

2nd Edition

Covers 802.11a, g, n & i

Creating & Administering Wireless Networks

802.11 Wireless Networks

The Definitive Guide

O'REILLY®

Matthew S. Gast



SECOND EDITION

802.11[®] Wireless Network

The Definitive Guide

E & S
TK
5103.2
.G385
2005

O'REILLY[®]

802.11 Wireless Networks: The Definitive Guide



Using a wireless network is a liberating experience. But underneath the experience lies a complex protocol, and even more complex issues arise when your data isn't limited to traveling on physical wires. How do you structure your network so mobile users can move around effectively? How do you extend wireless coverage so it's available everywhere you need it? What kinds of security issues do wireless networks raise? How do you tune your network for optimal performance? How do you provide enough capacity to support the users you expect initially, and how do you deal with the problems that arise as more users join the network?

802.11 Wireless Networks: The Definitive Guide, Second Edition discusses all these issues, and more. This book is for the serious system or network administrator who is responsible for deploying or maintaining a wireless network. It contains an extensive discussion of wireless security issues, including the problems with the WEP standard and a look at the alternatives: dynamic WEP, plus the 802.1X and the 802.11i security standards. Since network monitoring is essential to any serious network administrator, a chapter is devoted to network analysis and troubleshooting, using Ethereal and other tools.

802.11 Wireless Networks: The Definitive Guide brings you up-to-date on all the latest developments in wireless networking. In addition to 802.11b and 11a, this edition covers 802.11g and looks ahead to the 802.11n protocol, which is currently being standardized. This new edition greatly expands the discussion of network planning and architecture, paying special attention to mobility between access points, spectrum management, and power control. It's the only book available that discusses how to calculate the performance of your wireless network and how to tune your network for optimal performance.

Finally, *802.11 Wireless Networks: The Definitive Guide* shows you how to configure wireless cards and Linux, Windows, and Mac OS X systems, and how to work with access points. Few books in any field combine the theory you need to know with the practical experience and advice you need to get things working. *802.11 Wireless Networks: The Definitive Guide* is one of those books. If you are responsible for a wireless network, you need this book.

Matthew S. Gast is the leading author on the planning and deployment of wireless networks.

www.oreilly.com

ISBN 0-596-10052-3



9 780596 100520

US \$44.95

CAN \$62.95



6 36920 10052 2

Safari
BOOKS ONLINE
ENABLED

Includes
**FREE 45-Day
Online Edition**

University Libraries
Carnegie Mellon University
Pittsburgh PA 15213-3890

SECOND EDITION

802.11® Wireless Networks

The Definitive Guide

The Definitive Guide

Cover Designer: Elise Goldwasser

Matthew S. Coats

Matthew S. Gast

O'REILLY®

Beijing • Cambridge • Farnham • Köln • Paris • Sebastopol • Taipei • Tokyo

621.382
G25e2

802.11® Wireless Networks: The Definitive Guide, Second Edition

by Matthew S. Gast

Copyright © 2005 Matthew S. Gast. All rights reserved.
Printed in the United States of America.

Published by O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472.

O'Reilly books may be purchased for educational, business, or sales promotional use. Online editions are also available for most titles (safari.oreilly.com). For more information, contact our corporate/institutional sales department: (800) 998-9938 or corporate@oreilly.com.

Editor: Mike Loukides

Production Editor: Colleen Gorman

Cover Designer: Ellie Volckhausen

Interior Designer: David Futato

Printing History:

April 2002: First Edition.

April 2005: Second Edition.

Nutshell Handbook, the Nutshell Handbook logo, and the O'Reilly logo are registered trademarks of O'Reilly Media, Inc. 802.11® *Wireless Networks: The Definitive Guide*, Second Edition, the image of a horseshoe bat, and related trade dress are trademarks of O'Reilly Media, Inc.

802.11® and all 802.11-based trademarks and logos are trademarks or registered trademarks of IEEE, Inc. in the United States and other countries. O'Reilly Media, Inc. is independent of IEEE.

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and O'Reilly Media, Inc. was aware of a trademark claim, the designations have been printed in caps or initial caps.

While every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions, or for damages resulting from the use of the information contained herein.



This book uses RepKover™, a durable and flexible lay-flat binding.

