2nd Edition

Covers 802.11a, g, n & i

Creating & Administering Wireless Networks

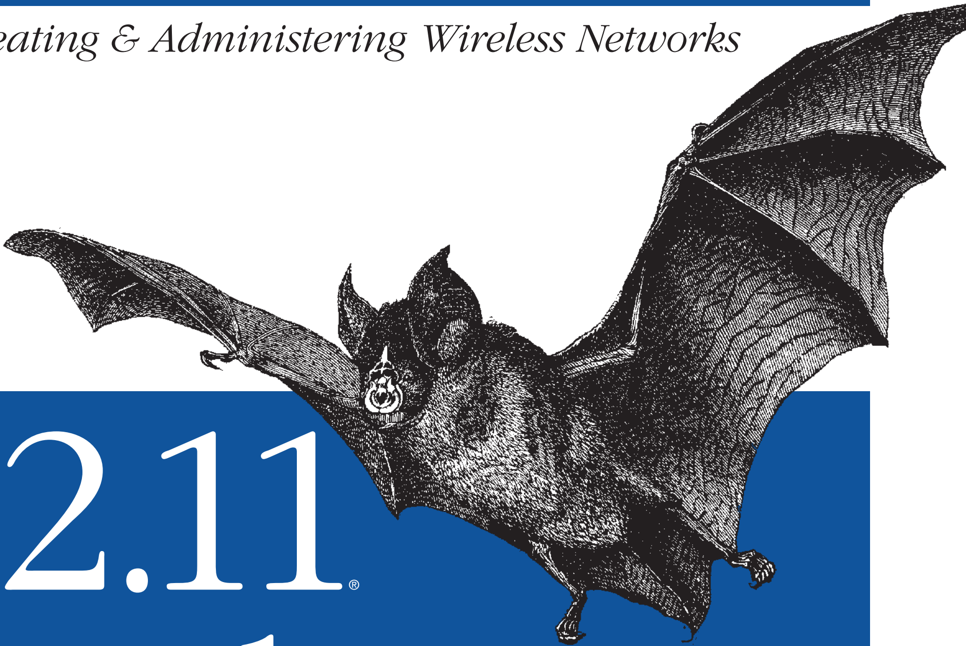

802.11® Wireless Networks

The Definitive Guide

O'REILLY®

Matthew S. Gast

CHAPTER 15

A Peek Ahead at 802.11n: MIMO-OFDM

802.11 task group N (TGn) has an interesting goal. Most IEEE task groups focus on increasing the peak throughput, making data fly as fast as possible during the time it is being transmitted. TGn's goal is to achieve 100 Mbps net throughput, after subtracting all the overhead for protocol management features like preambles, interframe spacing, and acknowledgments. Although the goal is 100 Mbps net throughput, the final proposal seems certain to blow past that number, and offer many times that throughput in maximum configurations. There are two roads to 100 Mbps: improve the efficiency of the MAC, increase the peak data rate well beyond 100 Mbps—or both.

Six complete proposals were made to the group creating the eventual 802.11n, but support has coalesced around two main proposals, from groups named TGnSync and WWiSE (short for "World-Wide Spectrum Efficiency"). Both camps have chipmakers. Atheros, Agere, Marvell, and Intel are part of TGnSync; Airgo, Broadcom, Conexant, and Texas Instruments are the core of WWiSE. However, quite a few manufacturers of electronic devices that might use 802.11 (Cisco, Nokia, Nortel, Philips, Samsung, Sanyo, Sony, and Toshiba) have also become part of the effort, and they are disproportionately represented in TGnSync.

At a very high level, both proposals are similar, though they differ in the emphasis on increasing peak data rates versus improving efficiency. Each of them makes use of multiple-input/multiple-output (MIMO) technology in several configurations and provides for backwards compatibility with installed systems in the same frequency band. Both support operation in the current 20 MHz channels, with provisions to use double-width 40 MHz channels for extra throughput.

As the standards war is fought across the globe at IEEE meetings, a "pre-N" access point has already hit the streets, based on Airgo's chipset. Purchasing it well before the standards process is underway is a roll of the dice. When most "pre-G" products were brought to market, the task group had begun to work in earnest on a single proposal. TGn is currently in the "dueling proposal" stage right now, and there is no guarantee that an early device will be firmware upgradeable to the final 802.11n standard. 802.11n is likely to be the last chance to standardize a PHY this decade. Developing a standard is as much political engineering as technical engineering. IEEE rules require that a proposal get a 75% supermajority vote before becoming the basis for a standard. As this book went to press, TGnSync was garnering a clear majority of support, but was still falling short of the necessary 75%. I expect that features from competing proposals will be incorporated into the working document to bring the vote count to the necessary level. As a result, this chapter describes both of the main competing proposals. Although TGnSync will probably be the basis for the 802.11n specification, some horse trading will likely result in a few WWiSE features being incorporated.

This chapter describes both the WWiSE and TGnSync proposals. The final standard will have some resemblance to both of them, and will likely pick and choose features from each. Fortunately, many basic concepts are shared between the two. As you read this, keep in mind that the proposals themselves may have changed quite a bit since the drafts upon which this chapter was based were written.

Common Features

Although the two proposals are different, there is a great deal of similarity between the two. Practically speaking, some features are required to reach the goal of 100 Mbps throughput.

Multiple-Input/Multiple-Output (MIMO)

Up until 2004, 802.11 interfaces had a single antenna. To be sure, some interfaces had two antennas in a diversity configuration, but the basis of diversity is that the "best" antenna is selected. In diversity configurations, only a single antenna is used at any point. Although there may be two or more antennas, there is only one set of components to process the signal, or *RF chain*. The receiver has a single input chain, and the transmitter has a single output chain.

The next step beyond diversity is to attach an RF chain to each antenna in the system. This is the basis of Multiple-Input/Multiple-Output (MIMO) operation.* Each RF chain is capable of simultaneous reception or transmission, which can dramatically improve throughput. Furthermore, simultaneous receiver processing has benefits in resolving multipath interference, and may improve the quality of the received signal far beyond simple diversity. Each RF chain and its corresponding antenna are responsible for transmitting a *spatial stream*. A single frame can be broken up and multiplexed across multiple spatial streams, which are reassembled at the receiver. Both the WWiSE and TGnSync proposals employ MIMO technology to boost the data rate, though their applications differ.

* MIMO is pronounced "MyMoe." I attended a symposium in which a standards committee attendee described the standardization vote on the acronym's pronunciation.

MIMO antenna configurations are often described with the shorthand "YxZ," where Y and Z are integers, used to refer to the number of transmitter antennas and the number of receiver antennas. For example, both WWiSE and TGnSync require 2x2 operation, which has two transmit chains, two receive chains, and two spatial streams multiplexed across the radio link. Both proposals also have additional required and optional modes. I expect that the common hardware configurations will have two RF chains on the client side to save cost and battery power, while at least three RF chains will be used on most access points. This configuration would use 2x3 MIMO for its uplink, and 3x2 MIMO on the downlink.

Channel Width

802.11a currently uses 20 MHz channels because that is the channel bandwidth allowed by all regulators worldwide. Doubling the channel bandwidth to 40 MHz doubles the theoretical information capacity of the channel. Although promising for the future, some regulators do not currently allow 40 MHz operation. Japan is the most notable exception.

MAC Efficiency Enhancements

As this book has repeatedly pointed out, the efficiency of the 802.11 MAC is often poor. In most usage scenarios, it is very difficult to exceed 50–60% of the nominal bit rate of the underlying physical layer. Every frame to be transmitted requires a physical-layer frame header, as well as the pure overhead of preamble transmission. The 802.11 MAC adds further overhead by requiring that each frame be acknowledged. Overhead is particular bad for small frames, when the overhead takes more time than the frame data itself. Figure 15-1 shows the efficiency, defined as the percentage of the nominal bit rate devoted to MAC payload data, for a variety of frame sizes. The values in the figure are exclusively for MAC payload data. Any network measurement would require additional LLC data, and networks that are encrypted would have additional overhead bytes. Furthermore, most network protocols provide their own acknowledgment facilities, which further reduces real-world efficiency. The point of Figure 15-1 is that small frames have particularly poor efficiency.

Both TGnSync and WWiSE adopt techniques to improve the efficiency of the radio channel. Concepts are similar, but the details differ. Both offer some form of *block ACKs* (sometimes called *frame bursting*). By removing the need for one acknowledgment frame for every data frame, the amount of overhead required for the ACK frames, as well as preamble and framing, is reduced. Block acknowledgments are helpful, but only if all the frames in a burst can be delivered without a problem. Missing one frame in the block or losing the acknowledgment itself carries a steep penalty in protocol operations because the entire block must be retransmitted.

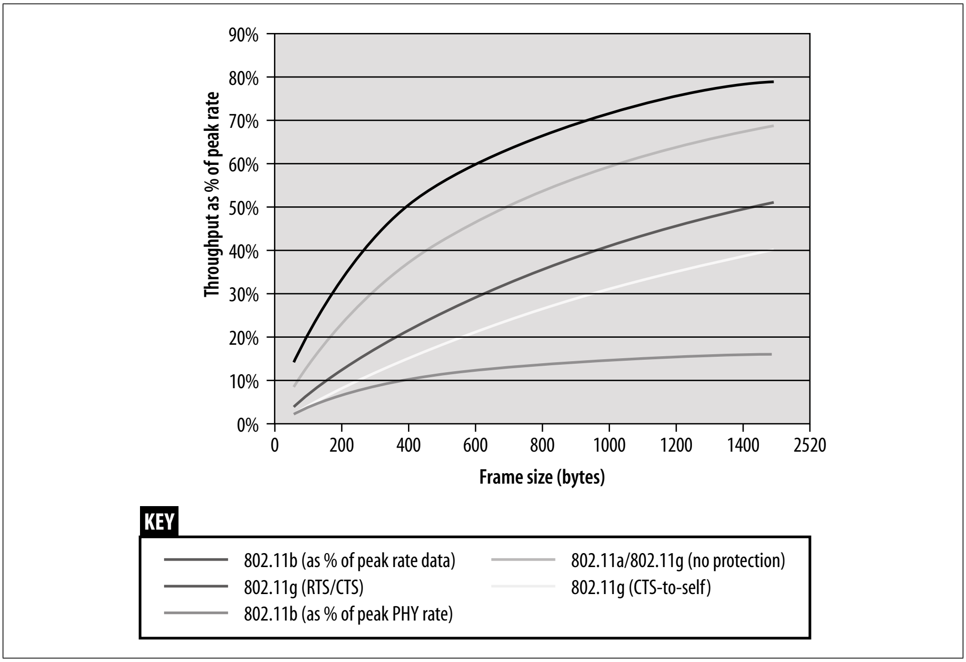

Figure 15-1. MAC Efficiency

Frame aggregation is also part of both proposals. Many of the packets carried by 802.11 are small. Interactive network sessions, such as telnet and SSH, make heavy use of rapid-fire small packets. Small packets become small frames, each of which requires physical-layer framing and overhead. Combining several small packets into a single relatively large frame improves the data-to-overhead ratio. Frame aggregation is often used with *MAC header compression*, since the MAC header on multiple frames to the same destination is quite similar.

