates greater than 1.5 Mb/s. Broadband Internet access became a mass-market industry reaching most of the 100 million homes and 3 million businesses in the United States. Broadband access reaches an even greater percentage of homes and businesses in some other industrialized countries, such as Korea and Canada. For a relatively long time now, it has been known that the major technologies for broadband Internet access are DSL, cable, wireless, and optical fiber communications. Of these four technologies, two—DSL and cable—are mature technologies and have reached the mass-market stage. Optical communications provides very high data rates, but the cost of laying optical fibers is extremely high, in fact, prohibitively high in urban areas where most of the customers are. Because of cost issues, it is not clear whether optical communications will ever become a mass-market technology. However, it is expected that broadband wireless access (BWA) systems will become a technology that is generally competitive with DSL and cable, connecting businesses and users' premises with core networks.

Clearly, if BWA systems are to compete with DSL and cable, their market growth will be significant for several reasons. First, in the long run, BWA systems are capable of delivering significantly higher data rates than DSL or cable. DSL technology requires customers to be close to the telephone company's central office. DSL can deliver up to 6 Mb/s at distances up to 18,000 feet. At present, most DSL subscribers get much lower speeds. Customers that are more than 18,000 feet away from the telephone company's central office cannot get DSL service. Similarly, cable broadband service delivers 1.5 Mb/s, but requires a TV

INTEL-1020
U.S. Patent 8,219,129

cable. Though most of businesses and residential users in North America have access to either DSL or cable, some do not. Furthermore, many more business and residential users around the world cannot get DSL or cable. These users represent a significant potential market for BWA. In addition, compared with DSL and cable networks, the use of wireless techniques results in a number of benefits for users and for service providers. For service providers, these benefits include relatively low up-front costs. For consumers, the benefits include the convenience of a wireless connection.

What is the architecture of BWA networks? Fixed broadband wireless access systems typically include at least one base station (BS) and a number of subscriber stations (SS). In addition there may be links between base stations, repeaters, and possibly other equipment. Base stations provide connections to core networks on one side and radio connection to other stations on the other. A reference fixed BWA system diagram is provided in Figure 4–1.

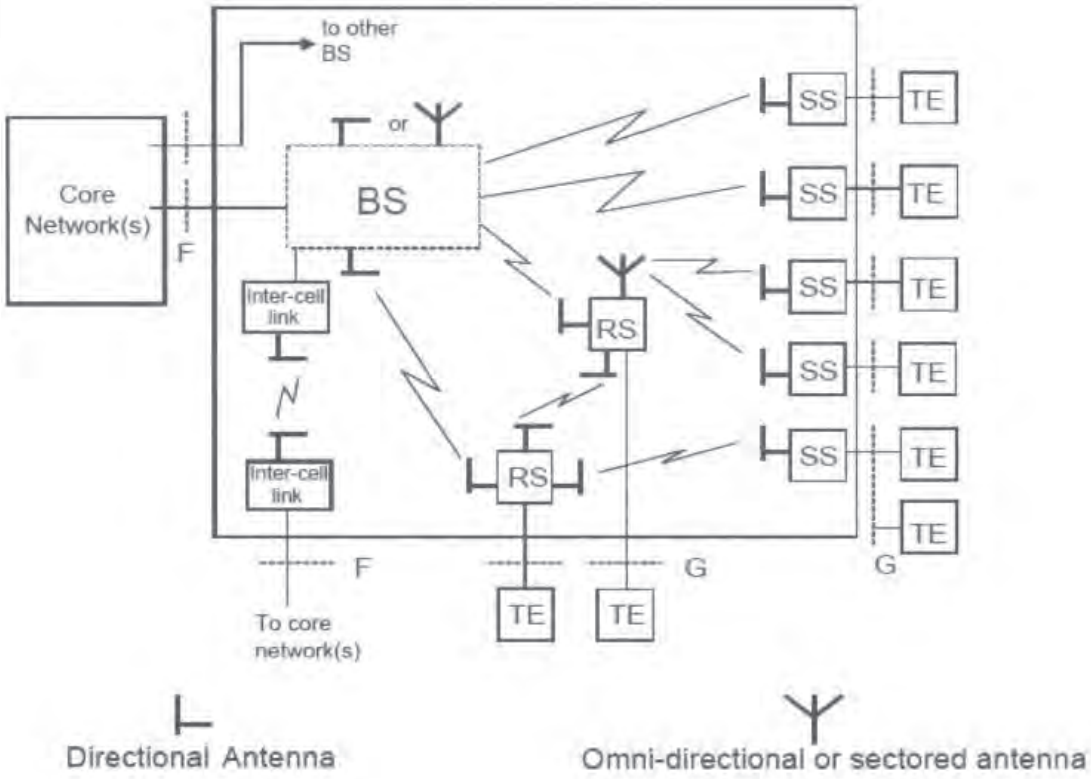

Figure 4–1: Fixed broadband wireless access (BWA)

INTEL-1020
U.S. Patent 8,219,129

Fixed BWA systems have multipoint architectures. The term multipoint includes point-to-multipoint (PMP) and multipoint-to-multipoint (MP-MP). PMP systems comprise base stations, subscriber stations and, in some cases, repeaters. On the uplink, the sending entity is the SS and the receiving entity is the BS. On the downlink, the sending entity is the BS and the receiving entity is the SS. Note that one difference compared with WLAN and WPAN is that antennas with a variety of radiation patterns are used. Base stations use relatively wide-beam antennas, divided into one or several sectors providing up to 360-degrees coverage with one or more antennas. The sectorized antenna of the base station is capable of handling multiple independent sectors simultaneously. Within a given frequency channel and antenna sector, all stations receive the same transmission. To achieve complete coverage of an area, more than one base station may be required. The connection between base stations is not part of the fixed BWA network itself. The connection is achieved by use of radio links, fiber optic cable, or equivalent means. Links between BSs may sometimes use part of the same frequency allocation as the fixed BWA itself.

Routing to the appropriate BS is a function of the core network. In general, subscriber stations use highly directional antennas that face the BS. Subscriber stations share use of the radio channel. This may be achieved by various access methods, including frequency division, time division, or code division. Some systems deploy repeaters. In a PMP system, repeaters are generally used to improve coverage to locations where the BSs have no line of sight within their normal coverage areas or, alternatively, to extend coverage of a particular BS beyond its normal transmission range. A repeater relays information from a BS to one or a group of SSs. A repeater may operate on the same downlink frequencies as those frequencies that it uses, facing the BS, or it may use different frequencies (i.e., demodulate and remodulate the traffic on different channels). MP-MP or mesh systems have the same functionality as PMP systems. In MP-MP systems, traffic may pass through one or more repeaters to reach a subscriber. Most stations are repeaters that also provide connections for local subscribers. Antennas are generally narrow-beam directional types with means for remote alignment.

These networks operate transparently, so users are not aware that services are delivered by radio. A BWA network provides connection to many user premises within a radio coverage area. It provides a pool of bandwidth, shared automatically among the users. Demand from different users is often statistically

INTEL-1020
U.S. Patent 8,219,129

of low correlation, allowing the network to deliver significant bandwidth-on-demand to many users with a high level of spectrum efficiency. Significant frequency reuse is employed. The range of applications is very wide and evolving quickly. It includes voice, data, and entertainment services of many kinds. Each subscriber may require a different mix of services; this mix is likely to change rapidly as connections are established and terminated. Traffic flow may be unidirectional, asymmetrical, or symmetrical, again changing with time. In some territories, systems delivering these services are referred to as Multimedia Wireless Systems (MWS) in order to reflect the convergence between traditional telecommunications services and entertainment services.

In response to these market needs, the IEEE 802.16 working group on Broadband Wireless Access was established in 1999. The industry interest was significant and the first standard (IEEE Std 802.16-2001) was approved in December 2001 and published in April 2002 [B82].

The IEEE 802.16 standard specifies the MAC and PHY of fixed broadband wireless access systems. Similar to IEEE 802.11, the MAC layer is structured to support multiple physical layer specifications to give the standard flexibility. IEEE 802.16 supports point-to-multipoint and multipoint-to-multipoint communications in the 2–66 GHz band. In addition to technological and business reasons, the existence of multiple physical layers reflects the fact that the electromagnetic propagation between 2 GHz and 66 GHz is not uniform. Each physical layer is appropriate to a particular frequency range and application. IEEE 802.16 is intended to operate in three different frequency bands.

The first group of frequency bands of interest is the licensed bands between 10 GHz and 66 GHz. In this frequency range, the wavelength is very short, and the attenuation of electromagnetic waves by various terrain and human-generated structures is severe. Therefore, line-of-sight between transmit and receive antennas is required. Multipath is not an issue, because only the main path from the transmitter to the receiver will have significant energy. While multipath is not a limiting factor, thermal noise and/or interference are the main limiting factors to the performance of wireless systems in this frequency band. These systems are used only in an outdoor setting, and rain will increase the attenuation experienced by the electromagnetic waves. Therefore, a sizable link margin must be reserved for rain loss. At the same time, the desire for high data rates means that high-order

INTEL-1020
U.S. Patent 8,219,129

modulation schemes must be used. In turn, these high-order modulation schemes require large signal-to-noise ratio (SNR) for satisfactory operation.

Another difference between IEEE 802.16 and IEEE 802.11 or IEEE 802.15 is that the range of interference typically exceeds the operating range. This is compounded by the fact that the rain cells producing the most severe rain losses are not uniformly distributed over the operational area. This creates the potential for scenarios in which the desired signal is severely attenuated but the interfering signal is not. While rain attenuation is a technical problem, operating in the 10–66 GHz band leads to two technical advantages. First, for the reasons outlined previously, multipath is negligible. Second, the frequency bands of operation in this physical environment are typically large, on the order of 25–28 MHz. This allows data rates in excess of 120 Mb/s.

The second group of frequency bands of interest is the licensed bands between 2 GHz and 11 GHz. In these bands, line of sight is not necessary. Multipath can be significant and appropriate measures must be taken. Because there may or may not be line of sight, the signal power can vary significantly. Retransmissions may be necessary because of the lossy behavior of the wireless medium. As a solution to these problems, the 802.16 standard provides advanced power management techniques, and the physical layer for operation between 2 GHz and 11 GHz can be based on OFDM.

The third group of frequency bands of interest is the unlicensed bands between 2 GHz and 11 GHz. While the physical characteristics of these unlicensed bands are the same as the licensed bands between 2 GHz and 11 GHz, there are two differences. Because these bands are unlicensed, first, there may be other users, which causes an interference problem. Second, regulations limit the output power. These problems require dynamic frequency selection and power management.

Figure 4–2 provides a block diagram of IEEE 802.16. The architecture follows general IEEE 802 guidelines. The MAC layer comprises three sublayers—convergence sublayer, common part sublayer, and security sublayer. The MAC includes a convergence sublayer to better handle the higher-layer protocols placed above the MAC. The central part of the MAC is the common part sublayer. Like in the other 802 standards, it handles channel access, connection establishment and maintenance, and QoS. The third sublayer is the security sublayer, providing authentication, secure key exchange, and encryption.

INTEL-1020
U.S. Patent 8,219,129

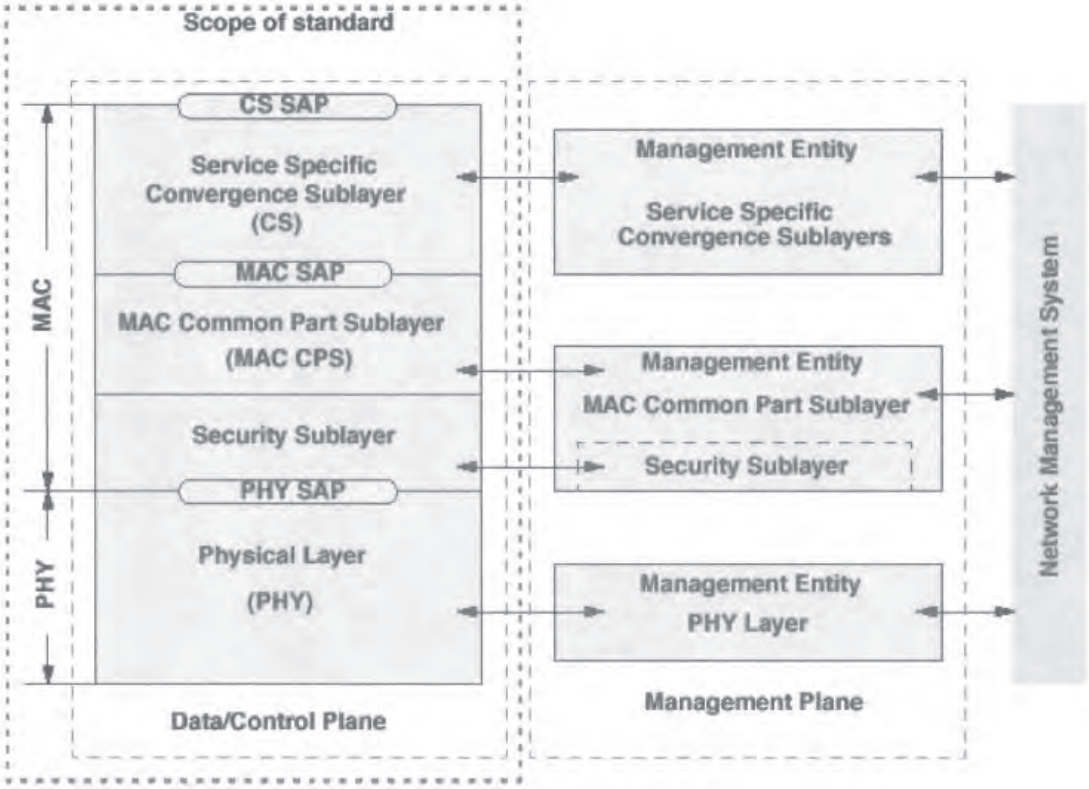

Figure 4–2: IEEE 802.16 Protocol structure

Although the MAC and PHY layers are independent, the integration between the MAC and the PHY is important for the overall system performance. For the MAC/PHY integration at a system level, two issues are important: the duplexing technique and whether the PHY is framed or nonframed. These choices should be made together with the PHY design. For a framed PHY, the MAC aligns its scheduling intervals with the underlying PHY framing. Better performance can be achieved with an unframed PHY, although it is not supported by 802.16. In this case, the MAC can choose the scheduling intervals to optimize system performance.

The upstream physical layer is based on the use of a combination of time division multiple access (TDMA) and demand assigned multiple access (DAMA). In particular, the upstream channel is divided into a number of "time slots." The number of slots assigned for various uses (registration, contention, guard, or user traffic) is controlled by the MAC layer in the base station and can vary in time for

INTEL-1020
U.S. Patent 8,219,129

optimal performance. The downstream channel can be either based upon continuous time division multiplexing (TDM), where the information for each subscriber station is multiplexed onto the same stream of data and is received by all subscriber stations located within the same sector or in an alternative method (defined for the burst mode of operation), which allows bursts to be transmitted to specific SSs in a similar fashion to the TDMA upstream bursts.

The next three sections of this chapter are devoted to the MAC layer: the convergence, common part, and security sublayers. The fourth section beginning on page 256 is devoted to the necessary MAC enhancements for when the physical layer operates in the 2 GHz–11 GHz frequency range. Then the physical layers of IEEE 802.16 are discussed in the fifth section on page 269. Coexistence is very important for wireless access systems and is considered in the sixth section on page 317. The chapter concludes with an overview of the technology and business trends for BWA systems.

MAC Convergence Sublayer

Above IEEE 802.16, the higher layers will predominantly be ATM [B5] or IEEE Std 802.3 (Ethernet) [B91]. Therefore, two convergence sublayer specifications are developed—ATM and packet. Because the convergence sublayer is a separate part of the MAC, a vendor that wants to use a higher layer different than ATM or Ethernet can develop a different convergence sublayer. The convergence sublayer is not a physical, but a logical interface. In transmit mode, the convergence sublayer receives data from a higher layer, and classifies the data. Based on the classification, it can perform additional processing before delivering the data to the MAC common part sublayer. This type of additional processing can be payload header suppression (PHS). In receive mode, the convergence sublayer accepts data from the MAC common part sublayer. If the peer convergence sublayer has modified the data frame, the receiver will restore the frame before passing it on to the higher layer.

The ATM convergence sublayer (CS) is designed to support the different ATM services. It accepts ATM cells, classifies these cells, and maps them to MAC frames. It may also perform Payload Header Suppression (PHS). An ATM connection can be either Virtual Path (VP) switched or Virtual Channel (VC) switched. The ATM convergence sublayer differentiates between these two types

INTEL-1020
U.S. Patent 8,219,129

of connections and assigns appropriate channel ID (CID) before delivering the data to the MAC common part sublayer. In this way, the ATM convergence sublayer guarantees the correct handling of the traffic by the MAC sublayer. The ATM header is thrown away, with the exception of the field indicating priority. Payload header suppression is the process of suppressing the repetitive portion of payload headers at the sender and restoring the headers at the receiver. It can be viewed as a compression method used to save bandwidth. A PHS Rule provides details on how a repetitive portion of the payload headers can be replaced with an appropriate index for transmission and subsequently regenerated at the receiving end. When payload header suppression is performed, a suppression mask can be used to allow select bytes not to be suppressed. This is used for sending bytes that change (such as IP sequence numbers) and still suppressing bytes that do not change. Payload header suppression is intended for unicast communications and is not defined in the multicast case. In addition to payload header suppression, multiple ATM cells, which share the same connection ID, may be packed together to further save bandwidth.

ATM interfaces are also characterized by the signaling mechanism that they use [B6]. The convergence sublayer does not support associated signaling, but supports common channel signaling (CCS). According to the common channel signaling mechanism for ATM networks [B6], signaling messages are carried over a connection completely independent of user connections. In addition, one ATM signaling channel can carry signaling messages for a number of user connections. Every IEEE 802.16 station must have a channel ID, corresponding to CCS messages. The questions regarding the mapping of ATM signaling messages are left to technology vendors to implement. This is one of the many issues that IEEE 802.16 leaves to vendors. In particular, designers of the base station MAC must map ATM signaling messages to corresponding MAC messages.

The second type of convergence sublayer is the packet convergence sublayer. It is used for the transport for all packet-based protocols such as Internet protocol (IP), point-to-point protocol (PPP) [B14], and IEEE 802.3/Ethernet. It performs similar function as the ATM convergence sublayer, including suppression of payload header information and rebuilding of any suppressed payload header information.

Both the ATM and the packet convergence sublayers perform classification of the higher-lay data. This classification is the process by which a MAC data is

INTEL-1020
U.S. Patent 8,219,129

associated with a particular connection. This process facilitates the delivery of MAC data with the appropriate QoS constraints. In general, a classifier is a set of matching criteria applied to each packet entering the IEEE 802.16 network. There are multiple classifiers, applied according to a certain priority. If more than one classifier matches the packet, the highest priority classifier is chosen. The destination MAC address, the source MAC address, and other parameters can be used for the classification. When Internet protocol is used, the IP headers can be included in the classification parameters. Each classifier also defines whether or not payload header suppression is performed. The output of the classification procedure is the connection identifier, or connection ID. The concept of a connection, or service flow, is central in IEEE 802.16. The connection ID can be considered a connection identifier even for nominally connectionless traffic like IP, because it serves as a pointer to destination and context information. This connection identifier is used to distinguish between multiple uplink channels, all of which are associated with the same downlink channel. In IEEE 802.16, the connection ID is 16 bits long. The use of a 16-bit connection ID permits a total of 64K connections within each downlink and uplink channel. It is possible that a packet does not match any of the classifiers. In this case, the output of the classification procedure is not standardized and is left at the discretion of the vendors. The packet can be associated with a default connection ID or discarded.

The operation of the convergence sublayer can be summarized as follows. After the convergence sublayer of the sending entity receives a packet from the higher layer, it applies its list of classifier rules to determine the service flow, connection ID, and whether to turn on or off payload header suppression. If suppression is enabled, the sending entity suppresses all the bytes except the bytes masked by the mask. The convergence sublayer then prefixes the frame with the payload header suppression index and sends the block of data to the MAC SAP for further processing. When the packet is received, the receiving entity will determine the associated connection ID by examination of the generic MAC header. The receiving convergence sublayer then reassembles the packet.

MAC Common Part Sublayer

The common part sublayer is the central part of the IEEE 802.16 MAC. It defines the multiple-access mechanism.

INTEL-1020
U.S. Patent 8,219,129

> Messages can be individually addressed, multicast, or broadcast. Multicast messages are addressed to a group of subscriber stations, and broadcast messages are addressed to all stations.

In the downlink, the base station is the only transmitter that is operating. Therefore, it does not have to coordinate its transmissions with other stations. In the downlink, the base station broadcasts to all stations. Stations check the address in the received messages and retain only those addressed to them. In the uplink direction, the user stations share the channel. To determine which subscriber station has the right to transmit, three basic principles are used—unsolicited bandwidth grants, polling, and contention procedures. The method for sharing the channel is dynamic and functions on an on-demand basis. Sharing is on demand to give the service providers opportunities to differentiate themselves in the marketplace. Depending on the class of service used, the SS may be issued continuing rights to transmit, or the right to transmit may be granted by the base station after receipt of a request from the user.

The duplexing techniques supported by IEEE 802.16 will be considered first. Duplexing here means how the downlink is separated from the uplink. Compared with the other wireless standards, IEEE 802.16 follows a much more centralized architecture. This is appropriate for a technology that does not have to support ad-hoc networks. There are two ways to separate downlink from uplink—in the frequency domain and in the time domain.

In frequency division duplexing (FDD), downlink and uplink transmissions are performed at different frequencies and therefore may overlap in time. Although theoretically the width of the frequency bands devoted to downlink and uplink transmissions can be adjusted, this is very difficult to accomplish in practice. As a result, in FDD the asymmetry between uplink and downlink is static. A fixed duration frame is used for both uplink and downlink transmissions. This facilitates the use of different modulation types. It also allows simultaneous use of both full duplex subscriber stations, which can transmit and receive simultaneously, and half duplex subscriber stations, which cannot. A half duplex subscriber station can listen to the downlink channel only when it is not transmitting in the uplink channel. If half duplex subscriber stations are used, the bandwidth controller does not allocate uplink bandwidth for a half duplex subscriber station at the same time that it is expected to receive data on the downlink channel. The fact that the uplink

INTEL-1020
U.S. Patent 8,219,129

and downlink channels use a fixed duration frame simplifies the bandwidth allocation algorithms. Figure 4–3 illustrates the bandwidth allocation in FDD.

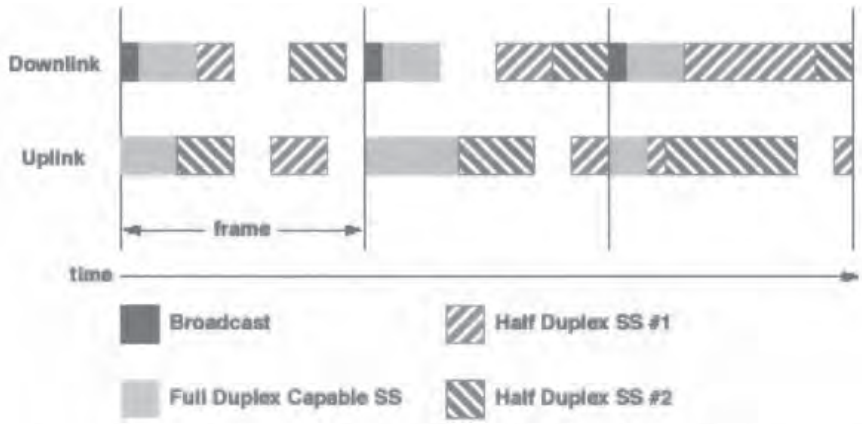

Figure 4–3: Example of burst FDD bandwidth allocation

Time-division duplexing (TDD) divides time into uplink and downlink transmission periods, and therefore downlink and uplink may be performed at the same frequency band. In the case of TDD, a frame also has a fixed duration and contains one downlink and one uplink subframe. The frame is divided into an integer number of physical slots, which help to partition the bandwidth easily. Because in TDD, the balance between the amount of time devoted to uplink and downlink can be dynamically adjusted, the asymmetry between downlink and uplink is dynamic. The split between uplink and downlink is a system parameter and is controlled at higher layers within the system. Figure 4–4 illustrates the TDD frame structure.

A MAC data frame in IEEE 802.16 consists of three parts—header, payload, and cyclic redundancy code (CRC). The payload and the CRC are optional. The payload information may vary in length. Two MAC header formats are defined—generic and bandwidth request. The generic header is used if the payload contains MAC management messages, or data from the higher layers that has passed through the convergence sublayer. Bandwidth request frames do not have payload and are identified by the bandwidth request header. A particular connection ID may require a CRC to be attached. If present, the CRC covers the generic MAC

INTEL-1020
U.S. Patent 8,219,129

header and the MAC payload. The CRC is calculated after encryption; i.e., the CRC protects the generic header and the ciphered payload.

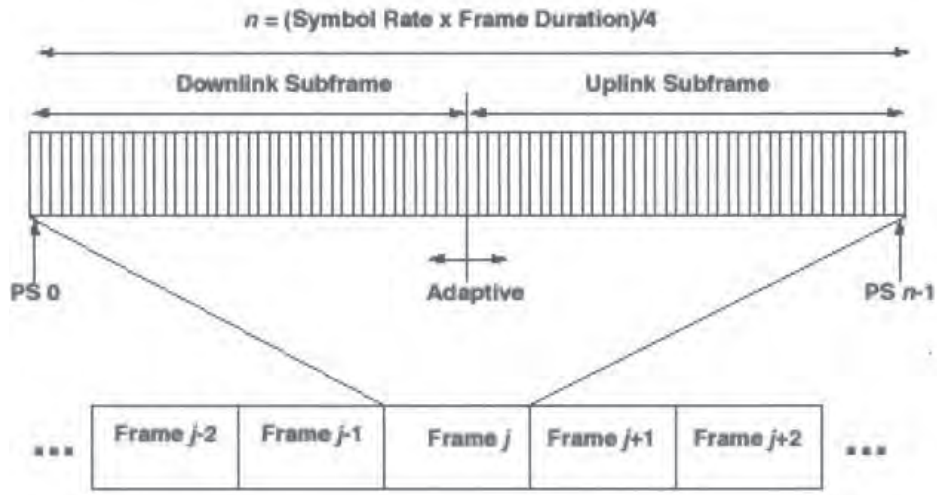

Figure 4–4: TDD frame structure

Multiple MAC frames may simply be concatenated into a single transmission in either the uplink or downlink directions. MAC management messages, user data, and bandwidth request frames may be concatenated into the same transmission.

Fragmentation and packing are dual processes. Fragmentation is the process by which one MAC frame is divided into more frames. This process allows efficient use of available bandwidth relative to the QoS requirements. Every frame contains a sequence number, which allows the receiver to reconstruct the original frames and to detect the loss of any intermediate packets.

> Frames can be concatenated, fragmented, or packed for efficient data transfer.

Packing is the opposite of fragmentation. If packing is turned on for a connection, the MAC may pack multiple frames into a single MAC frame. If ARQ is not being used, packing fixed-length frames is straightforward, because the receiving side can unpack simply by checking the length field in the MAC header and the entire length of the data. However, when packing variable-length MAC frames, where one MAC frame ends and another begins must be indicated. To accomplish this, the MAC attaches a packing subheader to each MAC frame.

INTEL-1020
U.S. Patent 8,219,129

The type field in the MAC header indicates the presence of packing subheaders. Packing frames on ARQ-enabled connections is very similar. The packing subheaders will be used by the ARQ protocol to identify and retransmit lost fragments.

Note that packing and concatenation are very similar.

The ultimate goal of fragmentation and packing is the efficient use of the air link. Fragmentation and packing also allow the MAC to tunnel various higher-layer traffic types without knowledge of the formats or bit patterns of those messages.

MAC management messages are used to describe the downlink and the uplink, and to handle ranging, registration, and privacy.

To describe the downlink and uplink, the BS transmits downlink channel descriptor (DCD) and uplink channel descriptor (UCD) messages at periodic intervals to define the physical layer characteristics of the downlink and uplink channels, correspondingly. The DCD and UCD messages contain the burst profile, which has information about the modulation type, forward error-correction type, preamble length, etc. The downlink map (DL-MAP) and the uplink map (UL-MAP) messages define the burst start times on the downlink and the uplink, respectively.

Another management function is ranging. Subscriber stations transmit ranging requests at initialization and periodically at the request of the base station to determine power and burst profile changes.

To efficiently support these management functions, at SS initialization, three management connections in each direction (uplink and downlink) are established between the SSs and the BS. These connection identifiers reflect the fact that there are inherently thee different QoS of management traffic between SSs and the BS. The basic connection is established during initial subscriber station registration and is used to transport MAC management messages with short delay. The primary management connection is for the exchange of longer, more delay-tolerant MAC management messages. Finally, the secondary management connection is used for delay tolerant, standards-based (DHCP, TFTP, SNMP, etc.) management messages.

INTEL-1020
U.S. Patent 8,219,129

Network entry and initialization

The procedure for initialization of a subscriber station can be divided into the following phases:

1) *Scan for downlink channel and establish synchronization with the BS.* Typically, subscriber station devices will have nonvolatile storage in which the last operational parameters are stored. In this case, these devices will first try to reacquire the downlink channel used previously. If this fails, the SS will begin to continuously scan the possible channels of the downlink frequency band of operation until it finds a valid downlink signal.
2) *Obtain downlink and uplink parameters.* The BS generates uplink channel descriptor (UCD) and downlink channel descriptor (DCD) messages on the downlink at periodic intervals. The BS generates UL-MAP and DL-MAP at intervals as specified in a particular PHY specification. These messages are addressed to all subscriber stations. First, to obtain downlink parameters, the MAC sublayer will search for the DL-MAP MAC management messages. The SS achieves MAC synchronization once it has received at least one DL-MAP message. An SS MAC remains in synchronization as long as it continues to successfully receive the DL-MAP and DCD messages. To obtain uplink parameters, the SS will wait for a UCD message from the BS in order to retrieve a set of transmission parameters for a possible uplink channel. These messages are transmitted periodically from the BS for all available uplink channels and are addressed to the MAC broadcast address. The SS can determine whether it may use the uplink channel from the channel description parameters. The SS collects all UCDs, which are different in their channel ID field, to build a set of usable channel IDs. The subscriber station will determine from the channel description parameters whether it may use the uplink channel. If the channel is not suitable, then the SS will try the next channel ID until it finds a usable channel. If the channel is suitable, the SS will extract the parameters for this uplink from the UCD. It will then wait for the next DL-MAP message and extract the time synchronization from this message. Before it can begin transmitting, the SS will wait for a bandwidth allocation map for the selected channel. Transmissions from the SS in the uplink will be in accordance with its MAC operation and the bandwidth allocation mechanism. The SS MAC is considered to have valid uplink parameters as long as it continues to successfully receive the UL-MAP and

INTEL-1020
U.S. Patent 8,219,129

UCD messages. If at least one of these messages is not received within certain time intervals, the SS will not use the uplink.

3) *Perform ranging.* Ranging is the process of acquiring the correct timing offset to align the SS's transmissions to a symbol marking the beginning of a mini-slot boundary. After the SS has received the UL-MAP message, it will scan the message to find an initial maintenance interval. The initial maintenance interval specified by the BS is large enough to account for the variation in delays between any two SSs. The maximum delay that one SS can experience is the maximum round-trip propagation delay due to cell radius plus the maximum allowable implementation delay. As a result, there could be significant variation in delay among subscriber stations. The ranging request message from the SS will be sent at the first initial maintenance transmission opportunity. The connection ID is set to zero in this case, because the SS does not have any connection IDs assigned to it. The first ranging request will be sent at minimum power level. This is an important step to minimize potential interference when a connection is being set up. If the first ranging request attempt is not successful, then at the next initial maintenance transmission opportunity, the SS will resend it at a power level, which is one step higher. This process will continue until the transmission is successfully received. Once the BS successfully receives the ranging request message, it will return a ranging response message addressed to the individual SS. The ranging response message contains the following information—basic and primary management CIDs assigned to this SS, RF power level adjustment, offset frequency adjustment, and timing offset. After receiving the ranging response from the BS, the SS will not transmit before it adjusts its RF signal in accordance with the ranging response. The SS must perform initial ranging at least once. Because operation in several channels must be supported, the subscriber station will attempt initial ranging on every suitable uplink channel before moving to the next available channel.

There are two types of ranging: initial ranging and periodic ranging. (Note that CDMA-based ranging is different.) After the initial ranging request, periodic ranging provides opportunities for other ranging requests and responses to adjust power levels, time, and frequency offsets. After the initial ranging request, further ranging requests by the SS can be transmitted using the basic CID. These ranging requests contain the power levels, time and

INTEL-1020
U.S. Patent 8,219,129

frequency offsets. The BS will return another ranging response message to the SS with any additional fine-tuning required. These ranging request/response steps are repeated until the response contains a Ranging Successful notification, or the BS aborts ranging. Once successfully ranged (ranging request is within tolerance of the BS), the SS can transmit data in the uplink.

The BS provides each SS a periodic ranging opportunity. The period is BS dependent. When bandwidth is granted per subscriber station, any allocation of uplink bandwidth constitutes a ranging opportunity. The goal of this periodic ranging is to adjust the quality of the wireless link. The SS monitors the carrier to noise plus interference ratio $C/(N + I)$ and compares the average value against the allowed range of operation. If the received $C/(N + I)$ goes outside of the allowed operating region, the SS requests a change to a new burst profile. The SS applies an algorithm to determine its optimal burst profile. The channel descriptors (UCD and DCD) are updated regularly by the BS.

4) *Immediately after ranging is completed, the SS informs the BS of its basic capabilities.* The BS responds with a message containing the intersection of the SS's and the BS's capabilities.

5) *The BS and SS perform authorization and key exchange, as specified in the MAC security sublayer.* Each SS contains the following unique information when shipped from the manufacturer: a 48-bit IEEE-style MAC address and a X.509 digital certificate. These are assigned during the manufacturing process. The MAC address is used during the authentication and registration processes to establish the appropriate connections for SSs. The security information contained in the X.509 digital certificate is used to authenticate the SS to the security server and authenticate the responses from the security and provisioning servers.

6) *After authorization and key exchange, registration occurs.* The main goal of the registration process is for the SS to receive its secondary management connection ID and, therefore, to become manageable. In this process, the SS sends a registration request to the BS. The BS responds with a registration response message, which includes the secondary management connection ID. The registration response specifically authorizes the SS to forward traffic to the network.