ISBN: 0-596-10052-3

[M]

Table of Contents

Foreword	xi
Preface	xiii
1. Introduction to Wireless Networking	1
Why Wireless?	1
What Makes Wireless Networks Different	5
A Network by Any Other Name.....	8
2. Overview of 802.11 Networks	12
IEEE 802 Network Technology Family Tree	13
802.11 Nomenclature and Design	14
802.11 Network Operations	23
Mobility Support	28
3. 802.11 MAC Fundamentals	32
Challenges for the MAC	34
MAC Access Modes and Timing	37
Contention-Based Access Using the DCF	41
Fragmentation and Reassembly	44
Frame Format	46
Encapsulation of Higher-Layer Protocols Within 802.11	54
Contention-Based Data Service	55
Frame Processing and Bridging	63

University Libraries
Carnegie Mellon University
Pittsburgh PA 15213-3890

4. 802.11 Framing in Detail	67
Data Frames	67
Control Frames	76
Management Frames	82
Frame Transmission and Association and Authentication States	111
5. Wired Equivalent Privacy (WEP)	114
Cryptographic Background to WEP	114
WEP Cryptographic Operations	117
Problems with WEP	123
Dynamic WEP	127
6. User Authentication with 802.1X	129
The Extensible Authentication Protocol	130
EAP Methods	135
802.1X: Network Port Authentication	141
802.1X on Wireless LANs	145
7. 802.11i: Robust Security Networks, TKIP, and CCMP	149
The Temporal Key Integrity Protocol (TKIP)	149
Counter Mode with CBC-MAC (CCMP)	159
Robust Security Network (RSN) Operations	163
8. Management Operations	170
Management Architecture	170
Scanning	171
Authentication	176
Preauthentication	180
Association	184
Power Conservation	188
Timer Synchronization	197
Spectrum Management	199
9. Contention-Free Service with the PCF	214
Contention-Free Access Using the PCF	214
Detailed PCF Framing	218
Power Management and the PCF	223

10. Physical Layer Overview	225
Physical-Layer Architecture	225
The Radio Link	226
RF Propagation with 802.11	232
RF Engineering for 802.11	238
11. The Frequency-Hopping (FH) PHY	241
Frequency-Hopping Transmission	241
Gaussian Frequency Shift Keying (GFSK)	246
FH PHY Convergence Procedure (PLCP)	249
Frequency-Hopping PMD Sublayer	250
Characteristics of the FH PHY	252
12. The Direct Sequence PHYs: DSSS and HR/DSSS (802.11b)	253
Direct Sequence Transmission	253
Differential Phase Shift Keying (DPSK)	260
The “Original” Direct Sequence PHY	263
Complementary Code Keying	267
High Rate Direct Sequence PHY	268
13. 802.11a and 802.11j: 5-GHz OFDM PHY	276
Orthogonal Frequency Division Multiplexing (OFDM)	277
OFDM as Applied by 802.11a	282
OFDM PLCP	290
OFDM PMD	293
Characteristics of the OFDM PHY	297
14. 802.11g: The Extended-Rate PHY (ERP)	298
802.11g Components	298
ERP Physical Layer Convergence (PLCP)	304
ERP Physical Medium Dependent (PMD) Layer	307
15. A Peek Ahead at 802.11n: MIMO-OFDM	311
Common Features	312
WWiSE	314
TGnSync	325
Comparison and Conclusions	340

16. 802.11 Hardware	343
General Structure of an 802.11 Interface	343
Implementation-Specific Behavior	348
Reading the Specification Sheet	350
17. Using 802.11 on Windows	353
Windows XP	353
Windows 2000	368
Windows Computer Authentication	369
18. 802.11 on the Macintosh	373
The AirPort Extreme Card	373
802.1X on the AirPort	378
19. Using 802.11 on Linux	387
PCMCIA Support on Linux	387
Linux Wireless Extensions and Tools	395
Agere (Lucent) Orinoco	402
Atheros-Based cards and MADwifi	404
802.1X on Linux with xsupplicant	407
20. Using 802.11 Access Points	415
General Functions of an Access Point	416
Power over Ethernet (PoE)	422
Selecting Access Points	424
Cisco 1200 Access Point	427
Apple AirPort	434
21. Logical Wireless Network Architecture	438
Evaluating a Logical Architecture	438
Topology Examples	452
Choosing Your Logical Architecture	474
22. Security Architecture	476
Security Definition and Analysis	477
Authentication and Access Control	482
Ensuring Secrecy Through Encryption	492