WWiSE

The WWiSE consortium includes several well-known chipmakers: Airgo (the manufacturer of the first "pre-N" devices on the market), Broadcom, Conexant, and Texas Instruments. Motorola joined the consortium in February 2005, just as this book headed to press.

MAC Enhancements

As would be expected from a name touting spectral efficiency, WWiSE is more the more heavily weighted towards improving the MAC efficiency of the two proposals. To get to 100 Mbps net payload throughput, 12,000 bytes (960,000 bits) need to be transmitted in 960 microseconds. WWiSE's PHY specification has a 135 Mbps data rate in a basic two-antenna configuration with two data streams, which can move the data in 711 microseconds. The remaining 249 microseconds are used for preambles, framing, interframe spacing, and the single block acknowledgment.

Channels and radio modes

WWiSE uses both 20 MHz and 40 MHz channels. 40 MHz operation may be through a single 40 MHz channel, or through a 20 MHz *channel pair* in which both channels are used simultaneously for data transmission. One channel is designated as the primary channel, and operates normally. The secondary channel is used only for channel aggregation, and does not have stations associated on it. The secondary channel is used for "overflow" from the primary; carrier sensing functions are performed only on the primary channel.

Although the use of two channels is really a physical layer operation, there are some housekeeping functions performed by the MAC. A new information element, the Channel Set element, is sent in the primary channel Beacon frames so that stations are informed of the secondary channel in the pair. Access points also send Beacon frames on the secondary channel; unlike most Beacon operations, though, the purpose is to discourage clients from associating, or other devices from choosing that channel for operation. A secondary channel Beacon frame is very similar to the primary channel Beacon, but the only supported rate is a mandatory MIMO PHY rate. To further discourage use of the channel, it may also include the contention-free information element.

Protection

Like 802.11g, the new PHYs require enhanced protection mechanisms to avoid interfering with existing stations. Naturally, the protection mechanisms specified in 802.11g are adopted for operation of 2.4 GHz stations that may have to avoid interfering with older direct sequence or 802.11b equipment. When access points detect the presence of older equipment, it will trigger the use of RTS-CTS or CTS-to-self protection as described in Chapter 14.

However, additional protection may be required to avoid having a MIMO station transmit at a rate not understood by 802.11a or 802.11g equipment. The WWiSE proposal contains an OFDM protection scheme to allow MIMO stations to appropriately set the NAV on older OFDM stations. The protection mechanism is identical to the one described in Chapter 14, but it takes place using OFDM data rates.

Finally, the WWiSE proposal uses two bits in the ERP information element in Beacon frames to indicate whether OFDM protection is needed. In some cases, OFDM protection may be needed to assist an older 802.11g network, but no protection is needed for 802.11b stations. Access points monitor the radio link to determine if OFDM protection is needed. To assist stations using channel pairs, they also report on whether a secondary channel is in use.

Aggregation, bursting, and acknowledgment

The WWiSE proposal increases the maximum payload size from 2,304 bytes to over 8,000 bytes. Increasing the payload increases the payload-to-overhead and the ratio can increase efficiency if the larger frames or bursts can be delivered successfully.

Aggregation bundles multiple higher-level network protocol packets into a single frame. Each packet gets a subframe header with source and destination addresses, and a length to delimit the packet, as shown in Figure 15-2. Aggregation can only be used when the frames bundled together have the same value for the Address 1 field, which is the receiver of the frame. Frames from an access point in an infrastructure network use Address 1 as the destination, so access points can only aggregate frames bound for a single station. A station in an infrastructure network can, however, aggregate frames to multiple destinations. Station transmissions use the Address 1 field for the AP, since all frames must be processed by the AP prior to reaching the backbone network. Upon aggregation, the destination address is the "next hop" processing station, and the source is the creator of the frame. Upon deaggregation, the individual subframes will be processed according to the sub-frame headers. Due to the requirement that the receiver address must be the same, it is not possible to aggregate a mixture of unicast, broadcast, and multicast data. The proposal contains no rules about when to use aggregation.

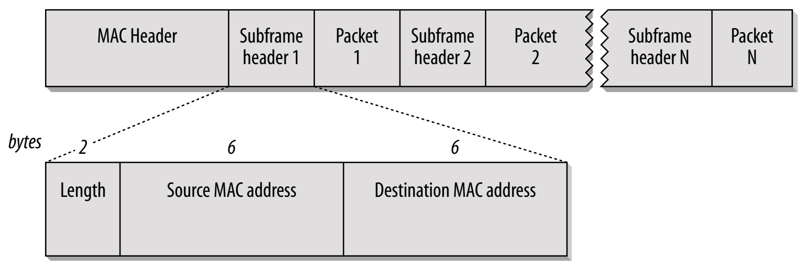

Figure 15-2. Aggregation in WWiSE

Bursting is a related, but slightly different concept. Frame aggregation glues higher-layer protocol packets together for transmission in larger lumps. Bursting does the same at the physical layer. Once a station has invested a significant amount of protocol overhead to obtain control of the channel, it can just keep on transmitting. One of the advantages of using multiple physical frames, as opposed to higher-layer frames, is that each physical frame has its own source and destination. A frame burst can consist of traffic intended for a variety of different destination addresses. In a frame burst, there are two additional interframe spaces defined, the Zero Interframe Space (ZIFS) and the Reduced Interframe Space (RIFS). Successive frames that use the same transmit power may use the ZIFS for immediate transmission. If the transmit power is changed between frames, the RIFS may be used. The RIFS is shorter than other interframe spaces, though, so it allows a station to retain control of the channel. In Figure 15-3, the first frame cannot be aggregated, and is transmitted after the transmitter gains control of the channel. Once it has gained control, it can hold on as long as allowed. The second and third frames use the same transmission power, and so are transmitted after the zero interframe space. Additionally, they share the Address 1 field and are therefore bundled into an aggregate frame. For transmission of the next frame, power needs to be changed, requiring the use of the reduced interframe space. The fourth and fifth frames can be aggregated, and are transmitted as a single aggregate frame. When the queued data has been transmitted, the station relinquishes control of the channel.

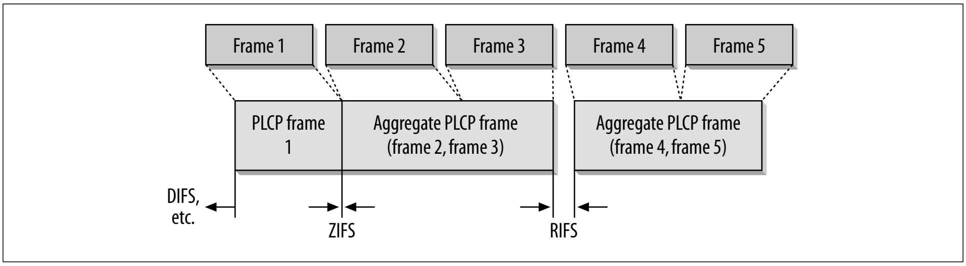

Figure 15-3. Bursting in WWiSE

In the initial version of the 802.11 MAC, a positive acknowledgment was required for every unicast data frame. WWiSE lifts this restriction, and allows for a more flexible acknowledgment policy. In addition to the "normal" policy, frames can be transmitted without an acknowledgment requirement, or with block acknowledgments instead.

The WWiSE MIMO PHY

The WWiSE proposal is a slight evolution of 802.11a, using MIMO technology. The basic channel access mechanisms are retained, as is the OFDM encoding. At a high level, the WWiSE PHY is mainly devoted to assigning bits to different antennas.

Structure of an operating channel

Like 802.11a, the radio channel is divided into 0.3125 MHz subcarriers. As in the 802.11a channel subdivisions, a 20 MHz channel in the WWiSE proposal is divided into 56 subcarriers. 40 MHz channels, which are optional, are divided into 112 subcarriers. In addition to being optional, 40 MHz channels are only supported in the 5 GHz band because it is not possible to squeeze multiple 40 MHz channels into the ISM band. (And if you thought network layout was hard with three channels, wait until you try with two!) Figure 15-4 shows the structure of both the 20 MHz and 40 MHz operating channels in the WWiSE proposal.

As in 802.11a, subcarriers are set aside as pilots to monitor the performance of the radio link. Fewer pilot carriers are needed in a MIMO system because the pilot carriers run through as many receiver chains. A 20 MHz 802.11a channel uses four pilot

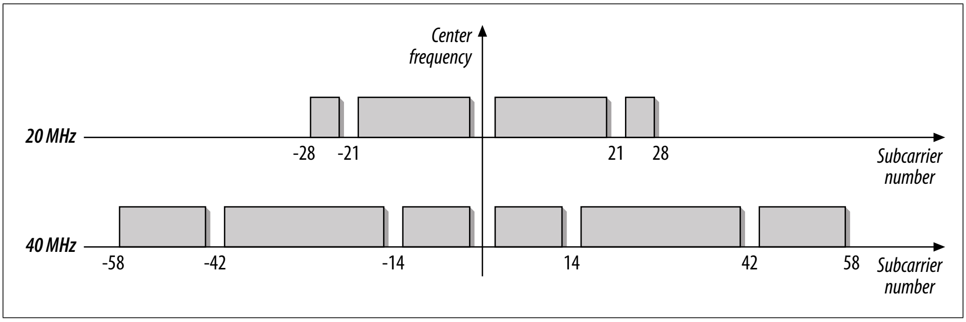

Figure 15-4. WWiSE pilot carrier structure

subcarriers. In the WWiSE proposal, a 20 MHz channel requires only two pilot carriers because each pilot is processed by two receiver chains, which has the same effect as four pilots processed by a single receiver chain. With fewer pilots, more subcarriers can be devoted to carrying data. 20 MHz WWiSE channels have 54 data subcarriers; 40 MHz channels have exactly twice as many at 108.

Modulation and encoding

The WWiSE proposal does not require new modulation rates. It uses 16-QAM (4 bit) and 64-QAM (6 bit) modulation extensively, but does not require finer-grained modulation constellations.