INTEL-1020
U.S. Patent 8,219,129

7) *At this point, IP connectivity can be established on the SS's secondary management connection.* The SS uses standard DHCP mechanisms [B26] to obtain an IP address and other configuration parameters.
8) *The subscriber station and base station need to have the current date and time.* This is necessary for time-stamping certain events.
9) *The SS and the BS exchange operational parameters.*
10) *If there are any preprovisioned service flows in the service contract of this SS, the BS sets up connections for these preprovisioned service flows.*

Channel access and QoS

The principal mechanism for providing QoS is to associate packets with a service flow. The service flow, the central concept of the MAC protocol, is a unidirectional flow of packets with a particular QoS. The SS and BS provide this QoS according to the QoS parameter set defined for the service flow. Service flows can be preconfigured or dynamically established. The QoS parameters of a service flow can be specified implicitly by specifying a service class name or explicitly. If a service class name is given, modifying parameters can also be included. Service flows exist in both the uplink and downlink direction and may exist without actually being activated to carry traffic.

When the higher-layer data is forwarded to the MAC, the convergence sublayer includes the connection ID identifying the connection across which the information is to be transmitted. The service flow for the connection is mapped to the Connection Identifier (CID). The service class is an optional object that may be implemented at the BS. A service class is defined in the BS to have a particular QoS parameter set. The QoS parameter sets of a service flow may contain a reference to the service class name. The service flow QoS parameter sets may augment and even override the QoS parameter settings of the service class, subject to authorization by the BS.

The service class serves two purposes. First, it allows operators to move the burden of configuring service flows from the provisioning server to the BS. Second, it allows higher-layer protocols to create a service flow by its service class name. Service classes are merely identifiers for a specific set of QoS parameter set values. A service identified by a service class is treated no

INTEL-1020
U.S. Patent 8,219,129

differently, once established, than a service that has the same QoS parameter set explicitly specified.

Service flows are distinguished by a 32-bit Service Flow Identifier (SFID); active service flows also have a 16-bit Connection Identifier (CID). A service flow is characterized by a set of QoS parameters such as bandwidth, latency, jitter, and throughput assurances.

To provide vendors with maximum flexibility, the way the QoS parameter set is provisioned is outside the scope of 802.16. Typically, the network management system will set up the QoS parameters. Following the centralized architecture, every BS must approve or deny every change to the QoS parameters. Within the BS, the authorization module approves or denies every change to the QoS parameters. This module performs the function of a QoS manager.

There is static and dynamic authorization. In the static authorization model, all possible services are defined in the initial configuration of each SS. The QoS parameters of provisioned service flows cannot be changed in the static model. Furthermore, the admitted QoS parameter set must be a subset of the provisioned QoS parameter set and the active QoS parameter set must be a subset of the admitted QoS parameter set.

In the dynamic authorization model, the authorization module also communicates through a separate interface to an independent policy server. This policy server may provide the authorization module with advance notice of upcoming admission and activation requests. The policy server also will specify the proper authorization action to be taken on those requests. Requests for admission and activation of service flows are checked by the authorization module to ensure that the QoS parameter set is a subset of the QoS set provided by the policy server. Before initial connection setup, the BS retrieves the provisioned QoS set for an SS. This is handed to the authorization module within the BS. The BS can implement mechanisms for overriding this automated approval process. Thus the dynamic model gives vendors the capability to implement any QoS policy desired. When modifying QoS parameters, care must be taken to prevent packet loss. In particular, any required bandwidth change is sequenced between the SS and BS. This sequencing means that if an uplink service flow's bandwidth is being reduced, the SS reduces its payload bandwidth first and then the BS reduces the bandwidth scheduled for the service flow. If an uplink service flow's bandwidth is

INTEL-1020
U.S. Patent 8,219,129

being increased, the BS increases the bandwidth scheduled for the service flow first and then the SS increases its payload bandwidth.

It is useful to think about three basic types of service flows:

- *Provisioned.* A service flow may be provisioned but not immediately activated. Such a service flow can be referred to as "deferred." The BS assigns a Service Flow ID for such a service flow but does not reserve resources. As a result of external action typically at higher layers, the SS or the BS may choose to activate a provisioned service flow. Service flows will be provisioned by the network management system.
- *Admitted.* The protocol supports a two-phase activation model. Admitted Service Flows that have resources assigned but not yet completely activated are in a transient state. The goal is to use resources efficiently, to perform admission control quickly, and to prevent potential theft-of-service situations.
- *Active.* A service flow is an active service flow if it is requesting and being granted bandwidth for transporting data packets. This completes the two-stage activation model.

Service flows may be created, changed, or deleted. The provisioning of service flows is done by means outside of the scope of 802.16, such as the network management system. The creation of dynamic service flows may be initiated by the subscriber station or by the base station. This is accomplished through a series of MAC management messages. In addition, protocols are defined for modifying and deleting service flows. When a service flow is deleted, all resources associated with it are released. However, if a basic, primary management or secondary management service flow of an SS is deleted, that SS is deregistered and will have to reregister. Also, if a service flow that was provisioned during registration is deleted, the provisioning information for that service flow is lost until the SS reregisters. However, deleting a provisioned service flow shall not cause an SS to reregister. Therefore, care should be taken before deleting such service flows.

The MAC is connection-oriented. All data communications are in the context of a connection. Service flows can be provisioned when a subscriber station is installed. After subscriber station registration, connections are associated with these service flows, one connection per service flow. As customer's service needs

INTEL-1020
U.S. Patent 8,219,129

change, new connections can be established or existing connections can be eliminated. The service flow defines the QoS parameters for the data that is exchanged on the connection. There are three connection management functions—creation, modification or maintenance, and deletion. The maintenance requirements vary depending upon the type of service connected. For example, connections that have a constant bandwidth requirement do not require maintenance. Connections that have dynamic (but relatively slowly changing) bandwidth requirements require some maintenance. IP services may require a substantial amount of ongoing maintenance due to their bursty nature and due to the high possibility of fragmentation. Existing connections may be changed in their characteristics on a dynamic basis in response to changing bandwidth requirements.

Uplink and downlink QoS is provided for each service flow.

The first parameter of QoS is bandwidth. A SS requests uplink bandwidth on a per-connection basis. However, the BS grants bandwidth on an aggregate basis for the entire SS. Requests for transmission are not associated with a subscriber station, but with a connection ID. Requests for transmission are always on a perconnection basis, because different connections can consume different amounts of bandwidth, even if they have the same service type. A subscriber station can have several connections. For example, a SS unit serving multiple tenants in an office building would make requests on behalf of all of them, though the service and other connection parameters may be different for each of them. Many higher-layer sessions may operate over the same wireless connection identifier. For example, many users within a company may be communicating with TCP/IP to different destinations, but bandwidth will be granted on an aggregate basis.

The main characteristic of the 802.16 standard is flexibility. This flexibility is extremely important for service providers, because it maximizes the revenue-generating potential of the technology. The multiple-access method enables the service to be tailored to the delay and bandwidth requirements of each user application. This flexibility is accomplished through five different types of uplink scheduling mechanisms. These are implemented using unsolicited bandwidth grants, polling, and contention procedures.

INTEL-1020
U.S. Patent 8,219,129

Mechanisms are defined in the protocol to allow vendors to optimize system performance using different combinations of these bandwidth allocation techniques. For example, contention may be used to avoid the individual polling of subscriber stations that have been inactive for a long period of time. The use of polling simplifies the access operation and guarantees that applications receive service on a deterministic basis if it is required. Scheduling services are designed to improve the efficiency of the poll/grant process. By specifying a scheduling service and its associated QoS parameters, the BS can anticipate the throughput and latency needs of the uplink traffic and provide polls and/or grants at the appropriate times. The basic services are Unsolicited Grant Service (UGS), real-time polling service (rtPS), non-real-time polling service (nrtPS) and Best Effort (BE) service. Each service is tailored to a specific type of data flow.

The bandwidth allocation and request mechanism are considered first. Increasing (or decreasing) bandwidth requirements is necessary for all services. Resources are assigned on a demand assignment basis as the need arises.

The Unsolicited Grant Service (UGS) is designed to support real-time service flows that generate fixed-size data packets on a periodic basis, such as T1/E1 and Voice over IP. The service offers fixed-size grants on a real-time periodic basis, which eliminates the overhead and latency of SS requests and assures that grants are available to meet the flow's real-time needs.

There are incompressible and compressible constant bit rate UGS connections. The needs of incompressible UGS connections do not change between connection establishment and termination. The requirements of compressible UGS connections may increase or decrease depending on traffic. In order for this service to work correctly, the subscriber station is prohibited from using any contention request opportunities and the BS is prohibited from providing any unicast request opportunities for that connection.

The real-time polling service (rtPS) is designed to support real-time data with variable-size packets on a periodic basis, such as MPEG video. The service offers real-time, periodic, unicast request opportunities, which meet the flow's real-time needs and allow the SS to specify the size of the desired grant. This service requires more request overhead than UGS but, because the grants are of variable size, the efficiency is higher. In order for this service to work correctly, the subscriber station is prohibited from using any contention request opportunities

INTEL-1020
U.S. Patent 8,219,129

for that connection. Also, the BS must provide opportunities for periodic unicast requests even if a grant is pending. Consequently, the SS uses only unicast request opportunities to obtain uplink transmission opportunities. (The SS could still use unsolicited data grant burst types for uplink transmission as well.)

The non-real-time polling service (nrtPS) is designed to support non real-time data that require variable size data grant burst types on a regular basis, such as high bandwidth FTP. The service offers unicast polls on a regular basis, which assures that the flow receives request opportunities even during network congestion. The BS typically polls nrtPS connection identifiers on an interval on the order of one second or less. In order for this service to work correctly, the SS is allowed to use contention request opportunities. Therefore, the SS uses contention request opportunities, as well as unicast request opportunities and unsolicited data grant burst types.

The intent of the best effort (BE) service is to provide efficient service to best effort traffic. In order for this service to work correctly, the SS must be set to allow use of contention request opportunities. This results in the SS using contention request opportunities as well as unicast request opportunities and unsolicited data grant burst types.

Requests refer to the mechanism that SSs use to indicate to the BS that they need uplink bandwidth allocation. Note that, in the request intervals, collisions can happen. The BS controls assignments on the uplink channel and determines which minislots are request intervals and therefore subject to collisions. The mandatory method of contention resolution is based on truncated binary exponential backoff. The initial backoff window and the maximum backoff window are controlled by the BS. When an SS has information to send and wants to enter the contention resolution process, it randomly selects a number within its backoff window. According to the CSMA protocol, the SS defers transmission for a contention transmission opportunity equal to the selected backoff counter.

> The backoff counter is decremented for only the contention transmission opportunities for which this SS is eligible.

This contention is only for the transmission of a bandwidth request. If the contention transmission is successful, the SS will receive a data grant. The subscriber station considers the contention transmission lost if no data grant has been given within a certain time interval. In this case, the SS

INTEL-1020
U.S. Patent 8,219,129

will increase its backoff window by a factor of two, as long as it is less than the maximum backoff window. Then the SS will randomly select a number within its new backoff window and repeat the deferring process. This retry process continues until the maximum number of retries has been reached. The maximum number of retries is independent of the initial and maximum backoff windows.

Note that the contention resolution process can be abandoned. If the SS receives a data grant at any time while deferring, it will abandon the contention resolution process and use the explicit transmission opportunity. The centralized architecture of IEEE 802.16 manifests itself in giving the BS flexibility in controlling the contention resolution by dynamically adjusting the values of Backoff Start and Backoff End. At one extreme, the BS may choose to set up the backoff window to emulate an IEEE 802.11-style backoff. This will ensure fairness and efficiency. At the other end, the BS may make the Backoff Start and Backoff End identical so that all SS are using the same backoff window.

A request may come as a standalone bandwidth request, or it may come as a piggyback request. Note that because the uplink burst profile can change dynamically, the SS cannot estimate the amount of time that is required. This amount of time will depend on the modulation format that is used. Therefore, all requests for bandwidth are in terms of the number of bytes needed to carry the MAC header and payload, but not the PHY overhead. There are two types of distinguishable bandwidth requests—incremental or aggregate. When the BS receives an incremental bandwidth request, it will add the quantity of bandwidth requested to its current perception of the bandwidth needs of the connection. When the BS receives an aggregate bandwidth request, it will replace its perception of the bandwidth needs of the connection with the quantity of bandwidth requested. Piggybacked bandwidth requests are always incremental. The self-correcting nature of the request/grant protocol requires that SSs periodically use aggregate bandwidth requests. The period may be a function of the QoS of a service and of the link quality. Although bandwidth is requested for the individual connections, it is granted for the entire SS. This allows more intelligent SSs to implement dynamic bandwidth allocation mechanisms of their own. This may be useful for real-time applications that require a faster response from the system. Subscriber stations can use bandwidth stealing. A station may use bandwidth that has already been granted to send another request. This is called "bandwidth stealing."

INTEL-1020
U.S. Patent 8,219,129

Polling is the process by which the BS allocates to the SS's bandwidth specifically for the purpose of making bandwidth requests. Note that polling is done on either an SS or connection basis. Furthermore, individual subscriber stations or groups of subscriber stations can be polled. Unicast polling is when a subscriber station is polled individually. In this case, no explicit message is transmitted to poll the subscriber station. Rather, the SS is allocated sufficient bandwidth in the uplink map to respond with a bandwidth request.

If insufficient bandwidth is available to individually poll many inactive SSs, some SSs may be polled in multicast groups or a broadcast poll may be issued. Certain connections are reserved for multicast groups and for broadcast messages. As with individual polling, the poll is not an explicit message but bandwidth allocated in the uplink map. The difference is that, rather than associating allocated bandwidth with an SS's basic connection, the allocation is to a multicast or broadcast connections.

SSs with currently active UGS connections may set the poll-me bit in the grant management subheader to indicate to the BS that they need to be polled to request bandwidth for non-UGS connections; otherwise SSs are not polled during active UGS connections.

MAC SECURITY SUBLAYER

The security sublayer has two component protocols. The first is the encapsulation protocol for data encryption. This encapsulation protocol defines (1) a set of supported cryptographic suites, i.e., pairings of data encryption and authentication algorithms, and (2) the rules for applying those algorithms to a MAC payload. Authentication is based only on digital certificates. The second component protocol of the security sublayer is called privacy key management (PKM). The PKM protocol describes how the BS distributes keys to client SS in a secure fashion. Through this key management protocol, the SSs and the BS synchronize keys.

> The security, or privacy, sublayer is designed with two goals in mind. First, it provides subscribers with privacy across the fixed broadband wireless network. Second, it provides operators with strong protection from theft of service.

INTEL-1020
U.S. Patent 8,219,129

The set of security information that a base station and a subscriber station share is called security association. This security information includes traffic encryption keys and cipher block chaining (CBC) initialization vectors. Security associations are uniquely identified by a security association identifier (SAID). For connections that are mapped to a security association encryption, and data authentication are performed as specified by that security association. Three types of security associations are defined—primary, static, and dynamic.

Each SS establishes an exclusive primary security association with its BS during the SS initialization process. Static security associations are provided within the BS. Dynamic security associations are established and eliminated on the fly in response to the initiation and termination of specific service flows. Both static and dynamic SAs can be shared by multiple SSs. A security association's shared information includes the cryptographic suite employed within the SA. In this case, a cryptographic suite is the security association's set of methods for data encryption, data authentication, and traffic encryption key (TEK) exchange. The shared information may also include traffic encryption keys and initialization vectors. The exact content of the SA depends the SA's cryptographic suite. Security associations are identified using security association IDs. The ID of the primary security association of a SS must be equal to the basic connection ID of that SS.

There are rules for mapping connections to security associations:

- All data connections are mapped to an existing security association.
- The basic and the primary management connections are not mapped to a security association, but the secondary management connection is mapped to the primary security association.
- All transport connections are mapped to a security association.

Encryption is always applied to the MAC payload; the header is not encrypted. The generic MAC header is not encrypted because it contains encryption information needed to decrypt a payload at the receiving station. There is a special bit field, called Encryption Control, which indicates whether the payload is encrypted. Any unencrypted MAC frame received on a connection mapped to a security association requiring encryption is discarded. To facilitate normal operation of the MAC sublayer, all MAC management messages are without encryption.

INTEL-1020
U.S. Patent 8,219,129

A SS uses the privacy key management (PKM) protocol to obtain authorization and traffic keying material from the BS, and to support periodic reauthorization and key refresh. The key management protocol uses X.509 digital certificates, the RSA public-key encryption algorithm, and strong symmetric algorithms to perform key exchanges between the SS and the BS. The privacy key management protocol adheres to a client/server model, where the SS is the client that requests keys, and the BS is the server that responds to those requests, ensuring that individual SS clients receive only keys for which they are authorized. The PKM protocol uses MAC management messaging. The key distribution mechanism has two steps. First, the PKM protocol uses public-key cryptography to establish a shared secret (i.e., an Authorization Key) between the SS and the BS. Then, the shared secret is used to secure subsequent exchanges of traffic encryption keys. This two-step mechanism for key distribution permits refreshing of traffic encryption keys without the significant overhead of computation-intensive public-key operations.

The keying material of a security association (e.g., DES key and CBC Initialization Vector) has a limited lifetime. When the BS delivers the keying material to an SS, it also provides the SS with that material's remaining lifetime. It is the responsibility of the SS to request new keying material from the BS before the current one expires at the BS. If the current keying material expires before a new set of keying material is received, the SS must perform network entry. Again, the PKM protocol specifies how the SS and the BS maintain key synchronization.

Using the PKM protocol, an SS requests from its BS the keying material for that security association. The BS ensures that each client SS only has access to the Security Associations it is authorized to access.

Authentication

The goal of the authentication process is for the BS to establish an authenticated identity of a client SS and the services (i.e., specific traffic encryption keys) the SS is authorized to access.

SS authorization, controlled by the authorization state machine, has three steps.

1) The BS authenticates a client SS's identity.

INTEL-1020
U.S. Patent 8,219,129

2) The BS provides the authenticated SS with an authorization key, from which a Key Encryption Key (KEK) and message authentication keys are derived.
3) The BS provides the authenticated SS with the IDs of the security association and the properties of primary and static security associations for which the SS is authorized to obtain keys.

Each SS carries a unique X.509 digital certificate issued by the manufacturer itself or by an external authority. The digital certificate contains the SS's public key and SS MAC address. A SS begins authorization by sending an Authentication Information message to its BS. The Authentication Information message contains the SS's X.509 certificate. The Authentication Information message is strictly informative; i.e., the BS may choose to ignore it. However, it does provide a mechanism for a BS to learn the manufacturer certificates of its client SS. The SS sends an Authorization Request message to its BS immediately after sending the Authentication Information message. This is a request for an authorization key, as well as for the security association IDs identifying any static security associations in which the SS is authorized to participate. Again, the Authorization Request includes again the X.509 certificate, a description of the cryptographic algorithms the requesting SS supports, and the SS's basic connection ID. The SS's cryptographic capabilities are presented to the BS as a list of cryptographic suite identifiers, each indicating a particular pairing of encryption and authentication algorithms the SS supports. Recall that the basic connection ID is the first static CID the BS assigns to an SS during initial ranging—the primary security association ID is equal to the basic connection ID.

In response to an Authorization Request message, the BS verifies the digital certificate. The use of the X.509 certificates prevents cloned SSs from passing fake credentials onto a BS. The BS determines whether the requesting SS is authorized for basic unicast services and if the SS's user has subscribed to any additional statically provisioned services (i.e., static security association IDs). The BS associates the SS's authenticated identity to a paying subscriber, and hence to the data services that subscriber is authorized to access.

INTEL-1020
U.S. Patent 8,219,129

> The protected services a BS makes available to a client SS can depend upon the particular cryptographic suites shared by the SS and the BS.

The BS then determines which encryption algorithm and protocol is common with the SS. From the list of cryptographic suites that the SS has provided, the BS selects one cryptographic suite to employ with the requesting SS's primary SA. The Authorization Reply that the BS sends back to the SS includes a primary security association descriptor, which, among other things, identifies the cryptographic suite the BS selected to use for the SS's primary SA. A BS rejects the authorization request if it determines that none of the offered cryptographic suites are satisfactory.

In addition to X.509, all SSs will either have factory-installed RSA private/public key pairs or provide an internal algorithm to generate such key pairs dynamically. If a SS relies on an internal algorithm to generate its RSA key pair, then the SS generates the key pair before its first authentication key (AK) exchange.

Initially the BS has a single active AK for each SS. After the BS has verified the digital certificate and selected a common cryptographic suite, it activates a second AK for this SS. The second AK is sent in the Authorization Reply message. The lifetime of this second Authorization key is finite and is set to be equal to the remaining lifetime of the first AK plus a predefined time period. In this fashion, for each authenticated client SS there are two active authorization keys with overlapping lifetimes. Then the BS activates an AK for the SS, uses the verified public key of the SS to encrypt this AK and then sends the encrypted key back to the requesting SS in an Authorization Reply message. The authorization reply includes the following:

1) An AK encrypted with the SS's public key
2) A four-bit key sequence number, used to distinguish between successive generations of authorization keys
3) A key lifetime
4) The identities (i.e., the IDs of the security association) and properties of the single primary and zero or more static security associations for which the SS is authorized to obtain keying information.

The Authorization Reply also contains an optional list of static security association descriptors; each static security association descriptor identifies the

INTEL-1020
U.S. Patent 8,219,129

cryptographic suite employed within that security association. The selection of a static security association's cryptographic suite is typically made independent of the requesting SS's cryptographic capabilities. In its Authorization Reply, a BS may include static SA descriptors identifying cryptographic suites that the requesting SS does not support. This completes initial authorization.

An authorization request can be rejected. Generally, there are two types of rejections—permanent or simple. Permanent authorization rejection indicates that the error is of a permanent nature. What is interpreted as a permanent error is subject to administrative control within the BS. Examples of permanent error conditions include unknown manufacturer of the SS, invalid signature on the SS certificate, inconsistencies between data in the certificate and data in accompanying PKM data, or incompatible security capabilities. Simple authorization rejection does not indicate if the failure was due to a permanent error condition. As a result, the SS can set a wait timer and can reattempt authorization after the timer expires.

However, authorization is not a one-time process. Authorization is always valid for a finite period of time. After achieving initial authorization, a SS periodically seeks reauthorization with the BS. A SS must periodically refresh its authorization key by reissuing an Authorization Request to the BS. Reauthorization is identical to authorization with the exception that the SS does not send Authentication Information messages during reauthorization cycles. To avoid service interruptions during reauthorization, successive generations of the SS's authorization keys have overlapping lifetimes. Both SS and BS should be able to support up to two simultaneously active authorization keys during these transition periods. If the SS does not complete the reauthorization before the expiration of the authorization keys, the SS is considered unauthorized.

An SS must maintain its authorization status with the BS in order to be able to refresh aging TEKs. The refreshing of TEKs is managed by state machines. Upon achieving authorization, a SS starts a separate TEK state machine for each of the security association IDs identified in the Authorization Reply message. Each TEK state machine operating within the SS is responsible for managing the keys associated with its respective security association ID. TEK state machines periodically send Key Request messages to the BS, requesting a refresh of keying material for their respective security association identifier. The BS responds to a

INTEL-1020
U.S. Patent 8,219,129

Key Request with a Key Reply message, containing the BS's active keying material. The traffic encryption key (TEK) in the Key Reply is triple DES (encrypt-decrypt-encrypt or EDE mode) encrypted, using a two-key, triple DES key encryption key (KEK) derived from the authorization key.

> A SS must maintain the two most recent generations of keys for each SA. This retention helps maintain uninterrupted service during the key transitions.

Since every key associated with a Security Association (SA) has a limited lifetime, the BS manages a two-bit key sequence number independently for each SA and distributes this key sequence number along with the SA's key to client SS. The BS increments the key sequence number with each new generation of keys. The MAC header includes this sequence number to identify the specific generation of the SA key used to encrypt the attached payload. Because the sequence number is a two-bit quantity, it wraps around to 0 when it reaches 3. Comparing a received MAC frame's key sequence number with what it believes to be the "current" key sequence number, an SS or BS can easily recognize a loss of key synchronization with its peer.

As discussed previously, at all times the BS maintains two active sets of keys per security association identifier. The lifetimes of the two generations overlap such that each generation becomes active halfway through the life of its predecessor and expires halfway through the life of its successor. A BS includes in its Key Replies both of the SAID's active generations of keying material. The Key Reply provides the requesting SS, in addition to the TEK and CBC initialization vectors, their remaining lifetimes. The receiving SS uses these remaining lifetimes to estimate when the BS will invalidate a particular TEK, and therefore when to schedule future Key Requests so that the SS requests and receives new keying material before its current keying material expires. The operation of the TEK state machine's Key Request scheduling algorithm, combined with the BS's regimen for updating and using an SAID's keying material, ensures that the SS will be able to continually exchange encrypted traffic with the BS.

A TEK state machine remains active as long as the SS is authorized to operate in the BS's security domain, i.e., it has a valid authorization key, and the SS is authorized to participate in that particular security association, i.e., BS continues to provide fresh keying material during rekey cycles.

INTEL-1020
U.S. Patent 8,219,129

There are two state machines—authorization state machine and TEK state machine. The authorization state machine will start an independent TEK state machine for each of its authorized SAIDs. The BS maintains two active TEKs per SAID. The BS includes in its Key Replies both of these TEKs, along with their remaining lifetimes. The BS encrypts downlink traffic with the older of its two TEKs. Note that the BS encrypts with a given TEK for only the second half of its lifetime. The BS decrypts uplink traffic with either the older or newer TEK, depending upon which of the two keys the SS was using at the time. The SS encrypts uplink traffic with the newer of its two TEKs and decrypts downlink traffic with either the older or newer TEK, depending upon which of the two keys the BS was using at the time.

Communication between authorization and TEK state machines occurs through the passing of events and protocol messaging. The authorization state machine generates events (i.e., Stop, Authorized, Authorization Pending, and Authorization Complete events) that are targeted at its child TEK state machines. TEK state machines do not target events at their parent authorization state machine. The TEK state machine affects the authorization state machine indirectly through the messaging a BS sends in response to a SS's requests.

Through operation of a TEK state machine, the SS attempts to keep its copies of the SAID's TEKs synchronized with those of its BS. A TEK state machine issues Key Requests to refresh copies of its SAID's keying material soon after the scheduled expiration time of the older of its two TEKs and before the expiration of its newer TEK. To accommodate for SS/BS clock mismatch and timing mismatches, the SS schedules its Key Requests several seconds before the newer TEK's estimated expiration in the BS. When it receives the Key Reply, the SS updates its records with the TEK Parameters from both TEKs contained in the Key Reply message.

The BS is responsible for maintaining keying information for all Security Associations. The BS must maintain two active TEKs and their associated initialization vectors per Security Association identifier.

Data encryption with DES

The encryption method used is the Cipher Block Chaining (CBC) mode of the U.S. Data Encryption Standard (DES). The CBC initialization vector is calculated

INTEL-1020
U.S. Patent 8,219,129

as follows. In the downlink, the CBC is initialized with the XOR of the initialization vector in the TEK keying information and the PHY synchronization field of the latest DL-MAP. In the uplink, the CBC is initialized with the XOR of the initialization vector in the TEK keying information and the PHY synchronization field of the DL-MAP that is in effect when the UL-MAP is received.

In the Key Reply message, TEK is encrypted by the BS using the two-key triple-DES in the encrypt-decrypt-encrypt (EDE) mode. This encryption operation can be described by

$$C = Ek1[Dk2[Ek1[P]]] \qquad \text{Eq. 4–1}$$

where P is 64-bit plaintext TEK, C is 64-bit ciphertext TEK, $E[]$ is 56-bit DES electronic codebook (ECB) mode encryption, and $D[]$ is 56-bit DES ECB decryption. The leftmost 64 bits of the 128-bit KEK are denoted $k1$, and the rightmost 64 bits of the KEK are denoted $k2$. The corresponding decryption operation is

$$P = Dk1[Ek2[Dk1[C]]] \qquad \text{Eq. 4–2}$$

The 56-bit DES keys are actually 64 bits long. In every byte, the seven most significant bits are independent and are used in a DES key. The least significant bit is a parity bit, so that every byte has odd parity.

The keying material for two-key triple DES consists of two distinct DES keys. The 64 most significant bits of the KEK are used in the encrypt operation. The 64 least significant bits are used in the decrypt operation.

MAC ENHANCEMENTS FOR 2–11 GHZ OPERATION

There is a significant interest in the wireless industry especially in the frequency bands between 2 GHz and 11 GHz. The advantages for operation in this frequency range include operation in a non-line of sight (NLOS) environment. Operation in this frequency band requires not only a different physical layer, but also additions (or, alternatively, enhancements) to the MAC layer. These enhancements to the MAC are necessary to support not only point-to-multipoint architectures but, in license-exempt bands, optional mesh topology. Because multipath is significant,

INTEL-1020
U.S. Patent 8,219,129

losses over the wireless medium will increase. To deal with these losses, the second area of additions to the MAC are related to introduction of ARQ on a per-connection basis. The third area of additions is related to the DFS and TPC mechanisms.

MAC enhancements for mesh systems

In unlicensed bands the topology of the network can be different from the point-to-multipoint type. In particular multipoint-to-multipoint (or mesh) networks can be used. Mesh architecture is a network where devices are capable of forwarding traffic from and to multiple other devices. Figure 4–5 is an illustration of the mesh concept.

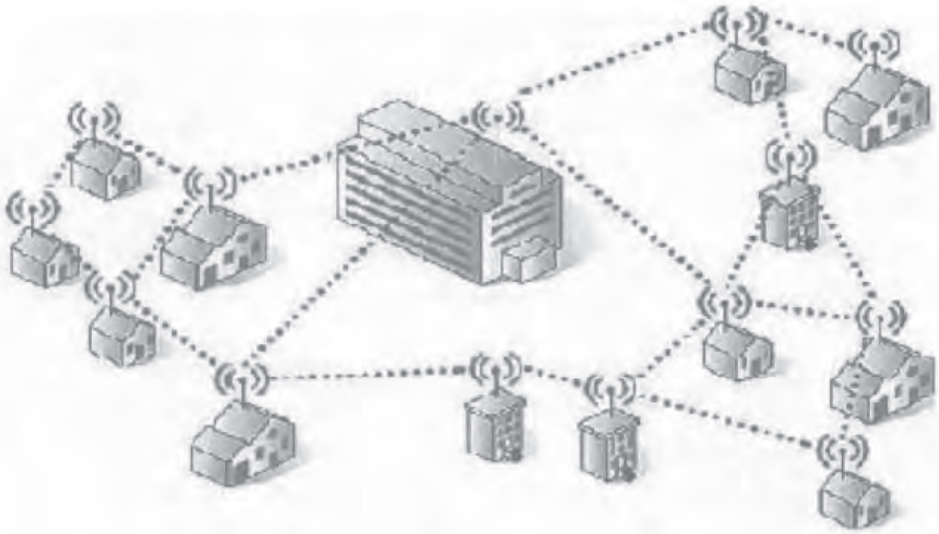

Figure 4–5: An example of a mesh network

Mesh mode is optional for operation in unlicensed bands between 2 GHz and 11 GHz. In mesh networks, a number of subscriber stations within a geographic area are interconnected and can act as repeater stations. This allows a variety of routes between the core network and any subscriber station. Clearly, at least one of the devices in a mesh network must be connected to a backhaul network. This device is termed a mesh BS. Devices that are not connected to the backhaul network are termed mesh SS. Uplink and downlink are different for mesh systems compared with point-to-multipoint networks. Downlink is defined as traffic away

INTEL-1020
U.S. Patent 8,219,129

from the mesh BS, and uplink is defined as traffic towards the mesh BS. The other three important elements of the architecture of mesh systems are neighbor, neighborhood, and extended neighborhood. The stations with which a node has direct links are called neighbors. All neighbors of a node form a neighborhood. The neighbors of a node are considered to be "one hop" away from the node. An extended neighborhood also contains all the neighbors of the neighborhood.

Mesh networks need a different multiple-access mechanism than point-to-multipoint systems. The reason is that mesh devices that are close to each other cannot transmit at the same time on the same channel. If a node is transmitting all nodes within a distance of two hops must remain silent. Therefore, even the BS must coordinate its transmission with other stations. The transmission activity of a node depends substantially on how much traffic it will have to forward to and from other nodes. In turn, the amount of traffic depends on the position of the node within the network. The main difference between mesh and point-to-multipoint mode is that the channel resources (e.g., the ability to transmit) are shared between the systems on demand basis. There are two principal ways to assign channel resources: distributed scheduling and centralized scheduling. A combination of both is also possible. Distributed scheduling implies distributed decision-making. In centralized scheduling, the mesh BS makes decisions as to who has the right to transmit at any given time instant.

Clearly, mesh networks need changes in the basic MAC protocol because of their architecture. The mesh capability support includes (1) mesh network configuration message, (2) mesh network entry message, and (3) messages for distributed and centralized mesh scheduling.

Mesh devices use mostly roof-mounted omnidirectional antennas, because coverage must be extended in all directions. However, mesh systems can use a variety of antenna types simultaneously. At the edge of the coverage area of the mesh network even highly directional antennae can be used.

In the coordinated distributed scheduling mode, all stations (BSs and SSs) coordinate their transmissions in their extended two-hop neighborhood. Every station (BS and SS) transmits a message at regular intervals to inform all the neighbors of its own schedule and proposed schedule changes. This message is used simultaneously to convey resource

INTEL-1020
U.S. Patent 8,219,129

requests to the neighbors. All the stations in a network must use this same channel to transmit schedule information in a format of specific resource requests and grants. Coordinated distributed scheduling ensures that transmissions are scheduled in a manner that does not rely on the operation of a BS, and that are not necessarily directed to or from the BS. Within the constraints of the coordinated schedules (distributed or centralized), uncoordinated distributed scheduling can be used for fast, ad-hoc setup of schedules on a link-by-link basis. Uncoordinated distributed schedules are established by directed requests and grants between two nodes These schedules must be set up to keep the resulting data transmissions (and the request and grant packets themselves) from colliding with the data and control traffic scheduled by the coordinated distributed methods or the centralized scheduling methods. The differences between the coordinated and uncoordinated distributed scheduling are that, in coordinated scheduling, management messages are scheduled in the control subframe in a collision-free manner. In uncoordinated scheduling, management messages can collide.

Centralized scheduling is performed by the BS. A Mesh Centralized Scheduling Configuration message is broadcasted in mesh systems using centralized scheduling. The BS shall broadcast this message to all its neighbors, and all the subscriber stations forward (rebroadcast) the message in a way specified in the message. The mesh BS gathers resource requests from all the mesh SSs within a certain threshold hop range (HR). The mesh BS will determine the amount of granted resources for each link in the network both in downlink and uplink, and communicates these grants to all the mesh SSs within the hop range. The grant messages do not contain the actual schedule, but each node can compute it by using the predetermined algorithm with given parameters.

The network connections and topology are the same as in the distributed scheduling mode. However, some portion of the scheduled transmissions for the SSs less than or equal to a certain number of hops from the BS will be either defined by the BS or computed by the SSs themselves. This hop range, HR threshold, may be determined at the system startup phase or may be dynamic, according to considerations such as network density, the proximity of other BSs, and/or the dynamic characteristics of the traffic streams.