Selecting Security Protocols	500
Rogue Access Points	510
23. Site Planning and Project Management	518
Project Planning and Requirements	519
Network Requirements	521
Physical Layer Selection and Design	532
Planning Access-Point Placement	537
Using Antennas to Tailor Coverage	546
24. 802.11 Network Analysis	553
Network Analyzers	553
Ethereal	556
802.11 Network Analysis Checklist	566
Other Tools	570
25. 802.11 Performance Tuning	573
802.11 Performance Calculations	573
Improving Performance	576
Tunable 802.11 Parameters	578
26. Conclusions and Predictions	585
Standards Work	585
Current Trends in Wireless Networking	587
The End	596
Glossary	597
Index	607

data frame, and the final acknowledgment are all considered part of the same atomic operation. Figure 3-3 illustrates the procedure.




Figure 3-3. RTS/CTS clearing

In Figure 3-3, node 1 has a frame to send; it initiates the process by sending an RTS frame. The RTS frame serves several purposes: in addition to reserving the radio link for transmission, it silences any stations that hear it. If the target station receives an RTS, it responds with a CTS. Like the RTS frame, the CTS frame silences stations in the immediate vicinity. Once the RTS/CTS exchange is complete, node 1 can transmit its frames without worry of interference from any hidden nodes. Hidden nodes beyond the range of the sending station are silenced by the CTS from the receiver. When the RTS/CTS clearing procedure is used, any frames must be positively acknowledged.

The multiframe RTS/CTS transmission procedure consumes a fair amount of capacity, especially because of the additional latency incurred before transmission can commence. As a result, it is used only in high-capacity environments and environments with significant contention on transmission. For lower-capacity environments, it is not necessary.

Hidden nodes have also become less of a problem as 802.11 has grown up. In small, quiescent networks with only a few stations associated to an access point, there is very little risk of simultaneous transmission, and plenty of spare capacity to be used for retransmission. In many larger environments, the coverage is dense enough that the clients are located physically close enough to an access point that they are all within range of each other. (In fact, the range of many client systems is probably too large for most networks, which will be explored in the planning phase of this book.)

You can control the RTS/CTS procedure by setting the *RTS threshold* if the device driver for your 802.11 card allows you to adjust it. The RTS/CTS exchange is performed for frames larger than the threshold. Frames shorter than the threshold are simply sent.

MAC Access Modes and Timing

Access to the wireless medium is controlled by coordination functions. Ethernet-like CSMA/CA access is provided by the *distributed coordination function* (DCF). If contention-free service is required, it can be provided by the *point coordination function* (PCF), which is built on top of the DCF. Between the free-for-all of the DCF and the precision of the PCF, networks can use the *hybrid coordination function* (HCF), a middle ground for quality of service between the two extremes. Contention-free services are provided only in infrastructure networks, but quality of service may be provided in any network that has HCF support in the stations. The coordination functions are described in the following list and illustrated in Figure 3-4:

DCF

The DCF is the basis of the standard CSMA/CA access mechanism. Like Ethernet, it first checks to see that the radio link is clear before transmitting. To avoid collisions, stations use a random backoff after each frame, with the first transmitter seizing the channel. In some circumstances, the DCF may use the CTS/RTS clearing technique to further reduce the possibility of collisions.

PCF

The point coordination function provides contention-free services. Special stations called *point coordinators* are used to ensure that the medium is provided without contention. Point coordinators reside in access points, so the PCF is restricted to infrastructure networks. To gain priority over standard contention-based services, the PCF allows stations to transmit frames after a shorter interval. The PCF is not widely implemented and is described in Chapter 9.

HCF

Some applications need to have service quality that is a step above best-effort delivery, but the rigorous timing of the PCF is not required. The HCF allows stations to maintain multiple service queues and balance access to the wireless medium in favor of applications that require better service quality. The HCF is not fully standardized yet, but is being produced as part of the eventual 802.11e specification.

Adding quality of service to the 802.11 MAC was a significant undertaking. Due to the added complexity in terms of framing, queue management, and signaling, the quality of service specifications were still subject to standards committee wrangling as this book was written, and discussion of them must be postponed to a future edition.

Carrier-Sensing Functions and the Network Allocation Vector

Carrier sensing is used to determine if the medium is available. Two types of carrier-sensing functions in 802.11 manage this process: the physical carrier-sensing and virtual carrier-sensing functions. If either carrier-sensing function indicates that the medium is busy, the MAC reports this to higher layers.