Coding is enhanced, however. A new convolutional code rate of 5/6 is added. Like the 2/3 and 3/4 code rates defined by 802.11a, the 5/6 code is defined by puncturing the output to obtain a higher code rate. WWiSE also defines the use of a low density parity check (LDPC) code.

Interleaver

In 802.11a, the interleaver is responsible for assigning bits to subcarriers. MIMO interleavers are more complex because they must assign bits to a spatial stream in addition to assigning bits to positions on the channel itself. The WWiSE interleaver takes bits from the forward error coder and cycles through each spatial stream. The first bit is assigned to the first spatial stream, the second bit is assigned to the second spatial stream, and so on. The interleaver is also responsible for scrambling the encoded bits within each spatial stream.

Space-time block coding

In most cases, an antenna will be used for each spatial stream. However, there may be cases when the number of antennas is greater than the number of spatial streams. If, for example, most APs wind up using three antennas while clients only use two, there is an "extra" transmit antenna, and the two spatial data streams need to be assigned to the three antennas. Transmitting a single spatial stream across multiple antennas is called *space-time block coding* (STBC).

The basic rule for splitting a spatial stream across multiple antennas is to transmit two related streams on different antennas. As discussed in Chapter 13 on 802.11a, the radio wave is composed of in-phase and quadrature components, where the quadrature wave is a quarter-cycle out of phase with the in-phase component. Phase shifts are represented mathematically by the imaginary part of the complex number in the constellation. The complex conjugate of a complex number has the same real part, but flips the sign on the imaginary part. Physically, the radio wave from the complex conjugate will have the same in-phase component, but the quadrature component will have the oppose phase shift. When there are extra antennas, the WWiSE proposal mandates that a spatial stream and its complex conjugate are transmitted on an antenna pair. Table 15-1 reviews the rules. The rules for splitting spatial streams are independent of the channel bandwidth, although 40 MHz spatial streams will carry more bits.

Table 15-1. WWiSE encoding rules when antennas outnumber spatial streams

Transmit antennas	Spatial streams	First spatial stream	Second spatial stream	Third spatial stream
2	1	Coded across antennas 1 and 2	N/A	N/A
3	2	Coded across antennas 1 and 2	Transmitted normally on third antenna	N/A
4	2	Coded across antennas 1 and 2	Coded across antennas 3 and 4	N/A
4	3	Coded across antennas 1 and 2	Third antenna	Fourth antenna

Modulation rates

There are 24 data rates defined by the WWiSE PHY, with 49 different modulation options. Rather than take up a great deal of space in a table, here is a basic formula for the data rates:

Data rate (Mbps) = 0.0675 × channel bandwidth × number of spatial streams × coded bits per subcarrier × code rate

Channel bandwidth
: Either 20 for 20 MHz channels, or 40 for 40 MHz channels or channel pairs.

Number of spatial streams
: The number of spatial streams can be equal to 1, 2, 3, or 4. It must be less than or equal to the number of transmission antennas. Support for at least two spatial streams is mandatory.

Coded bits per subcarrier
: In most cases, this will either be 6 for 64-QAM or 4 for 16-QAM. BPSK (1 coded bit per subcarrier) and QPSK (2 coded bits per subcarrier) are only supported in the 20 MHz channel mode with one spatial stream.

Code rate
: The code rate may be 1/2 or 3/4 when used with 16-QAM, and 2/3, 3/4, or 5/6 when used with 64-QAM.

There may be multiple ways to get to the same data rate. As an example, there are four ways to get 108 Mbps:

- Four spatial streams in 20 MHz channels, using 16-QAM with R=1/2.
- Two spatial streams in 20 MHz channels, using 64-QAM with R=2/3.
- One spatial stream in a 40 MHz channel, using 64-QAM with R=2/3.
- Two spatial streams in 40 MHz channels, using 16-QAM with R=1/2.

In a basic mode with a single spatial stream, channel capacity is slightly higher than with 802.11a because fewer pilot carriers are used. Single-channel modulation tops out at 60.75 Mbps, rather than the 54 Mbps in 802.11a. By using all the highest throughput parameters (four 40 MHz spatial streams, with 64-QAM and a 5/6 code), the WWiSE proposal has a maximum throughput of 540 Mbps.

MIMO and transmission modes

Previous 802.11 PHY specifications had fairly simple transmission modes. The WWiSE proposal has 14 transmission modes, depending on 3 items:

- The number of transmit antennas, noted by xTX, where x is the number of transmit antennas. It ranges from 1 to 4, although a single antenna is only supported for 40 MHz channels. All 20 MHz channels must use at least two transmit antennas, though they may have only one spatial stream.
- Whether the frame is used in a greenfield (GF) or mixed mode (MM) environment. Mixed mode transmissions use physical headers that are backwards-compatible with other OFDM PHYs, while greenfield transmissions use a faster physical header.
- The channel bandwidth, which may be 20 MHz or 40 MHz.

Table 15-2 shows the resulting 14 transmission modes. There are several physical layer encodings defined for each of these modes, and they will be discussed in the PLCP section. The number of active antennas is only loosely related to the number of spatial streams. A system operating in the 4TX40MM mode has four transmit antennas, but it may have two or three spatial streams.

Table 15-2. WWiSE transmission modes

	20 MHx channels	40 MHz channels
Greenfield	2TX20GF 3TX20GF 4TX20GF	1TX40GF 2TX40GF 3TX40GF 4TX40GF
Mixed mode	2TX20MM 3TX20MM 4TX20MM	1TX40MM 2TX40MM 3TX40MM 4TX40MM

WWiSE PLCP

The PLCP must operate in two modes. In Greenfield mode, it operates without using backwards-compatible physical headers. Greenfield access is simpler: it can operate without backwards compatibility. As a starting point, consider Figure 15-5; it shows the PLCP encapsulation in the 1TX40GF, 2TX20GF, and 2TX40GF modes.

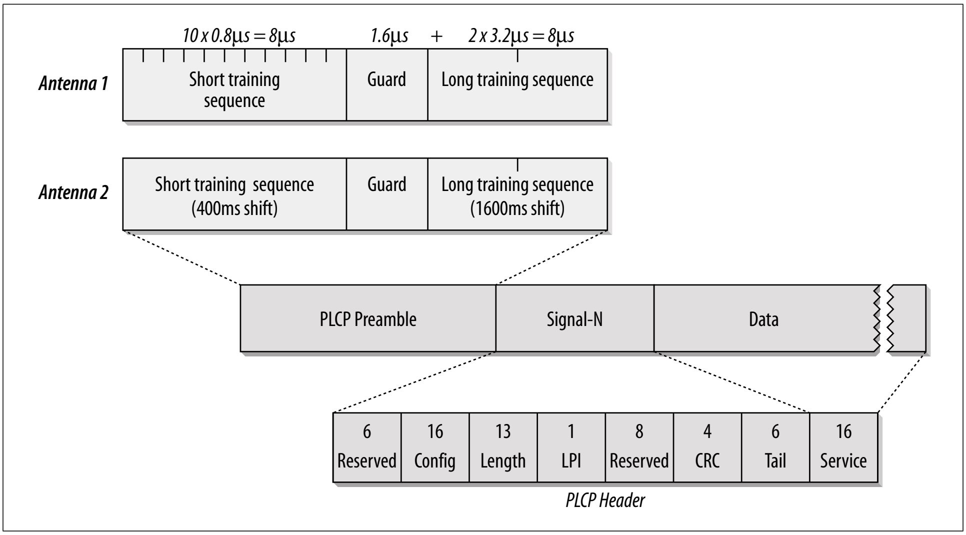

Figure 15-5. Greenfield 1TX40 and 2TX20/2TX40 modes

The fields in the frame are similar in name and purpose to all of the other PLCP frames discussed in this book.

MIMO-OFDM PLCP Preamble
: The preamble consists of well-known bit sequences to help receivers lock on to the signal. Depending on the transmission mode, the preamble may be split into multiple parts. It generally consists of both short and long training sequences. In the WWiSE proposal, the same preamble is transmitted on all the antennas, but with small time shifts relative to the others. Figure 15-5 shows the training sequences used by two-antenna transmission modes. Although the training sequences consist of different bits, the shift is the same. Naturally, the single antenna 40 MHz mode would only have one active antenna transmitting a preamble.

SIGNAL-N
: The SIGNAL-N field contains information that helps to decode the data stream. It is always sent using QPSK, R=1/2, and is not scrambled. It contains information on the number of spatial streams, channel bandwidth, modulation, and coding, and a CRC. More detail on the SIGNAL-N field follows this section.

SERVICE

The SERVICE field is identical to its usage in 802.11a. Unlike the other components of the PLCP header, it is transmitted in the Data field of the physical protocol unit at the data rate of the embedded MAC frame. The first eight bits are set to 0. As with the other physical layers, MAC frames are scrambled before transmission; the first six bits are set to 0 to initialize the scrambler. The remaining nine bits are reserved and must set to 0 until they are adopted for future use.

Data

The final field is a sequence of four microsecond symbols that carry the data. Data bits have six zero tail bits to ramp down the error correcting code, and as many pad bits as are required to have an even symbol block size.

The SIGNAL-N field

The SIGNAL-N field is used in all transmission modes. It has information to recover the bit stream from the data symbols. The SIGNAL-N field is shown in Figure 15-6.

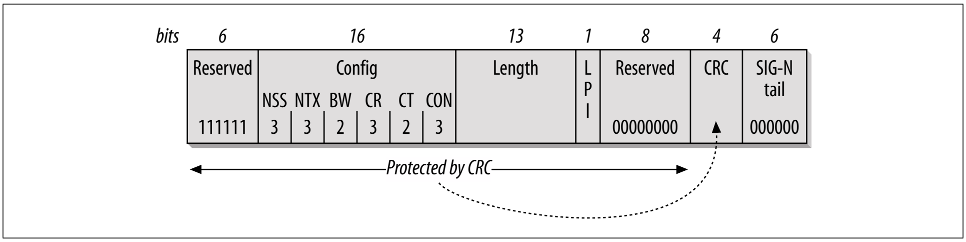

Figure 15-6. WWiSE SIGNAL-N field

CONFIG

Six fields are grouped into the Configuration subfield.

NSS (number of spatial streams)

Three bits are used to indicate how many spatial streams are used. The value is zero-based, so it ranges from zero to three.

NTX (number of transmission antennas)

Three bits are used to indicate how many antennas are used to carry the number of spatial streams. The value is zero-based, so it ranges from zero to three.