In the basic form, the BS provides the schedule for all the SSs less than or equal to HR threshold hops from the BS. The BS determines the flow assignments from

INTEL-1020
U.S. Patent 8,219,129

the resource requests from the SSs within the HR threshold hop range. Subsequently, the SSs themselves determine the actual schedule from these flow assignments by using a common algorithm that divides the frame proportionally to the assignments. Thus, the BS acts just like the BS in a PMP network except that not all the SSs have to be directly connected to the BS and the assignments determined by the BS extends also to those SSs not directly connected to the BS but less than HR threshold hops from it. The SS resource requests and the BS assignments are both transmitted during the control portion of the frame. Centralized scheduling ensures that transmissions are coordinated to ensure collision-free scheduling over the links in the routing tree to and from the BS. Centralized scheduling is more efficient than the distributed scheduling for traffic streams (or collections of traffic streams that share links) that persist over a duration greater than the cycle time to relay the new resource requests and distribute the updated schedule.

In the absence of clearly defined uplink and downlink, only TDD is supported in mesh mode.

The first portion of a downlink (MAC) frame contains unencrypted frame control information destined for all SSs. This control section must contain a DL-MAP message for the downlink channel followed by one UL-MAP message for each associated uplink channel. The DL-MAP message does not describe the locations of the payload bursts that immediately follow; it describes the locations of the downlink bursts at some prescribed time in the future. This delay is necessary because of MAC and PHY processing delays.

Since all the basic functions like scheduling and network synchronization are based on the neighbor information, each node (BS and SS) in the mesh network maintains a physical neighborhood list containing the MAC address and the distance (measured in hops) of this neighbor from the present node. A SS that has a direct link to the BS will synchronize to the BS, while a SS that is at least two hops from the BS will synchronize to its neighbor SSs that are closer to the BS. Special packets, called network configuration and network entry packets, provide a basic level of communication between nodes in different nearby networks. These packets are used to synchronize both centralized and distributed control mesh networks. These packets are also used to synchronize the transmissions among multiple co-located BSs.

INTEL-1020
U.S. Patent 8,219,129

Node initialization and network entry procedures in mesh mode are in some aspects different from those in PMP mode. A new node entering the mesh network obeys the following procedure. First, it goes through scanning and synchronization. After scanning and synchronization, a station listens for mesh configuration packets to build its physical neighbor table. Then it selects a sponsor node from the neighbor table based on receive signal quality, and gains a coarse network synchronization using a received network configuration packet from the selected sponsor. Then the new node transmits a packet, which includes the selected sponsor node's address. After a sponsor node is selected, the new station can perform authentication. The sponsor node will forward authentication frames to a node in the mesh, called authentication node. This authentication node will act as a BS, verifying the digital certificate of the new station. MAC management messages are forwarded to the new station through its sponsor node. A node can exchange TEKs with other nodes, not just the BS. A TEK state machine is identified by a SAID. Every node will maintain separate TEK state machines per SAID for all nodes it exchanges TEKs with. If operation with a selected sponsor fails, a different sponsor is selected and the procedure repeated. The sponsor will also include the link information for this new node in its neighbor list. After authorization the new station receives its node ID, establishes IP connectivity, time of day, and exchanges the operational parameters.

Advanced antenna systems (AAS)

Systems that work in the 2–11 GHz band can use advanced antenna systems (AAS). Advanced antenna systems are based on a technology that uses more than one antenna element. A significant advantage of this technology is that it improves range and system capacity. These are very important advantages, especially when due to a combination of technical and regulatory reasons that cannot be easily achieved in other ways. It is expected that advanced antenna systems will be used in the future, and 802.16 as an option includes functionality to support them.

The first advantage of AAS is that their spectral efficiency can be increased linearly with the number of antenna elements. This is achieved by steering beams to multiple users simultaneously so as to realize an intercell frequency reuse of one and an in-cell reuse factor proportional to the number of antenna elements. An additional benefit is the SNR gain realized by coherently combining multiple

INTEL-1020
U.S. Patent 8,219,129

signals. Another possible benefit is the reduction in interference achieved by steering nulls in the direction of co-channel interferers. It must be noted that these techniques do not require multiple antennas at the SS. However, it is still possible to use multiple antennas at the SS to achieve further benefits.

Initially only the BS will have AAS. The design of the AAS option provides a mechanism to migrate from a non-AAS system to an AAS-enabled system, in which the initial replacement of the non-AAS-capable BS by an AAS-capable BS should not cause service interruption. The goal of 802.16 is to make possible systems where the BS, and possibly some SS, use AAS, while other SSs do not. This is achieved by dedicating only part of any frame to AAS traffic (Figure 4–6). Part of the frame is dedicated to non-AAS traffic and part is dedicated to AAS traffic. The allocation is performed dynamically by the BS. Non-AAS SSs can ignore AAS traffic, which they can identify based on the DL-MAP/UL-MAP messages. AAS-enabled SSs will use dedicated DL-MAP/ UL-MAP messages and are therefore prevented from colliding with non-AAS traffic. Non-AAS subscriber stations are hence able to operate whether or not they receive any signal in the AAS part of the frame, and AAS subscriber stations are able to operate whether or not they receive any signal in the non-AAS part of the frame.

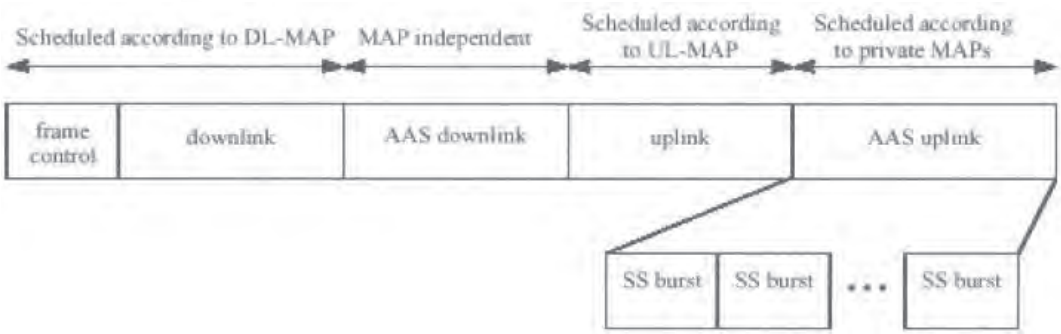

Figure 4–6: Frame structure for systems with advanced antennas

Additional functionality required to support AAS in the MAC layer can be divided into two categories—control functionality and utility functionality. The control functionality is responsible for the DL-MAP and UL-MAP distribution, and channel description. The utility function is responsible for the PHY-related information provided by the MAC.

INTEL-1020
U.S. Patent 8,219,129

Furthermore the following additional logical channels are required for AAS support:

1) *Downlink Synchronization Channel:* This channel is used for time and frequency synchronization by AAS subscriber stations.
2) *Downlink Polling Channel:* This channel is required for systems where the array adaptation using signals received in the Uplink Contention Channel may not be adequate to open a Downlink Traffic Channel to a subscriber station.
3) *Downlink Traffic Channels:* This is the part of the physical channel that carries downlink user traffic, as scheduled by the BS. Their availability requires array adaptation. In contrast to similar channels in the non-AAS systems, multiple simultaneous Downlink Traffic Channels can be open to spatially separated subscriber stations.
4) *Uplink Contention Channel:* This channel is used for SS-initiated random access, including ranging and bandwidth request. Its availability does not require array adaptation.
5) *Uplink Traffic Channels:* This is the part of the physical channel that carries uplink user traffic, as scheduled by the BS. Their availability requires array adaptation.

The process of registration in AAS systems is different from the regular registration process. This is because the adaptive array operating in the PHY cannot be effective until the MAC and PHY of the BS identify the registering subscriber station. The adaptation of BS antenna array can be accomplished only after the BS has identified the SS. On the other hand, if the BS cannot adapt the array to the new subscriber, there is a chance that the SS will not receive any valid signal from the BS at all. Because the regular registration process requires that the SS receive a valid message first, it is not always possible to rely on this process. If the SS decodes the DL-MAP and DCD messages, it can proceed with the network entry process, just like in the non-AAS case (power adjustment, rate adaptation, authentication, encryption, etc.). In this case, the BS can adapt its adaptive array during the ranging process.

INTEL-1020
U.S. Patent 8,219,129

> Note also that AAS subscribers might not be able to request bandwidth using the usual contention mechanism. This happens because the adaptive array may not have a beam directed at the SS when it is requesting BW, and the BW request will be lost. There are two ways out of this situation. The SS can broadcast the bandwidth request or the BS can poll the SS.

If the SS cannot decode the DL-MAP and DCD messages, it can use the following procedure: An AAS BS must reserve a fixed, predefined part of the frame as initial-ranging contention slots for AAS SS. These contention slots are located at a well-known location relative to the downlink preamble, so even an SS that can only identify the DL preamble can locate it. These contention slots are called AAS-alert-slots and their number and location in the frame is PHY specific. The AAS SS that cannot decode the DL-MAP and DCD messages will attempt ranging on the AAS-alert-slots. Then the SS will wait for a private DL-MAP and private DCD messages from the BS. If these messages are received, the SS continues the network entry process like a non-AAS SS. If the private DL-MAP and DCD messages are not received, the SS will use an exponential backoff algorithm for selecting the next frame in which to attempt alerting the BS for its presence.

Adaptive Arrays require knowledge of the channel at both downlink and uplink. The BS can obtain channel information in two ways. The first way is to perform uplink channel estimation and assume that the downlink channels are identical to the uplink channels. The second way is to use feedback from the SS, thus transmitting the estimated channel state from the SS to BS. The first method is simpler and is well suited for TDD systems. The second method is more suitable for FDD systems, where reciprocity does not apply (because uplink and downlink are in different frequency bands). The second method can also be used in TDD systems. In the second method, the BS must periodically send a message to the SS to signal that channel state information should be updated. The period with which this message is sent is called Channel Estimation Interval (CEI). It is vendor-specific and is an internal parameter of the BS. a. The value of CEI is defined according to the channel stability over time (a typical value is 20 ms). The BS is responsible to determine the actual value of CEI, and for the distribution of this value to all SSs. The SS must perform channel estimations on a regular time basis

INTEL-1020
U.S. Patent 8,219,129

to be able to provide up-to-date estimations upon request. In response to a request from the BS, the SS will reply with the most up-to-date estimation of the channel.

Automatic repeat request (ARQ)

Multipath in the range between 2 GHz and 11 GHz leads to intersymbol interference, reduced signal-to-noise ratio, and increased BER and PER. To deal with these significant issues, 802.16 provides an ARQ mechanism. The ARQ mechanism is an optional part of the MAC layer and can be enabled on a per-connection basis. The per-connection ARQ and associated parameters are specified and negotiated during connection creation or change. A connection cannot have a mixture of ARQ and non-ARQ traffic. Similar to other properties of the MAC protocol, the scope of a specific instance of ARQ is limited to one unidirectional connection. The ARQ feedback information can be sent as a standalone MAC management message on the appropriate basic management connection or piggybacked on an existing connection. ARQ feedback cannot be fragmented. The receiver provides positive or negative ARQ acknowledgments.

Recall that the format of a MAC frame is a single header and one or more whole or fragmented payloads. In addition to the single MAC header, there may be subheaders. The MAC header contains information whether subheaders are present. If ARQ is used, a MAC frame can contain ARQ feedback payload, preceded by an appropriate subheader.

Transmitter and receiver state machine operations include comparing fragment sequence numbers and taking actions based on which is larger or smaller. In this context, it is not possible to compare the numeric sequence number values directly to make this determination. The ARQ protocol is responsible for choosing the right fragment size on a per-fragment basis. The fragment size is not a fixed value that remains constant for a particular connection. Once defined, the size of a fragment cannot be changed between retransmissions. An ARQ fragment may be in one of the following four states—not-sent, outstanding, discarded, and waiting-for-retransmission. Any ARQ fragment begins as not-sent. After it is sent, it becomes outstanding for a period of time termed the retry timeout. While a fragment is in outstanding state, either it is acknowledged and then discarded, or transitions to state waiting-for-retransmission after a timeout. An ARQ fragment can become waiting-for-retransmission before the timeout expires if it is

INTEL-1020
U.S. Patent 8,219,129

negatively acknowledged. An ARQ fragment may also change from waiting-for-retransmission to discarded when a positive acknowledgment for it is received or after another timeout expires. The transmitter policy is that if any waiting-for-retransmission ARQ fragments exist, they should be given precedence over not-sent packets for the same connection. When fragments are retransmitted, the fragment with the lowest frame sequence number is retransmitted first. The ARQ transmit fragment state sequence is shown in Figure 4–7.

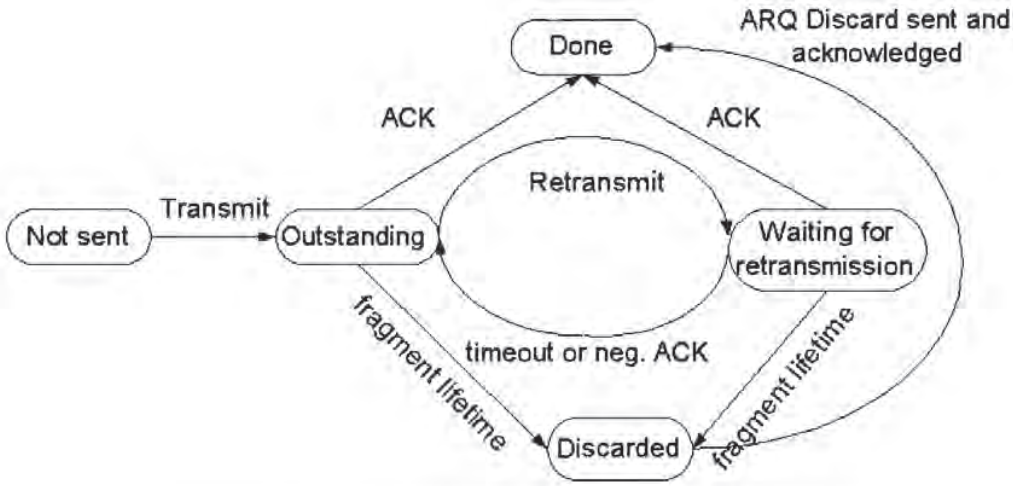

Figure 4–7: ARQ transmitter state machine

When an acknowledgment is received, the transmitter checks the validity of the frame sequence number. To be valid, a frame sequence number should be in a specified range. If the frame sequence number is not valid, the transmitter will ignore the acknowledgment. There is cumulative acknowledgment for multiple fragments. The cumulative acknowledgment contains a bitmap. All fragments where the corresponding bit is 1 in the bitmap are positively acknowledged, and those with the corresponding bit set to 0 are negatively acknowledged.

Under abnormal error conditions, as a last resort, a station may send a Reset message. Exactly when a Reset is sent is outside the scope of 802.16.

INTEL-1020
U.S. Patent 8,219,129

When a frame is received, its integrity is determined based on the CRC-32 checksum. If the frame passes the checksum, it is unpacked and defragmented, if necessary. The receiver maintains a sliding window. The start of the window is the lowest fragment that has not been received or has been received in error. When an ARQ fragment with a number in the range defined by the sliding window is received, the receiver will accept it. ARQ fragment numbers outside the sliding window are rejected as out of order.

For each ARQ fragment accepted fully and without errors, an acknowledgment message is sent to the transmitter. Acknowledgments may be either for specific ARQ fragments (i.e., contain information on the acknowledged ARQ fragment numbers), or cumulative (i.e., contain the highest ARQ fragment number below which all ARQ fragments have been received correctly) or a combination of both (i.e., cumulative with selective). Acknowledgments are sent in the order of the ARQ fragment numbers they acknowledge. How often acknowledgments are generated depends on implementation. The interval is not specified in 802.16. The ARQ receiver state machine has four states—not received, received, received correctly, and done. Initially a fragment is not received. If the CRC check is correct and the frame sequence number is in range a fragment is received correctly. If the fragment sequence number is out of range, the fragment is considered not received.

DFS for license-exempt operation

For operation in unlicensed bands, DFS is required. Unlicensed bands can be allocated for more than one use. In addition to being available on unlicensed basis, some bands have primary users. These primary users of the bands legally have the right of way. Wireless data communication systems can operate in unlicensed bands only if they are capable of avoiding the primary users by employing a dynamic frequency selection (DFS). If a frequency band contains a primary user, the wireless data communications equipment must detect the primary user and move to another band. Devices can move to another band for other reasons such as low RSSI or interference from other systems. Therefore, DFS is used when devices choose frequency band for operation. Each downlink channel will be

INTEL-1020
U.S. Patent 8,219,129

characterized in terms of its RSSI. Valid channels are only those channels that have a high enough SNR. DFS procedures provide for the following:

- Requesting and reporting measurements
- Testing channels for the presence of primary users
- Detecting primary users
- Ceasing operation on a channel after primary users have been found
- Selecting and advertising a new channel

Note that several of these procedures are outside of the scope of 802.16.

The process also requires channel state monitoring by the SS. Channels are assigned by the BS. The initial ranging mechanism consists of a series of transactions where the SS send a ranging request message, and the BS answers with a ranging response message and the SS corrects its transmission power or timing parameters as instructed by the BS. During periodic maintenance, the SS can be either polled by the BS to transmit a ranging request message, or can send it independently. The DFS mechanism messages are appended to the ranging messages. In the ranging response message, the BS can request the DFS info. The SS shall reply with an appropriate ranging request message containing the requested information. The ranging request/response transactions continue until the BS is satisfied with the results.

MAC enhancements for OFDM and OFDMA PHYs

The WirelessMAN-OFDM physical layer supports two contention-based bandwidth request mechanisms. In addition to the mandatory bandwidth request mechanism, the MAC provides another bandwidth request mechanism. The station requesting bandwidth can send what is called a contention code. This contention code is modulated on a contention channel. The contention channel consists of four carriers. The contention code and the contention channel are randomly chosen with uniform distribution. When the BS receives this contention code' it will provide an predefined UL allocation for the SS to transmit a bandwidth request and, optionally, additional data. If the BS does not provide an uplink allocation, the SS will assume that the ranging code transmission resulted in collision and will follow the contention resolution procedure.

INTEL-1020
U.S. Patent 8,219,129

The OFDMA physical layer supports two mandatory contention-based bandwidth request mechanisms. In addition to the standard mechanism for bandwidth requests, there is a CDMA-based mechanism. The OFDMA layer specifies a Ranging Subchannel and a set of Ranging codes. These ranging codes are special PN codes used for making bandwidth requests. The SS selects with equal probability a ranging code from the code subset. The selected ranging code is modulated onto the ranging subchannel and transmitted during the appropriate uplink allocation. When the BS receives this contention code, it will provide a predefined uplink allocation for the SS. If the BS does not provide an uplink allocation, the SS will assume that the ranging code transmission resulted in collision and will follow the contention resolution procedure.

IEEE 802.16 PHYSICAL LAYERS

Every physical layer has two sublayers—a transmission convergence sublayer and a Physical Medium Dependent (PMD) sublayer. Each PMD sublayer may require the definition of a unique transmission convergence sublayer. In addition, a physical layer may be accompanied by a physical layer management entity. IEEE 802 standards do not specify the exact functions of the management entity; the functions are always left to vendors. In general, the management entity is responsible for such functions as gathering of layer-dependent status and interaction with the general system management. Physical layer management functions include frequency adjustment, power management, propagation delay compensation, etc. The management information specific to PHY layer is represented as a management information base (MIB) for this layer. The generic model of the device management is that it either reads or sets the value of a parameter in the MIB.

Recalling the architecture of IEEE 802.16 shown in Figure 4–2, the PHY can be considered as providing services to the MAC layer. These services are accessed through the PHY service access point (SAP). Therefore the communication between the MAC and the PHY is described using a set of primitives. There are three types of primitives between the MAC and the PHY—primitives for data transfer, primitives for management functions, and service primitives that have local significance and support sublayer-to-sublayer interactions related to layer control.

INTEL-1020
U.S. Patent 8,219,129

IEEE 802.16 defines several physical layers, for business and other reasons. The main reason is that there was not significant support within the working group for a single physical layer. By standardizing several physical layers at once, vendors are left free to implement whichever they want. Ultimately this flexibility allows service providers the ability to optimize system deployments with respect to cell planning, cost, radio capabilities, services, and capacity.

Physical layer for 10–66 GHz

The characteristics of outdoor wireless channels are different. In outdoor wireless communication at frequencies over 10 GHz, the multipath components will be absorbed rather than reflected. The energy of the non-line-of-sight components that reach the receiver will be very small. This is a typical scenario for the operation of wireless Internet access equipment. The IEEE 802.16 working group has studied the characteristics of wireless channels and has come up with three types of channels.

The first type of channel is type 0, where no multipath components are present and the frequency response of the channel is equal to a constant. Channel type 1 has one weak non-line-of-sight component, and its frequency response can be expressed as:

$$H(j\omega) = 0.995 + 0.0995e^{-j0.75}e^{-j\omega 400/B} \qquad \text{Eq. 4–3}$$

where B is a parameter that can assume values between 15 and 25. This parameter is the baud rate that is used for communication over this channel, measured in Mbaud.

Channel type 2 has the strongest multipath components. Its frequency response can be expressed as:

$$H(j\omega) = 0.286e^{-j0.75} + 0.953e^{-j\omega 400/B} + 0.095e^{-j\omega 800/B} \qquad \text{Eq. 4–4}$$

At frequencies between 10 GHz and 66 GHz, the predominant fade mechanism in outdoor wireless communication is that resulting from rain attenuation. Given certain parameters such as frequency of operation, desired link availability and BER, the rain attenuation will determine the maximum cell radius of operation. An internationally accepted method for computation of rain fade attenuation

INTEL-1020
U.S. Patent 8,219,129

probability is that defined by ITU-R 530 [B92]. As an example, for operation at 28 GHz and link availability of 99.995%, the maximum cell radius is about 3.5 km in ITU rain region *K*.

A significant mitigation technique for the control of both intrasystem and intersystem interference is the antenna radiation pattern envelope (RPE) discrimination provided by system antennas. The RPE discrimination has significance for both clear sky and rain faded propagation conditions. The RPE requirements for aggressive intrasystem frequency reuse plans may exceed the RPE requirements for the control of intersystem coexistence. Recommended antenna RPE characteristics are described in "Coexistence " on page 317.

Rain attenuation will be different in general for the desired signal and for the interfering signal, i.e., there will be a differential rain fading loss. This differential rain fading loss between the two transmission paths can have a significant impact on the system performance. For example, at operational frequencies around 28 GHz, the estimated rain cell diameter is approximately 2.4 km.

The physical layer for operation between 10 GHz and 66 GHz must allow for flexible spectrum usage and support both time division duplexing (TDD) and frequency division duplexing (FDD). The wireless channel is different for the different SSs. The burst transmission format of the PHY is framed so that it supports adaptive burst profiling in which transmission parameters, including the modulation and coding schemes, can be adjusted individually to each subscriber station (SS) on a frame-by-frame basis. The FDD case supports full duplex SSs as well as half duplex SSs, which do not transmit and receive simultaneously. Note that because the architecture of IEEE 802.16 is substantially different from the architecture of IEEE 802.11 and IEEE 802.15, it uses adaptive modulation. Adaptive modulation means that systems communicate using multiple burst profiles, and furthermore the BS communicates with multiple subscriber stations using different burst profiles.

The uplink physical layer is based on a combination of time division multiple access (TDMA) and demand-assigned multiple access (DAMA). In particular, according to the TDMA part the uplink channel is divided into a number of time slots. The number of slots assigned for various uses (registration, contention,

INTEL-1020
U.S. Patent 8,219,129

guard, or user traffic) is controlled by the MAC layer in the base station and may vary over time for optimal performance.

The downlink channel is time division multiplexed (TDM), with the information for each subscriber station multiplexed onto a single stream of data and received by all subscriber stations within the same sector. To support half duplex FDD subscriber stations, a provision is also made for a TDMA portion of the downlink. The downlink physical layer includes a transmission convergence sublayer that inserts a pointer byte at the beginning of the payload to help the receiver identify the beginning of a MAC frame. The PMD layer performs randomization, FEC encoding, and modulation according to QPSK, 16-QAM, or 64-QAM (optional).

The uplink physical layer is based upon TDMA burst transmission. Each burst is designed to carry variable-length MAC frames. The PMD layer in the uplink performs randomization of the incoming data, FEC encoding, and modulation using QPSK, 16-QAM (optional), or 64-QAM (optional) The parameters the 10–66 GHz PHY are summarized in Table 4–1.

Table 4–1: Parameter values for the 10–66 GHz PHY

Parameter	Value
Symbol rate	16 to 40 Mbaud
Modulation density	2, 4, or 6
FEC block size	0–511 bytes
FEC payload	0–255 bytes
Uplink preamble length	16 or 32 symbols
PHY overhead	0–256 symbols

Power can adjusted in steps of 0.5 dB, k 0.5, where $k = -128 \ldots 127$.

This physical layer specification operates in a framed format. The supported frame durations are 0.5 ms, 1 ms, and 2 ms. Each frame contains a downlink subframe and an uplink subframe. The downlink subframe begins with

INTEL-1020
U.S. Patent 8,219,129

information necessary for frame synchronization and control. In the TDD case, the downlink subframe comes first, followed by the uplink subframe. In the FDD case, uplink transmissions occur concurrently with the downlink frame. Each SS attempts to receive all portions of the downlink except for those bursts whose burst profile is either not implemented by the SS or is less robust than the SS's current operational downlink burst profile. Half duplex SSs do not listen to portions of the downlink coincident with their allocated uplink transmission.

In FDD operation, the uplink and downlink channels are on separate frequencies. The capability of the downlink to be transmitted in bursts facilitates the use of different modulation types and allows the system to simultaneously support full duplex subscriber stations (which can transmit and receive simultaneously) and half duplex subscriber stations (which do not).

In the case of TDD, the uplink and downlink transmissions share the same frequency but are separated in time. A TDD frame also has a fixed duration and contains one downlink and one uplink subframe. The TDD framing is adaptive in that the link capacity allocated to the downlink versus the uplink may vary. Between the downlink burst and the subsequent uplink burst there is a gap, referred to as the Tx/Rx transition gap. This gap allows time for the BS to switch from transmit to receive mode and SSs to switch from receive to transmit mode. Similarly, there is a gap between the uplink burst and the subsequent downlink burst, called the Rx/Tx transition gap. Both transition gaps are an integer number of physical slots durations and start on a physical slot boundary.

Downlink physical layer

The available bandwidth in the downlink direction is defined in terms of physical slots. The available bandwidth in the uplink direction is defined in terms of mini-slots, where the minislot length is $2m$ physical slots (m ranges from 0 through 7). The number of physical slots with each frame is a function of the symbol rate. The symbol rate is selected in order to obtain an integral number of physical slots within each frame. For example, with a 20 Mbaud symbol rate, there are 5000 physical slots within a 1 ms frame.

The downlink frames can be TDD and FDD. The TDD downlink subframe contains preamble, frame control section, and data. The preamble is used for synchronization and equalization. The frame control section contains downlink

INTEL-1020
U.S. Patent 8,219,129

and uplink maps stating the physical slots at which bursts begin. The frame control section is not encrypted and is destined for all SSs. It also contains information about the modulation type and FEC code length for each burst profile. The data is organized into bursts with different burst profiles and therefore different level of transmission robustness. The bursts are transmitted in order of decreasing robustness. For example, with the use of a single FEC type with fixed parameters, data begins with QPSK modulation, followed by 16-QAM, followed by 64-QAM. At the end of the frame there is a Tx/Rx transition gap. Each SS receives and decodes the control information of the downlink and looks for MAC headers indicating that the remainder of the downlink subframe contains data for that SS.

In the FDD case, similar to the TDD case, the downlink subframe begins with a preamble followed by a frame control section and a TDM portion organized into bursts transmitted in decreasing order of burst profile robustness. This TDM portion of the downlink subframe contains data transmitted to full duplex SSs, half duplex SSs scheduled to transmit later in the frame than they receive, and half duplex SSs not scheduled to transmit in this frame. The FDD downlink subframe continues with a TDMA portion used to transmit data to any half duplex SSs scheduled to transmit earlier in the frame than they receive. This allows an individual SS to decode a specific portion of the downlink without the need to decode the entire downlink subframe. In the TDMA portion, each burst begins with the downlink TDMA burst preamble for phase resynchronization. Bursts in the TDMA portion need not be ordered by burst profile robustness. The FDD frame control section includes a map of both the TDM and TDMA bursts.

The TDD downlink subframe, which inherently contains data transmitted to SSs that transmit later in the frame than they receive, is identical in structure to the FDD downlink subframe for a frame in which no half duplex SSs are scheduled to transmit before they receive.

The frame control section contains the DL-MAP message, followed by the UL-MAP message for each associated uplink channel. In addition, it may contain DCD and UCD messages.

Two downlink burst preambles are used. The Frame Start Preamble is at the beginning of each downlink frame. The Downlink TDMA Burst Preamble is at the beginning of each TDMA burst in the TDMA portion of the downlink

INTEL-1020
U.S. Patent 8,219,129

subframe. Both preambles use QPSK modulation and are based upon +45 degrees rotated constant amplitude zero autocorrelation (CAZAC) sequences. These sequences can be normalized, because the amplitude of the preamble depends on the downlink power adjustment rule. In the case of the constant-peak power scheme, the preamble is transmitted so that its constellation points coincide with the modulation scheme's outermost constellation points. In the case of the constant mean power scheme, the preamble is transmitted with the mean power of the constellation points of the modulation scheme in use.

A burst corresponds to either a TDM burst beginning with the Frame Start Preamble or a TDMA burst beginning with a Downlink TDMA Burst Preamble.

The Frame Start Preamble consists of a 32-symbol sequence generated by repeating the 16-symbol QPSK CAZAC sequence $1+j, -1+j, -1-j, 1-j, 1+j, -1-j, 1+j, -1-j, 1+j, 1-j, -1-j, -1+j, 1+j, 1+j, 1+j, 1+j$. The Downlink TDMA Burst Preamble consists of a 16-symbol sequence generated by repeating the 8-symbol CAZAC sequence $1+j, 1+j, 1+j, -1+j, -1-j, 1+j, -1-j, -1+j$. In the case of TDMA downlink, a burst includes the Downlink TDMA Burst Preamble of length p physical slots, and the downlink map entry points to its beginning.

The first portion of the message after the preamble is the frame control section. The frame control section consists of a DL-MAP message followed by an UL-MAP message. The DL-MAP message contains information at which physical slots the burst profile changes. In general, the number of physical slots i (which is an integer) allocated to a particular burst can be calculated from the downlink MAP, which indicates the starting position of each burst as well as the burst profiles. In the frame control section, there may be DCD and UCD messages. The DCD message contains information about the burst profile, such as modulation type and FEC code type and parameters.

In the TDM portion, downlink data is transmitted in order of decreasing burst robustness. In the case of TDMA, data are grouped into bursts, but are not necessarily grouped in the order of robustness.

The downlink physical medium dependent layer is summarized in the block diagram in Figure 4–8.

INTEL-1020
U.S. Patent 8,219,129

The downlink channel supports adaptive burst profiling on the user data portion of the frame. Up to 12 burst profiles can be defined. The parameters of each are communicated to the SSs via MAC messages during the frame control section of the downlink frame.

Because there are optional modulation and FEC schemes that can be implemented at the subscriber station, a method for identifying the capability to the base station is required (i.e., including the highest order modulation supported, the optional FEC coding schemes supported, and the minimum shortened last codeword length supported). This information is communicated to the base station during the subscriber registration period.

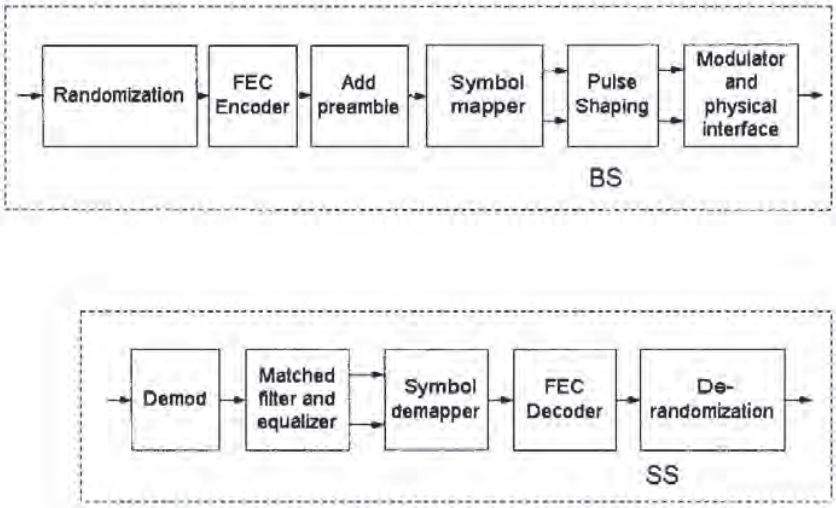

Figure 4–8: Block diagram of the downlink physical medium dependent sublayer

The operation of each block in Figure 4–8 will be described next.

The first signal-processing operation that the transmitted data is subject to is randomization. Randomization prevents long sequences of ones or zeros. An adequate number of bit transitions facilitates clock recovery. The downlink data is randomized by modulo-2 addition (alternatively, XOR) of every data bit with the output of a pseudorandom binary sequence generator, specified as a shift register. The preambles are not randomized. The randomizer sequence is applied only to information bits. The generator polynomial for the shift register is $x^{15} + x^{14} + 1$

INTEL-1020
U.S. Patent 8,219,129

and the initialization sequence is 100101010000000. The rightmost bit has weight 15, the next one has weight 14, and so on. At the beginning of each burst, the shift register is cleared and initialized. The first randomization bit is the modulo-2 addition of the two rightmost bits of the initialization sequence. Clearly, the first six bits of the randomization sequence are all zeros.

For the same business and technical reasons that were outlined in the introduction of this chapter, four forward error correction (FEC) schemes are defined in 802.16. Defining four FEC schemes allows the performance of IEEE 802.16 equipment to improve over time, together with improvements in semiconductor and signal processing technologies. Notably these improvements do not require going through the standardization process again.

The 802.16 standard defines concatenated coding schemes. Concatenated coding has been studied for a long time and is well known. Concatenated coding methods employ two constituent codes: an inner code and an outer code. The reason for using two types of codes is that in general there are two types of errors—random errors and burst errors.

In code type 1, the outer code is a systematic Reed-Solomon code over GF(256) and no inner code is used. The block length is $N = 255$ bytes, and the information block length K varies from 6 to 255 bytes. As a result, this code is able to correct errors ranging from 16 to 0 bytes. Therefore, it can be represented as RS $(255, 255 - R)$ where R is the number of redundancy bytes and ranges between 0 and 32, inclusive. The value of K is a parameter specified for each burst profile by the MAC. Code type 1 is useful either for a large data block or when high coding rate is required. This code has two modes of operation—fixed codeword and shortened last codeword operation.