Figure 3-4. MAC coordination functions

Physical carrier-sensing functions are provided by the physical layer in question and depend on the medium and modulation used. It is difficult (or, more to the point, expensive) to build physical carrier-sensing hardware for RF-based media, because transceivers can transmit and receive simultaneously only if they incorporate expensive electronics. Furthermore, with hidden nodes potentially lurking everywhere, physical carrier-sensing cannot provide all the necessary information.

Virtual carrier-sensing is provided by the Network Allocation Vector (NAV). Most 802.11 frames carry a duration field, which can be used to reserve the medium for a fixed time period. The NAV is a timer that indicates the amount of time the medium will be reserved, in microseconds. Stations set the NAV to the time for which they expect to use the medium, including any frames necessary to complete the current operation. Other stations count down from the NAV to 0. When the NAV is non-zero, the virtual carrier-sensing function indicates that the medium is busy; when the NAV reaches 0, the virtual carrier-sensing function indicates that the medium is idle.

By using the NAV, stations can ensure that atomic operations are not interrupted. For example, the RTS/CTS sequence in Figure 3-3 is atomic. Figure 3-5 shows how the NAV protects the sequence from interruption. (This is a standard format for a number of diagrams in this book that illustrate the interaction of multiple stations with the corresponding timers.) Activity on the medium by stations is represented by the shaded bars, and each bar is labeled with the frame type. Interframe spacing is depicted by the lack of any activity. Finally, the NAV timer is represented by the bars on the NAV line at the bottom of the figure. The NAV is carried in the frame headers on the RTS and CTS frames; it is depicted on its own line to show how the NAV relates to actual transmissions in the air. When a NAV bar is present on the NAV line, stations should defer access to the medium because the virtual carrier-sensing mechanism will indicate a busy medium.

To ensure that the sequence is not interrupted, node 1 sets the NAV in its RTS to block access to the medium while the RTS is being transmitted. All stations that hear the RTS defer access to the medium until the NAV elapses.




Figure 3-5. Using the NAV for virtual carrier sensing

RTS frames are not necessarily heard by every station in the network. Therefore, the recipient of the intended transmission responds with a CTS that includes a shorter NAV. This NAV prevents other stations from accessing the medium until the transmission completes. After the sequence completes, the medium can be used by any station after distributed interframe space (DIFS), which is depicted by the contention window beginning at the right side of the figure.

RTS/CTS exchanges may be useful in crowded areas with multiple overlapping networks. Every station on the same physical channel receives the NAV and defers access appropriately, even if the stations are configured to be on different networks.

Interframe Spacing

As with traditional Ethernet, the interframe spacing plays a large role in coordinating access to the transmission medium. 802.11 uses four different interframe spaces. Three are used to determine medium access; the relationship between them is shown in Figure 3-6.




Figure 3-6. Interframe spacing relationships

We've already seen that as part of the collision avoidance built into the 802.11 MAC, stations delay transmission until the medium becomes idle. Varying interframe spacings create different priority levels for different types of traffic. The logic behind this

is simple: high-priority traffic doesn't have to wait as long after the medium has become idle. Therefore, if there is any high-priority traffic waiting, it grabs the network before low-priority frames have a chance to try. To assist with interoperability between different data rates, the interframe space is a fixed amount of time, independent of the transmission speed. (This is only one of the many problems caused by having different physical layers use the same radio resources, which are different modulation techniques.) Different physical layers, however, can specify different interframe space times.

Short interframe space (SIFS)

The SIFS is used for the highest-priority transmissions, such as RTS/CTS frames and positive acknowledgments. High-priority transmissions can begin once the SIFS has elapsed. Once these high-priority transmissions begin, the medium becomes busy, so frames transmitted after the SIFS has elapsed have priority over frames that can be transmitted only after longer intervals.

PCF interframe space (PIFS)

The PIFS, sometimes erroneously called the priority interframe space, is used by the PCF during contention-free operation. Stations with data to transmit in the contention-free period can transmit after the PIFS has elapsed and preempt any contention-based traffic.

DCF interframe space (DIFS)

The DIFS is the minimum medium idle time for contention-based services. Stations may have immediate access to the medium if it has been free for a period longer than the DIFS.

Extended interframe space (EIFS)

The EIFS is not illustrated in Figure 3-6 because it is not a fixed interval. It is used only when there is an error in frame transmission.