BW (bandwidth)

Two bits carry the channel bandwidth. 20 MHz is represented by zero, and 40 MHz is represented by one.

CR (code rate)

Three bits indicate the code rate. 1/2 is zero, 2/3 is one, 3/4 is two, and 5/6 is three.

CT (code type)

Two bits indicate the type of code. Zero is a convolutional code, and one is the optional LDPC.

CON (constellation type)

Three bits indicate the type of constellation: zero for BPSK, one for QPSK, two for 16-QAM, and three for 64-QAM.

LENGTH

A 13-bit identifier for the number of bytes in the payload of the physical frame. It ranges from zero to 8,191.

LPI (Last PSDU indicator)

When multiple physical frames are sent in a burst, the LPI bit is set on the last one to notify other stations that the burst is coming to an end.

CRC

The CRC is calculated over all the fields except for the CRC and the tail bits.

Tail

Six bits are used as tail bits to ramp down the convolutional coder.

In the other transfer modes, shown in Figure 15-7, the preamble is split into chunks. In between the chunks, there may be Signal fields. SIGNAL-N fields are defined by the 802.11n proposal and are only decoded by 802.11n stations; the SIGNAL-MM field is used to retain backwards compatibility in a mixed mode with older OFDM stations. It is identical to the Signal field used by 802.11a, and is shown in Figure 13-16.

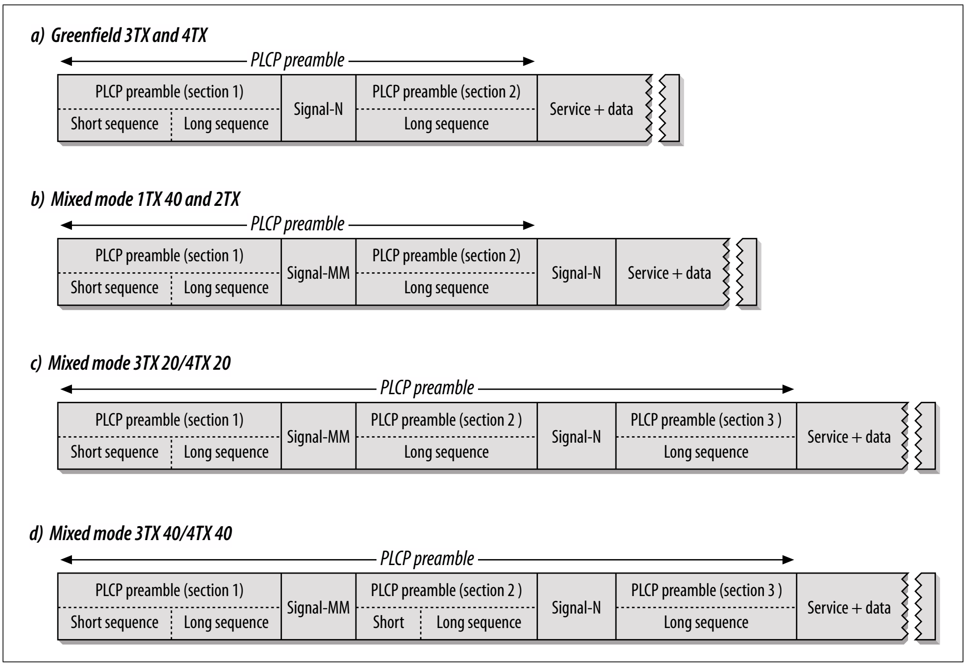

Figure 15-7. PLCP frame format for other transfer modes

WWiSE PMD

Figure 15-8 shows the basic layout of the WWiSE transmitter. It is essentially the same as the 802.11a transceiver, but it has multiple transmit chains. The interleaver is responsible for dividing coded bits among the different transmit chains and spatial streams.

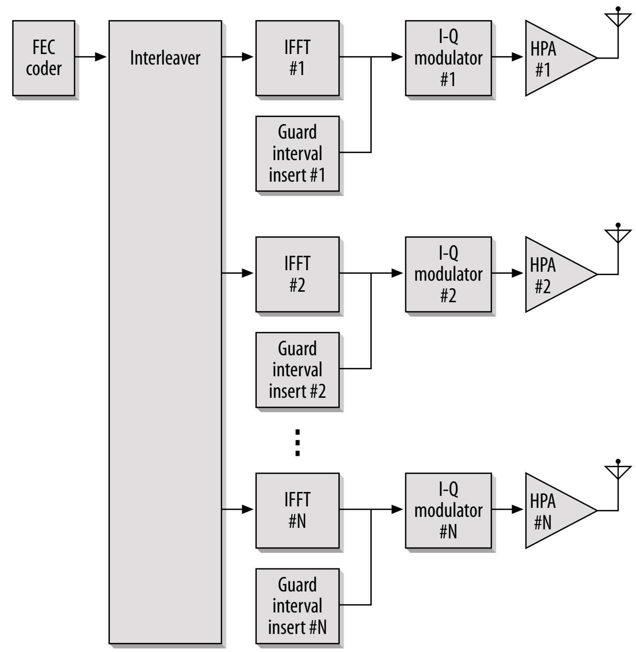

Figure 15-8. WWiSE transceiver

Sensitivity is specified by the proposal, and it is identical to what is required of 802.11a receivers. Table 15-3 shows the required sensitivity. The proposal does not have any adjacent channel rejection requirements.

Table 15-3. WWiSE receiver sensitivity

Constellation	Rate	Sensitivity (dBm)	802.11a Sensitivity (dBm), for reference
BPSK	1/2	–82	–82
BPSK	3/4	–81	–81
QPSK	1/2	–79	–79
QPSK	3/4	–77	–77
16-QAM	1/2	–74	–74
16-QAM	3/4	–70	–70
64-QAM	2/3	–66	–66

Table 15-3. WWiSE receiver sensitivity (continued)

Constellation	Rate	Sensitivity (dBm)	802.11a Sensitivity (dBm), for reference
64-QAM	3/4	–65	–65
64-QAM	5/6	–64	N/A

Characteristics of the WWiSE PHY

Parameters specific to the WWiSE PHY are listed in Table 15-4. Like the other physical layers, it also incorporates a number of parameters to adjust for the delay in various processing stages in the electronics.

Table 15-4. WWiSE MIMO PHY parameters

Parameter	Value	Notes
Maximum MAC frame length	8,191 bytes	
Slot time	9 μs	
SIFS time	16 μs	The SIFS is used to derive the value of the other interframe spaces (DIFS, PIFS, and EIFS).
RIFS time	2 μs	
Contention window size	15 to 1,023 slots	
Preamble duration	16 μs	
PLCP header duration	4 μs	
Receiver sensitivity	-64 to -82 dBm	Depends on speed of data transmission

TGnSync

The TGnSync consortium is composed of a wider array of companies. In addition to the chipmakers that one would expect to find (Atheros, Agere, and Intel, and Qualcomm), TGnSync notably includes manufacturers of other electronic devices. Network equipment manufacturers and even consumer electronics companies are represented. One of the goals of TGnSync is to support new networked devices in the home; promotional materials refer to sending HDTV or DVD video streams across wireless networks. The goal of streaming video probably accounts for some of the emphasis placed on high peak data rates.

TGnSync MAC Enhancements

Although the TGnSync proposal has a higher peak data rate, the group did not completely neglect the development of MAC enhancements to improve efficiency and operation. Efficiency is improved through the development of frame aggregation and bursting, as well as changes to acknowledgment policies. Some protection of older transmissions is performed at the MAC layer. Notably, several MAC enhancements are designed to save battery power, which is likely a reflection of the group's membership.

Channels, radio modes, and coexistence

Although some regulators do not allow them, the TGnSync proposal makes 40 MHz channel support mandatory. If it were adopted without change, a TGnSync chipset would support both 20 MHz and 40 MHz channels, even in regulatory domains that did not allow the latter channel bandwidth. The TGnSync proposal also has MAC features that enable the use of networks with both 20 MHz- and 40 MHz-capable stations. When stations have large amounts of data to transmit, it is possible to negotiate a temporary use of a wider channel before falling back to 20 MHz operation.

MAC operational modes can also be classified based on the types of stations in the network. *Pure mode* networks consist only of 802.11n stations. No protection is necessary to account for older 802.11a and 802.11g stations. Alternatively, TGnSync 802.11n devices may operate in the *legacy mode* just like an 802.11a or 802.11g station. Most operation, though, will be in *mixed mode*, where a TGnSync network must co-exist with a legacy network on the same channel, and may accept associations from older 802.11a or 802.11g stations.

Association requests are handled differently in each mode. Pure mode networks stay pure by ignoring association requests from older stations, and sending Beacon frames with an information element that directs associated stations to use only the new 802.11n transmission modes. Pure mode networks also transmit Beacon frames using the TGnSync high throughput PLCP, which makes them unreadable by legacy devices. Mixed mode access points are visible to legacy devices because they transmit Beacons using the legacy format.

Mixed mode is required to coexist with older devices. (If the experience of 802.11g deployment is any guide, most 802.11n devices are likely to operate in the mixed mode for quite some time.) Mixed mode is a broad classification with several subdivisions. *Mixed capable* networks will allow association from legacy devices, but do not divide time between legacy and high-throughput transmissions. Access points in *managed mixed* networks do actively divide the time between high-throughput transmissions and legacy transmissions. Much like the division between the contention-free period and the contention period (see Chapter 9), an AP operating in managed mixed mode will allow legacy stations their timeslice, while using mechanisms similar to the protection mechanism to reserve some timeslice for MIMO stations only.

Aggregation and bursting

Initial 802.11 stations typically send frames in the order they are received. For throughput purposes, it is highly desirable to reorder frames so that they can coalesce into larger aggregated frames. Aggregation in TGnSync is a MAC-layer function that bundles several MAC frames into a single PLCP frame for transmission.

Figure 15-9 shows the basic format of a single physical-layer frame containing several MAC layer frames. Several MAC frames are put into the same PLCP frame, with an appropriate delimiter between them. The delimiter has a small reserved field, a length field for the following MAC frame, a CRC to protect the delimiter, and a unique pattern to assist in recovering individual frames from the aggregate. MAC frames are put into the aggregate without modification, and contain the full header and MAC CRC. Even if one frame out of an aggregate is lost, it may be possible to successfully receive all the remaining frames.