In fixed codeword operation, the number of information bytes K is the same in each Reed–Solomon codeword. If the MAC messages in a burst require fewer bytes than are carried by an integral number of codewords, stuff bytes (FF in hexadecimal notation) are added so that the total message length is an integral multiple of K bytes. The SS and the BS can determine the number of codewords in its downlink burst. Using the burst length, both the SS and the BS calculate the number of full-length RS codewords that can be carried by each burst.

INTEL-1020
U.S. Patent 8,219,129

In the shortened last codeword mode of operation, the number of information bytes in the final Reed–Solomon block of each burst is reduced from the normal number K, while the number of parity bytes R remains the same. The BS reduces the number of information bytes in the last codeword to minimize the number of stuff bytes to add to the end of the MAC message. The length of the burst is then set to the minimum number of physical slots required to transport all of the burst's bytes, which include preamble, information, and parity bytes. The BS implicitly communicates the number of bytes in the shortened last codeword to the SS via the Downlink Map message, which defines the starting physical slot of each burst. The SS uses the Downlink Map information to calculate the number of full-length RS codewords and the length of the shortened last codeword that can be carried within the specified burst size. The BS performs a similar calculation as the SS for its encoding purposes. To allow the receiving hardware to decode the previous Reed–Solomon codeword, in this mode of operation a codeword will have at least 6 information bytes. The number of information bytes carried by the shortened last codeword is between 6 and K bytes, inclusive. Only if the number of information bytes to be sent by the BS is less than 6 bytes will stuff bytes be added at the end.

In code type 2 the outer code is the same RS code as in code type 1, with one minor restriction. When using code type 2 in the shortened last codeword mode of operation, the number of information bytes in the shortened last codeword must be an even number. Consequently, if an odd number of information bytes must be sent, a stuff byte is added because the inner code requires a pair of bytes on which to operate.

The inner code is a (24,16) block convolutional code (BCC). This inner code in code type 2 can be considered as an equivalent nonsystematic block code. The input to the block code is 16 bits, and the output (the codeword) is 24 bits long. The encoder of this equivalent block code can be described by the following equations:

$$c23 = b15 \oplus b0 \oplus b1$$

$$c22 = b15 \oplus b1$$

$$c21 = b14 \oplus b15 \oplus b0$$

$$c20 = b13 \oplus b14 \oplus b15$$

INTEL-1020
U.S. Patent 8,219,129

$c19 = b13 \oplus b15$

$c18 = b12 \oplus b13 \oplus b14$

$c17 = b11 \oplus b12 \oplus b13$

$c16 = b11 \oplus b13$

$c15 = b10 \oplus b11 \oplus b12$

$c14 = b9 \oplus b10 \oplus b11$

$c13 = b9 \oplus b11$

$c12 = b8 \oplus b9 \oplus b10$

$c11 = b7 \oplus b8 \oplus b9$

$c10 = b7 \oplus b9$

$c9 = b6 \oplus b7 \oplus b8$

$c8 = b5 \oplus b6 \oplus b7$

$c7 = b5 \oplus b7$

$c6 = b4 \oplus b5 \oplus b6$

$c5 = b3 \oplus b4 \oplus b5$

$c4 = b3 \oplus b5$

$c3 = b2 \oplus b3 \oplus b4$

$c2 = b1 \oplus b2 \oplus b3$

$c1 = b1 \oplus b3$

$c0 = b0 \oplus b1 \oplus b2.$

Code type 2 is useful for low-to-moderate coding rates that provide good performance.

While code types 1 and 2 must be implemented by all BSs and SSs, code types 3 and 4 are optional. In code type 3, the outer code is still the same, but the inner code is (9, 8) parity check code. The code itself is a simple bit-wise parity check.

INTEL-1020
U.S. Patent 8,219,129

This code adds one parity bit to every eight bits. The parity bit is the sum modulo 2 of all eight bits.

In code type 4, the outer code is a block turbo code (BTC), and no inner code is used. A block turbo code has the properties of a concatenated code. Code type 4 has significant advantages. It can be used to either extend the range of a base station or increase the data rate at the same range. The BTC idea is illustrated in Figure 4–9.

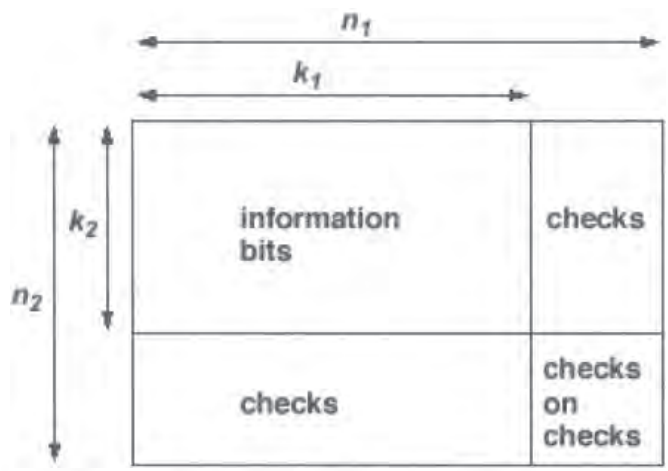

Figure 4–9: Two-dimensional product code matrix

Essentially, the idea of turbo codes is to encode the data twice. First, the k_1 information bits in the rows are encoded into n_1 bits, by using an (n_1,k_1) block code. Then the k_2 information bits in the columns are encoded into n_2 bits by using the same or possibly different block code. The parity bits of the first code are also encoded by the second code. The overall block size of such a product code is $n = n_1 \times n_2$, the total number of information bits is $k = k_1 \times k_2$, and the code rate is k/n. As a constituent block coding scheme the standard specifies (31,26) and (63,57) extended Hamming block codes. Shorter codes can be obtained by deleting an appropriate number of rows and columns.

When using the Block Turbo Coding, bit interleaving can be used. At present, two modes of interleaving are defined. The first mode of interleaving, type 1, means actually no interleaving. In this mode, the encoded bits are read from the encoder row by row, in the order that they were written. Type 2 interleaver is a block interleaver. In this mode, the encoded bits are read from the encoder after the first

INTEL-1020
U.S. Patent 8,219,129

k_2 rows are written into the encoder memory. The bits are read column by column, proceeding from the top position in the first column. Currently, a type 3 interleaver is not defined. It is possible that in the future an interleaver that has advantages in terms of performance and/or implementation can be defined.

It is important that the frame control section of the downlink frame be encoded with a fixed set of parameters known to the SS at initialization in order to ensure that all subscriber stations can read the information. The control portion of the frame is encoded with a type 2 FEC, where the outer code is a (46, 26) Reed–Solomon code and an inner code is a (24, 16) BCC. Following FEC, QPSK modulation is used.

In communication systems, digital modulation specifies how bits are mapped to symbols. The downlink physical layer of IEEE 802.16 supports several modulation formats. The constellation can be selected per subscriber based on the quality of the RF channel. In a high SNR environment, a more complex modulation scheme can be used. If the air-link degrades, the system can revert to the less complex constellations. The goal is to maintain a constant error rate. In the downlink, the mandatory modulation formats are QPSK and 16-QAM. 64-QAM is an option. Figures 4–10 through 4–12 describe the bit mappings for QPSK, 16-QAM, and 64-QAM modulation.

The distance from the origin determines the power that sends the signal. There are two power adjustment rules—constant constellation peak power and constant constellation mean power. In the constant peak power scheme, corner points are transmitted at equal power levels regardless of modulation type. In the constant mean power scheme, the signal is transmitted at equal mean power levels regardless of modulation type. The power adjustment rule is specified in the DCD management message.

Prior modulation, the I and Q signals are filtered by industry-standard square-root raised cosine pulse-shaping filters with excess bandwidth factor of 0.25. Finally, at the antenna port the transmitted waveform is given by:

$$S(t) = I(t)\cos(2\pi f_c t) - Q(t)\sin(2\pi f_c t) \qquad \text{Eq. 4–5}$$

where $I(t)$ and $Q(t)$ are the signals after pulse-shaping, and f_c is the RF carrier frequency.

INTEL-1020
U.S. Patent 8,219,129

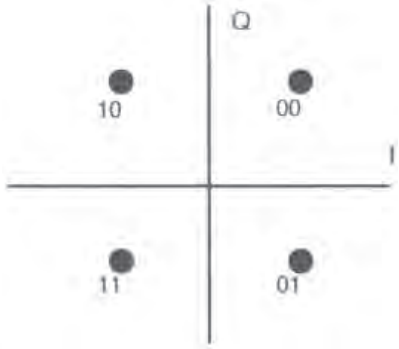

Figure 4–10: Bit-to-symbol mapping for QPSK in the downlink

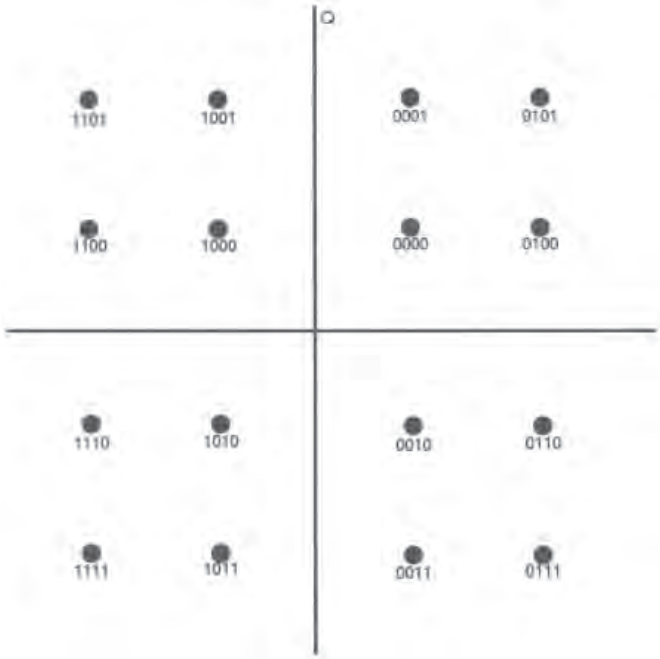

Figure 4–11: Gray-coded 16-QAM in the downlink

INTEL-1020
U.S. Patent 8,219,129

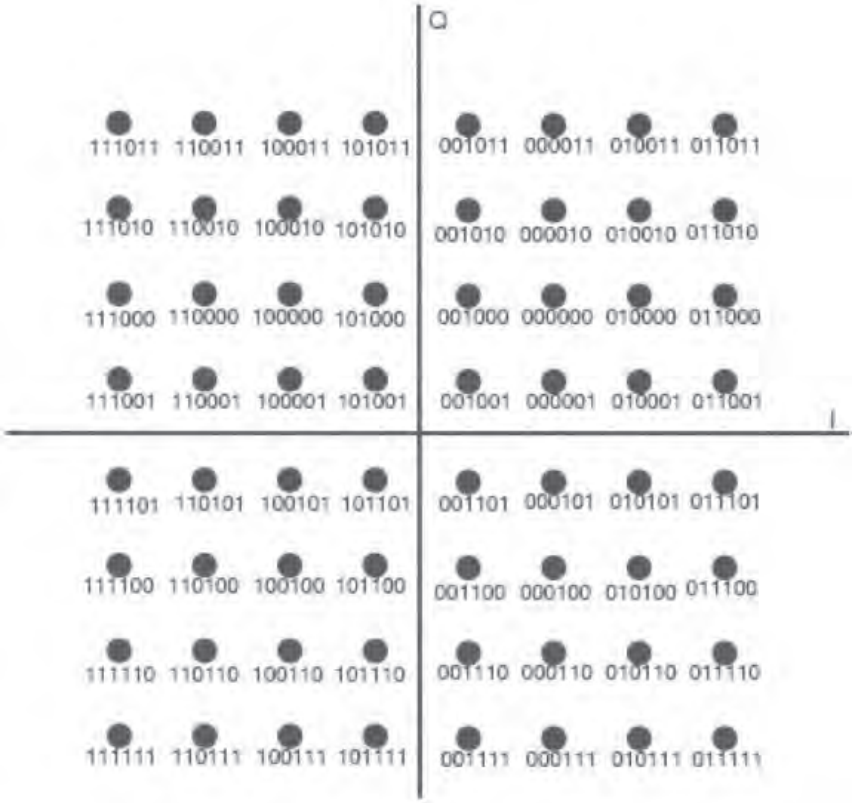

Figure 4–12: Gray-coded bit-to-symbol mapping for the optional 64 QAM in the downlink

Uplink physical layer

Three classes of bursts may be transmitted by the SS during the uplink subframe—bursts transmitted in contention opportunities reserved for initial maintenance, bursts transmitted in contention opportunities provided by multicast and broadcast polls, and bursts transmitted in intervals specifically allocated to individual SSs. One frame can contain any of these bursts. Bursts can occur in any order and in any quantity (limited by the number of available physical slots) within the frame, at the discretion of the BS uplink scheduler. All uplink transmissions are made according to the uplink burst profiles, specified by the BS, as indicated in the UL-MAP message in the frame control section of the downlink.

INTEL-1020
U.S. Patent 8,219,129

The burst profiles include information such as burst profile, FEC type, and preamble length.

SS transition gaps separate the transmissions of the various SSs during the uplink subframe. The gap allows some finite time for the previous burst to subside, followed by a preamble allowing the BS to synchronize to the new SS. The preamble and gap lengths are broadcast periodically in the UCD message.

Each uplink burst begins with an uplink preamble. This preamble is based upon a repetition of a +45 degrees rotated constant amplitude zero autocorrelation (CAZAC) sequence. Two preambles can be used—one of 16-symbol length or one of 32-symbol length. The base station defines the preamble length and communicates it to the SS in the UCD message. The 16-symbol preamble is formed by repeating twice the following CAZAC sequence:

$$\{1+j, -1+j, 1+j, 1+j, -1-j, -1+j, -1-j, 1+j\}.$$

The 32-symbol preamble is formed by repeating twice

$$\{1+j, 1+j, 1-j, -1-j, -1-j, 1+j, 1+j, -1+j, 1+j, 1+j, -1+j, 1+j, -1-j, 1+j, -1-j, 1-j\}$$

These are QPSK sequences. The amplitude of the preamble depends on the uplink power adjustment rule—constant peak power or constant mean power. In the case of constant peak power, the constellation points of the preamble coincide with the outermost constellation points of the modulation scheme in use. In the case of the constant mean power, the mean power of the constellation points of the preamble is the same as the mean power of the modulation scheme in use.

The uplink transmission convergence sublayer is identical to the downlink transmission convergence sublayer. The uplink physical medium dependent layer coding and modulation are summarized in the block diagram shown in Figure 4–13.

Randomization and FEC encoding in the uplink are identical to the corresponding operations in the downlink.

The type of modulation and the power adjustment rule are set by the base station and are communicated in the UCD message. QPSK is mandatory, and the 16-QAM and 64-QAM are optional. The mappings of bits to symbols are identical

INTEL-1020
U.S. Patent 8,219,129

to those in the downlink. An SS should not produce an EIRP spectral density exceeding either +30 dBW/MHz or the local regulatory requirements.

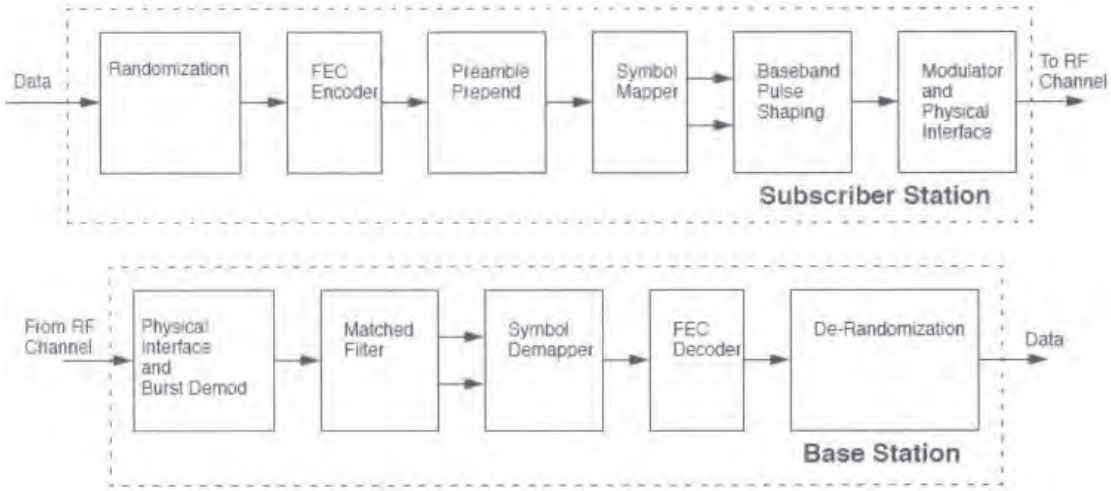

Figure 4–13: Uplink PHY

The band between 10 GHz and 66 GHz is very wide, and regulatory requirements vary significantly between different regions. Due to wide variations in local regulations, no frequency plan is specified in 802.16. No single plan can accommodate all cases. To ensure that products built to 802.16 have interoperability over the air interface, the parameters in Table 4–2 should be followed. Systems shall use Nyquist square-root raised cosine pulse shaping with a roll-off factor of 0.25. Note that baud rates are chosen to provide an integer number of physical slots per frame. A frame duration of 1 ms is recommended as a compromise between transport efficiency (with lower frame overhead) and latency.

Note that for 802.16, synchronization and carrier frequency offset estimation and compensation, among the subscriber stations and the BS are more important than in the other wireless standards. Following the centralized architecture, the BS will make periodic measurements of the carrier frequency offset and will communicate it to the subscriber stations by a MAC message.

INTEL-1020
U.S. Patent 8,219,129

Table 4–2: Baud rates and channel sizes for a roll-off factor of 0.25

Channel size (MHz)	Symbol rate (Mbaud)	Bit rate (Mb/s) QPSK	Bit rate (Mb/s) 16 QAM	Bit rate (Mb/s) 64 QAM
20	16	32	64	96
25	20	40	80	120
28	22.4	44.8	89.6	134.4

An important part of the physical layer that is left to vendors is the power control algorithm. To provide power control in the uplink, the base station must be capable of providing accurate power measurements of the received burst signal. Then the BS will send information about the power level back to the subscriber stations in calibration messages. The power control algorithm should also take into account various power fluctuations due to the characteristics of the transmission medium, the performance of the RF power amplifier in different burst profiles, and other factors. The exact implementation of power control is vendor-specific.

Physical layers for 2–11 GHz operation

The physical layer in the previous section is for the 10–66 GHz band. Operating in licensed bands requires a license, which is very expensive, if at all possible to obtain. IEEE 802.16a specifies a physical layer for operation at frequencies between 2 GHz and 11 GHz. Clearly the band between 2 GHz and 11 GHz also includes unlicensed bands. The protocol for operation in the unlicensed frequencies is called WirelessHUMAN. The 2–11 GHz bands provide a physical environment where, due to the longer wavelength, line of sight is not necessary and multipath may be significant. The channel bandwidths used in this physical environment typically vary from 1.5 MHz to 14 MHz. The ability to support non-line-of-sight requires additional PHY functionality, such as the support of advanced power management techniques, interference mitigation/coexistence, and smart antennas.

Both single-carrier and multiple-carrier modulation can be used. In general, the goal of a modulation technique is to transfer data over a prescribed channel bandwidth, within transmit power, reliability, and receiver complexity constraints.

INTEL-1020
U.S. Patent 8,219,129

Single-carrier systems require an equalizer at the receiver to compensate for any distortion resulting from the channel, because the multipath creates intersymbol interference among the received symbols. This equalization can be performed in time, frequency, and mixed (time and frequency) domain. In multicarrier systems, data is mapped to symbols and then multiplexed using an orthogonal transformation into a number of simultaneous lower-speed streams. The discretetime signal after the orthogonal transformation is converted to an analog waveform and mixed to a RF frequency. The nonlinear distortions introduced as part of the RF conversion can create significant out-of-band interference and must be reviewed in the context of:

- Deployment in specific frequency bands and out-of-band requirements,
- Coexistence with adjacent (in-band and out-of-band) systems (in particular TDD/FDD),
- Trade-offs in terms of guard bands, system performance, and system complexity.

In multicarrier systems, the total available channel bandwidth is subdivided among these multiple streams. Although the frequency response over the entire channel bandwidth may not be ideal (i.e., nonconstant), the spacing between the modulated carriers is small, so that the frequency response over the signaling bandwidth of one carrier is approximately constant. Equalization for multicarrier systems is still required, but it becomes a simple multiplication per carrier.

IEEE 802.16's specification for unlicensed bands (known as WirelessHUMAN for Wireless High-Speed Unlicensed Metropolitan Area Networks) is specifically designed for operation in the 5 GHz band, but it can be applicable to other license-exempt bands in the 2–11 GHz range. In the U.S. and Canada, the 5.15–5.25 GHz lower U-NII band is restricted to indoor use only. Because in BWA the BS is outdoors, this band is not available for BWA.

Operation in the unlicensed bands introduces additional interference and coexistence issues. Government regulations limit the allowed radiated power. IEEE 802.16a introduces mechanisms to detect and avoid interference such as dynamic frequency selection (DFS). Here again, DFS means the ability of a system to switch to different physical RF channels based on channel measurement criteria. DFS is a feature that all wireless technologies operating in unlicensed bands must have.

INTEL-1020
U.S. Patent 8,219,129

All frequency layers between 2 GHz and 11 GHz have some common characteristics. Frequency division duplex (FDD), half duplex frequency division (H-FDD), and time division duplex (TDD) modes provide for bidirectional operation, except for operation in the license-exempt band, where provision is made for TDD operation only. TDD flexibility permits efficient allocation of the available bandwidth and hence is capable of efficiently allocating the available traffic transport capacity for applications with varying ratios of uplink-to-downlink traffic transport demand. TDD operates in single, paired, or noncontiguous blocks of frequencies. FDD/H-FDD can be used by applications that require fixed asymmetric allocation between their uplink and downlink traffic transport demand. FDD and H-FDD operate in paired downlink/uplink sub-bands.

Table 4–3 gives a summary of the different BWA systems.

Table 4–3: BWA systems

Standard	Targeted frequency band	PHY	MAC	Duplexing
WirelessMAN-SC	10–66 GHz	SC	Basic	TDD, FDD
WirelessMAN-SC2	2–11 GHz	SC2	Basic+ARQ+STC+AAS	TDD, FDD
WirelessMAN-OFDM	2–11 GHz licensed	OFDM	Basic+ARQ+STC+DFS+AAS	TDD, FDD
WirelessMAN-OFDM	2–11 GHz unlicensed	OFDM	Basic+ARQ+STC+DFS+mesh+AAS	TDD
WirelessMAN-OFDMA	2–11 GHz licensed	OFDMA	Basic+ARQ+STC+DFS+AAS	TDD, FDD
WirelessMAN-OFDMA	2–11 GHz unlicensed	OFDMA	Basic+ARQ+STC+DFS+mesh+AAS	TDD

INTEL-1020
U.S. Patent 8,219,129

WirelessMAN-SC2 Physical layer for 2–11 GHz frequencies

This physical layer is based on single-carrier technology and designed for NLOS operation in the 2–11 GHz band. To enable NLOS operation, this physical layer supports channel estimation and equalization. The main characteristics of this physical layer are:

- TDMA uplink
- TDM downlink
- Concatenated FEC using Reed-Solomon and pragmatic TCM codes
- Optional BTC and convolutional TC
- No FEC option using ARQ, optional STC
- Optional AAS support

The WirelessMAN-SC2 is similar to the 10–66 GHz physical layer. As shown in Figure 4–14, the first step is randomization (or scrambling), performed as in WirelessMAN-SC. Data randomization is performed on the source bits, before FEC encoding, on both the downlink (DL) and uplink (UL).

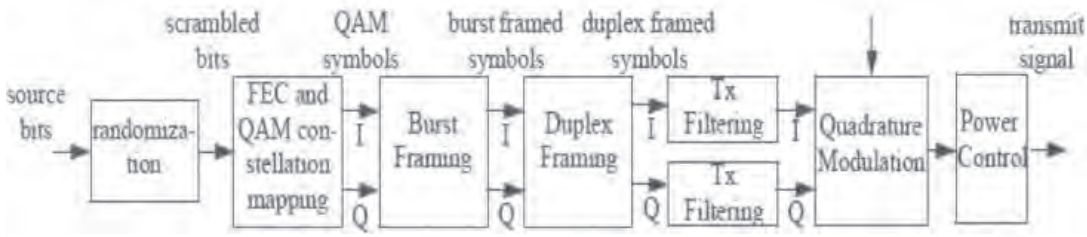

Figure 4–14: WirelessMAN-SC2 transmitter signal processing

The bits from the randomizer are passed to the channel encoder. The FEC mechanism is different from WirelessMAN-SC, because it is designed for a multipath environment. It is well-known that the performance of codes in channels with frequency selective fading is different from the performance in AWGN channels. Note that in general the amount of coding gain required may differ between uplink and downlink due to their different propagation characteristics.

INTEL-1020
U.S. Patent 8,219,129

In WirelessMAN-SC2, a concatenated FEC scheme is used, based on the serial concatenation of a Reed-Solomon outer code and a rate-compatible TCM inner code, as shown in Figure 4–15. Between the outer and the inner code, byte interleaving can be used as an option. The outer code is RS (255, 239) over GF (256). This code can be shortened and punctured to enable variable block sizes and variable error-correction capability. Shortening is when the number of information bytes that are encoded, K, is less than 239. In this case, the first $239-K$ data bytes of the block are set to zero but are not transmitted. When the data cannot be divided into an integer number of K-byte blocks, a shortened RS code will be used for the last block. The capability to shorten the code is mandatory, while puncturing is optional. Puncturing is when the code has fewer than 16 parity bytes. Puncturing may be used, depending on the burst profile.

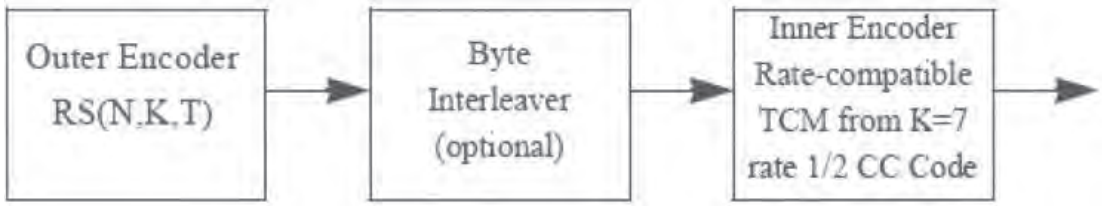

Figure 4–15: Concatenated coding scheme

The inner code is a rate-compatible pragmatic TCM code, derived from the rate 1/2 constraint length $K = 7$, binary convolutional code, which is identical to the convolutional code in IEEE 802.11a, with generator polynomials designated by 171 and 133 in octal notation. In addition to rate 1/2, this convolutional code can also support other rates through puncturing. Through puncturing, rates of 2/3, Ω, 5/6, and 7/8 can be obtained. Table 4–4 summarizes the mandatory and optional modulations and coding rates for WirelessMAN-SC2.

The pragmatic TCM code is constructed from both nonsystematic coded bits taken from the rate 1/2 binary convolutional encoder outputs and systematic uncoded bits taken directly from the encoder input. The resulting coded bits are then mapped to symbol constellations. The choice of a particular code rate and modulation is made based on burst profile parameters.

INTEL-1020
U.S. Patent 8,219,129

Table 4–4: Mandatory (M) and optional (O) modulations and coding rates for WirelessMAN-SC2

Modulation	Supported coding rates	Support	
		Uplink	Downlink
BPSK	1/2, 3/4	M	O
QPSK	1/2,2/3,3/4,5/6,7/8	M	M
16-QAM	1/2, 3/4	M	M
64-QAM	2/3, 5/6	O	M
256-QAM	3/4, 7/8	O	O

Interleaving between the inner and outer code is optional. Interleaving is not used in broadcast burst profiles. When interleaving is used, its usage and parameters will be specified within a burst profile. The interleaver changes the order of bytes from the Reed Solomon (RS) encoder output. Interleaving is used to spread consecutive bits into separate symbols after modulation. In this way, the interleaver prevents a series of consecutive bad symbols, which may occur due to multipath. A de-interleaver in the receiver restores the order of the bytes before RS decoding. The interleaver is a block interleaver, where a table is written a byte at a time row-wise (one row per RS code word) and read a byte at a time column-wise. The number of columns is equal to the codeword length of the Reed–Solomon code. The number of rows used by the interleaver is a burst parameter and is equal to the number of RS codewords that the interleaver can store. So that bursts exceeding an intended receiver's capabilities are not generated, the largest number of rows supported by a station is communicated during SS basic capability negotiation. The product of the number of rows and the number of columns is the interleaver block size.

The bits exiting the inner code (after puncturing, if it is used) are mapped to I and Q symbol coordinates using either a pragmatic TCM or Gray code symbol map, depending on the code rate and modulation scheme. All BPSK and QPSK code rates and rate 1/2 16-QAM shall use the Gray-coded constellation maps depicted in Figure 4–22 on page 297. Rate 3/4 16-QAM and all code rates for 64-QAM use the pragmatic TCM constellation map depicted in Figure 4–23 on page 297. All

INTEL-1020
U.S. Patent 8,219,129

code rates for 256-QAM use the pragmatic TCM constellation map depicted in Figure 4–24 on page 298.

Figure 4–16 illustrates the rate 1/2 pragmatic TCM encoder for 16-QAM. The baseline rate 1/2 binary convolutional encoder first generates a two-bit constellation index, b_3b_2, associated with the *I* symbol coordinate. Then it generates a two-bit constellation index, b_1b_0, for the *Q* symbol coordinate. The *I* index generation precedes the *Q* index generation. Note that this encoder should be interpreted as a rate 2/4 encoder, because it generates one four-bit code symbol per two input bits. For this reason, input bit sequence that feeds this encoder must have length divisible by two.

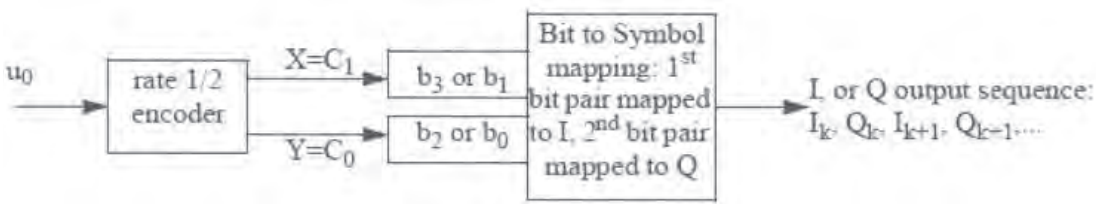

Figure 4–16: TCM encoder for rate 1/2 16 QAM

Figure 4–17 illustrates the rate 3/4 pragmatic TCM encoder for 16-QAM. This encoder uses the same rate 1/2 binary convolutional encoder, along with two systematic bits that are passed directly from the encoder input to the encoder output. With this structure, the encoder is capable of simultaneously generating four output bits per three input bits. The sequence of arrival for the $u_2u_1u_0$ input into the encoder is u_2 arrives first, u_1 second, u_0 last. During the encoding process, the encoder generates a two-bit constellation index, (b_3b_2, for the *I* symbol coordinate) and simultaneously generates another two bit constellation index, designated b_1b_0, for the *Q* symbol coordinate. The length of the input bit sequence must be divisible by three. Figure 4–23 depicts the bits-to-symbol-constellation map that is applied to the rate 3/4 16-QAM encoder output.

INTEL-1020
U.S. Patent 8,219,129

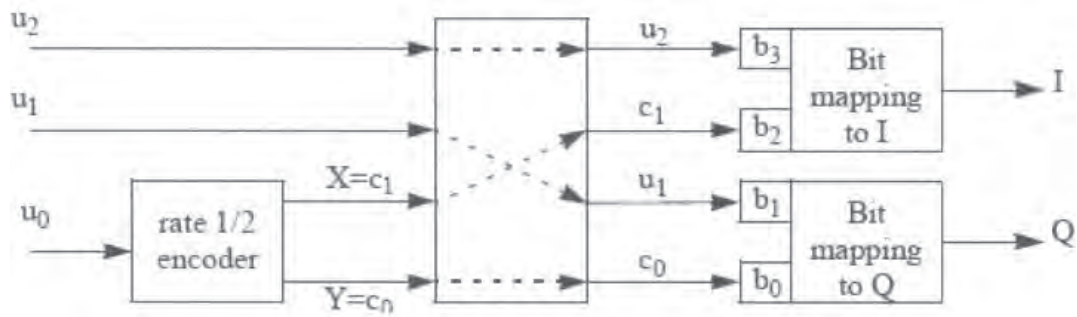

Figure 4–17: Rate 3/4 pragmatic TCM encoder for 16 QAM

Figure 4–18 illustrates the rate 2/3 pragmatic TCM encoder for 64-QAM. This encoder uses again the same rate 1/2 binary convolutional encoder, along with one systematic bit that is passed directly from the encoder input to the encoder output. The sequence of arrival for the u_1u_0 input into the encoder is u_1 arrives first, u_0 last. The encoder (as a whole) then generates a three bit constellation index, $b_5b_4b_3$, which is associated with the *I* symbol coordinate. From the next two bits, the encoder generates another three-bit constellation index, $b_2b_1b_0$, which is associated with the *Q* symbol coordinate. The *I* index generation should precede the *Q* index generation. Figure 4–23 depicts the bits-to-symbol-constellation map that is applied to the rate 2/3 64-QAM encoder output.

> This encoder should be interpreted as a rate 4/6 encoder, because it generates one six-bit code symbol per four input bits. For this reason, the length of the input bit sequence must be divisible by four.

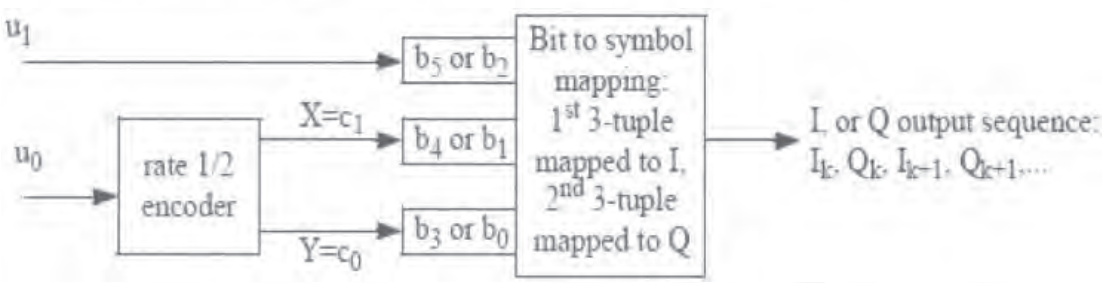

Figure 4–18: Pragmatic TCM encoder for rate 2/3 64 QAM

INTEL-1020
U.S. Patent 8,219,129

Figure 4–19 illustrates the rate 5/6 pragmatic TCM encoder for 64-QAM. This encoder uses a rate 3/4 punctured version of the rate 1/2 binary convolutional encoder, along with two systematic bits that are passed directly from the encoder input to the encoder output. The rate Ω punctured code is generated from the baseline rate 1/2 code using an appropriate puncture mask. Puncture samples are sequenced c_3 first, c_2 second, c_1 third, and c_0 last. The sequence of arrival for the $u_4u_3u_2u_1u_0$ input into the encoder is u_4 arrives first, u_3 arrives second, u_2 arrives third, u_1 arrives next to last, and u_0 arrives last. During the encoding process, the pragmatic encoder generates a three-bit constellation index, $b_5b_4b_3$, for the *I* symbol coordinate, and simultaneously generates another three-bit constellation index, $b_2b_1b_0$, for the *Q* symbol coordinate. Note that the length of the input bit sequence must be divisible by five. The bits-to-symbol-constellation map that is applied to the rate 5/6 64-QAM encoder output is shown in Figure 4–23 on page 297.