Interframe spacing and priority

Atomic operations start like regular transmissions: they must wait for the relevant interframe space, typically the DIFS, to end before they can begin. However, the second and any subsequent steps in an atomic operation take place after the SIFS, rather than after the DIFS. Because the SIFS is shorter than the other interframe spaces, the second (and subsequent) parts of an atomic operation will grab the medium before another type of frame can be transmitted. By using the SIFS and the NAV, stations can seize the medium for as long as necessary.

In Figure 3-5, for example, the short interframe space is used between the different units of the atomic exchange. After the sender gains access to the medium, the receiver replies with a CTS after the SIFS. Any stations that might attempt to access the medium at the conclusion of the RTS would wait for one DIFS interval. Partway through the DIFS interval, though, the SIFS interval elapses, and the CTS is transmitted.

Contention-Based Access Using the DCF

Most traffic uses the DCF, which provides a standard Ethernet-like contention-based service. The DCF allows multiple independent stations to interact without central control, and thus may be used in either IBSS networks or in infrastructure networks.

Before attempting to transmit, each station checks whether the medium is idle. If the medium is not idle, stations defer to each other and employ an orderly exponential backoff algorithm to avoid collisions.

In distilling the 802.11 MAC rules, there is a basic set of rules that are always used, and additional rules may be applied depending on the circumstances. Two basic rules apply to all transmissions using the DCF:

1. If the medium has been idle for longer than the DIFS, transmission can begin immediately. Carrier sensing is performed using both a physical medium-dependent method and the virtual (NAV) method.
 - a. If the previous frame was received without errors, the medium must be free for at least the DIFS.
 - b. If the previous transmission contained errors, the medium must be free for the amount of the EIFS.
2. If the medium is busy, the station must wait for the channel to become idle. 802.11 refers to the wait as *access deferral*. If access is deferred, the station waits for the medium to be idle for the DIFS and prepares for the exponential backoff procedure.

Additional rules may apply in certain situations. Many of these rules depend on the particular situation “on the wire” and are specific to the results of previous transmissions.

1. Error recovery is the responsibility of the station sending a frame. Senders expect acknowledgments for each transmitted frame and are responsible for retrying the transmission until it is successful.
 - a. Positive acknowledgments are the only indication of success. Atomic exchanges must complete in their entirety to be successful. If an acknowledgment is expected and does not arrive, the sender considers the transmission lost and must retry.
 - b. All unicast data must be acknowledged. Broadcast data is not acknowledged. (As a result, unicast data has an inherently higher service quality than broadcast data, even though the radio link is inherently a broadcast medium.)
 - c. Any failure increments a retry counter, and the transmission is retried. A failure can be due to a failure to gain access to the medium or a lack of an acknowledgment. However, there is a longer congestion window when transmissions are retried (see next section).
2. Multiframe sequences may update the NAV with each step in the transmission procedure. When a station receives a medium reservation that is longer than the current NAV, it updates the NAV. Setting the NAV is done on a frame-by-frame basis and is discussed in much more detail in the next chapter.




Figure 3-8. Fragmentation burst

Fragments and their acknowledgments are separated by the SIFS, so a station retains control of the channel during a fragmentation burst. The NAV is also used to ensure that other stations do not use the channel during the fragmentation burst. As with any RTS/CTS exchange, the RTS and CTS both set the NAV from the expected time to the end of the first fragments in the air. Subsequent fragments then form a chain. Each fragment sets the NAV to hold the medium until the end of the acknowledgment for the next frame. Fragment 0 sets the NAV to hold the medium until ACK 1, fragment 1 sets the NAV to hold the medium until ACK 2, and so on. After the last fragment and its acknowledgment have been sent, the NAV is set to 0, indicating that the medium will be released after the fragmentation burst completes.

Frame Format

To meet the challenges posed by a wireless data link, the MAC was forced to adopt several unique features, not the least of which was the use of four address fields. Not all frames use all the address fields, and the values assigned to the address fields may change depending on the type of MAC frame being transmitted. Details on the use of address fields in different frame types are presented in Chapter 4. Figure 3-9 shows the generic 802.11 MAC frame. All diagrams in this section follow the IEEE conventions in 802.11. Fields are transmitted from left to right.




Figure 3-9. Generic 802.11 MAC frame

802.11 MAC frames do not include some of the classic Ethernet frame features, most notably the type/length field and the preamble. The preamble is part of the physical layer, and encapsulation details such as type and length are present in the header on the data carried in the 802.11 frame.