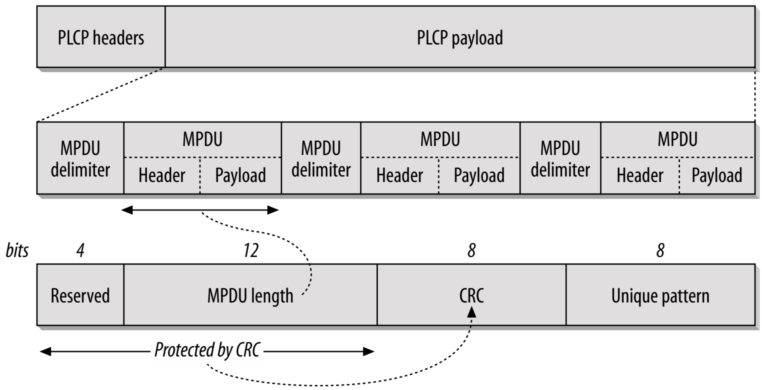

Figure 15-9. TGnSync frame aggregation

Exchanging aggregated frames is only possible once the channel has been configured for it. Figure 15-10 illustrates the process. The sender of an aggregate, called the *initiator*, must send an Initiator Aggregation Control (IAC) frame. IAC frames work much like RTS frames, but have additional fields to assist with channel control. Initiators can request channel measurements, offer different types of coding on the aggregate frame, and accept aggregates in the return direction. Upon receiving the IAC, the destination system, called the *responder*, generates a Responder Aggregation Control (RAC) frame. RAC frames work much like CTS frames: they close the loop by notifying the sender that an aggregate will be accepted, and finishing the parameter negotiation. When aggregate frames are received, an acknowledgment is required. TGnSync defines a new acknowledgment type, the BlockACK, which can be used to acknowledge all the MAC frames contained in an aggregate.

To further improve MAC efficiency, TGnSync defines a MAC header compression algorithm for use in conjunction with aggregate frames. It works in the same manner as Van Jacobsen header compression on serial dial-up lines. Frames between two destinations share most of the fields in the MAC header, most notably the MAC addresses inside the packet. Therefore, a one-byte Header ID (HID) is assigned to a unique set of the three MAC addresses inside a MAC frame. The Header ID can also save the Duration field, since the aggregate will have its own Duration, as well as the two bytes for QoS control. When frames are transmitted between the same sender

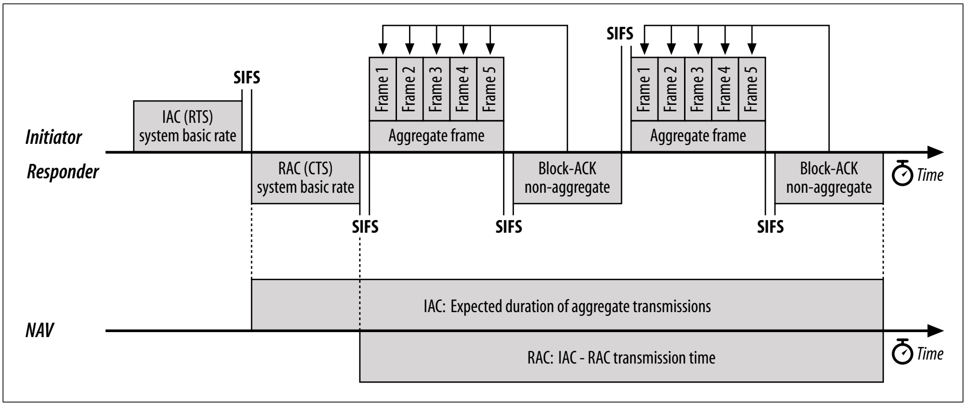

Figure 15-10. TGnSync block acknowledgment

and destination, rather than repeating the same 22 bytes of header information, it is replaced by the corresponding single byte Header ID. Figure 15-11 (a) shows the use of the header compression MAC frames. First, a header frame containing the full header is transmitted, and a header ID is assigned. The header ID can be used to reference the prior full header by sending a single byte to reference previously-transmitted information about the Duration, addressing, and QoS data.

Figure 15-11 (b) illustrates the use of header compression. Five frames for two destinations from the MAC have been aggregated into a single frame for physical transmission. Rather than transmit a complete header on each constituent MAC frame in the aggregate, the system uses header compression. There are two destinations and therefore two unique MAC headers. They are each transmitted and assigned a header ID number. The five data frames following the aggregate each refer to the appropriate header number. A header ID number is unique only within the context of a single aggregate frame. Compared to transmitting full headers on all five frames, the overhead due to MAC framing is cut by more than half.

Header compression is useful when a single aggregate contains multiple frames between the same source and destination pair. However, the benefits of aggregation in TGnSync are not confined to pairs. Single-receiver aggregation is required; an optional extension allows aggregate frames to contain MAC frames for multiple receivers, in which case they are called Multiple Receiver Aggregate (MRA) frames. Inside the single transmitted aggregate frame, there are multiple Initiator Access Control frames. Each IAC specifies an offset to transmit the response to the aggregated frames, which will usually be a block acknowledgment response. To distinguish multiple receiver aggregate frames from single-receiver aggregate frames, multiple-receiver frames start with a control item called the Multiple Receiver Aggregate Descriptor (MRAD). Figure 15-12 shows the operation of multiple-receiver aggregation. The initiator's aggregate frame starts with the aggregate descriptor, and is followed by the aggregated frames for each destination. An IAC frame is used to divide them.

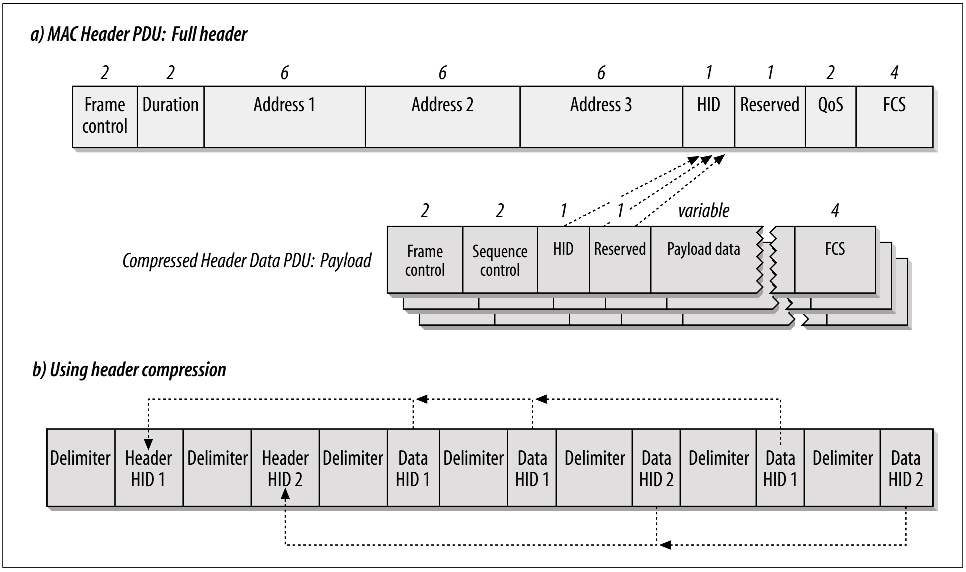

Figure 15-11. TGnSync MAC header compression

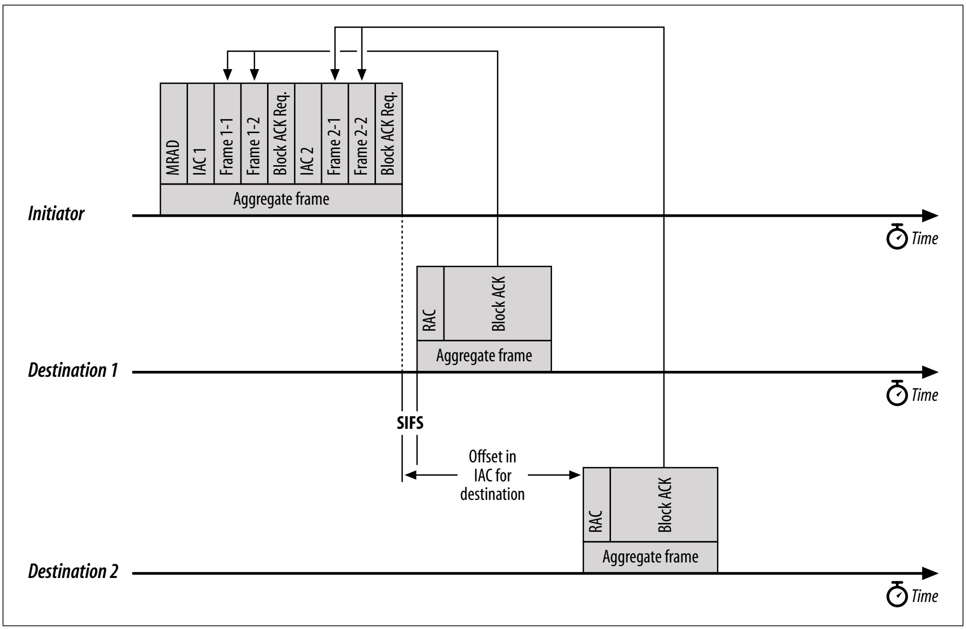

Figure 15-12. TGnSync MRA

Protection

As with all PHYs that have followed existing hardware on to the market, the TGnSync proposal implements protection to avoid having the new PHY step on transmissions from the old PHY. Protection in the TGnSync proposal can take one of two main forms. The first class is based on the MAC's virtual carrier sensing mechanism with the network allocation vector. The second class is based on "spoofing," which uses the existing PLCP header format to carry duration information. Each station may make its own determination as to the appropriate mechanism.

Setting long NAV values to protect the duration of a frame exchange is a small adaptation of the 802.11g protection mechanism. At the start of a frame exchange, the RTS frame will contain a NAV long enough to protect the entire frame exchange. The RTS frame is sent using a "legacy" rate, and can be understood by existing OFDM receivers. In response, the target station sends a CTS message back, also with a long NAV value. According to the basic access rules of the MAC, other stations defer access to the medium due to the RTS/CTS clearing, and the two stations are free to exchange frames at higher data rates using modulations that would not be understood by older stations. LongNAV intervals may be terminated early by using a CF-End frame. When this protection is used for aggregate frames, the RTS is replaced by an Initiator Access Control (IAC) frame, and the CTS is replaced by a Responder Access Control (RAC) frame; the principle of operation, however, remains identical. See Figure 15-13.