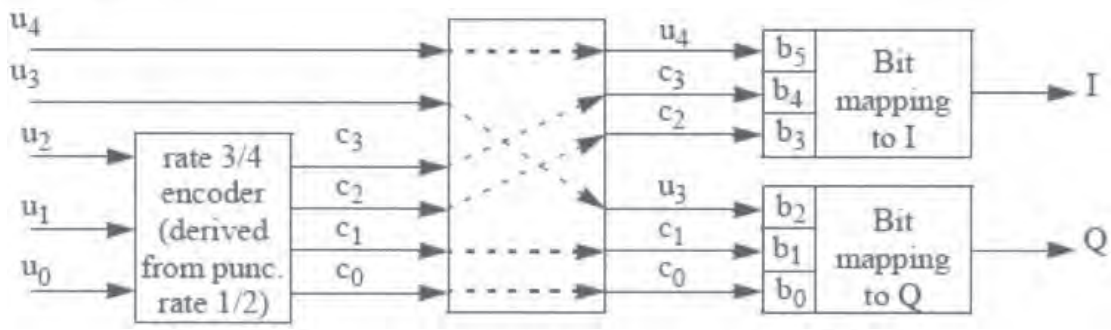

Figure 4–19: TCM encoder for rate 5/6 64 QAM

Figure 4–20 illustrates the rate 3/4 pragmatic TCM encoder for 256-QAM. This encoder uses the rate 1/2 binary convolutional encoder, along with two systematic bits that are passed directly from the encoder input to the encoder output. The sequence of arrival for the $u_2u_1u_0$ input into the encoder is u_2 arrives first, u_1 next, u_0 last. Note that the encoder first generates a four-bit constellation index, $b_7b_6b_5b_4$, which is associated with the *I* symbol coordinate. Provided another four bit encoder input, it generates a four-bit constellation index, $b_3b_2b_1b_0$, which is associated with the *Q* symbol coordinate. The *I* index generation precedes the *Q* index generation. Note that this encoder should be interpreted as a rate 6/8 encoder, because it generates one eight-bit code symbol per six input bits. For this reason, the length of the input bit sequence should be divisible by six. Figure 4–24

INTEL-1020
U.S. Patent 8,219,129

(on page 298) depicts the bits-to-symbol-constellation map that is applied to the rate 3/4 256-QAM encoder output.

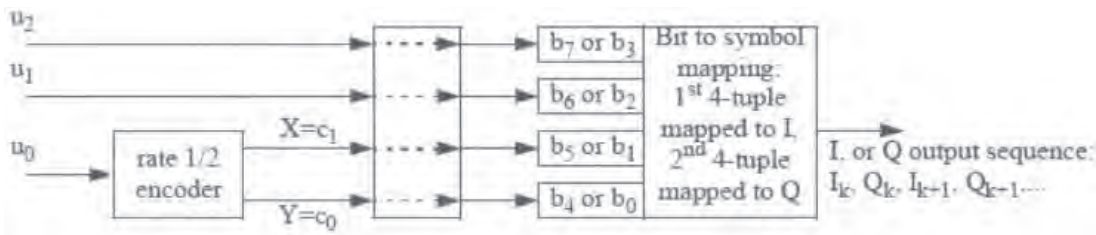

Figure 4–20: Optional TCM encoder for rate 3/4 256 QAM

Next, Figure 4–21 illustrates the rate 7/8 pragmatic TCM encoder for 256-QAM. This encoder uses a rate 3/4 punctured version of the baseline rate 1/2 binary convolutional encoder, along with two systematic bits that are passed directly from the encoder input to the encoder output. The rate 3/4 punctured code is generated from the baseline rate 1/2 code using an appropriate puncture mask. Puncture samples are sequenced c_3 first, c_2 second, c_1 third, and c_0 last. The sequence of arrival for the $u_6u_5u_4u_3u_2u_1u_0$ input into the encoder (as a whole) is u_6 arrives first, u_5 arrives second, u_4 arrives third, u_3 arrives fourth, u_2 arrives fifth, u_1 arrives next to last, and u_0 arrives last. During the encoding process, the encoder generates a four-bit constellation index, $b_7b_6b_5b_4$, for the *I* symbol coordinate, and simultaneously generates another four-bit constellation index, $b_3b_2b_1b_0$, for the *Q* symbol coordinate. The length of the input bit sequence must be divisible by seven. Figure 4–24 depicts the bits-to-symbol-constellation map that is applied to the rate 7/8 256-QAM encoder output.

To make the average transmitted power identical and equal to 1 for all modulation formats, the constellation points are multiplied by an appropriate constant. The QAM symbols are next framed within a message burst, which typically introduces additional framing symbols. The burst symbols are then multiplexed into a duplex frame, which may contain multiple bursts. The *I* and *Q* components are injected into pulse-shaping filters, quadrature modulate a certain carrier frequency, and amplified with power control so that the proper output power is transmitted.

INTEL-1020
U.S. Patent 8,219,129

The WirelessMAN-SC2 standard envisions the possibility that FEC would not be necessary if ARQ is enabled. In this case, the bits after scrambling are directly mapped to a QAM constellation using appropriate Gray coding. In this case, for BPSK, QPSK, and 16 QAM, the mapping in Figure 4–22 is used, but for 64 QAM and 256-QAM the mappings in Figure 4–25 and Figure 4–26 are used. No FEC operation is mandatory for QPSK; no FEC is optional for other modulation methods.

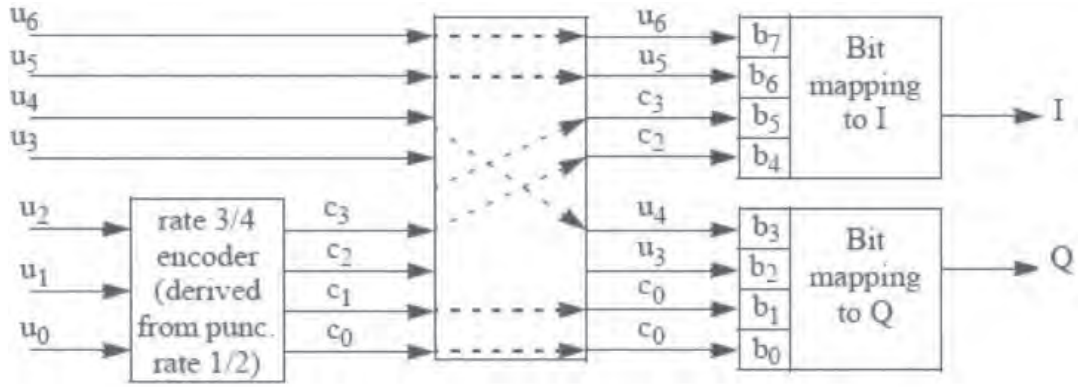

Figure 4–21: Optional rate 7/8 TCM encoder for 256 QAM

INTEL-1020
U.S. Patent 8,219,129

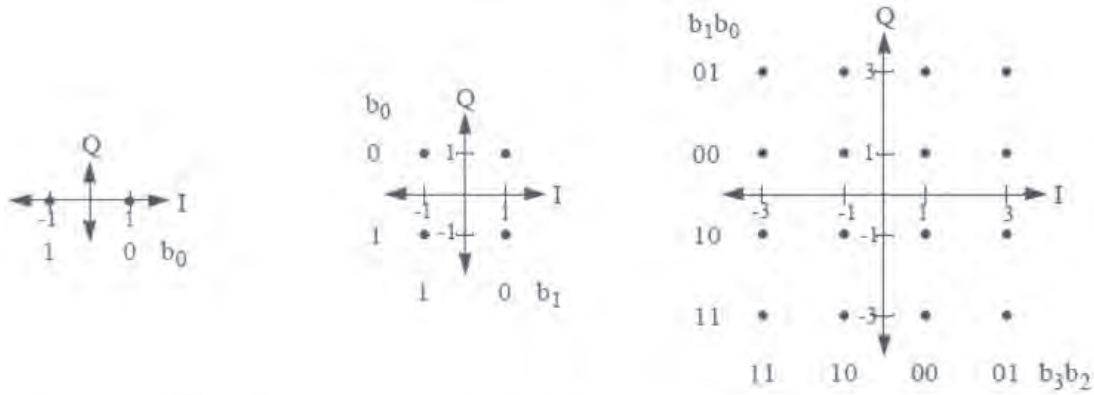

Figure 4–22: Gray maps for BPSK, QPSK, and 16 QAM constellations

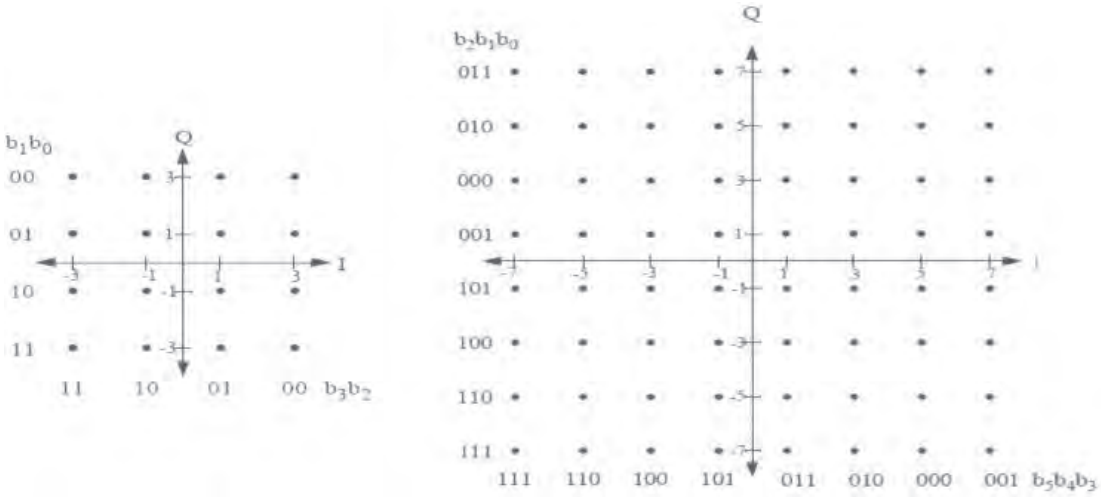

Figure 4–23: Gray-coded symbol maps for 1/2 16 QAM and 64 QAM

INTEL-1020
U.S. Patent 8,219,129

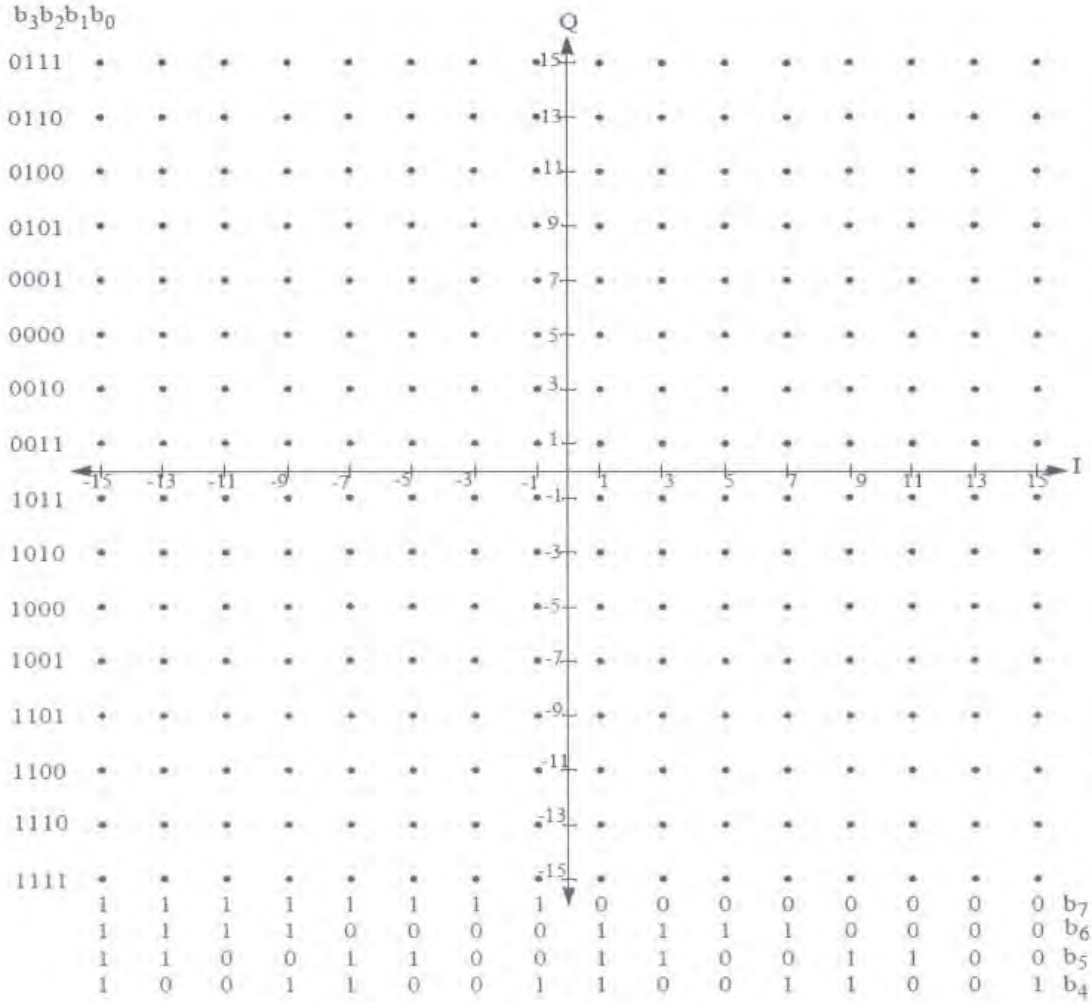

Figure 4–24: Bit-to-symbol mapping map for 256 QAM

INTEL-1020
U.S. Patent 8,219,129

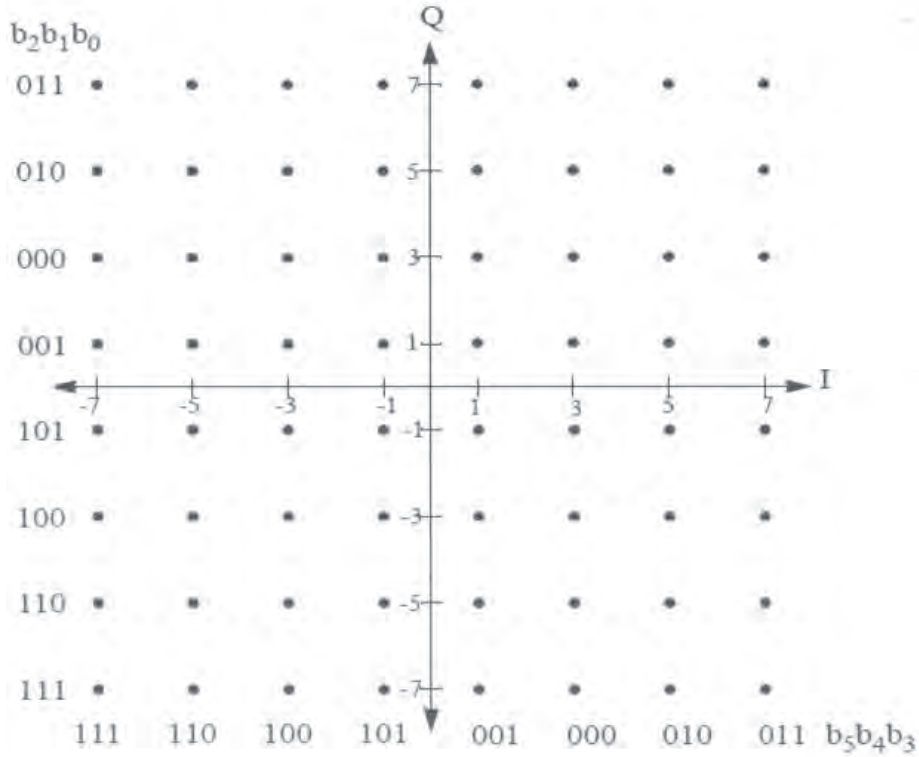

Figure 4–25: Gray coding for 64 QAM when no FEC is used

INTEL-1020
U.S. Patent 8,219,129

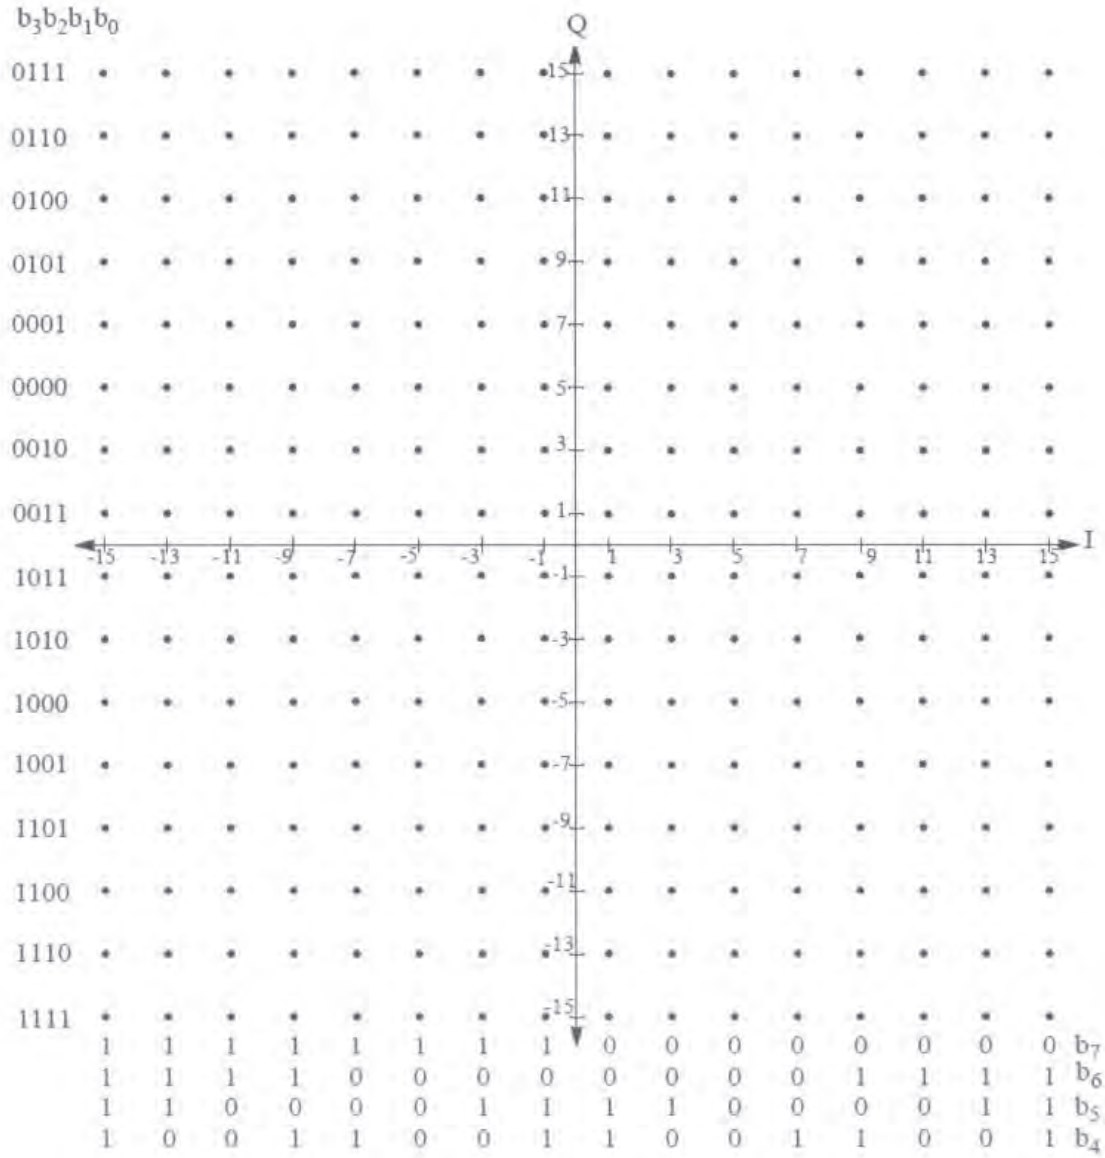

Figure 4–26: Gray coding for 256 QAM when no FEC is used

To achieve higher coding gain, optional turbo product coding (TPC) can be used for all modes. The constituent codes in the turbo coding case are extended Hamming codes or parity check codes. Many soft decoding algorithms can be used for turbo decoding. To match an arbitrary required packet size, BTCs can be

INTEL-1020
U.S. Patent 8,219,129

shortened by removing symbols from the BTC array. Rows, columns, or parts thereof can be removed until the appropriate size is reached. For transmission, an encoded BTC block shall be read out of the encoded matrix as a serial bit stream, starting with the least significant bit and ending with the most significant bit. This bit stream is sent to a symbol mapper the same as the symbol mappers in the no-FEC case. If not enough encoded bits are available to fill the last symbol of an allocation, sufficient zero-valued bits (unscrambled) will be appended to the end-of-the-serial stream to complete the symbol.

Another option in 802.16 is using convolutional turbo codes. The convolutional turbo code encoder and its constituent encoder are depicted in Figure 4–27 and Figure 4–28, respectively. It uses a double binary circular recursive systematic convolutional code.

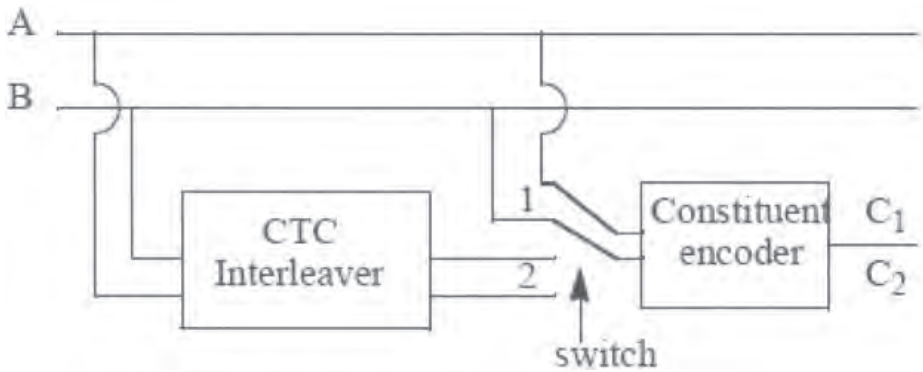

Figure 4–27: A convolutional turbo encoder

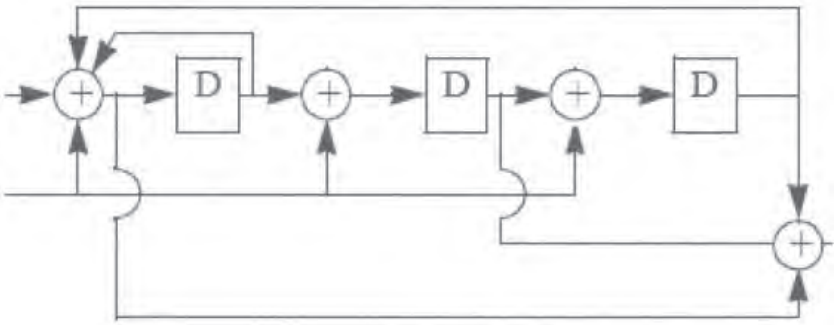

Figure 4–28: A constituent convolutional encoder

INTEL-1020
U.S. Patent 8,219,129

During the encoding process, the bits of the data are alternately fed to A and B, starting with the msb of the first byte being fed to A. The encoder is fed by blocks of k bits or N couples ($k = 2N$ bits). For all frame sizes, N is a multiple of 4 and limited between 8 and 256. After initialization, the encoder is fed the sequence in the natural order (position 1) with the incremental address $i = 0,...,N - 1$. This first encoding is called C1 encoding. Then the encoder is initialized differently and is fed by the interleaved sequence (switch in position 2) with incremental address $j = 0,\ldots,N - 1$. This second encoding is called C2 encoding. Code-rates of 1/2, 2/3, 3/4, 5/6, and 7/8. are achieved by selectively deleting the parity bits (puncturing). The puncturing patterns are identical for both codes C1 and C2.

Both DL and UL data are formatted into bursts that use the framed burst format. The DL supports the most general case of TDM bursts, while the UL supports TDMA bursts. TDMA burst and continuous DL operational modes are subclasses of the TDM burst DL mode. The burst frame consists of a burst preamble and burst payload. Pilot words can be inserted in the payload. A pilot word is a unique word, known to the receiver, and used to help signal processing operations at the receiver. The length in symbols, U, of a unique word (UW) is a burst profile parameter. For best performance, U should be at least as long as the intended channel's delay spread. 802.16 specifies that unique words have lengths at powers of two. These unique words have I and Q components, where

$$I[n] = \cos(\theta[n]) \qquad \text{Eq. 4–6}$$

$$Q[n] = \sin(\theta[n]) \qquad \text{Eq. 4–7}$$

When the length L is 16, 64, or 256, the unique words are derived from Frank–Zadoff sequences, $\theta_{FZ}[n] = (2\pi pq)/\sqrt{L}$. It is clear that when the length is 16, the unique word consists of 4-PSK (equivalent to QPSK) symbols, and when the length is 64 or 256, the unique word consists of 8-PSK and 16-PSK symbols respectively. The sequence length of 64 is mandatory. Length of 16 is also mandatory for symbol rates below 1.25 Msymb/s, and a length of 256 is mandatory for symbol rates above 20 Msymb/s.

Optionally, unique words with lengths of 8, 32, 128, or 512 derived from Chu sequences can also be used, $\theta_{Chu}[n] = \pi n^2/L$. All of these unique words are chosen

INTEL-1020
U.S. Patent 8,219,129

because they can be easily detected due to their constant amplitude zero autocorrelation (CAZAC) properties [B108].

A downlink burst may contain time division multiplexed messages that are adaptively modulated for the intended message recipients. When a MAC frame control message is to be transmitted within a downlink burst, it will always appear first and must use QPSK. Subsequent messages within the burst shall be sequenced in decreasing order of modulation robustness, beginning with the most robust modulation that is supported at the transmitter. An uplink burst contains a single message and uses a single modulation format within a burst.

Prior to modulation, the I and Q signals are filtered by a square-root raised cosine filter with a roll-off factor of 0.25. Roll-off factors of 0.18 and 0.15 are optional. The transmitted waveform is obtained using quadrature modulation.

802.16 does require power control of the transmitted waveform, but leaves the details of the algorithm to vendors. It also requires channel quality measurements (RSSI and CINR) to be performed.

Antenna diversity is another technique that is likely to find application in BWA systems. To provide another dimension of flexibility, the WirelessMAN-SC2 physical layer supports antenna diversity as well.

With two-way delay transmit diversity, two transmit antennas are used, and the output of the second antenna is delayed with respect to the first. Receive diversity does not require any changes in the physical layer. The receiver equalization is expected to handle the extra delay spread introduced in the system due to the delayed output of the second antenna.

However, when the optional STC transmit diversity scheme [B4] is used, a different framing mechanism is required. The STC transmit diversity scheme logically pairs blocks of data separated by delay spread guard intervals. These paired blocks are jointly processed at both the transmitter and receiver. The technique to be described is particularly amenable to a frequency domain equalizer implementation. Suppose that $s_0[n]$ and $s_1[n]$ are the blocks to be transmitted, each of length F symbols. A two antenna transmitter can transmit these two sequences as follows. The first antenna will transmit $s_0[n]$ and $s_1[n]$, separated by guard intervals (Figure 4–29). The length of these guard intervals is equal to the length of the unique word, which, as has been mentioned, should be

INTEL-1020
U.S. Patent 8,219,129

equal to the length of the channel impulse response. Instead of guard intervals composed of zeros, it is better to use cyclic prefixes or unique words.

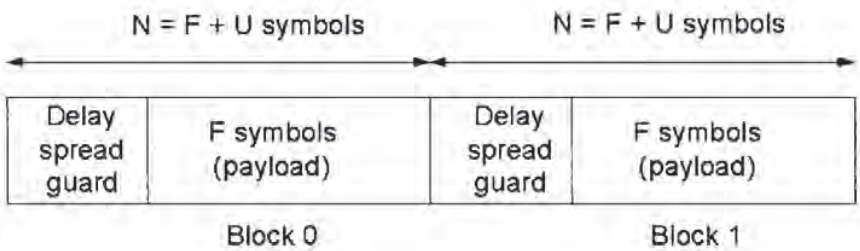

Figure 4–29: Paired blocks in STC transmit diversity combining

The second antenna will reverse the order in which sequences are transmitted, and will conjugate and time-reverse each sequence. Therefore the first block transmitted by the first antenna will be $-s_1^*[L-n]$ and the second block will be $-s_0^*[L-n]$. At the receiver, frequency-domain equalization can be used to recover the transmitted signal, provided that the channel is known.

A STC burst consists of a preamble, followed by a payload. The preamble consists of several unique words, and the number of such unique words is a burst profile parameter. The payload may consist of multiple pairs of blocks as in Figure 4–29. The preamble structure may also be inserted within a transmission as a group of contiguous pilot words to assist in the channel estimation. In this case, the group of contiguous pilot words is considered external to the paired data blocks. The pilot word repetition interval and the number of unique words comprising one pilot are burst profile parameters.

STC-encoded data and conventionally encoded data will not be time-division multiplexed within the same burst.

WirelessMAN-OFDM PHY

The WirelessMAN-OFDM physical layer is designed for NLOS operation between 2 GHz and 11 GHz and is based on OFDM modulation. It can be used in unlicensed and licensed bands.

INTEL-1020
U.S. Patent 8,219,129

The OFDM symbol duration, or the related carrier spacing in frequency, is a major design parameter of an OFDM system. The symbol duration is composed of the FFT interval and of the cyclic prefix (CP). The presence of a cyclic prefix leads to an equivalent loss of SNR, because the cyclic prefix at the receiver is discarded and therefore part of the transmitter energy is wasted. The advantages of having a cyclic prefix are many (multipath immunity, better detection and synchronization performance) and outweigh this energy loss. Note that in the downlink, the BS is not allowed to change the length of the cyclic prefix, because the SS will lose synchronization.

One difference in IEEE 802.16 with respect to IEEE 802.11a is that the OFDM symbol is divided into logical subchannels to support scalability, multiple access, and advanced antenna array-processing capabilities.

An OFDM symbol consists of data carriers, pilot carriers, and null carriers (Figure 4–30). The null carriers, the DC carrier and the guard carriers, are not used. Four parameters describe the OFDM symbol. The number of FFT points is 256, out of which 200 are used, with 55 reserved as guard bins. The DC carrier is also unused. Within the 200 used carriers, 8 carriers are used for pilots. The cyclic prefix can be of lengths 64, 32, 16, or 8 carriers.

Scrambling, FEC, and interleaving are performed prior to modulation. The same randomizer is used as in the physical layer for frequencies between 10 GHz and 66 GHz, described in the previous section. On the downlink, the scrambler is initialized at the start of every frame with 100101010000000. On the uplink, the scrambler is initialized with the vector in Figure 4–31.

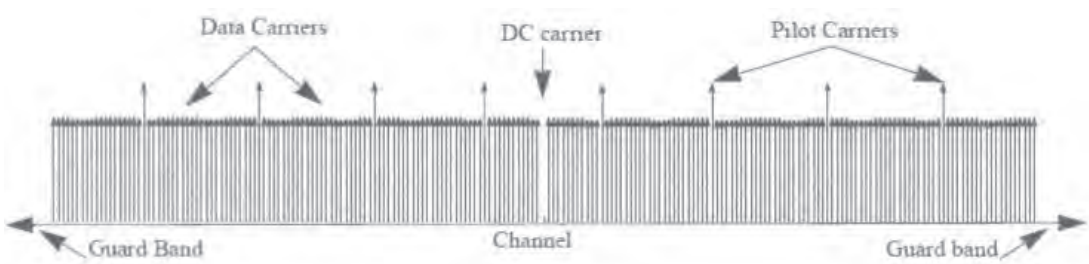

Figure 4–30: An OFDM symbol.

INTEL-1020
U.S. Patent 8,219,129

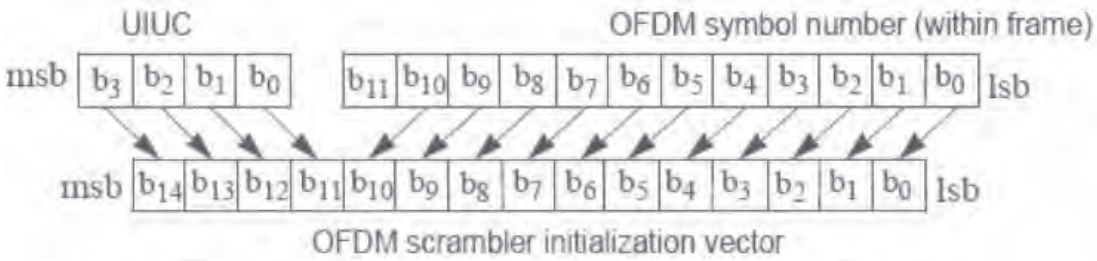

Figure 4–31: Scrambler initialization on the uplink

A FEC, consisting of the concatenation of a RS (255, 239) outer code over GF (256) and a rate-compatible convolutional inner code, is supported on both the uplink an downlink. Support of block turbo codes with extended Hamming or parity check codes is an option. Convolutional turbo codes are also optional. The block turbo code and the convolutional turbo code are the same as the codes in WirelessMAN-SC2. The Reed–Solomon-Convolutional coding rate 1/2 must always be used when requesting access to the network and in the FCH burst. The encoding is performed by first passing the data in block format through the RS encoder and then passing it through a zero-terminating convolutional encoder. Shortening and puncturing can be performed. After bit interleaving, the data bits are entered serially to the constellation mapper. Gray-mapped QPSK and 16 QAM, as shown in Figure 4–32, are mandatory for this physical layer, while 64 QAM is optional. The constellations shown in Figure 4–32 must be normalized by multiplying the constellation point with the indicated factor c to achieve equal average power. Adaptive modulation and coding must be supported in the downlink per SS. In the uplink, each SS can use different modulation schemes based on the MAC burst configuration messages coming from the BS. After modulation the symbols are mapped onto the carriers. The coding and modulation parameters are summarized in Table 4–5.