Frame Control

Each frame starts with a two-byte Frame Control subfield, shown in Figure 3-10. The components of the Frame Control subfield are:

Protocol version

Two bits indicate which version of the 802.11 MAC is contained in the rest of the frame. At present, only one version of the 802.11 MAC has been developed; it is assigned the protocol number 0. Other values will appear when the IEEE standardizes changes to the MAC that render it incompatible with the initial specification. So far, none of the revisions to 802.11 have required incrementing the protocol number.

Type and subtype fields

Type and subtype fields identify the type of frame used. To cope with noise and unreliability, a number of management functions are incorporated into the 802.11 MAC. Some, such as the RTS/CTS operations and the acknowledgments, have already been discussed. Table 3-1 shows how the type and subtype identifiers are used to create the different classes of frames.




Figure 3-10. Frame control field

In Table 3-1, bit strings are written most-significant bit first, which is the reverse of the order used in Figure 3-10. Therefore, the frame type is the third bit in the frame control field followed by the second bit (b3 b2), and the subtype is the seventh bit, followed by the sixth, fifth, and fourth bits (b7 b6 b5 b4).

Table 3-1. Type and subtype identifiers

Subtype value	Subtype name
Management frames (type=00) ^a	
0000	Association request
0001	Association response
0010	Reassociation request
0011	Reassociation response
0100	Probe request
0101	Probe response

Table 3-1. Type and subtype identifiers (continued)

Subtype value	Subtype name
1000	Beacon
1001	Announcement traffic indication message (ATIM)
1010	Disassociation
1011	Authentication
1100	Deauthentication
1101	Action (for spectrum management with 802.11h, also for QoS)
Control frames (type=01) ^b	
1000	Block Acknowledgment Request (QoS)
1001	Block Acknowledgment (QoS)
1010	Power Save (PS)-Poll
1011	RTS
1100	CTS
1101	Acknowledgment (ACK)
1110	Contention-Free (CF)-End
1111	CF-End+CF-Ack
Data frames (type=10)	
0000	Data
0001	Data+CF-Ack
0010	Data+CF-Poll
0011	Data+CF-Ack+CF-Poll
0100	Null data (no data transmitted)
0101	CF-Ack (no data transmitted)
0110	CF-Poll (no data transmitted)
0111	CF-Ack+CF-Poll (no data transmitted)
1000	QoS Data ^c
1001	QoS Data + CF-Ack ^c
1010	QoS Data + CF-Poll ^c
1011	QoS Data + CF-Ack + CF-Poll ^c
1100	QoS Null (no data transmitted) ^c
1101	QoS CF-Ack (no data transmitted) ^c
1110	QoS CF-Poll (no data transmitted) ^c
1111	QoS CF-Ack+CF-Poll (no data transmitted) ^c
(Frame type 11 is reserved)	

^a Management subtypes 0110–0111 and 1110–1111 are reserved and not currently used.

^b Control subtypes 0000–0111 are reserved and not currently used.

^c Proposed by the 802.11e task group, but not yet standardized. Note that these frames all have a leading one, which has caused some to refer to the first bit as the QoS bit.

ToDS and FromDS bits

These bits indicate whether a frame is destined for the distribution system. All frames on infrastructure networks will have one of the distribution system bits set. Table 3-2 shows how these bits are interpreted. As Chapter 4 will explain, the interpretation of the address fields depends on the setting of these bits.

Table 3-2. Interpreting the ToDS and FromDS bits

	To DS=0	To DS=1
From DS=0	All management and control frames Data frames within an IBSS (never infrastructure data frames)	Data frames transmitted from a wireless station in an infrastructure network.
From DS=1	Data frames received for a wireless station in an infrastructure network	Data frames on a "wireless bridge"

More fragments bit

This bit functions much like the "more fragments" bit in IP. When a higher-level packet has been fragmented by the MAC, the initial fragment and any following nonfinal fragments set this bit to 1. Large data frames and some management frames may be large enough to require fragmentation; all other frames set this bit to 0. In practice, most data frames are transmitted at the maximum Ethernet size and fragmentation is not often used.

Retry bit

From time to time, frames may be retransmitted. Any retransmitted frames set this bit to 1 to aid the receiving station in eliminating duplicate frames.