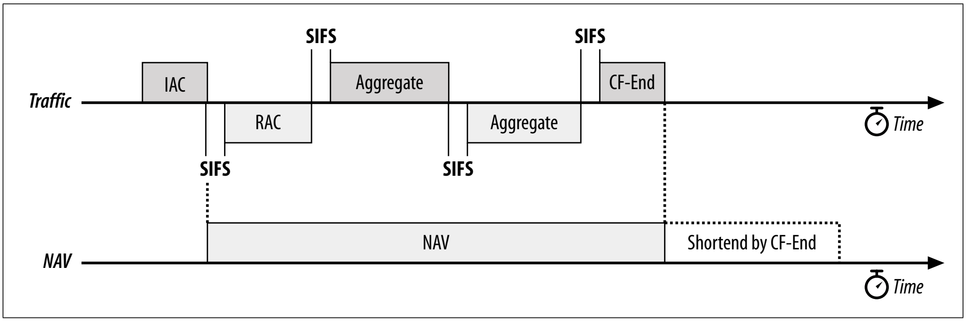

Figure 15-13. TGnSync protection: LongNAV

The second class of protection is called "spoofing," and depends on setting the length field in the PLCP header. TGnSync retains the existing OFDM header described in Chapter 13. Because it is identical to the 802.11a/g format, spoofing is effective with all stations. The OFDM PLCP header contains two numbers that are used by receivers to determine how long the transmission will take. The SIGNAL field, which is shown in Figure 13-16, encodes both the transmission rate for the body and its length in bits. Stations decode the signal field and divide the number of bits by the rate to come up with an approximate transmission time.* To maximize

* There are some slight offsets to account for interframe spacing, but the concept remains identical.

the amount of time that can be spoofed, the data rate in the legacy SIGNAL field is always set to the lowest possible value of six Mbps.

In *pairwise spoofing*, two stations will each send an incorrect length and rate so that older stations will be in receiving mode for the duration of the current frame and its next response. Newer TGnSync stations ignore the older SIGNAL field, and use an 802.11n SIGNAL field instead. Figure 15-14 illustrates pairwise spoofing. When Frame 1 is transmitted, pairwise spoofing is used to lengthen the receiving time of the frame through the end of Frame 2. TGnSync stations will interpret the spoofing as a longer NAV, and will therefore act as if the NAV were set for the duration of Frame 2. Naturally, a station within range of the responder will be set to the receiving state; if, however, there is a hidden node, the NAV will protect transmission over Frame 2. 802.11a/g stations will interpret the spoofed time as receiving time, even if they are out of range of the second frame. When Frame 2 is transmitted, it also employs pairwise spoofing to protect Frame 2 and Frame 3.

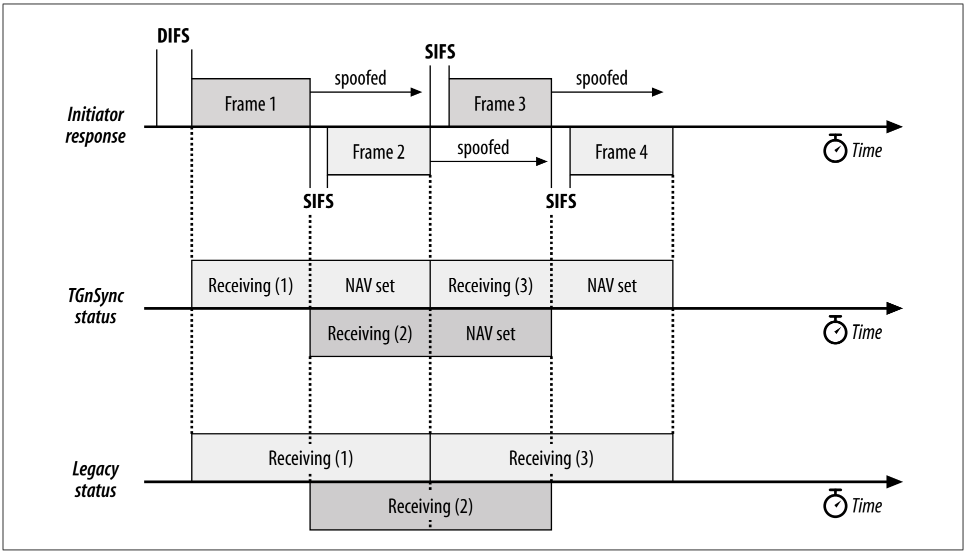

Figure 15-14. TGnSync protection: pairwise spoofing

If a long lock-out period is needed for multiple responses to a single frame, *single-ended spoofing* may also be used. With single-ended spoofing, the first frame in the exchange uses spoofing to protect the entire exchange, allowing all the responses to come in during the protected period. Figure 15-15 illustrates single-ended spoofing with frame aggregation. The first aggregate frame is a multiple-receiver aggregate, allowing responses from two other stations. It sets spoofed duration equal to the time expected for the entire exchange. TGnSync stations will go into receiving mode for the duration of the first frame, and then act as if the NAV were set for the spoofed duration. Legacy devices go into receiving mode for the entire spoofed duration.

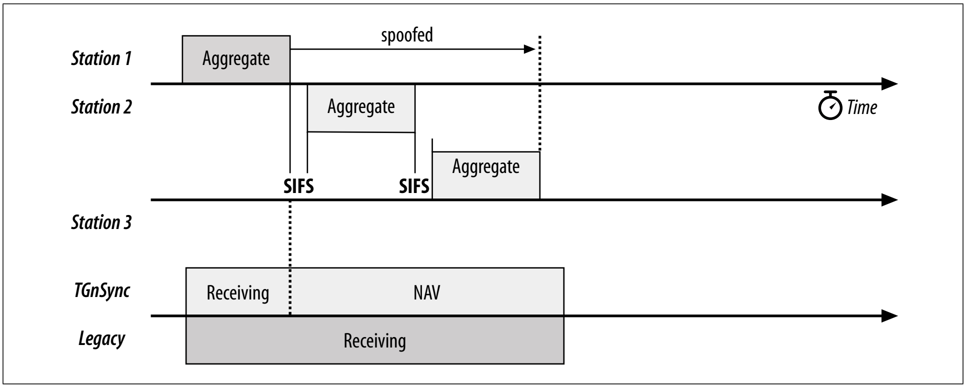

Figure 15-15. TGnSync protection: single-ended spoofing

Powersaving

TGnSync defines the Timed Receive Mode Switching (TRMS) protocol to conserve energy and extend battery life. Traditional 802.11 powersaving works by completely shutting down an interface and requiring buffering at the AP. In single-input/single-output radios, there is only one RF chain to shut down. With MIMO systems, however, there can be significant power savings by shutting down unused RF chains, but retaining a single active chain to monitor the radio link. The two states of the system are called *MIMO enabled* for full receive capability, and *MIMO disabled* when all but one RF chain is shut down.

Stations activate TRMS power saving by including an information element in the association request. The basic parameter in TRMS powersaving is the hold time. After a station transmits a frame, it stays awake for the duration of the hold time. Any transmitted frame resets the hold timer to its maximum value. Setting the hold time to zero indicates that the station will remain fully operational for one slot time before sleeping.

In infrastructure networks, the AP is responsible for maintaining the TRMS hold timer for every station. If the timer elapses, the AP must conclude it has entered the MIMO disabled state, and trigger it to power on sleeping receive chains. In an independent BSS, each station must maintain a hold timer for all the other stations.

The timer is a tunable parameter. If it is set to a larger value, stations will use more power to keep the receiver fully operational. Throughput is likely to be better, but at the cost of some battery life. In some cases, network capacity may not be affected much at all. Networks that use the NAV lengthening procedure for protection must transmit the initial RTS/CTS exchange in single-antenna mode manner compatible with all OFDM stations, and will cause stations to enable MIMO operation without additional frames. Advanced transmission modes may suffer, however, because many of them require multiantenna frame exchanges to become fully operational.

TGnSync PHY Enhancements

To develop a higher peak data rate, the TGnSync proposal depends on technology similar to WWiSE. Frames are divided into spatial streams that can be multiplexed across antennas in a MIMO configuration. More aggressive coding, including a larger constellation, higher convolutional code rate, and a reduced guard interval are present to improve the data rate. Wider channels are also required by TGnSync where supported. Support for 40 MHz channels must be built in to TGnSync-compliant devices, whereas it is optional in WWiSE.

Structure of a channel

Both 20 MHz and 40 MHz channels are divided into 0.3125 MHz subcarriers, just as in 802.11a. The 20 MHz channel is identical to an 802.11a channel, and is shown in Figure 15-16 (a). The 40 MHz channel proposed in TGn Figure 15-16 (b) is a modification of the 20 MHz structure. Two 20 MHz channels are bonded together, and the resulting spectral band is divided into 128 subchannels. The center frequencies of the old 20 MHz channels are located at +/–32. The legacy channels apply a spectral mask from –6 to +6 and roll off the amplitude of transmissions at the end of the bands. With a single continuous channel, however, there is no need to use a spectral mask, and the middle of the band can be used at full strength. Full-strength transmissions in the middle of the band allow for eight new subcarriers. Using a single contiguous 40 MHz block of spectrum reclaims subcarriers that would have otherwise been wasted. Thus, in TGnSync, a 40 MHz channel provides throughput equal to 2.25 times the 20 MHz channel, rather than simply doubling the throughput. To further boost throughput, one of the pilot carriers from the 20 MHz channel is removed, so a 40 MHz channel has 6 pilot carriers instead of 8.

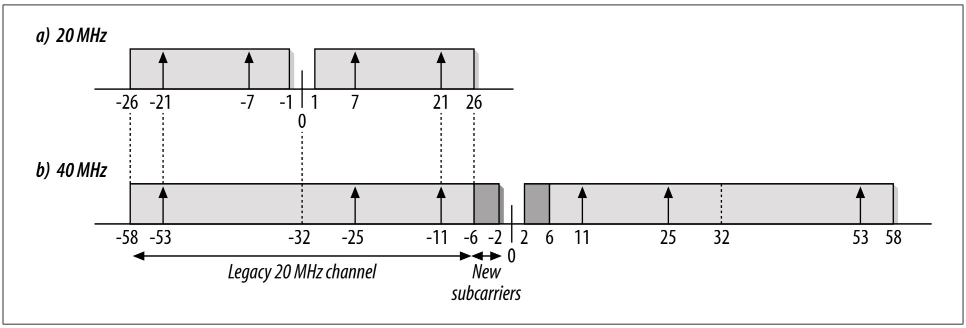

Figure 15-16. TGnSync channel structure

Basic MIMO rates

There are 32 modulation and coding pairs defined by the TGnSync PHY. In the basic MIMO mode, every spatial stream must use an identical modulation technique, so the data rate is simply a multiple of the single spatial stream data rate. Rather than take up a great deal of space in a table, here is a basic formula for the data rates:

> Data rate (Mbps) = 12 × channel bandwidth factor × number of spatial streams × coded bits per subcarrier × code rate × guard interval factor

Channel bandwidth factor
: 20 MHz channels are the baseline, and are assigned a channel bandwidth factor of 1. 40 MHz channels carry more than twice the data, and are assigned a channel bandwidth factor of 2.25.