INTEL-1020
U.S. Patent 8,219,129

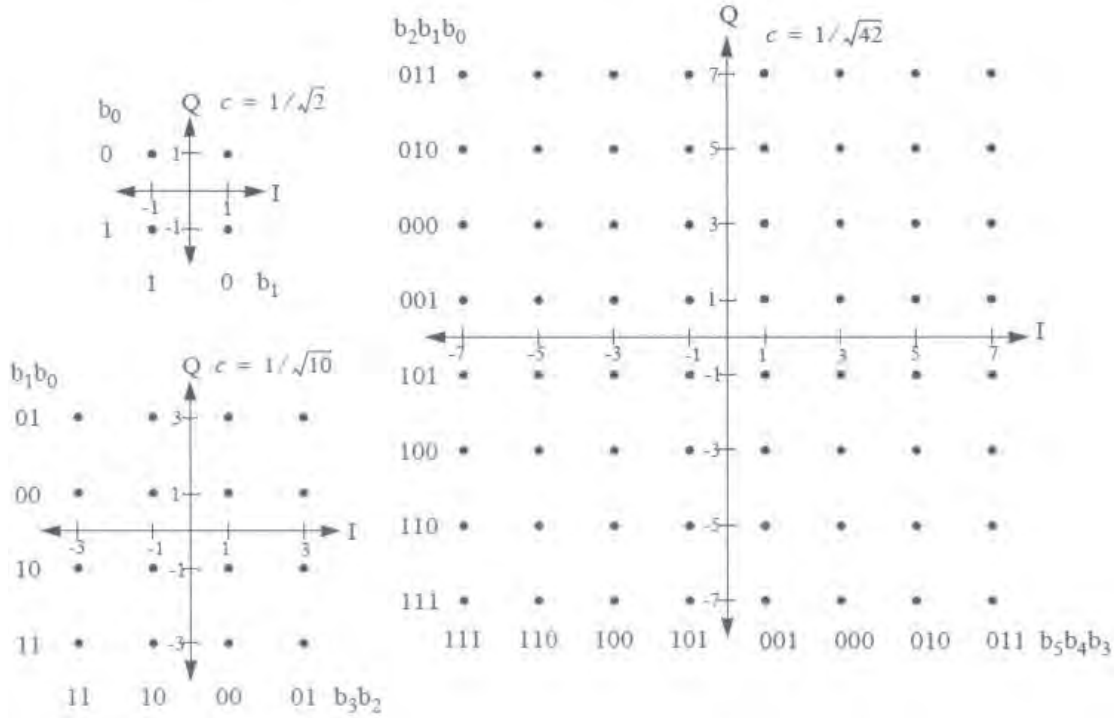

Figure 4–32: QPSK, 16 QAM, and 64 QAM for WirelessMAN-OFDM

Table 4–5: Modulation and coding parameters for WirelessMAN-OFDM

Modulation	Uncoded block size (bytes)	Coded block size (bytes)	Overall coding rate	RS code	CC code rate
QPSK	24	48	1/2	(32,24,4)	2/3
QPSK	36	48	3/4	(40,36,2)	5/6
16 QAM	48	96	1/2	(64,48,8)	2/3
16 QAM	72	96	3/4	(80,72,4)	5/6
64 QAM	96	144	2/3	(108,96,6)	Ω
64 QAM	108	144	3/4	(120,108,6)	5/6

INTEL-1020
U.S. Patent 8,219,129

Pilots are inserted in the OFDM symbols in the following way. A pseudorandom binary sequence generator of length 11 is used, which is described by the polynomial $1 + x^9 + x^{11}$. The initialization sequences are all ones in the downlink and 10101010101 in the uplink. Starting with the third bit, the output of this random sequence generator, w_k, is used to obtain the pilot tones in the first OFDM symbol following the frame preamble. The pilot tones are modulated using BPSK. On the downlink $c_{-84} = c_{-36} = c_{60} = c_{84} = 1 - 2w_k$ and $c_{-60} = c_{-12} = c_{12} = c_{36} = 1 - 2\overline{w_k}$. On the uplink $c_{-84} = c_{-36} = c_{12} = c_{36} = c_{60} = c_{84} = 1 - 2w_k$ and $c_{-60} = c_{-12} = 1 - 2\overline{w_k}$.

Two preambles are defined—short and long. The short preamble is used in the uplink and consists of 128 samples repeated twice and preceded by a cyclic prefix (CP). The length of the cyclic prefix is the same as the length of the cyclic prefix of the payload-carrying OFDM symbols. The long preamble is used in the downlink. It consists of a CP, followed by 64 samples, repeated four times, followed by a CP and 128 samples, repeated twice. (See Figure 4–33 and Figure 4–34.)

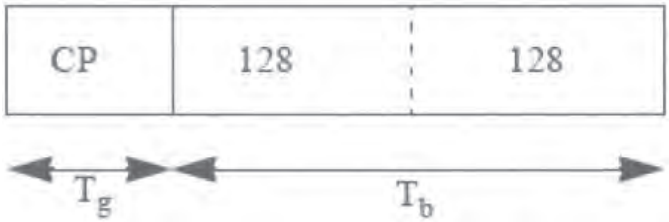

Figure 4–33: Short preamble for uplink data and downlink AAS

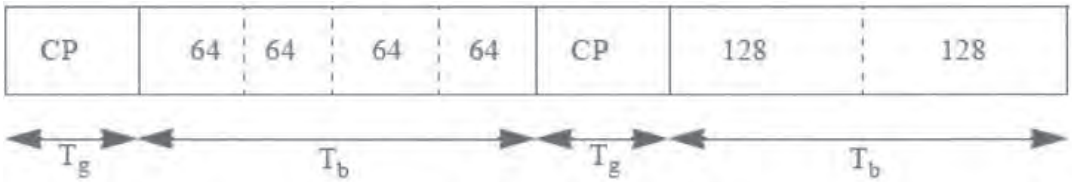

Figure 4–34: Long preamble for downlink and initial ranging

In licensed bands, TDD and FDD can be used. In unlicensed bands, only TDD can be used. The physical layer supports frame-based transmission. One frame

INTEL-1020
U.S. Patent 8,219,129

consists of a downlink subframe and an uplink subframe. The uplink subframe consists of contention interval for ranging and bandwidth requests, followed by packets from one or more SSs. Every packet consists of a short preamble, followed by an integer number of OFDM symbols. In the downlink, transmissions start with the long preamble, followed by one special OFDM symbol, called FCH symbol. The FCH OFDM symbol conveys the burst parameters and other control information such as the DCD and UCD, and is transmitted using QPSK and rate 1/2 coding. The FCH symbol is followed by one or more bursts, each transmitted with a different burst profile. Each burst contains an integer number of OFDM symbols. A TDD frame is shown in Figure 4–35.

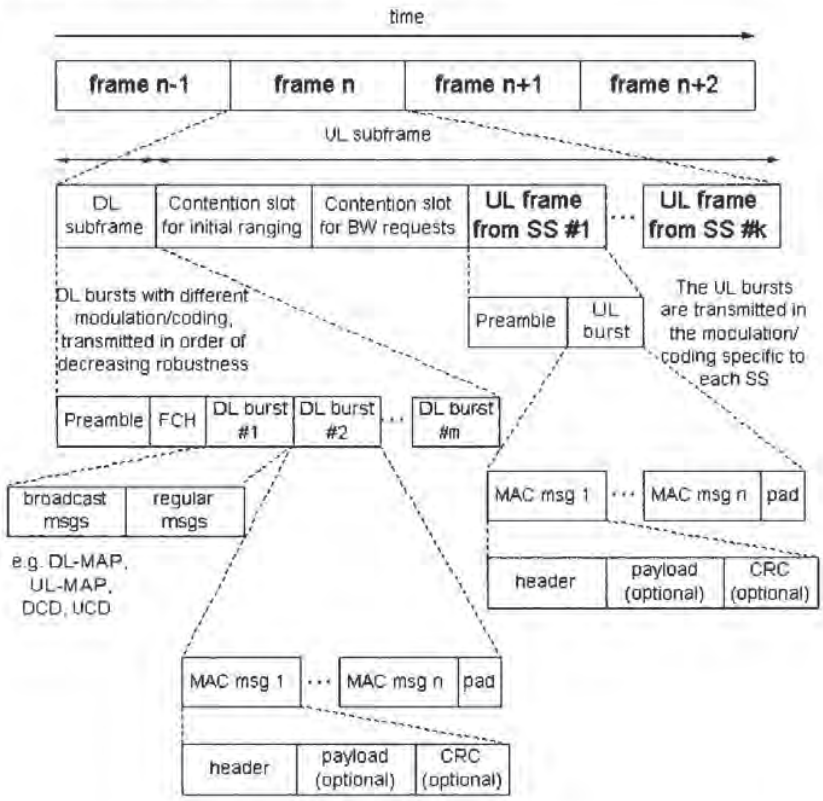

Figure 4–35: An example of an OFDM frame structure with TDD

For FDD, the downlink subframe and the uplink subframe have the same structure as in the TDD case. If FDD is used, instead of time gaps between subframes and frames, the BS must not schedule uplink transmission from a half duplex SS during a downlink allocation.

INTEL-1020
U.S. Patent 8,219,129

Mesh networks are an option, and have an optional frame structure, shown in Figure 4–36.

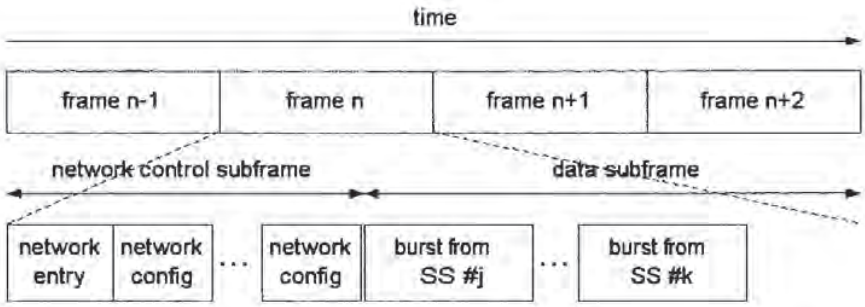

Figure 4–36: Frame structure for mesh networks

A mesh frame consists of a control subframe and a data subframe. The control subframe is used for network creation, maintenance, and coordinated scheduling. All transmissions in the control subframe are done using QPSK and the mandatory coding method. All preambles in the mesh frame are long preambles. Note that distributed scheduling messages can be transmitted in the data subframe as well.

AAS and STC are also supported by this physical layer as options. Power control is required, but not specified and left to vendors.

As an option, STC can be used on the downlink to provide second order spatial transmit diversity. There are two transmit antennas on the BS side and one reception antenna on the SS side. This scheme requires Multiple Input Single Output (MISO) channel estimation. Decoding is very similar to maximum ratio combining. Figure 4–37 shows a block diagram of STC using an OFDM physical layer.

Both antennas transmit two different OFDM data symbols at the same time. The first time antenna 0 transmits s_0, and antenna 1 transmits s_1. The second time antenna 0 transmits $-s_1^*$ and antenna 1 transmits $-s_0^*$. STC is applied independently on each carrier with respect to pilot tone positions. Figure 4–38 serves as illustration, where only one pilot carrier is shown.

INTEL-1020
U.S. Patent 8,219,129

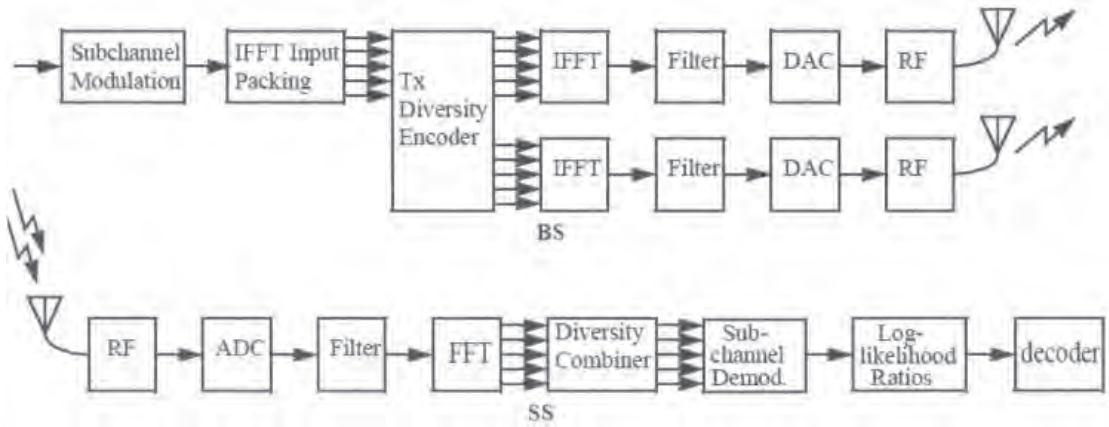

Figure 4–37: A block diagram of a transmitter and receiver constructed according to Alamouti

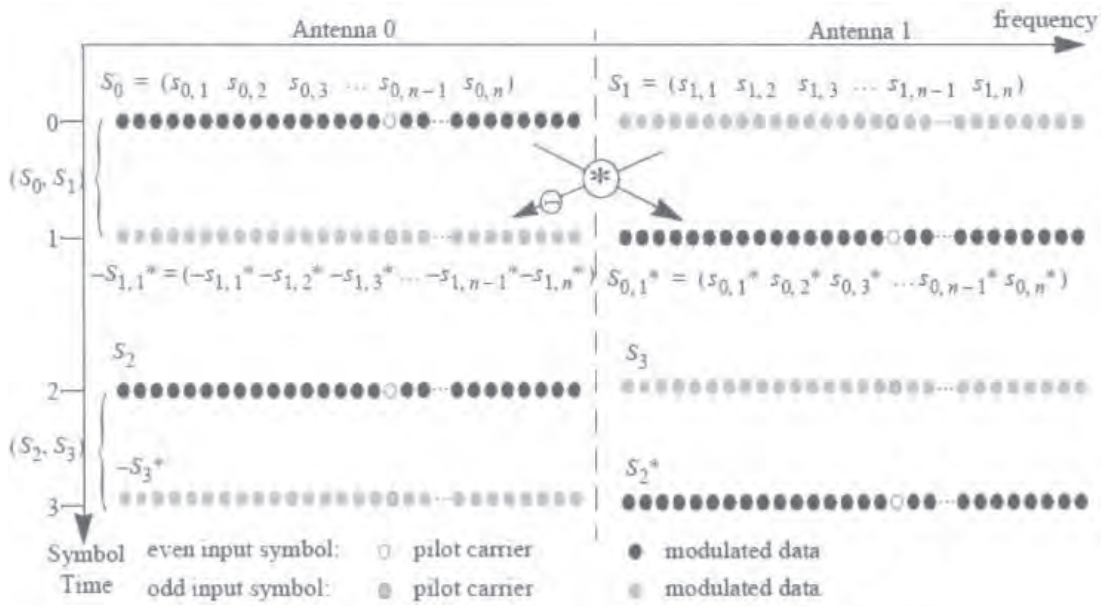

Figure 4–38: STC usage with WirelessMAN-OFDM

Compared with the other wireless standards, in IEEE 802.16 channel quality measurements are much more important. A BS is responsible for making assignments and/or reassignments based on both upstream and downstream channel quality assessments. It can directly determine channel quality on upstream channels from the messages that it receives from SSs. However, to

INTEL-1020
U.S. Patent 8,219,129

determine the channel quality of its downstream channels, a BS requires channel quality reports from SSs. The BS might also decide on the adaptive modulation type and FEC to be used when communicating with a SS. RSSI and CINR computation is mandatory. RSSI measurements do not require demodulation and can be simply done even at low signal levels. On the other hand, CINR measurements require demodulation. However, CINR is more useful because it provides information on the actual operating condition of the receiver, including interference, noise levels, and signal strength.

Since NLOS subscriber locations may experience Rayleigh or Ricean fading, signal quality (and signal quality measures) may vary with time. For this reason, both mean and standard deviation statistics are defined for the RSSI and CINR measurements. The mean statistics convey the average signal quality level for a channel. The standard deviation statistics are intended to convey the variation in that average value.

The receiver signal processing includes frame detection, fine synchronization, coarse and fine frequency offset estimation and compensation, and channel estimation. This is followed by an orthogonal transformation, phase estimation from the pilot subcarriers, derotation the subcarrier values according to estimated phase, and equalization by dividing each subcarrier value with a complex estimated channel response coefficient. Finally, for each data-carrying subcarrier, find the closest constellation point, compute the Euclidean distance from it and compute the RMS average of all errors in a packet. All of the receiver signal processing operations are similar to IEEE 802.11a. Receiver sensitivity is defined for BER of 10^{-6}. The sensitivity ranges from –65 dBm for 64 QAM with a rate 3/4 code to –91 dBm for QPSK with a rate 1/2 code.

WirelessMAN-OFDMA

WirelessMAN-OFDMA is based on OFDM and designed for NLOS operation in the 2–11 GHz band.

Similar to OFDM, an OFDMA symbol consists of data carriers, pilot carriers, and null carriers. However, OFDM and OFDMA are slightly different. For OFDM, a channel is defined as consisting of all carriers residing within the full signaling bandwidth. In OFDM mode, the basic resource allocation quantum is an OFDM symbol. The amount of data that fits into an OFDM symbol depends on the

INTEL-1020
U.S. Patent 8,219,129

constellation and the coding method used within this symbol as well as the number of data carriers per symbol. For OFDMA, subchannels are defined as a fraction of the available carriers within the full signaling bandwidth. In OFDMA mode, the basic resource allocation quantum is a subchannel. Furthermore, the carriers forming one subchannel do not have to be adjacent. For all FFT sizes, each OFDMA symbol contains an integer number of subchannels, both on downlink and on uplink. The amount of data that fits into a subchannel depends on the constellation and the coding method used within this subchannel as well as the number of data carriers per subchannel.

Therefore, OFDMA is a multiple access method simultaneously in the time and frequency domains. In OFDMA one slot, also called data region, has two dimensions—time and subchannel. The time dimension denotes consecutive OFDM symbols in the time domain, and the subchannel dimension denotes blocks within the OFDM symbols. Figure 4–39 shows three OFDMA subchannels.

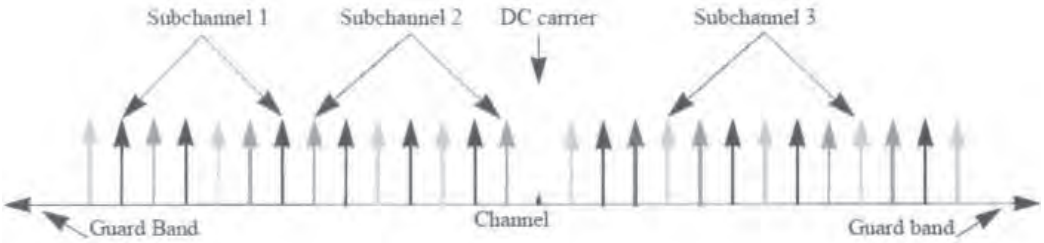

Figure 4–39: An example of three OFDMA subchannels

All used carriers are partitioned into fixed-location pilots, variable-location pilots, and data. In the downlink, the variable-location pilots shift their location every symbol repeating every four symbols. After mapping the pilots, the rest of the used carriers are the data subchannels. Because the variable location pilots change location in each symbol, repeating every fourth symbol, the locations of the carriers in the data subchannels must change also. To allocate the data subchannels, the remaining carriers are partitioned into groups of contiguous carriers. Each subchannel consists of one carrier from each of these groups. The number of groups is therefore equal to the number of carriers per subchannel. The number of the carriers in a group is equal to the number of subchannels.

INTEL-1020
U.S. Patent 8,219,129

Regarding the subcarrier allocations in OFDMA, there is a subtle difference between uplink and downlink. This difference is that, in the downlink, the pilot tones are allocated first; what remains are subchannels that are used exclusively for data. In the uplink, however, the set of used carriers is first partitioned into subchannels, and then the pilot carriers are allocated from within each subchannel. Thus, in the downlink, there is one set of common pilot carriers, but in the uplink, each subchannel contains its own set of pilot carriers. This is necessary because, in OFDMA, the BS downlink is broadcast to all SS, but in the uplink, each subchannel may be transmitted from a different SS.

In OFDMA, the FFT size is 2048, out of which 1702 are used in the downlink, and 1696 are used in the uplink. On both the downlink and the uplink there are 32 subchannels, each with 48 data carriers. The cyclic prefix is 1/32, 1/16, 1/8, and 1/4 of the FFT length, or alternatively 64, 128, 256, or 512 samples. In the downlink, there are 142 variable-location pilots, 32 fixed-location pilots, eight of which coincide with the variable location pilots, so the total number of pilots is 166. In the uplink within each subchannel, there are 48 data carriers, one constant-location pilot carrier at index 26, and four variable-location pilot carriers. The variable-location pilots change with each symbol, repeating every 13 symbols. The variable-location pilot locations will never coincide with the constant-location pilot at index 26. The remaining 48 carriers are data carriers.

The signal processing operations are as follows. Each FEC block spans one OFDMA subchannel in the subchannel axis and three OFDM symbols in the time axis. The FEC blocks are mapped in such a way that the lowest numbered FEC block occupies the lowest numbered subchannel in the lowest numbered OFDM symbol. When the edge of a slot is reached, the mapping continues from the lowest-numbered OFDMA subchannel in the next OFDM symbol.

In licensed bands, the duplexing can be TDD or FDD, and in unlicensed bands it can be only TDD.

In PMP TDD operation (Figure 4–40), each frame consists of downlink and uplink, separated by transition gaps. Each burst transmission consists of multiples of three (or six, if STC used) OFDMA symbols. On the downlink subframe, first a FCH is transmitted in the most robust burst profile. The FCH is a data structure that contains the DL-MAP message and can additionally include the UL-MAP, DCD or UCD messages. The DL-MAP message is preceded by a DL frame

INTEL-1020
U.S. Patent 8,219,129

prefix, which has information where the FCH message is and how it can be demodulated. Since different modulations and coding rates can be used, the transitions between modulations and coding can take place not only on OFDMA symbol boundaries in time domain, but on subchannels within an OFDM symbol in frequency domain. A BS supporting the AAS option will allocate (at the end of the uplink frame) an initial ranging slot for AAS SS that have to initially alert the BS for their presence.

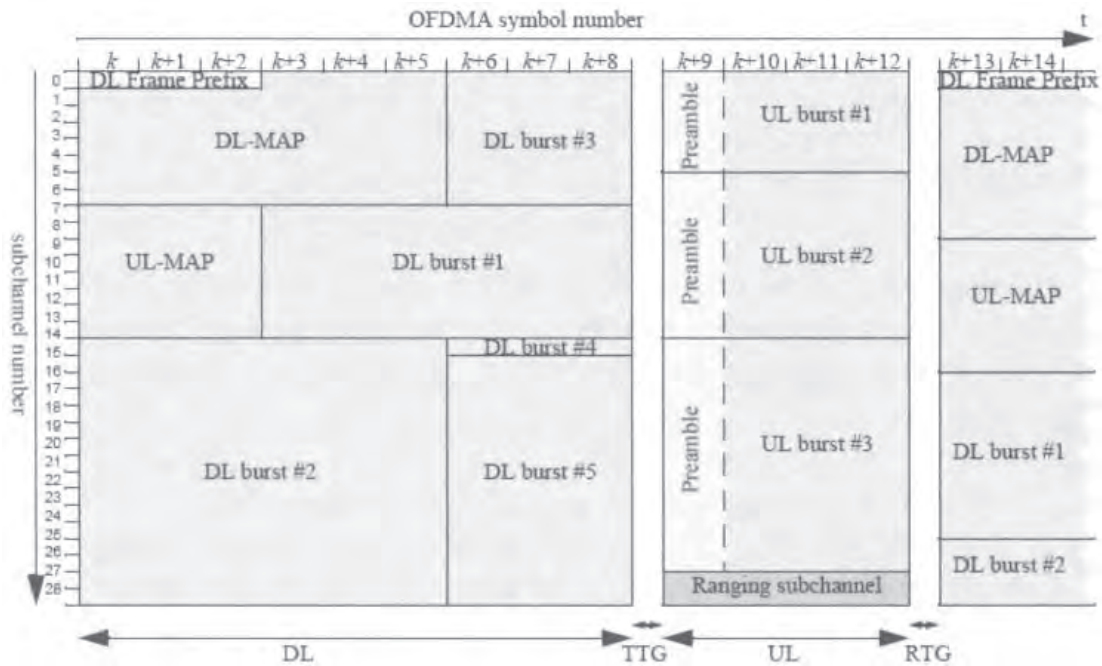

Figure 4–40: An example of a TDD OFDMA frame

When used with the WirelessMAN-OFDMA physical layer, the MAC will define a single ranging channel. There is one ranging channel, composed of one or more pairs of adjacent subchannels, where the index of the lowest-numbered subchannel is even. The indices of the subchannels that compose the ranging channel are specified in the UL-MAP message. Users are allowed to collide on this ranging channel. To effect a ranging transmission, each user randomly chooses one ranging code from a bank of specified binary codes. These codes are obtained from a simple random binary sequence generator, described by the polynomial $1 + x^1 + x^4 + x^7 + x^{15}$ and initialized by the sequence 100101010000000. The first code is formed by the first 106 bits in this sequence;

INTEL-1020
U.S. Patent 8,219,129

the next is formed by taking the 107–212 bits; etc. There are 48 codes in all, divided into three usage groups—codes for initial ranging, periodic ranging, and bandwidth requests. The BS dynamically assigns codes to each usage group. These codes are then BPSK modulated onto the carriers in the ranging channel, one bit per carrier. This process is used for initial ranging, periodic ranging, and bandwidth requests. Initial ranging is used by SSs that want to synchronize for the first time. Initial ranging is performed during two OFDM symbols, with the same ranging code transmitted during the two symbols. Periodic ranging and bandwidth requests use one OFDM symbol.

As an option, similar to the OFDM case, STC [B4] can be used on the downlink to provide second-order space transmit diversity. The only difference with WirelessMAN-OFDM is that variable-location pilots are kept constant for two symbols.

Information bits are randomized with the same randomizer as for OFDM. This randomizer is initialized in the DL with same sequence as in the OFDM case. In the downlink it is initialized in a different way, with the sequence obtained as shown in Figure 4–41.

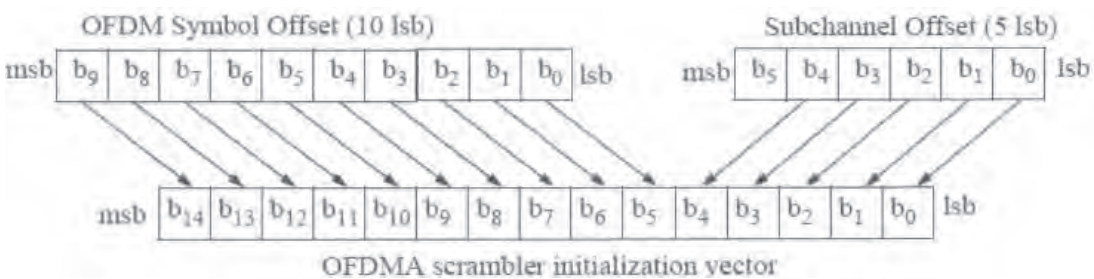

Figure 4–41: Scrambler initialization vector for OFDMA

The FEC mechanism and data modulation formats are similar to those used in the WirelessMAN-OFDM physical layer. Pilot modulation is also similar. The difference in pilot modulation is that each pilot is transmitted with a 2.5 dB boost over the average power of each data tone. This boost is achieved by using $c_k = 8/3\left(\frac{1}{2} - w_k\right)$, instead of $c_k = 2\left(\frac{1}{2} - w_k\right)$, as in BPSK.

The OFDMA receiver sensitivity is the same as in the OFDM case.

INTEL-1020
U.S. Patent 8,219,129

In the licensed bands, channelization is performed according to the regulatory requirements. The minimum width of one channel is 1.25 MHz. In the middle and upper UNII bands, the same channelization scheme as for IEEE 802.11a is used, with the difference that channels of width 10-MHz can also be used. In unlicensed bands, mesh networks can be used in the same fashion as for WirelessMAN-OFDM.

The same types of channel quality measurements are required (RSSI and CINR), as in the OFDM case.

COEXISTENCE

Radio waves propagate in unbounded medium. Interference among different wireless systems operating in the same frequency band is inevitable. Therefore, it is better if measures are taken to minimize the interference. As discussed in the previous chapter, coexistence is a new issue for wireless communications. Compared with IEEE 802.11 and IEEE 802.15.1, 802.15.3, and 802.15.4, the resolution of coexistence issues is more important for IEEE 802.16 because of the usage scenarios and business model of IEEE 802.16 service providers.

What are the interference mechanisms that can reduce the performance of fixed BWA systems? First, although intrasystem interference is often a significant source of performance degradation, it is not considered. The reason is that intrasystem interference is under the control of the operator. Thus, only intersystem interference mechanisms, where interoperator coordination is appropriate, are considered here. Several coexistence problems should be discussed.

First, there are coexistence issues with other (non-IEEE 802.16) systems, mainly in the unlicensed bands. In the 5-GHz region, satellite services and radar services are two primary users of the bands. The RF center frequency of these satellites is 5.305 GHz; they use antennas with very high gain, and their peak radiated power can be significant. Other devices that also operate in these frequency bands are road transport and traffic telemetric (RTTT) devices [B128]. They are allocated in the 5795–5805 MHz band, with an extension band of 5805–5815 MHz, which in the United States may be used on a national basis at multilane road junctions. These devices are split into the roadside unit (RSU) and the onboard unit (OBU).

INTEL-1020
U.S. Patent 8,219,129

Second, there are coexistence issues also among IEEE 802.16 operators, both in the licensed and unlicensed bands. IEEE 802.16 also provides recommendations directed at equipment manufacturers and different recommendations aimed at companies providing services using this technology.

Looking at the sources of interference will be helpful here. Figure 4–42 shows the main sources of interference to a fixed BWA base station, with a sectoral-coverage antenna.

The victim BS is shown as a black triangle on the left, with its radiation pattern represented as ellipses. The desired SS transmitter is shown on lower right of figure. The letters in Figure 4–42 illustrate several cases of interference to a base station.

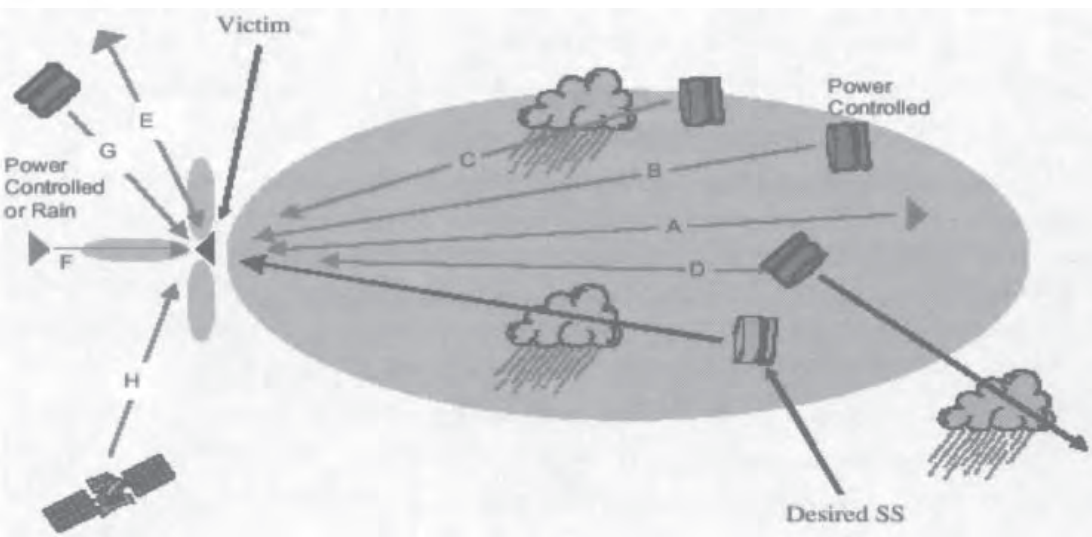

Figure 4–42: Interference sources

Case A is BS-to-BS interference in which each BS antenna is in the main beam of one or more other BSs. This can be a significant problem, because BS antennas tend to be elevated, with a high probability of line-of-sight path to each other. BS-to-BS interference can be reduced by ensuring that there is no BS transmission on frequencies being used for reception at other BSs.

INTEL-1020
U.S. Patent 8,219,129

Case B is SS-to-BS interference in which each antenna is in the main beam of the other. As SS antenna gain is much higher than the BS antenna gain, this might appear to be the worst possible case. However, fixed BWA PMP systems can safely be assumed to employ uplink adaptive power control at subscriber stations. This power control equalizes the received signal strength arriving at a BS from near and far SSs on adjacent channels. Note that active control of downlink power from BS transmitters is usually not employed. The net effect is that Case B may not be more severe than Case A. In addition, the narrow beamwidth of a SS antenna ensures that Case B is much less common an occurrence than Case A. However, band planning cannot eliminate Case B interference.

Case C is just like Case B, but the interfering signal is subject to rain attenuation. Therefore it is less of a problem than Case B.

Case D is also like Case B, but the interference is from a SS's sidelobe. However, because this SS sees rain in its main transmission, it does not turn down its power. Case D can be more serious than Case B if the interference does not come from a BWA system, but from a satellite communication system.

Case E is similar to Case A, it is also BS-to-BS interference, but the interference comes from another BS's sidelobe or backlobe. Case E is less serious than Case A.

Case F covers BS-to-BS backlobe-to-backlobe or sidelobe-to-sidelobe. The low gains involved here ensure that this is a problem only for co-deployment of systems on the same rooftop. Like all sources of BS-to-BS interference, this can be virtually eliminated in FDD via a coordinated band plan.

Case G is interference from an SS to the BS's sidelobe or backlobe. It is less severe than Case B and is not considered further.

Finally, **Case H** covers interference from a satellite downlink or stratospheric downlink. This case is not considered in the Recommended Practice.

With the previous simplifying assumptions, the dominant sources of interference that require detailed modeling are shown in Figure 4–43. Case A will tend to dominate unless there is a harmonized band plan for the use of FDD. Case B is always a concern. Case D is less significant than Case B if the source is another BWA system, but could be significant if it comes from a transmitter with high

INTEL-1020
U.S. Patent 8,219,129

output power. Case F is a concern only for co-sited BSs and can be largely mitigated by the use of a harmonized band plan with FDD.