Power management bit

Network adapters built on 802.11 are often built to the PC Card form factor and used in battery-powered laptop or handheld computers. To conserve battery life, many small devices have the ability to power down parts of the network interface. This bit indicates whether the sender will be in a powersaving mode after the completion of the current atomic frame exchange. 1 indicates that the station will be in powersave mode, and 0 indicates that the station will be active. Access points perform a number of important management functions and are not allowed to save power, so this bit is always 0 in frames transmitted by an access point.

More data bit

To accommodate stations in a powersaving mode, access points may buffer frames received from the distribution system. An access point sets this bit to indicate that at least one frame is available and is addressed to a dozing station.

Protected Frame bit

Wireless transmissions are inherently easier to intercept than transmissions on a fixed network. If the frame is protected by link layer security protocols, this bit is set to 1, and the frame changes slightly. The Protected Frame bit was previously called the WEP bit.

Order bit

Frames and fragments can be transmitted in order at the cost of additional processing by both the sending and receiving MACs. When the “strict ordering” delivery is employed, this bit is set to 1.

Duration/ID Field

The Duration/ID field follows the frame control field. This field has several uses and takes one of the three forms shown in Figure 3-11.




Figure 3-11. Duration/ID field

Duration: setting the NAV

When bit 15 is 0, the duration/ID field is used to set the NAV. The value represents the number of microseconds that the medium is expected to remain busy for the transmission currently in progress. All stations must monitor the headers of all frames they receive and update the NAV accordingly. Any value that extends the amount of time the medium is busy updates the NAV and blocks access to the medium for additional time.

Frames transmitted during contention-free periods

During the contention-free periods, bit 14 is 0 and bit 15 is 1. All other bits are 0, so the duration/ID field takes a value of 32,768. This value is interpreted as a NAV. It allows any stations that did not receive the Beacon* announcing the contention-free period to update the NAV with a suitably large value to avoid interfering with contention-free transmissions.

PS-Poll frames

Bits 14 and 15 are both set to 1 in PS-Poll frames. Mobile stations may elect to save battery power by turning off antennas. Dozing stations must wake up periodically. To

* Beacon frames are a subtype of management frames, which is why “Beacon” is capitalized.

ensure that no frames are lost, stations awaking from their slumber transmit a PS-Poll frame to retrieve any buffered frames from the access point. Along with this request, waking stations incorporate the association ID (AID) that indicates which BSS they belong to. The AID is included in the PS-Poll frame and may range from 1–2,007. Values from 2,008–16,383 are reserved and not used.

Address Fields

An 802.11 frame may contain up to four address fields. The address fields are numbered because different fields are used for different purposes depending on the frame type (details are found in Chapter 4). The general rule of thumb is that Address 1 is used for the receiver, Address 2 for the transmitter, and Address 3 field for filtering by the receiver. In an infrastructure network, for example, the third address field is used by the receiver to determine whether the frame is part of the network it is associated to.*

Addressing in 802.11 follows the conventions used for the other IEEE 802 networks, including Ethernet. Addresses are 48 bits long. If the first bit sent to the physical medium is a 0, the address represents a single station (unicast). When the first bit is a 1, the address represents a group of physical stations and is called a *multicast* (or *group*) address. If all bits are 1s, then the frame is a *broadcast* and is delivered to all stations connected to the wireless medium.

48-bit addresses are used for a variety of purposes:

Destination address

As in Ethernet, the destination address is the 48-bit IEEE MAC identifier that corresponds to the final recipient: the station that will hand the frame to higher protocol layers for processing.

Source address

This is the 48-bit IEEE MAC identifier that identifies the source of the transmission. Only one station can be the source of a frame, so the Individual/Group bit is always 0 to indicate an individual station.

Receiver address

This is a 48-bit IEEE MAC identifier that indicates which wireless station should process the frame. If it is a wireless station, the receiver address is the destination address. For frames destined to a node on an Ethernet connected to an access point, the receiver is the wireless interface in the access point, and the destination address may be a router attached to the Ethernet.

* 802.11 specifies that stations should ignore frames that do not have the same BSSID, but most products do not correctly implement BSSID filtering and will pass any received frame up to higher protocol layers.

Transmitter address

This is a 48-bit IEEE MAC address to identify the wireless interface that transmitted the frame onto the wireless medium. The transmitter address is used only in wireless bridging.

Basic Service Set ID (BSSID)

To identify different wireless LANs in the same area, stations may be assigned to a BSS. In infrastructure networks, the BSSID is the MAC address used by the wireless interface in the access point. Ad hoc networks generate a random BSSID with the Universal/Local bit set to 1 to prevent conflicts with officially assigned MAC addresses.