Number of spatial streams
: The number of spatial streams can be equal to 1, 2, 3, or 4. It must be less than or equal to the number of transmission antennas. Support for at least two spatial streams is mandatory.

Coded bits per subcarrier
: This is either 6 for 64-QAM, 4 for 16-QAM, 2 for QPSK, or 1 for BPSK.

Code rate
: The code rate may be 1/2 when used with BPSK; 1/2 or 3/4 when used with QPSK or 16-QAM; or 2/3, 3/4, or 7/8 when used with 64-QAM.

Guard interval factor
: The basic guard interval is 800 ns, and is assigned the factor 1. 400 ns guard intervals increase throughput slightly, and are assigned a factor of 1.11.

In a basic mode with a single spatial stream, channel capacity is identical to 802.11a, with the exception that a 7/8 rate code may be used for a 63 Mbps data rate. By using the highest capacity parameters (four 40 MHz channels, 64-QAM with a 7/8 code rate, and the short guard interval), the TGnSync proposal has a maximum throughput of 630 Mbps.

Transmit modes

There are three MIMO modes that the TGnSync proposal calls for. In the mandatory *basic MIMO* mode, the number of spatial streams is equal to the number of antennas. Each spatial stream is modulated and transmitted identically. Each channel is coded using the same modulation, and sent with the same transmission power. Any changes in transmission rate are based on the implicit feedback of lost acknowledgments.

Two optional modes take advantage of information learned about the radio channel, which is referred to as "closed-loop" operation. TGnSync devices send "sounding" frames to each other to measure the performance of the link. Based on the information gleaned from sounding and calibration, *beamforming* can be used to boost signal quality. Higher signal quality means that a given data rate can be used at longer range. For a given signal-to-noise ratio, a beamformed transmission can carry more data. Beamforming is an optional protocol feature. It is unlikely that most client devices will be unable to send beamformed transmissions. However, client devices must be able to receive beamformed frames to receive the benefits.

In the *basic MIMO with beamforming* mode, every channel must be coded the same way. Before beginning transmission, a sounding exchange is required to calibrate the radio channel. Based on the information from the sounding exchange, the power and coding for the spatial streams is selected. Basic beamforming mode requires that all spatial streams be transmitted at the same power with the same coding. Basic beamforming can be used whenever the number of spatial streams is less than or equal to the number of transmission antennas, but its signal processing advantages are most evident when the number of antennas transmitting a signal is greater than the number of receiving antennas. If the number of spatial streams is less than the number of transmit antennas, a *spatial steering* matrix is used to assign bits to transmission antennas.

An optional *advanced beamforming MIMO* (ABF-MIMO) *mode* is also defined. It works in a manner similar to the basic beamforming mode, but with the additional capability of using different transmission power on each transmit stream, as well as the possibility of using a different modulation and code on each spatial stream. Like the basic beamforming mode, it requires the gathering of radio status information to calibrate the channel. An optional mode in the advanced beamforming mode allows beamforming to occur in both directions if it is supported in both directions. The advanced beamforming MIMO mode also includes one new constellation: 256-QAM, which transmits 8 coded bits per subcarrier.

To obtain the throughput for the advanced beamforming mode, use the equation in the previous section for each spatial stream, and add the spatial streams together. For 256-QAM, use 8 coded bits per subcarrier. 256-QAM is only used with a rate R=3/4 code rate.

Optional coding

In addition to the convolutional code supported by the original OFDM specification, the TGnSync proposal also includes two optional additional error correction codes. The first technique uses the Reed-Solomon block code, which was developed in 1960. It is widely used in many digital applications, most notably as the error-correction code on CDs and DVDs. The TGnSync proposal combines the Reed-Solomon code with the existing convolutional code in a conventional manner. First, the data stream is encoded with the Reed-Solomon code, and then the output of that encoding process is handed to a convolutional code.* Both codes have complementary properties. Convolutional codes work by spreading errors out over time, and can deal with relatively isolated

* The combination of a block encoder followed by a convolutional encoder is especially popular, and has been used extensively by deep-space probes. Galileo, Cassini, the Mars Pathfinder, and the Mars Rover all used Reed-Solomon/convolutional code combinations.

errors; the Reed-Solomon code works well at correcting error bursts. An alternative to the Reed-Solomon/convolutional combination is a low-density parity check (LDPC) code.

Optional short guard interval

To further improve MAC efficiency, the TGnSync proposal allows the use of a short guard interval. In the 802.11a and 802.11g standards as well as the WWiSE proposal, the guard interval is 800 ns. In Chapter 13, it was discussed that that the guard interval should be two to four times the delay spread. An 800 ns guard interval allows a 200 ns delay spread, which is much higher than was observed in many environments. Most offices and homes have a much smaller delay spread, on the order of 50–100 ns. In that case, using a 400 ns guard interval can boost throughput by approximately 10%.

TGnSync Physical Transmission (PLCP and PMD)

The basic frame format of the PLCP in the TGnSync proposal is shown in Figure 15-17. It uses the same header as the existing OFDM, and therefore does not require the use of high-overhead protection to avoid interfering with 802.11a or 802.11g networks. Fields prefaced with "L-" in the figure are legacy fields common to 802.11a and 802.11g; see Chapters 13 and 14 for details.

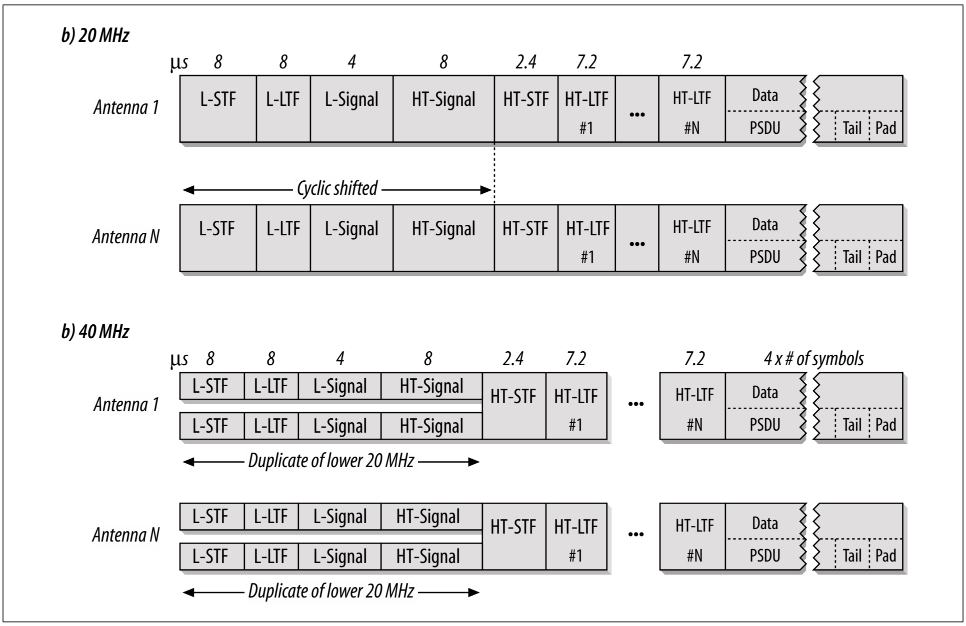

Figure 15-17. TGnSync PLCP frame format

Legacy header

The first three fields in the TGnSync PLCP frame are identical to the 802.11a/g PLCP header.

L-STF (Legacy Short Training Field)
: The legacy short training field is identical to its definition in 802.11a. It lasts eight microseconds.

L-LTF (Legacy Long Training Field)
: The legacy long training field is identical to its definition in 802.11a. It also lasts eight microseconds.

L-SIG (Legacy Signal)
: These three fields are identical to the corresponding fields used by 802.11a and adopted by 802.11g. For details, see Figure 13-16. The L-SIG field is transmitted using BPSK, R=1/2 modulation and coding.

The contents of the L-SIG field is ignored by TGnSync stations. When spoofing protection is employed, the contents of the legacy signal field are not descriptive. TGnSync stations will look in the high-throughput headers following the legacy headers for the actual length and coding rate of the enclosed MAC frame.

To take advantage of the spatial diversity that is the result of having multiple antennas transmitting the same legacy header, TGnSync offers an optional cyclic delay. Each antenna transmits its legacy frame with a slight alteration in its cyclic prefix length, such that the total delay shift is 50 ns.

When 40 MHz channels are used, as in Figure 15-17 (b), the legacy header is transmitted on each 20 MHz subchannel. That is, subcarriers –58 to –6 in the lower 20 MHz subchannel and subcarriers +6 to +58 in the upper 20 MHz subchannel both transmit the older 802.11a-style header.

High Throughput header

Immediately following the legacy preamble is a “high throughput” header specific to the TGnSync proposal. The main component of the high throughput header is the high throughput signal (HT-SIG) field, which is shown in Figure 15-18. The HT-SIG field is used to detect whether a frame carries TGnSync-encoded data at high data rates, or if it is merely a legacy data frame. The HT-SIG field is modulated conservatively, using Q-BPSK, R=1/2 modulation and coding. Q-BPSK uses two data points in its constellation, but they are present on the in the quadrature component. The Q-BPSK constellation is compared to the BPSK constellation in Figure 15-18 (b).

The three-byte high-throughput header is composed of several fields, and is transmitted in the order shown. Least significant bits in each field are transmitted first.