Figure 4–44 shows the main sources of interference to a subscriber station having a narrow beamwidth antenna. The victim subscriber station is shown along with its radiation pattern (ellipses). The BS and several interferers are also shown. The victim SS cases are fundamentally different from the victim BS cases because the antenna pattern is very narrow. The letters in Figure 4–44 illustrate several cases of interference to a subscriber station:

Case A is SS-to-SS interference, where the beams are collinear. This case rarely happens in practice.

Case B covers BS-to-SS interference.

Case C covers the case of a narrow-beam transmitter (fixed BWA or point-to-point) or satellite uplink at full power, due to rain in its path, but radiating from its sidelobe towards the victim. This case is more likely to occur than Case A because it could occur with any orientation of the interferer.

Case D covers BS-to-SS interference picked up by a sidelobe or backlobe of the victim. This case could be common because BSs radiate over wide areas, and this case could occur for any orientation of the victim.

Case E covers SS-to-SS interference picked up by a sidelobe or backlobe of the victim station. Similar to the reasoning in the victim BS Cases B and C, the worst case can be assumed to be clear air in the backlobe with the interferer having turned its power down.

The source of interference in **Case F** is a satellite downlink. Case F is not considered in the Recommended Practice.

INTEL-1020
U.S. Patent 8,219,129

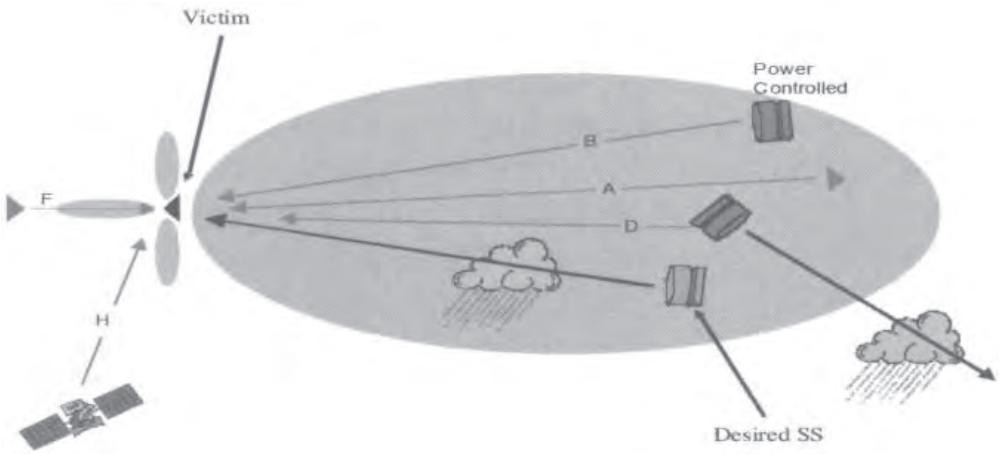

Figure 4–43: Simplified model of the interference to a BS

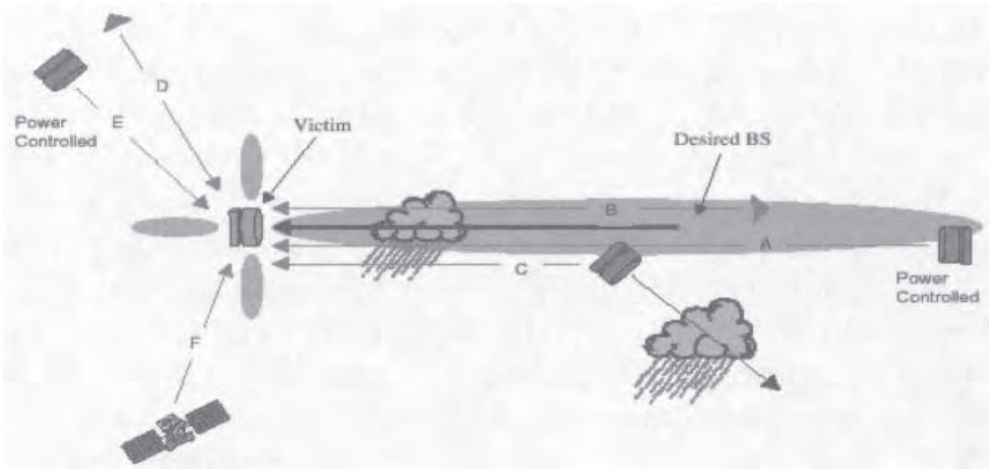

Figure 4–44: Interference to SS

How is the interference calculated? A number of techniques can be used, such as worst-case analysis, Monte Carlo simulations, or the by calculating the interference area. The most appropriate method depends on the interference mechanism. Worst-case analysis is more appropriate when there is a single dominant source of interference. An example is the interference from a single dominant BS into the victim BS of an adjacent system. There are many cases where a simple worst-case analysis is of limited use. Where there are many possible interference paths, the worst case could be very severe but may also be

INTEL-1020
U.S. Patent 8,219,129

very improbable. An example is the interference between subscriber stations of different operators in the same geographical area. Monte Carlo simulations provide a means of assessing the probability of occurrence of a range of interference levels at victim stations. The recommended geographical or frequency spacing is then a compromise in which an acceptably small proportion of cases suffer interference above the recommended limit. In some scenarios, it can be shown that specific parts of the coverage area will suffer high levels of interference while other areas are not affected. The interference area (IA) is the proportion of the sector coverage area where interference is above the target threshold. This is equivalent to the probability that a randomly positioned station (within the nominal coverage area) will experience interference above the threshold. In several scenarios, the interference area value is a small percentage and the locations are predictable. Although high levels of interference do occur, they are sufficiently localized to be acceptable.

The two types of interference—co-channel and adjacent channel— are shown in Figure 4–45. Note that the channel bandwidth of the co-channel interferer may be wider or narrower than the desired signal. In the case of a wider co-channel interferer (as shown), only a portion of its power will fall within the receive filter bandwidth. In this case, the interference can be estimated by calculating the power arriving at the receive antenna and then multiplying by a factor equal to the ratio of the filter's bandwidth to the interferer's bandwidth. An out-of-channel interferer is also shown. Here a portion of the interferer's spectral sidelobes or transmitter output noise floor falls co-channel to the desired signal; i.e., within the receiver filter's passband. This can be treated as co-channel interference. It cannot be removed at the receiver; its level is determined at the interfering transmitter. By characterizing the power spectral density of sidelobes and output noise floor with respect to the main lobe of a signal, this form of interference can be approximately computed in a manner similar to the co-channel interference calculation, with an additional attenuation factor due to the suppression of this spectral energy with respect to the main lobe of the interfering signal.

There are two types of unwanted emissions. First, out-of-band emissions are within 200% of the bandwidth from the edge of the authorized emissions. Second, there are spurious emissions, which are beyond this point. The spectral density of unwanted emissions at the input to the antenna port should be attenuated

INTEL-1020
U.S. Patent 8,219,129

significantly. Spurious emissions should not exceed an absolute level of –70 dBW/MHz in any 1 MHz band.

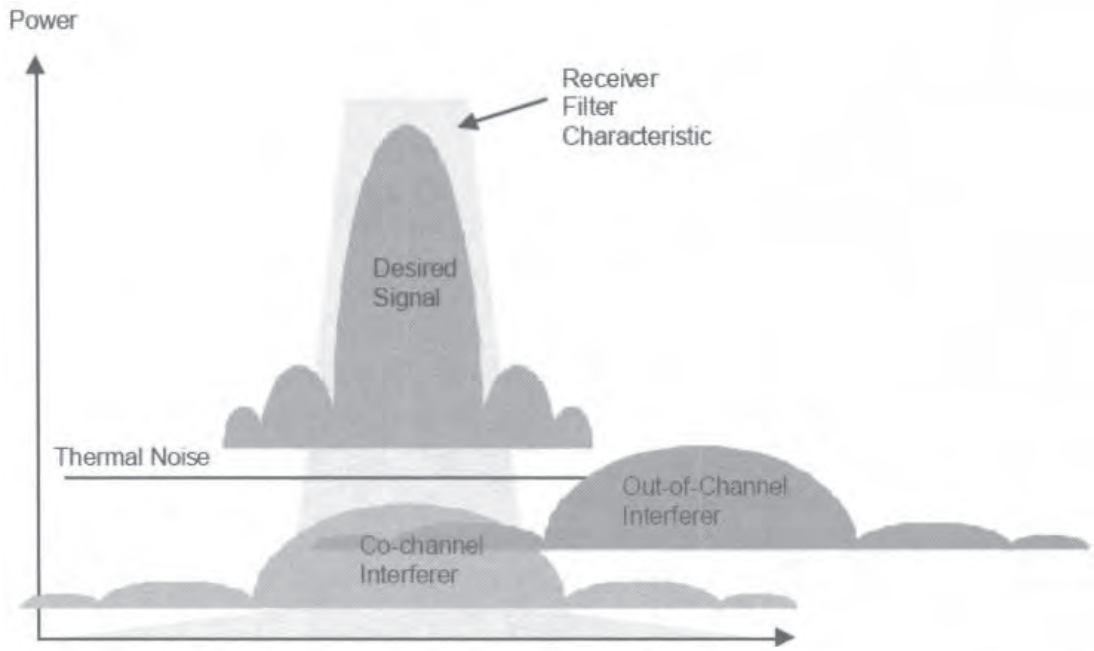

Figure 4–45: An illustration of co-channel and out-of-channel interferer

In the unlicensed portion of the 2–11 GHz band, coexistence issues are especially important, because any wireless system can operate in this portion of the band. Two categories of coexistence mechanisms are considered—methods that fall within the scope of the WirelessMAN standard and those that fall outside that scope. Within the scope of the WirelessMAN standard, two methods are identified—DFS and TPC. For unlicensed use, DFS and TPC are required both by the standard and by regulatory bodies in the United States and Europe. Outside the scope of the WirelessMAN standard, two methods are identified—antenna directivity and antenna polarization. Here is a summary of the four methods.

1) *DFS*. As frequency planning is not practical in licensed-exempt bands, DFS can be used to avoid assigning a channel to a channel occupied by another system. DFS is generally based on comparison of a C/I threshold against idle time RSSI measurements. DFS is predominantly effective to combat interference from and to ground based systems, such as WLANS, RTTT, radar

INTEL-1020
U.S. Patent 8,219,129

and other wireless systems. It is generally ineffective to combat interference from and to airborne systems, such as airborne radars and satellites.

2) *TPC.* This is a technique to adaptively adjust the transmit power of a transmitter to maintain the received signal level within some desired range. With power control, the transmitter EIRP is reduced according to the link margin. Shorter link ranges, hence, result in lower transmitted power levels. For PMP systems, the average EIRP will typically be several decibels below the legal limit. For mesh systems, the EIRP values decrease rapidly as customer deployment density increases. As power control is also influenced by C/I levels, the use of adaptive power control with DFS, where possible, tends to result in the most effective interference mitigation.

3) *Antenna directivity.* Antenna directivity, in horizontal but especially in vertical direction can significantly reduce a BWA's interference potential and resilience. Vertical directivity especially reduces the interference caused to satellite systems, which are designated primary users of part of the addressed bands. It also can significantly help reduce interference to and from indoor WLAN systems. Horizontal directivity significantly reduces the probability of interference to other systems (assuming interference is mainly caused in the main lobe), but tends to increase the severity of the interference, as the energy in the main lobe is generally higher.

4) *Antenna polarization.* Antenna cross-polarization in the 5GHz band can achieve an isolation of up to 15 dB in LOS, but the isolation reduces significantly in near-LOS and NLOS environments. Most deployments use both horizontal and vertical polarization to maximize spectral reuse. Polarization hence has the potential to provide some isolation between differently polarized systems, especially in LOS. However, given the operational needs and implementation of most systems in the targeted spectrum, the effectiveness will be mostly marginal.

While in the unlicensed bands there will be interference from other wireless systems, in the licensed bands coexistence issues among multiple operators will dominate. It is noted that in general some service providers may deploy systems not conforming to IEEE 802.16. The coexistence problem in this case becomes very complex and is not considered. Because interference among operators can be a contentious issue, IEEE 802.16 produced a Recommended Practice called IEEE Std 802.16.2. It should be noted that the coexistence issue is so complex in general

INTEL-1020
U.S. Patent 8,219,129

that it cannot be addressed in a standard. These recommendations can be considered as a system model. The input to the model is the system architecture, transmitted power, spectral masks and antenna patterns, and receiver parameters, such as noise floor degradation and blocking performance. The outputs of the model are recommendations for the deployment of BWA systems. These recommendations include band plans, separation distances, and power spectral flux density limits. The information is therefore valuable as a first step in planning the deployment of systems. In the model calculations, the geographical or frequency spacing between systems is being varied until the interference is below an acceptable threshold.

There are two main interference scenarios. First, there is a co-channel scenario, in which two operators are in either adjacent territories or territories within radio line of sight of each other and have the same spectrum allocation. Second, there is an adjacent channel scenario in which the licensed territories of two operators overlap and they are assigned adjacent frequency bands. Coexistence issues may arise simultaneously from both scenarios as well as from these scenarios involving multiple operators. These, however, are very complex, and are not considered. First, the general recommendations, applicable in all cases, are considered. Then, the recommendations for the co-channel and adjacent channel cases are considered separately.

The general recommendations are:

1) Each operator should accept a 1-dB degradation [the difference in dB between C/N and $C/(N + I)$] due to cumulative interference from other operators. Operators should include 1-dB margin even if the first to deploy in a region. Thus interference that produces a degradation of 1 dB to the system's C/N is considered acceptable. The logic here is that there is no such thing as interference-free wireless transmission. Therefore, if the interference is 6 dB below the noise, it can be said that systems coexist. For the noise floor to increase by 1 dB, the interference power level must be 6 dB below the receiver's thermal noise floor. According to the link budget calculations for wireless system (previously presented), the level of the received signal will decrease depending on the distance and the atmospheric conditions (rain, etc.). The receiver will receive the signal only when the signal-to-noise ratio exceeds a certain threshold. A way to account for interference is to determine

INTEL-1020
U.S. Patent 8,219,129

$C/(N + I)$, the ratio of carrier level to the sum of noise and interference. Receiver sensitivity determines the minimum detectable signal and is a key factor in any link design. However, as the level of receiver noise floor increases, the sensitivity degrades. This, in turn, causes reduction in cell coverage, degradation in link availability, and loss of revenues. The factors contributing to the increase in noise power divide into two groups—internal and external. The internal factors include the thermal noise generated by various components within the receiver, intermodulation noise, and intranetwork co- and adjacent-channel interference. The external factor is internetwork interference. The amount of degradation in receiver sensitivity is directly proportional to the total noise power added to the thermal noise, I, consisting of intra- and internetwork components. In order to reduce the internetwork contribution, it is recommended that the effect of any fixed BWA network on any other coexisting BWA network should not degrade the receiver sensitivity of that fixed BWA network by more than 1 dB.

For example, consider a receiver with 6 dB noise figure. The receiver thermal noise is –138 dBW/MHz. Interference of –138 dBW/MHz would double the total noise, or degrade the link budget by 3 dB. Interference of –144 dBW/MHz, 6 dB below the receiver thermal noise, would increase the total noise by 1 dB to –137 dBW/MHz, degrading the link budget by 1 dB.

2) Each operator should collaborate in good faith with other known operators before initial deployment and before every relevant system modification. In general, analyses to evaluate the potential for interference and any possible mitigation solution should be performed before system implementation. Coordination with adjacent operators could significantly lower the potential for interference. Best results may be obtained if full cooperation and common deployment planning is achieved. It is legitimate to consider the capital investment an incumbent operator has already made and the additional investment required by an incumbent operator for a change enhancing the ability to coexist versus the capital investment costs that the new operator would incur.

3) The degree of coexistence between systems depends on the emission levels of the various transmitters. Thus it is important to recommend an upper limit on transmitted power or, more accurately, a limit for the equivalent isotropically radiated power (EIRP). Because point-to-multipoint systems span very broad

INTEL-1020
U.S. Patent 8,219,129

frequency bands and use many different channel bandwidths, a better measure of EIRP for coexistence purposes is in terms of *power spectral density* (psd) expressed in dBW/MHz rather than simply power in dBW. Note that the regulatory limits are significantly higher (e.g., 15 dB) than supported by most currently available equipment. If systems are deployed using the maximum regulatory limits, they should receive a detailed interference assessment unless they are deployed in isolated locations, remote from adjacent operators. In the coexistence simulations, reasonable cell sizes and link budgets are considered. Typical parameters used for the BS and in coexistence simulations for the recommended practice are as follows—Tx Power: +24 dBm (–6 dBW), SS Antenna Gain: +34 dBi, BS Antenna Gain: +19 dBi, and Carrier Bandwidth: 28 MHz (+14.47 dB-MHz). (These design parameters apply to frequency range 2 (23.5–43.5 GHz.) According to 802.16.2, a BS should not produce an EIRP power spectral density exceeding +14 dBW/MHz. However, it is strongly recommended that a maximum EIRP power spectral density of 0 dBW/MHz be used. A SS should not produce an EIRP spectral density exceeding +30 dBW/MHz. However, it is strongly recommended that a maximum EIRP power spectral density of +15 dBW/MHz. (For the specific sub-band 25.25–25.75 GHz, a subscriber station should not exceed +24 dBW in any 1 MHz bandwidth [ITU-R F.1509]. Repeaters with directional antennas should not produce an EIRP spectral density exceeding +30 dBW/MHz, but a value of +15 dBW/MHz is recommended. Repeaters employing omnidirectional or sectored antennas should not produce an EIRP spectral density exceeding +14 dBW/MHz, with a recommended value of 0 dBW/MHz. Furthermore, a SS should employ uplink power control with at least 15 dB of dynamic range.

4) The antennas play an important role in the coexistence. Independently of polarization, the operator needs to consider the antenna radiation pattern in the azimuth and elevation planes relative to the required coverage footprint. Two linear polarization orientations, horizontal and vertical, are recommended. The required polarization purity is captured in the specification of antenna cross-polar discrimination (XPD). The XPD of an antenna for a given direction is the difference in dB between the peak co-polarized gain of the antenna and the cross-polarized gain of the antenna in the given direction. Cross-polarization can be effective in mitigating interference between adjacent systems. A typical cross-polarization isolation of 25–30 dB can be

INTEL-1020
U.S. Patent 8,219,129

achieved with most antennas today. This is sufficient to counter co-channel interference for QPSK and 16-QAM modulation schemes. Also, the Radiation Pattern Envelopes (RPEs) are independent of polarization. The performance of BS antennas is here divided into two electrical classes. Class 1 represents the minimum recommended performance. Class 2 antennas have enhanced RPEs and represent more favorable coexistence performance. Class 2 antennas are recommended for operation in environments in which interference levels could be potentially significant. Fixed BWA systems employ SS antennas that are highly directional, narrow-beam antennas. Although it is not as important for coexistence as the BS RPE, the RPE of the SS antenna is a factor in determining intersystem interference.

The performance of SS antennas is divided into three electrical classes. Class 1 is defined with moderate sidelobe characteristics and represents the minimum recommended performance. Class 2 and Class 3 antennas have enhanced RPEs and represent increasingly favorable coexistence performance. Most of the coexistence problems are the result of main-beam interference. The side lobe levels of the base station antennas are of a significant but secondary influence. The side lobe levels of the subscriber antenna are of tertiary importance.

Furthermore, mechanical parameters of the antenna also can affect coexistence. The mechanical characteristics that are considered are wind and ice loading, water ingress, temperature and humidity, and vibration. For deployments in hurricane-prone areas, more robust antenna systems may be required.

Factors related to the antennas can also be efficient mitigation techniques. These are antenna-to-antenna isolation, orientation, tilting, directivity, and polarization. Out of these changing the antenna orientation is especially effective in the case of interference arising from main-beam coupling. Near the service area boundaries, coexistence is improved with a high-performance antenna with high directivity as opposed to a broader range sectorized antenna or omnidirectional antenna.

5) It is recommended that the subscriber and base station equipment have the ability to detect and react to failures, either software or hardware, in a manner

INTEL-1020
U.S. Patent 8,219,129

to prevent unwanted emissions and interference. For example, in the absence of a correctly received downlink signal, the SS transmitter should be disabled.

These five recommendations are valid in all cases. Consider specifically the following co-channel and adjacent channel scenarios.

1) *The co-channel case.* The most problematic interference occurs between base stations. Careful deployment of fixed BWA systems is important in minimizing interference problems. Coordination is recommended between licensed service areas where both systems are operating co-channel (i.e., over the same fixed BWA frequencies) and where the service areas are in close proximity, (e.g., the shortest distance between the respective service boundaries is less than 60 km). The operators are encouraged to arrive at mutually acceptable sharing agreements that would allow for the provision of service by each licensee within its service area to the maximum extent possible. Under the circumstances where a sharing agreement between operators does not exist or has not been concluded, and where service areas are in close proximity, a coordination process should be employed. As described previously, distance is suggested as the first trigger mechanism for coordination between adjacent licensed operators. If the boundaries of two service areas are within 60 km of each other, then the coordination process is recommended. The number 60 km is obtained by calculating the radio horizon by $R = 4.12(\sqrt{h_1} + \sqrt{h_2})$, where h_1 and h_2 are in the heights of the antennas above clutter in meters, and R is in kilometers. Assuming that the BS antennas are 55 m above the radio clutter, the radio horizon will be 60 km. If the base stations are separated by a distance larger than the radio horizon, coexistence problems among them will not arise. Other forms of interference are still possible (for example, SS to BS), but they will be generally less severe. It is suggested that service area demarcation is used. IEEE 802.16 also recommends that government regulators specify trigger values of for each frequency band. If the trigger values are exceeded, the operator should try to modify the deployment or should coordinate with the affected operator. The coordination trigger values of –114 (dBW/m)/MHz (24, 26, and 28 GHz bands) and –111 (dBW/m 2)/ MHz (38 and 42 GHz bands) are recommended. These "triggers" should be applied before deployment and before each relevant system modification. If these trigger values are exceeded, there are several possible solutions. Increasing the distance between the interfering

INTEL-1020
U.S. Patent 8,219,129

transmitter and receive would obviously result in decreased interference. If the distance cannot be increased, the same effect can be achieved by lowering the transmitter power. As a mitigation technique in the co-channel case, frequency exclusion is a possible, albeit very undesirable, approach for avoiding interference. This practice involves dividing or segregating the spectrum so that neighboring licensees operate in exclusive frequencies, thus avoiding any possibility for interference. Instead of frequency exclusion, by retaining spare frequencies to use only when interference is detected, some potential co-channel and adjacent channel problems can be eliminated.

2) *Adjacent channel case in which the served territories overlap.* When operators are in the same area but on adjacent frequency bands, analysis and simulation indicate that deployment may require a guard frequency. In most co-polarized cases, the guard frequency should be equal to one equivalent channel of the widest bandwidth system. Where administrations do not set aside guard channels, the affected operators would need to reach agreement on how the guard channel is apportioned between them. This situation has one interesting special case—when multiple operators share an antenna site by using harmonized transmissions. For FDD systems, this implies that each operator's base station transmits in the same frequency sub-block (typically on a different channel) and that their terminals transmit in the corresponding paired sub-block. For TDD systems, harmonization implies frame, slot, and uplink/downlink synchronization. It is clear that planning is required for co-sited antennas of BSs. When tackling coexistence between FDD systems operating in adjacent frequency blocks in the same or overlapping areas with defined uplink and downlink frequency bands, co-siting of base station transmitters actually facilitates coexistence.

Alternative schemes might be available in the future. For example, technologies such as adaptive arrays or beam-steering antennas can focus a narrow beam towards individual users throughout the service area in real-time to avoid or minimize coupling with interfering signals. Beam-shaping arrays, which create a null in the main beam towards the interfering source, represent another possible approach towards addressing interference. Using more robust modulation and enhanced signal processing techniques can help in deployment scenarios where the potential for interference is high.

INTEL-1020
U.S. Patent 8,219,129

BWA BUSINESS AND TECHNOLOGY TRENDS

IEEE 802.16 is currently interested in the following technologies. First, all physical layers at present work above 2 GHz. However, it is possible and economically justified to have operation below 2 GHz. The MAC of IEEE 802.16 may be modified to support point-to-point operation [B133], and not only point-to-multipoint. Another initiative within IEEE 802.16 is to support mobility. Task Group 802.16e was established to support this technology [B30], [B97], [B99], [B131]. The solution that will be developed by 802.16e will likely be devoting one part of the superframe for fixed stations and one part for mobile stations. The further technical development of IEEE 802.16 will likely take place in these areas.[11]

There will be one billion Internet users in the year of 2005, and even more in future years. A relatively small, but growing percentage of them are accessing the Internet using broadband technology. Therefore, the potential market for WMAN is significant. To promote interoperability among IEEE 802.16 devices and to promote WMAN technology based on IEEE 802.16, several companies, including Intel, Fujitsu®, and Nokia, have formed a trade organization called WiMAX Forum. This trade organization will certify the compliance of IEEE 802.16 equipment and will promote it in the marketplace. Some analysts expect the first IEEE 802.16 products to be available in 2004–2005 and expect the total market to reach $1 billion by 2008. It remains to be seen whether these market predictions will be correct. IEEE 802.16 is very suitable for developing countries that lack telephone and cable infrastructure. In developed countries, the market opportunity in rural and suburban areas lacking DSL and cable is also significant. Even in urban areas, where DSL and cable services are available, many users will choose wireless because of its convenience. For example, while DSL works at distances up to 18 000 feet from a telephone company's central office, IEEE 802.16 works at distances up to 30 km from a base station.

One reason the broadband wireless Internet access market in the United States and other countries has remained small has been the lack of a standard, especially during the years of high growth in the telecommunications industry.

11. For the latest information on the 802.16 Working Group and ongoing projects, see http://grouper.ieee.org/groups/802/16/.

INTEL-1020
U.S. Patent 8,219,129

BWA business is different from the WLAN and WPAN businesses in one important aspect. The use of WPAN and WLAN devices does not require monthly service fees. As a result, service providers have been generally struggling to devise ways to raise revenue from WLAN and WPAN service. On the other hand, because BWA devices are used for Internet access, their use will certainly involve a monthly service fee. In addition, BWA can enable other revenue-generating services, such as interactive games, video on demand, etc. To make BWA systems even more competitive for the "first mile" connection to services, IEEE 802.16 provides maximum flexibility and relatively low "up-front" costs.

Overall, it seems plausible that, one way or another, wireless Internet access will succeed in the marketplace. At present in the U.S., UK, Sweden, and other countries, wireless broadband service is available in isolated areas. Examples in the U.S. are small service providers such as Mesh Networks Inc., in Florida, BroadBand Solutions in Utah, and other companies. These networks are not based on IEEE 802.16, but they support mobility, providing coverage at speeds of about 40–50 mph.

One question that remains unanswered is whether having licensed spectrum and unlicensed spectrum versions of the same service makes business sense. Some telecommunication service providers that bought spectrum at significant cost soon afterwards found themselves in an untenable business environment and ceased operation altogether.

In general, IEEE 802.16 makes the competition and cooperation with 3G and 4G cellular technologies much more interesting. Communications anytime, anywhere clearly requires wireless networks to seamlessly connect with other wireless networks. Flarion Technologies of Bedminster, NJ, and SOMA Networks™ of San Francisco, CA, are some of the companies that offer so-called 4G technologies, which integrate cellular communications and wireless data.

IEEE 802.16's major weakness is that it does not support mobility. While originally mobility was not required, recently demand has been increasing for mobile broadband wireless access. Within IEEE 802. there are two efforts to provide mobile BWA systems. First, IEEE 802.16 has created a task group to explore how to modify IEEE 802.16 so that it supports mobility. Also, a separate working group was established, IEEE 802.20, with the goal to develop an air interface for a technology achieving data rate of 1 Mb/s at speeds up to

INTEL-1020
U.S. Patent 8,219,129

250 km/hour at cell ranges of up to 15 km. The typical application is Internet access in trains and other fast-moving vehicles. IEEE 802.20 will work in licensed bands below 3.5 GHz. The modulation technologies will achieve spectral efficiencies of more than 1 bit/second/Hz/cell, which is more than double than that of today systems. IEEE 802.20 will support voice, video, and data services. VOIP, online gaming, and financial transactions are some of the possible applications. By providing ubiquitous mobile broadband connectivity the standard will be strongly competitive with DSL and cable, as well as 3G cellular technologies.

INTEL-1020
U.S. Patent 8,219,129

Copyright © 2004 IEEE

Chapter 5 Concluding Remarks

We are at the beginning of a significant change in wireless systems and services as IEEE 802 wireless standards continue to evolve. Future standards will require further innovations in architectures, spectrum allocation, and protocols [B12], [B50], [B101]. Ultimately, these areas are related, and innovations in one area will lead to significant developments in other areas. Convergence among the different wireless technologies is still in its infancy. Due to technological and business issues, convergence among the different wireless standards and truly ubiquitous connectivity will require a major change. This convergence includes wireless networks seamlessly operating with other wireless networks and even with wired networks. Future standards will have highly efficient use of the wireless spectrum. High-performance physical layers are needed to provide data rates on the order of 1 Gb/s. At these data rates, the coherence bandwidths become a smaller percentage of the channel bandwidths, and there will be a greater number of resolvable paths. Clearly, OFDM is very attractive at these higher data rates. Space-time coding and smart antennas provide higher spectral efficiencies through better spatial reuse. Furthermore, advanced signal processing (such as multiuser detection) can also help attain higher spectral efficiencies. Clearly, parametric control over the modulation and coding will be available, together with power levels and other related functions. The adaptation must be quick. Both centralized and ad-hoc networking must be efficiently supported. The MAC may need to include signaling information and the protocol for this seems an attractive research topic. Note that variations of OFDM include MC-CDMA and other coded schemes. MC-CDMA applies a spreading sequence on the frequency domain. Higher-level modulations can be used in some carriers and not in others. It is possible to combine all multiple access schemes—TDMA/FDMA/CDMA is feasible. The research into the best combination of multiple access scheme, modulation, and code is expected to continue. Note that to use spectrum efficiently, in addition to the multiple-access scheme, the duplexing scheme must be considered. QoS adaptability requires the QoS to be negotiated between the stations.

INTEL-1020
U.S. Patent 8,219,129

In the process of communication, resources need to be assigned and managed by a multilayer adaptive radio resource management (RRM). This RRM will include not only dynamic spectrum management, but also dynamic management of the multiple access schemes. Not only TDMA and CSMA schemes will be used, but more complex time/slot/space schemes will be provided where the resource will be assigned dynamically to achieve efficient multiple access for a variety of traffic types.

A technology that allows ultimate flexibility in many respects is software radio [B109]. Software-defined radio (SDR) is a radio with enough programmability, i.e., where enough is loosely defined. The SDR Forum is the main industry group promoting SDR and developing standards in SDR technology. While current technology limitations do not allow the full potential of software radio to be realized, the availability of inexpensive SDR solutions will be one of the key enabling technologies for the future wireless systems.

INTEL-1020
U.S. Patent 8,219,129

Copyright © 2004 IEEE

Bibliography

[B1] Aad, I. and Casteluccia, C., "Differentiation mechanisms for IEEE 802.11," IEEE INFOCOM 2001.

[B2] Aboba, B., Beadles, M., "The Network Access Identifier," RFC 2486, January 1999.

[B3] Aiello, R., Brethur, V., and Kareev, U., "UWB, a solution for location awareness in TG4 applications," IEEE 802.15-03/050r0, January 2003.

[B4] Alamouti, Siavash M., "A Simple Transmit Diversity Technique for Wireless Communications," *IEEE Journal on Select Areas in Communications*, Vol. 16, No. 8, pages 1451–1458, October 1998.

[B5] ATM Forum Specification af-uni-0010.002, ATM User-Network Interface Specification, Version 3.1, September 1994.

[B6] ATM Forum Specification af-sig-0061.000, ATM User-Network Interface (UNI) Signaling Specification, Version 4.0, July 1996.

[B7] ARIB STD-T66 Second Generation Low Power Data Communication System/Wireless LAN System 1999.12.14 (H11.12.14) Version 1.0. Association of Radio Industries and Businesses (ARIB). Japan.

[B8] Baldwin, K., Webster, M., and Halford, S., "Topics in MIMO channel modeling," IEEE 802.11-03-566r0, July 2003.

[B9] Bing, Benny, *Broadband Wireless Access*, Kluwer, Norwell, MA, 2000.

[B10] Benveniste, M., "Self-configurable wireless LAN systems," IEEE 802.11-02-211r0, March 2002.

[B11] Benveniste, M., "Wireless LANs and neighborhood capture," preprint.

[B12] Berezdivin, R., Breinig R., and Topp, R., "Next-generation wireless communication concepts and technologies," *IEEE Communications Magazine*, pp. 108-116, March 2002.

INTEL-1020
U.S. Patent 8,219,129

[B13] Bluetooth Special Interest Group, "Core Specification Version 1.1," Specification of the Bluetooth System, February 22, 2001. [Bluetooth_11_Specifications_Book.pdf]

[B14] Blunk, L., Vollbrecht, J., "PPP Extensible Authentication Protocol (EAP)," RFC 2284, March 1998.

[B15] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119, March, 1997.

[B16] Cali, F., Conti, M., and Gregori, E., " Dynamic tuning of the 802.11 protocol to achieve a theoretical throughput limit," IEEE/ACM Transactions on Networking, vol. 8, pp. 785-799, 2000.

[B17] Callaway, E., "TG4 battery life extension," IEEE 802.15-03/037r0, January 2003.

[B18] Cantoni, A. and Godara, L. C., "Fast Algorithm for Time Domain Broad Band Adaptive Array Processing," IEEE Trans. Aerops. Electr. Syst., AES-18, 682, September 1982.

[B19] Chayat, N., Tentative criteria for comparison of modulation methods, IEEE P802.11-97/96, Sept. 1997.

[B20] Chiasserini, C. F. and Rao, R. R., "Coexistence mechanisms for interference mitigation between 802.11 WLANs and Bluetooth," IEEE INFOCOM 2002.

[B21] Choi, W. Y. and Lee, S. K., "MAC throughput enhancement by dynamic dot11RTSThreshold," IEEE 802.11-03-0509r0, July 2003.

[B22] Clark, M., Leung, K., McNair, B., and Kostic, Z., "Outdoor IEEE 802.11 cellular networks: radio link performance," IEEE International Conference Communications, 2002.

[B23] Congdon, P., et al. "IEEE 802.1X Usage Guidelines," Internet draft, draft-congdon-radius-8021x-15.txt, July 2001.

[B24] Cooklev, T., "Dynamic bandwidth allocation and channel coding in providing QoS for wireless local area networks," International Conference Telecommunications, Papeete, Tahiti, February 2003.

INTEL-1020
U.S. Patent 8,219,129

[B25] Cooklev, T., Tzannes, M., Lee, D., and Lanzl, C., "Extended data rate 802.11a," IEEE 802.11-02/231r0, March 2002.

[B26] Droms, R., "Dynamic Host Configuration Protocol," RFC 2131, March 1997.

[B27] Dworkin, M., Draft recommendation for block cipher modes of operation: The CCM mode for authentication and confidentiality, NIST special publication 800-38C, Sept. 2003.

[B28] ECC "Strategic Plans for the Future. Use of the Frequency Bands 862-870 MHz and 2400-2483.5 MHz for Short Range Devices," May, 2002.

[B29] Edmonson, P. J., "Wearable computers and wireless connectivity," Wireless 2000, pp. 260-267, Calgary, Canada.

[B30] Eidson, B. and McKown, R., "WirelessMAN-SCa PHY extensions for mobility," IEEE C802.16e-03/27, May 2003.

[B31] Elders-Boll, H. and Dawid, H., "Implementation of a 480 Mb/s Viterbi decoder for 802.15.3a," IEEE 802.15-03/349r0, Sept. 2003.

[B32] ERC Recommendation 70-03, Relating to the use of Short Range Devices (SRD), April 2002.

[B33] ETSI, "Criteria for comparison," ETSI BRAN #9, July 1998.

[B34] ETSI EN 300 220-1: Electromagnetic Compatibility and Radio Spectrum Matters (ERM); Short Range Devices (SRD); Radio equipment to be used in the 25 MHz to 1000 MHz frequency range with power levels ranging up to 500 mW; Part 1: Technical characteristics and test methods.

[B35] ETSI EN 300 328-1: 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; Part 1: Technical characteristics and test conditions.

INTEL-1020
U.S. Patent 8,219,129

[B36] ETSI EN 300 328-2: 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; Part 2: Harmonized EN covering essential requirements under article 3.2 of the R&TTE Directive.

[B37] Faulhaber, G. and Farber, D. J., "Spectrum management: property rights, markets, and the commons," presented at the FCC, June 2002.

[B38] Fahmy, N., Todd, T., and Kezys, V., "Ad hoc networks with smart antennas using 802.11-based protocols," IEEE Int. Conference Communications 2002.

[B39] Federal Communications Commission (FCC), Document: 47 CFR, Part 15, section 15.247, USA

[B40] Federal Communications Commission (FCC), Part 15 regulations, March 13, 2003.

[B41] Federal Communications Commission (FCC), " Unlicensed NII devices in the 5 GHz frequency range," Federal register, vol. 62, No. 21, pp. 4649-4657, January 31, 1997.

[B42] Federal Communications Commission (FCC), "Unlicensed NII devices in the 5 GHz frequency range," Federal register, vol. 63, No. 147, pp. 40831-40837, July 31, 1998.

[B43] Fettweis, G. and Nitsche, G., "1/4 Gb/s WLANs," IEEE 802.11-02-320r0, May 2002.

[B44] FIPS Pub 198, The Keyed-Hash Message Authentication Code (HMAC), Federal Information Processing Standards Publication 198, US Department of Commerce/N.I.S.T., March 6, 2002.

[B45] Fixed wireless access system using quasi-millimeter-wave and millimeter-wave band frequencies—point-to-multipoint system, ARIB STD-T59, rev. 1, March 2000.

[B46] Ghosh, M., Ouyang, X., Dolmans, G., "On the use of multiple antennae for 802.11," IEEE 802.11-02-180r0, March 2002.

INTEL-1020
U.S. Patent 8,219,129

[B47] Hartsen, J. C., "The Bluetooth radio system," *IEEE Personal Communications Magazine,* pp. 28–36, February 2000.

[B48] Heegard, Chris et. al., "High-Performance Wireless Ethernet," *IEEE Communications Magazine,* Nov. 2001, pp. 64–73.

[B49] Herrera, J., et. al., "Switched-beam antennas in millimeter-wave band broadband wireless access networks," IEEE C802.16-03/09, July 2003.

[B50] Hirano, J., Yew, T., and Aramaki, T., "Global area network concept," IEEE 802.11-02-149r0, March 2002.

[B51] Holland, G., Vaidya, N., and Bahl, P., "A rate-adaptive MAC protocol for multi-hop wireless networks," ACM SIGMOBILE, pp. 236–250, 2001.

[B52] http://www.arib.or.jp/english

[B53] http://www.bluetooth.com

[B54] http://www.etsi.org

[B55] http://www.fcc.gov

[B56] http://www.ic.gc.ca

[B57] http://www.ieee802.org

[B58] http://www.ieee802.org/11

[B59] http://www.ieee802.org/15

[B60] http://www.ieee802.org/16

[B61] http://www.wi-fi.org

[B62] http://www.x9.org

[B63] http://www.zigbee.org

[B64] Iachella, S., Deployment engineering of WLAN systems, M.S. thesis, San Francisco State University, 2003.

INTEL-1020
U.S. Patent 8,219,129

[B65] IEEE 802™-2001, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.[12]

[B66] IEEE 802.1Q™, 2003 Edition, IEEE Standards for Local and metropolitan area networks—Virtual Bridged Local Area Networks.

[B67] IEEE 802.1X™-2001 IEEE Standards for Local and Metropolitan Area Networks: Port-Based Network Access Control.

[B68] IEEE Std 802.11™, 1999 Edition (R2003) (ISO/IEC 8802-11: 1999), IEEE Standard for Information Technology—Telecommunications and Information Exchange between Systems—Local and Metropolitan Area Network—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.

[B69] IEEE Std 802.11a™-1999 (R2003) (ISO/IEC 8802-11:1999/Amd 1:2000(E)), 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—Amendment 1: High-speed Physical Layer in the 5 GHz band.

[B70] IEEE 802.11b™-1999 (R2003) Supplement to 802.11-1999, Wireless LAN MAC and PHY specifications: Higher speed Physical Layer (PHY) extension in the 2.4 GHz band.

[B71] IEEE 802.11e™/D8, IEEE Draft Amendment to Standard for Information Technology—Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements—Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications: Medium Access Control (MAC) Quality of Service (QoS) Enhancements.

[B72] IEEE 802.11F™-2003, IEEE Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation.

12. 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/.

INTEL-1020
U.S. Patent 8,219,129

[B73] IEEE 802.11g™-2003, 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—Amendment 4: Further Higher-Speed Physical Layer Extension in the 2.4 GHz Band.

[B74] IEEE 802.11h™-2003, IEEE Standard for Information technology—Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Spectrum and Transmit Power Management Extensions in the 5GHz band in Europe.

[B75] IEEE 802.11i™-2004, IEEE Standard for Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control Security Enhancements.

[B76] IEEE P802.11k™, Draft Specification for radio resource management, 2003.

[B77] IEEE 802.12™, 1998 Edition (ISO/IEC 8802-12:1998), IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area network—Specific requirements— Part 12: Demand-priority access method, physical layer and repeater specifications for 100 Mb/s Operation.

[B78] IEEE Std 802.15.1™-2002—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).

[B79] 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.

INTEL-1020
U.S. Patent 8,219,129

[B80] IEEE Std 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).

[B81] IEEE Std. 802.15.4™-2003, IEEE Standard for Information Technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 15.4: Wireless Medium Access Control and Physical Layer Specifications for Wireless Personal Area Networks.

[B82] IEEE 802.16™-2001, IEEE Standard for Local and Metropolitan Area Networks—Part 16: Air Interface for Fixed Broadband Wireless Access Systems.

[B83] IEEE Std 802.16a™-2003, IEEE Standard for Local and metropolitan area networks—Part 16: Air Interface for Fixed Broadband Wireless Access Systems—Amendment 2: Medium Access Control Modifications and Additional Physical Layer Specifications for 2-11 GHz.

[B84] IEEE 802.16.2™-2004, IEEE Recommended Practice for Local and metropolitan area networks—Coexistence of Fixed Broadband Wireless Access Systems.

[B85] IETF RFC 2459: Internet X.509 Public Key Infrastructure Certificate and CRL Profile, Internet Request for Comments 2459, January 1999.

[B86] IETF RFC 3279: Algorithms and Identifiers for the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) profile, Internet Request for Comments 3279, April 2002.

[B87] IETF RFC 3280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, Internet Request for Comments 3280, April 2002.

[B88] Industry Canada, IC. Document: GL36. Ottawa, Canada.

[B89] ISO/IEC 10038 Information technology—Telecommunications and information exchange between systems—Local area networks—Media Access Control (MAC) Bridges, (also ANSI/IEEE Std 802.1D-1993), 1993.

INTEL-1020
U.S. Patent 8,219,129

[B90] ISO/IEC 15802-3: 1998, Information Technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 3: Media Access Control (MAC) Bridges.

[B91] ISO/IEC 8802-3 Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, (also ANSI/IEEE Std 802.3-1996), 1996.

[B92] ITU-R Recommendation P.530-8, Propagation data and prediction methods required for the design of terrestrial line-of-sight systems.

[B93] Jai, R., Multipath Fading Channel Modeling for High Speed PHY, IEEE P802.11-97/135, Nov. 1997.

[B94] Jeon, T., et. al., "Optimal combining of STBC and spatial multiplexing for MIMO OFDM," IEEE 802.11-03-513r0, July 2003.

[B95] Johnson, D. and Dudgeon, D., *Array Signal Processing,*Prentice-Hall, Englewood Cliffs, NJ. 1993

[B96] Jun Li, Weinstein, S., Zhang, J., and Tu, N., "Public access mobility LAN: extending the wireless Internet into the LAN environment," *IEEE Trans. Wireless Communications,* Jun. 2002, pp. 22–30.

[B97] Kaitz, T., et. al., "Initial PHY proposal for 802.16e," IEEE C802.16e-03/03, January 2003.

[B98] Kamerman, A. and Monteban, L., "WaveLAN-II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal,* pp. 118–133, 1997.

[B99] Kitroser, I., "802.16e handoff description," IEEE C802.16e-03/55, Sept. 2003.

[B100] Kleindl, G., "Signaling for adaptive modulation," IEEE 802.11-03-283r0, May 2003.

INTEL-1020
U.S. Patent 8,219,129

[B101] Khun-Jush, J., "Integration of WLAN and wide area mobile networks," IEEE 802.11-02-106, January 2002.

[B102] Kowalski, J., "Link measurement results give thumbs-up to MAC FEC!," IEEE 802.11-01/422r0, July 2001.

[B103] Lewis, B. "Mesh networks for fixed broadband wireless access," IEEE C802.16-03/10r1, July 2003.

[B104] Liberti, J.C. and Rappaport, T.S., *Smart Antenna for Wireless Communications,* Prentice-Hall, Englewood Cliffs NJ, 1999.

[B105] Lien, J., Wu, C., Liu, H., and Lung, Y., "An adaptive multirate IEEE 802.11 wireless LAN," International Conf. On Inf. Networking, pp. 411–418, 2001.

[B106] McCann, S., "Proposed wireless interworking group (WIG) baseline document," IEEE 802.11-03-004r0, January 2003.

[B107] McCorkle, J., "DS-CDMA—The technology of choice for UWB," IEEE 802.15-03/277r0SG3a, July 2003.

[B108] Milewski, A., "Periodic Sequences with Optimal Properties for Channel Estimation and Fast StartUp Equalization," IBM J. Res. Develop., Vol. 37, No.5, Sept. 1983, pp. 426–431.

[B109] Mitola, J., *Software radio architecture,* John Wiley, New York, NY, 2002.

[B110] Mulhens, O., "QoS-enabled multihop meets high data rate," IEEE 802.11-02-333r0, May 2002.

[B111] Mulhens, O., "Cluster-based multihop networking with WPANs," IEEE 802.15-02/312r0, July 2002.

[B112] Nasipuri, A., et al, "A MAC protocol for mobile ad-hoc networks using directional antennas," Proc. IEEE Wireless Communications and Networking Conf (WCNC), Sept. 2000.

[B113] NIST FIPS Pub 197: Advanced Encryption Standard (AES), Federal Information Processing Standards Publication 197, US Department of Commerce/N.I.S.T., November 2001.

INTEL-1020
U.S. Patent 8,219,129

[B114] O'Hara, Robert and Petrick, Al, *IEEE 802.11 Handbook: A Designer's Companion,* IEEE Standards Information Network/IEE Press, 1999.

[B115] Okamoto, G. and Xu, G., "Throughput multiplication of wireless LANs for multimedia services: spread-spectrum with SDMA," IEEE Vehicular Technology Conference, April 1996.

[B116] Pahlavan, A. Levesque, *Wireless Information Networks,* John Wiley, New York, NY, 1995.

[B117] Pazhyannur, R., "WLAN-cellular interworking," IEEE 802.11-03-193r1, March 2003.

[B118] PKCS#1 v2.1, RSA Cryptography Standard, RSA Laboratories, June 14, 2002.9

[B119] Porcino, D., Hirt, W., "Ultra-wideband radio technology: potential and challenges ahead," *IEEE Communications Magazine,* vol. 41, pp. 66–74, July 2003.

[B120] Prabhu, B. J. and Chockalingam, A., "A routing protocol and energy efficient techniques in Bluetooth scatternets," IEEE Int. Conference Communications, 2002.

[B121] Rappaport, T., *Wireless Communications: Principles and Practice, Prentice-Hall,* Upper Saddle River, NJ, 2002.

[B122] Recommendation ITU-R P.1238-1, Propagation data and prediction methods for the planning of indoor radiocommunication systems and radio local area networks in the frequency range 900 MHz to 100 GHz. Question ITU-R 211/3.

[B123] Regnier, J., "Benefits of smart antennas in 802.11 networks," IEEE 802.11-03-025r0, 2003.

[B124] Rigney, C., "RADIUS Accounting," RFC 2866, June 2000.

[B125] Rigney, C., Rubens, A., Simpson, W., Willens, S., "Remote Authentication Dial In User Service (RADIUS)," RFC 2865, June 2000.

[B126] Rigney, C., Willats, W., Calhoun, P., "RADIUS Extensions," RFC 2869, June 2000.

INTEL-1020
U.S. Patent 8,219,129

[B127] Rivest, R., Dusse, S., "The MD5 Message-Digest Algorithm," RFC 1321, April 1992.

[B128] Standard specification for telecommunications and information exchange between roadside and vehicle systems—5 GHz band Dedicated Short-Range Communications (DSRC) Medium-Access Control (MAC) and Physical Layer (PHY) specifications, American Society for Testing and Materials, ASTM E2213, Nov. 2002, Aug. 2003.

[B129] SEC 1: Standards for Efficient Cryptography, Elliptic Curve Cryptography, Version 1.0, Certicom Research, September 20, 2000.

[B130] Siep, Tom, *An IEEE Guide: How to Find What You Need in the Bluetooth Spec.*, IEEE Standards Information Network/IEE Press, 2000

[B131] Seggal, Y., et. al., "OFDMA modification for mobility," IEEE C802.16e-03/25r2, May 2003.

[B132] Stallings, W., *Cryptography and Network Security: Principles and Practice*, Prentice-Hall, Upper Saddle River, NJ, 2002.

[B133] Stanwood, K. and Marks, R., "A proposal to add point-to-point option to 802.16 MAC," IEEE 802.16-03/11, July 2003.

[B134] Unbehaun, M., "Coverage Planning for Indoor Wirelss LAN Systems," ICC 2001, Helsinki, Finland.

[B135] Verdu, S., *Multiuser Detection*, Cambridge University Press, 1998.

[B136] Viterbi, A. J., Wolf, J. K., Zehavi, E., and Padovani, R., "A pragmatic approach to trellis-coded modulation," *IEEE Communications Magazine*, July 1989, pp. 11–19.

[B137] Veeraraghavan, M., Cocker, N., and Moors, T., "Support of voice services in IEEE 802.11 wireless LANs," IEEE INFOCOM 2001.

[B138] Veeraraghavan, M., Cocker, N., Moors, T., "Support of voice services in IEEE 802.11 wireless LANs," IEEE INFOCOM 2001.

[B139] Welborne, M., et. al. "Merger #2 proposal DS-CDMA," IEEE 802.15-03/334r3, Sept. 2003.

INTEL-1020
U.S. Patent 8,219,129

[B140] Whiting, D., Housely, R., Ferguson, N., "Counter with CBC-MAC (CCM)," RFC 3610, Sept. 2003.

[B141] Win, M. Z., Cassioli, D., and Molisch, A., "The ultra-wide bandwidth indoor channel: from statistical modeling to simulations," IEEE 802.15-02/284r0SG3a, June 2002.

[B142] Win, M. Z., Cassioli, D., and Molisch, A., "A statistical model for the UWB inddor channel," IEEE 802.15-02/285r0SG3a, June 2002.

[B143] Wolf, J. K. and Zehavi, E., "P2 codes: Pragmatic trellis codes utilizing punctured convolutional codes," *IEEE Communications Magazine,* February 1995, pp. 94–99.

[B144] Wu, H., et. al. "IEEE 802.11 distributed coordination function (DCF) analysis and enhancement," IEEE Int. Conf. Communications, 2002.

[B145] Xao, Y. and Rosdahl, J., "Throughput analysis for 802.11a higher data rates," IEEE 802.11-02-138r0, March 2002.

[B146] Yanover, V., et. al., "Some thoughts for the New Concepts forum," IEEE C802.16-03/13, July 2003.

[B147] Yeh, J.-Y. and Chenm, C., "Support of multimedia services with the IEEE 802.11 MAC protocol," IEEE ICC 2002.

[B148] Yergeau, F., "UTF-8, a transformation format of Unicode and ISO 10646," RFC 2044, October 1996.

[B149] Yu, H., et. al., "MIMO OFDM with antenna selection," IEEE 802.11-03-514r0, July 2003.

[B150] Yu, H., "HDR 802.11a solution using MIMO-OFDM," IEEE 802.11-02-294 r0, May 2002.

[B151] Zorn, G., Leifer, D., Rubens, A., Shriver, J., Holdrege, M., Goyret, I., "RADIUS Attributes for Tunnel Protocol Support," RFC 2868, June 2000.

[B152] Zorn, G., Mitton, D., Aboba, B., "RADIUS Accounting Modifications for Tunnel Protocol Support," RFC 2867, June 2000.

INTEL-1020
U.S. Patent 8,219,129

Copyright © 2004 IEEE

Glossary

access point: A special station that allows a network to connect with another network.

acknowledgment field: Bit present in every frame that will, if set, require acknowledgment.

active service flow: A service flow requesting and being granted bandwidth for transporting data packets.

adaptive modulation: Different modulation for each subscriber station.

ad-hoc traffic indication message (ATIM) window: The amount of time after every beacon that a device must stay awake.

admitted service flow: A service flow with resources assigned but not yet activated.

alternating wireless medium access: Collaborative coexistence method where two networks with separate links have their transmissions scheduled in different time segments.

authentication key: Secret key derived in initialization and shared between two or more parties as the base for all their security transactions. The four kinds are unit, combination, temporary, and initialization. Also called link key.

authentication node: In mesh mode, a node that verifies the digital certificate of new stations, acting as a BS.

authorization module: BS module acting as QoS manager, approving or denying changes to the QoS parameters.

bandwidth stealing: When a station uses bandwidth already given for another request.

beacon frame: Management frame that maintains the synchronization of the local timers in the stations and delivers protocol-related parameters.

INTEL-1020
U.S. Patent 8,219,129

bridges: Layer 2 devices used to connect network segments at the MAC layer.

broadband: In ITU terminology, transmission rates greater than 1.5 Mb/s.

bursty traffic: Signal traffic with messages of arbitrary length separated by intervals of random duration, such as office data communication (e.g., e-mail, Internet access, file sharing).

CAT5: Special cables required for Ethernet wiring.

Channel Estimation Interval (CEI): Vendor-specific interval, typically 20 ms, that defines how often the BS must ask an SS for updated channel state information.

channel time assignment (CTA): Dedicated time slot from the parent piconet that determines the piconet identifier of child or neighbor piconets.

ciphertext: In a cryptographic system, an encrypted message.

contention window (CW): The maximum extent of the backoff time.

contention-free period: Time period for WLANs in which a point coordinator is active.

data region: OFDMA slot with time and subchannel dimensions. The time dimension denotes consecutive OFDM symbols in the time domain, and the subchannel dimension denotes blocks within the OFDM symbols.

delay spread: Difference between the arrival times of the first and last paths of signals.

digital modulation: As defined by the FCC, when the combination of data rate, coding, and modulation method has a 6 dB bandwidth greater than 500 kHz and a maximum transmitted spectral density of less than +8 dBm/3 KHz.

distribution services: Services that manage traffic when an access point is present: association, disassociation, distribution, integration, and reassociation.

donut pattern: Omnidirectional pattern radiating equally in all directions around an axis or azimuth in a horizontal plane with a circular pattern outward.

Doppler spread: Measure of spectral broadening caused by relative mobility.

INTEL-1020
U.S. Patent 8,219,129

encapsulation protocol: One of two protocols of MAC security sublayer. Defines a set of supported cryptographic suites and the rules for applying them to a MAC payload.

encapsulation: The way cryptographic data is constructed from plaintext data.

Encryption Control: Bit field in the MAC header indicating whether the payload is encrypted.

fading: Fluctuation in the received signal, the result of the randomness with which the multipath components add at the receiver. Fading depends on both the spatial position of the receiver, and on the frequency of transmission.

FCH OFDM symbol: Symbol following the long preamble on downlink that conveys the burst parameters and other control information.

Fresnel zone: Area around the optical line of sight that must be free of obstacles to avoid random signal loss.

guaranteed time slots (GTSs): Dedicated slots allocated by the network-coordinator that comprise the contention-free period. Used for periodic traffic.

HiperAccess: ETSI technology for wireless broadband access, much like IEEE 802.16.

HiperLAN: ETSI technology for wireless LAN, much like IEEE 802.11.

interframe space (IFS): Time interval between frames. The five types are short IFS (SIFS), point coordination function IFS (PIFS), distributed IFS (DIFS), arbitration IFS (AIFS), and extended IFS (EIFS).

interleaving: A technique to reduce the statistical dependence of errors in which the symbols in one code block are not transmitted in consecutive order, but instead are interspersed among other symbols.

interworking: The ability to perform a task by two systems using different set of rules.

jitter: Delay variation.

link key: See *authentication key.*

INTEL-1020
U.S. Patent 8,219,129

link-level authentication: Open-system and shared-key authentication, defined in 802.11 in 1997.

log-normal shadowing: An effect showing random path loss according to a Gaussian (or normal) distribution for the free-space path loss model, which provides the average value.

multipath: Phenomenon by which a signal will reach the receiver by traveling multiple paths.

multipoint: Architecture of fixed BWA system that includes base stations, subscriber stations and, in some cases, repeaters. May be either point-to-multipoint (PMP) and multipoint-to-multipoint (MP-MP).

neighbor: (1) Stations with which a node has direct links, "one hop" away from node. Architectural element of mesh systems. (2) Dependent piconet whose coordinator does not participate in the parent piconet.

nonce: Abbreviation for "number once," the random number sent by a device in an authentication effort.

offline classification: Channel classification when a piconet is not established or is on hold. May involve RSSI measurements.

packet traffic arbitration: Collaborative coexistence method where 802.11b and 802.15.1 networks can decide what type of packet to send to minimize the interference between them.

periodic traffic: Signal traffic with an approximately constant data rate and a small interarrival variance between messages, such as voice and video.

per-packet key: A concatenation of an initialization vector (IV) and a private encryption key. Also called a seed.

piconet: Basic ad-hoc network of a wireless personal area network.

plaintext: In a cryptographic system, a message to be encrypted.

poll interval: Time interval equal to the maximum time between subsequent transmissions from the master to a particular slave on an ACL link, used to provide dynamic bandwidth allocation and latency.

INTEL-1020
U.S. Patent 8,219,129

power control: Adjustment of power levels in a monotonic sequence, required for class 1 equipment.

privacy key management (PKM): One of two protocols of the security sublayer. Describes how the BS distributes and synchronizes keys in a secure fashion.

provisioned service flow: Service flow with an assigned Service Flow ID but no reserved resources. A flow not immediately activated. Also called a deferred service flow.

puncturing: A method to reduce the number of codeword bits and increase the rate of the code.

receiver sensitivity: Certain minimum value of the power of a received signal. Above this point, communication is possible.

repeaters: Layer 2 devices used to connect network segments at the PHY layer.

routers: Devices used to connect network segments at the network layer.

Rx/Tx transition gap: Interval between the uplink burst and the subsequent downlink burst that allows time for the BS and SS to switch modes (BS from receive to transmit, SS from transmit to receive). Begins on physical slot boundary.

scatternet: A collection of overlapping piconets.

security association: Set of security information shared by a BS and SS. Includes traffic encryption keys and cipher block chaining (CBC) initialization vectors. May be primary, static, or dynamic.

seed: See *per-packet key.*

service access point: The interface between any two protocol layers.

service flow: Unidirectional flow of packets provided to a particular QoS according to the QoS parameter set.

session: The finite time interval in which units participate in a piconet.

slot time: The time unit of the backoff timer.

INTEL-1020
U.S. Patent 8,219,129

sniff mode: Mode used to reduce the duty cycle of the slave, where the master can transmit in regularly spaced, specified time slots.

snooze mode: Low-power mode to be used when there are no pending messages and no activity from the devices. The coordinator stops transmitting but keeps track of superframe boundaries and listens for messages.

station services: Services that are part of every station: authentication, deauthentication, key distribution, data authentication, replay protection, privacy, and delivery of data.

station: A logical device that participates in a network. Consists of a physical layer and a medium access control layer.

superframe: Frame with two periods: contention-free (when DCF is active) and contention (when PCF is active).

targeted beacon transmission time (TBTT): The intended arrival point of the next beacon frame, announced in every beacon frame.

transition security network: Inherently insecure network in which some equipment supports advanced security algorithms and some does not.

transmission opportunity: The interval of time when a particular station has the right to initiate transmissions onto the wireless medium.

turnaround time: The time for a device to switch from receive to transmit mode (RxTx turnaround time), or from transmit to receive mode (TxRx turnaround time).

Tx/Rx transition gap: Interval between the downlink burst and the subsequent uplink burst that allows time for the BS and SS to switch modes (BS from transmit to receive, SS from receive to transmit). Begins on physical slot boundary.

U-NII bands: Unlicensed national information infrastructure bands, 5.15–5.825 GHz.

unscheduled service periods: Service periods in response to a poll.

WirelessHUMAN: Protocol for operating in the unlicensed frequencies in the 2–11 GHz bands.

INTEL-1020
U.S. Patent 8,219,129

Copyright © 2004 IEEE

Index

Numerics

802.11 3, 19, 20, 22, 24, 45, 107, 110, 133, 134, 135, 144, 165, 166, 167, 168, 169, 170, 171, 176, 177, 178, 179, 183, 188, 195, 199, 207, 214, 216, 217, 228, 247, 271, 290, 304, 311, 316
MAC 70–96
Physical layer 46, 96
Security 49–66
802.11a 107, 110
802.15 19, 21, 24, 70, 135–173, 176, 177, 178, 179, 180, 181, 182, 183, 187, 188, 192, 195, 199, 201, 202, 204, 207, 214, 215, 216, 217, 229, 271, 316
MAC layer 181–195
Physical layer 137–138, 195–201
802.16 19, 24, 70, 228–331
802.20 331, 332

A

Access point 48, 52, 55, 57, 66, 67, 68, 70, 77, 78, 79, 81, 82, 112, 117, 122, 233
ACL 145, 146, 147, 148, 150, 151, 153, 163, 168, 169, 171, 172, 174, 186, 213
Advanced antenna systems 32, 261
Advanced encryption standard 61, 214, 215, 216
Antenna 120–123, 261–265, 302, 323, 326
Antenna diversity 302
Antenna polarization 120, 323
Attenuation 29
Authentication 50, 52, 54, 55, 57, 61, 70, 158, 186, 248, 250, 251, 252, 253, 261
Automatic repeat request 40, 142, 149, 171, 236, 257, 265, 266, 267, 289, 296

B

Backoff 75, 208, 247
Bandwidth 235, 247, 326
Beacon 91, 94, 213
BER 71, 169, 170, 171, 172, 176, 177, 270, 311
Binary phase-shift keying (BPSK) 109, 110
Bit 71
Bluetooth 18, 24, 135–167, 171–179, 201
Connection modes 151–152
Physical layer 137–138, 195–201
Security 154–162
BPSK 101, 106, 109, 110, 195, 207, 291, 296, 307, 315
BS, base station 226, 227, 237, 238, 239, 240, 241, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 254, 255, 256, 257, 258, 259, 260, 261, 262, 263, 264, 268, 269, 271, 273, 277, 278, 283, 284, 285, 286, 287, 309, 310, 313, 314, 317, 318, 319, 320, 326, 327, 328
BWA 4, 225, 226, 227, 287, 316, 317, 318, 319, 323, 324, 325, 327, 328

C

CCA 110

INTEL-1020
U.S. Patent 8,219,129

CCK 99
CINR, carier-to-interference-plus-noise-ratio 302, 311
Coexistence 117, 164–178, 216, 316, 324
Collision 20
Comparing WPAN and WLAN 133–135
Connection 134, 145, 237, 241, 249, 251
Contention window 73, 86, 87
Contention-free 77, 79, 80, 84, 91, 187, 188, 189, 191, 192, 193, 212
Convolutional coding 300, 305
CSMA 20, 37, 46, 70, 73, 78, 112, 117, 134, 144, 187, 188, 202, 207, 208, 209, 212, 246, 336
CTS 76, 121

D

DFS 47, 111, 257, 267, 268, 287, 322, 323
Downlink 257, 263, 273, 274, 275, 278
DSSS 8, 15, 97, 124, 207

E

EIRP (effective radiated power relative to isotropic) 25, 177, 285, 323, 325
Encryption 50, 60, 66, 159, 249, 251, 255, 256
Enhanced distributed coordination function 85, 89, 90, 91
Error 71, 74, 177
ETSI 14, 15, 24

F

Fading 29
Fast Fourier transform (FFT) 107, 109, 110
FCC 5, 6, 8, 10, 11, 14, 15, 45, 47, 201
FDD 234, 235, 264, 271, 272, 273, 274, 287, 288, 307, 308, 313, 318, 329
FDMA 335
FEC, forward error control 142, 143, 144, 147, 148, 161, 171, 172, 272, 274, 275, 276, 277, 281, 284, 289, 296, 304, 305, 311, 313, 315
FFT 106, 107, 109, 110, 304, 312, 313
Frequency band 116, 195, 206
Frequency hopping 9, 15, 41, 173

G

Gray coding 282, 291, 296, 298, 299, 305

H

Hybrid coordination function 84, 89, 90, 91, 93, 94

I

IETF 52, 54, 68
Interference 165, 316, 317, 320, 325
Internet 1, 37, 48, 52, 54, 133, 178, 225, 232, 233, 270, 330, 331
ISM 7, 8, 10, 13, 14, 20, 21, 103, 116, 137, 165, 166, 195, 202, 204, 205
ITU 24, 149, 225, 271, 326

M

MAC 19, 20, 21, 36, 37, 39, 40, 41, 42, 43, 46, 47, 49, 52, 55, 61, 66, 83, 99, 110, 136, 137, 144, 165, 167, 176, 178, 179, 180, 199, 209, 215, 233, 272, 274, 276, 277, 278, 285, 302, 305, 335
802.11 70–96
Fixed broadband wireless (802.16) 228, 231–268, 269
High-rate WPAN (802.15.3) 181–195
Low-rate WPAN (802.15.4) 203, 207–216
Security 213–216, 240, 248–256
Medium-access control (MAC) 110

INTEL-1020
U.S. Patent 8,219,129

MIC 60, 61, 62, 214
Multimedia 228

O

OFDM 9, 32, 41, 104, 106, 107, 109, 110, 113, 195, 229, 287, 303, 304, 307, 308, 309, 311, 312, 313, 314, 315, 335
Orthogonal frequency-division multiplexing (OFDM) 107, 109, 110

P

Packet 37, 148, 178, 232
Packet binary convolutional coding 47, 99, 101, 113
PER 166, 170, 172
Periodic traffic 37
Physical layer 19, 21, 40, 41, 42, 46, 72, 86, 99, 104, 136, 137, 167, 176, 178, 179, 182, 184, 188, 193, 194, 195, 196, 199, 200, 204, 206, 207, 228, 230, 238, 247, 256, 260, 262, 263, 269, 270, 271, 272, 285, 286
802.11 96
802.15.3 195–201
Bluetooth (802.15.1) 137–138
Fixed broadband wireless (802.16) 269–316

Q

QAM 106, 109, 110, 138, 195, 196, 197, 198, 272, 274, 281, 282, 284, 291, 292, 293, 294, 295, 296, 297, 298, 299, 311, 327
QPSK 100, 101, 106, 195, 196, 199, 200, 205, 272, 274, 275, 281, 284, 291, 296, 301, 302, 305, 308, 309, 311, 327
Quadrature amplitude modulation (QAM) 109, 110
Quality of Service 46, 47, 57, 71, 80, 83, 84, 89–94, 96, 135, 141, 144, 151, 174, 179, 181, 187, 190, 194, 202, 212, 218, 229, 233, 236, 237, 241, 244, 245, 247, 335

R

Radio frequency 7, 25, 72, 122, 141, 149, 167, 204, 205, 239, 281, 282, 286, 287, 316
Randomization, scrambling 276, 284
Receiver 311, 325
Reed-Solomon 277, 278, 281, 289, 290, 305
RTS 76, 92, 93, 121, 193

S

SCO 145, 146, 147, 149, 150, 151, 153, 163, 168, 169, 171, 172, 174
SNR, signal-to-noise ratio 166, 229, 261, 268, 281, 304
Software radio 336
Spectrum 5, 6
Speech 2, 84, 149, 245
Spread-spectrum 8
SS, subscriber station 226, 227, 234, 237, 238, 239, 240, 241, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 254, 255, 257, 258, 260, 262, 263, 264, 268, 269, 271, 273, 274, 277, 278, 281, 283, 284, 285, 291, 309, 311, 313, 314, 317, 318, 319, 320, 326, 327, 328
Standards 13, 17, 18, 19, 20, 21, 133

T

TCM, trellis coded modulation 199, 289, 290, 291, 292, 293, 294, 295
TCP/IP 1, 244

INTEL-1020
U.S. Patent 8,219,129

TDD 144, 154, 235, 236, 260, 264, 271, 273, 274, 287, 288, 307, 308, 313, 329
TDMA 37, 38, 187, 189, 230, 271, 272, 274, 275, 289, 301, 335, 336
TKIP 54, 61
TPC 47, 111, 257, 299, 322, 323
Turbo coding 280, 300

U

UDP 68, 69
UNII 47, 104, 112
Unlicensed 8, 13
Uplink 244, 257, 263, 283, 285

V

Video 84

W

WEP 58, 59
WLAN 4, 19, 20, 45, 46, 58, 67, 72, 112, 113, 117, 122, 124, 133, 134, 165, 167, 168, 171, 227, 323, 331
WPAN 4, 19, 21, 133, 227, 331
Coexistence 164–178
High Rate 179–201
Low Rate 201–218

INTEL-1020
U.S. Patent 8,219,129