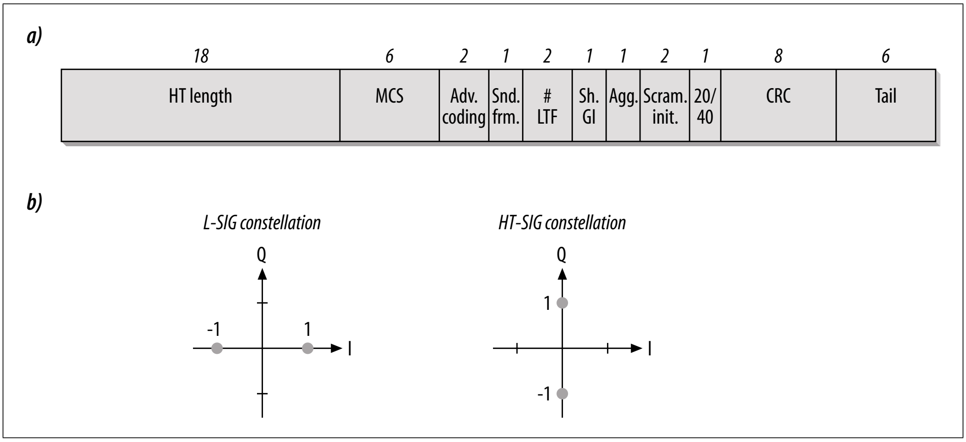

Figure 15-18. TGnSync HT-SIG field

HT-Length (18 bits)
: This field is the number of bytes in the payload of the PLCP frame. When aggregation is used, it may be quite large if several full-size MAC frames are included in the payload.

Modulation and Coding Set (MCS) (6 bits)
: One of the disadvantages to using MIMO is that there are myriad options when you account for different modulation schemes, differing numbers of spatial streams, and different code rates. The MCS field selects the modulation and coding scheme, along with the number of spatial streams. Values of 0–31 are used for basic MIMO modes, and values of 33–63 are used for advanced MIMO modes.

Advanced Coding (2 bits)
: This two-bit field is used to indicate whether the optional advanced coding is used. No advanced coding is indicated by zero. LDPC is indicated by 1. Reed-Solomon coding is indicated by 2. The value 3 is not used.

Sounding packet (1 bit)
: Requests and responses used to measure channel performance set this bit. When set, it indicates that every antenna is transmitting its own spatial stream. If it is not set, then the frame should not be used to measure channel information.

Number of HT-LTFs (2 bits)
: Following the HT-Signal field are high-throughput training fields. Each spatial stream requires a training field.

Short Guard Interval (1 bit)
: When set to one, this bit flag indicates that the short 400 ns guard interval is used on MIMO symbols in the Data field of the frame.

Aggregation (1 bit)

If this bit is set to one, it indicates that the PLCP frame carries several MAC frames in an aggregate burst.

Scrambler initialization (2 bits)

These two bits are used to seed the scrambler.

20/40 BW (1 bit)

If set to one, this bit indicates that a 40 MHz channel is used. When set to zero, it indicates a 20 MHz channel.

CRC (8 bits)

The CRC is used to protect the legacy signal field, and all the fields in the HT-SIG field before the CRC.

Tail (6 bits)

The HT-SIG field is protected by a convolutional code, and as always, six bits are needed to ramp down the convolutional coder.

High-Throughput training fields

Following the high-throughput headers are high-throughput short and long training fields. A single short training field spans the entire operating channel. In 20 MHz channels, the bandwidth of the high-throughput short training field (HT-STF) is 20 MHz. When the wider 40 MHz channels are used, the HT-STF has a bandwidth of 40 MHz. The short training field fine-tunes the receivers in MIMO operation.

When several spatial streams are transmitted over several chains, finer control of the amplification applied to the incoming signal is important. Long training fields (HT-LTFs) are used to further tune each receiver chains. One HT-LTF is used for each spatial stream. In basic MIMO mode, there is one receiver chain for each spatial stream; in the advanced mode, there may be more receiver chains than spatial streams.

Data, tail, and padding

Data bits are encoded according to modulation and coding methods that are defined by the high-throughput header. Like other OFDM physical layers, data is scrambled before transmission, using the scrambler initialization bits in the high-throughput header. Following the data, there is a six-bit tail that ramps down the convolutional code, and enough padding bits to make the data to be transmitted equal to the symbol block size.

TGnSync PMD

The basic design of a TGnSync transceiver is shown in Figure 15-19, which depicts the design of a beamforming transmitter, rather than a basic MIMO transmitter. An incoming scrambled frame is handed to the forward-error correction coder, which is usually a convolutional coder. Coded bits from the FEC coder are then sent to different spatial streams by the spatial parser, which is responsible for dividing the unified bit stream into subsidiary streams for transmission. Each spatial stream is punctured up to the desired rate. In the beamforming mode, the puncturing occurs for each spatial stream and may occur at different rates. Basic transmitters must puncture every stream to the same rate. (Logically, the puncturing in the basic MIMO transmitter can occur before the spatial parser.) Each spatial stream now consists of a sequence of coded bits, ready for mapping on to OFDM carriers by the interleaver. After the interleaver, each block of bits can be mapped on to a single symbol by the constellation mapper. In the basic MIMO mode, each interleaved spatial stream is processed by a single transmission chain; the advanced mode uses a spatial steering matrix to assign symbols to any transmit chain. The spatial steering matrix shown in the figure could be replaced by a one-to-one interface between the spatial stream processors and the transmit antennas for basic MIMO operation. Each transmit chain takes its symbol sequence and modulates it on to the airwaves. There is no mention in the specification of required channel rejection or sensitivity performance.

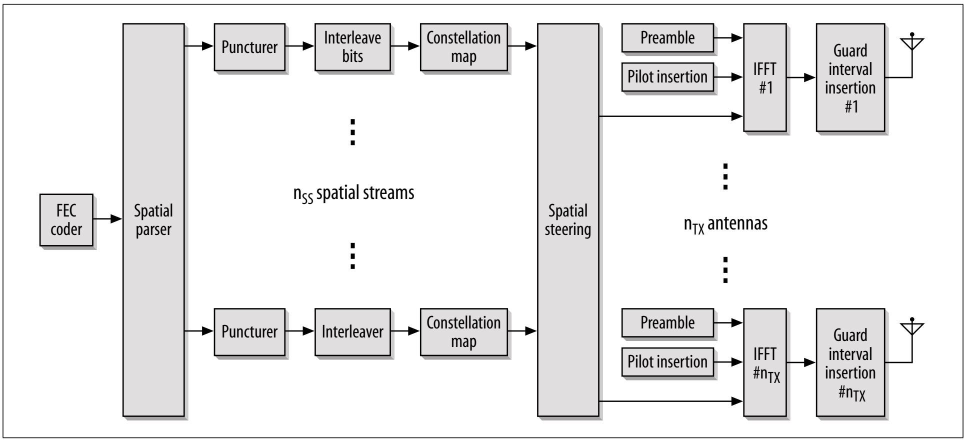

Figure 15-19. 2x2 TGnSync MIMO transceiver

Comparison and Conclusions

Both proposals are essentially MIMO evolutions of the 802.11a PHY. Both require support for a 2x2 mode, where both the sender and receiver have two transceivers active. However, it is likely that most products that are based on the eventual 802.11n standard (which may or may not resemble either of the proposals) will support at least some of the optional modes. It is likely that cost constraints on client devices will restrict them to operating in a two-transceiver mode, while APs will have more transceivers. Basic APs may have only two, while the most expensive enterprise-class APs have three or four transceivers.

Table 15-5 shows the data rates for the two spatial stream modes for each proposal. Higher data rates are possible with additional spatial streams, albeit at the extra cost of silicon. TGnSync's proposed peak rates are higher, although at the cost of more aggressive coding. WWiSE's 135 Mbps data rate is accomplished by using 64-QAM at R=5/6; TGnSync gets to 140 Mbps by using a 7/8 rate code and cutting the guard interval to half. (Without the short guard interval, TGnSync's data rate is only 126 Mbps.) The advanced beamforming mode uses the much larger 256-QAM constellation. Though TGnSync's data rates are higher, I expect that the more aggressive coding will lead to shorter range.

Table 15-5. Top speed for major 802.11n proposals (two spatial streams)

	20MHz channels	40 MHz channels
WWISE	135 Mbps	270 Mbps
TGnSync		
Basic mode	140 Mbps (+3.7%)	315 Mbps (+16.7%)
Advanced beamforming mode	160 Mbps (+18.5%)	360 Mbps (+33.3%)

Spectral usage is a major point of contention between the two groups. WWiSE has focused much more heavily on improving MAC efficiency than on the data rate, even going so far as to argue that using 40 MHz channels to improve the data rate before improving efficiency is a waste of scarce unlicensed spectrum. While there may be merit to that point, the 40 MHz channelization approach used by TGnSync has the advantage of being able to reclaim spectrum in the middle of a wider channel. WWiSE merely doubles throughput in their 40 MHz mode, while TGnSync squeezes more than double the capacity out of the wider channel. Both approaches have their drawbacks. TGnSync's proposal would probably lead to chipsets that are always capable of 40 MHz operation, adding extra cost and complexity, even though regulators may not allow them. In countries where 40 MHz channels are allowed, the extra speed would be welcome. In areas where 40 MHz channels are only a pipe dream of chipset manufacturers, the extra cost may not be welcome. On the other hand, WWiSE's denial of the need for high-speed channels seems to be denying the five-fold leap in data rates that occurs with every new 802.11 PHY.

For maximum speed, TGnSync requires closed-loop operation, which would be a major undertaking to implement in silicon. Sounding frames must be used to measure the channel, and responses must be collected to calibrate the radio channel. WWiSE uses only open loop operation, which is simpler to implement. The WWiSE proposal also offers the ability to spread a single encoded stream across multiple antennas without using closed-loop operation. If closed-loop operation were to be problematic to implement in silicon, 802.11n could be delayed unacceptably.

Frame aggregation is an important part of meeting the larger goal set for the eventual 802.11n standard. However, taking full advantage of aggregation opportunities requires more intelligent queuing than is currently implemented. Whether 802.11n offers a huge increase in speed is likely to depend a great deal on how well improved queuing algorithms are able to coalesce collections of small packets into large aggregates. Neither proposal specifies queueing, so the performance boost may vary between vendors.

Aggregation as designed by the protocol is a bit more intelligent in TGnSync, although this is only a minor advantage. WWiSE's proposal only allows aggregation when the Address 1 field in the MAC header is the same. In an infrastructure network, the Address 1 field is the BSSID. All frames from a station to an AP can be aggregated, so the two proposals are identical in the upstream direction. In the downstream direction, WWiSE must use a physical-layer frame burst to change directions. Each new direction must have a new PLCP header. TGnSync can reduce overhead by using a multiple-receiver aggregate frame, and collecting responses from each receiver in the aggregate.

Powersaving modes in 802.11 have been neglected, and are not very sophisticated. TGnSync attempts to come up with extended powersaving operations for some of the new MAC structures, while WWiSE does not. This is likely due to the presence of equipment vendors that use chips in the TGnSync consortium, while WWiSE is comprised only of chip vendors. While the trade-off between high speed and battery life is application-specific and may not always make sense, it is always good to see the standards bodies thinking ahead about problems that may occur.