The number of address fields used depends on the type of frame. Most data frames use three fields for source, destination, and BSSID. The number and arrangement of address fields in a data frame depends on how the frame is traveling relative to the distribution system. Most transmissions use three addresses, which is why only three of the four addresses are contiguous in the frame format.

Sequence Control Field

This 16-bit field is used for both defragmentation and discarding duplicate frames. It is composed of a 4-bit fragment number field and a 12-bit sequence number field, as shown in Figure 3-12. Sequence numbers are not used in control frames, so the Sequence Control field is not present.




Figure 3-12. Sequence Control field

Higher-level frames are each given a sequence number as they are passed to the MAC for transmission. The sequence number subfield operates as a modulo-4096 counter of the frames transmitted. It begins at 0 and increments by 1 for each higher-level packet handled by the MAC. If higher-level packets are fragmented, all fragments will have the same sequence number. When frames are retransmitted, the sequence number is not changed.

What differs between fragments is the fragment number. The first fragment is given a fragment number of 0. Each successive fragment increments the fragment number by 1. Retransmitted fragments keep their original sequence numbers to assist in reassembly.

Stations that implement the QoS extensions use a slightly different interpretation of the sequence control field because multiple transmit queues need to be maintained.

Frame Body

The frame body, also called the Data field, moves the higher-layer payload from station to station. As originally specified, 802.11 can transmit frames with a maximum payload of 2,304 bytes of higher-level data. Implementations must support larger frame bodies to accommodate additional headers for security and QoS. 802.2 LLC headers use 8 bytes for a maximum network protocol payload of 2,296 bytes. Preventing fragmentation must be done at the protocol layer. On IP networks, Path MTU Discovery (RFC 1191) will prevent the transmission of frames with Data fields larger than 1,500 bytes.

802.11 differs from other link layer technologies in two notable ways. First, there is no higher-level protocol tag in the 802.11 frame to distinguish between higher-layer protocol types. Higher-level protocols are tagged with a type field by an additional header, which is used as the start of the 802.11 payload. Second, 802.11 does not generally pad frames to a minimum length. Many frames used by 802.11 are short, and the chips and electronics used in network interfaces has progressed to the point where a pad is no longer necessary.

Frame Check Sequence

As with Ethernet, the 802.11 frame closes with a frame check sequence (FCS). The FCS is often referred to as the cyclic redundancy check (CRC) because of the underlying mathematical operations. The FCS allows stations to check the integrity of received frames. All fields in the MAC header and the body of the frame are included in the FCS. Although 802.3 and 802.11 use the same method to calculate the FCS, the MAC header used in 802.11 is different from the header used in 802.3, so the FCS must be recalculated by access points.

When frames are sent to the wireless interface, the FCS is calculated before those frames are sent out over the wireless link. Receivers can then calculate the FCS from the received frame and compare it to the received FCS. If the two match, there is a high probability that the frame was not damaged in transit.

On Ethernets, frames with a bad FCS are simply discarded, and frames with a good FCS are passed up the protocol stack. On 802.11 networks, frames that pass the integrity check may also require the receiver to send an acknowledgment. For example, data frames that are received correctly must be positively acknowledged, or they are retransmitted. 802.11 does not have a negative acknowledgment for frames that fail the FCS; stations must wait for the acknowledgment timeout before retransmitting.

Encapsulation of Higher-Layer Protocols Within 802.11

Like all other 802 link layers, 802.11 can transport any network-layer protocol. Unlike Ethernet, 802.11 relies on 802.2 logical-link control (LLC) encapsulation to carry higher-level protocols. Figure 3-13 shows how 802.2 LLC encapsulation is used to carry an IP packet. In the figure, the “MAC headers” for 802.1H and RFC 1042 might be the 12 bytes of source and destination MAC address information on Ethernet or the long 802.11 MAC header from the previous section.




Figure 3-13. IP encapsulation in 802.11

Two different methods can be used to encapsulate LLC data for transmission. One is described in RFC 1042, and the other in 802.1H. Both standards may go by other names. RFC 1042 is sometimes referred to as *IETF* encapsulation, while 802.1H is sometimes called *tunnel* encapsulation.

As you can see in Figure 3-13, though, the two methods are quite similar. An Ethernet frame is shown in the top line of Figure 3-13. It has a MAC header composed of source and destination MAC addresses, a type code, the embedded packet, and a frame check field. In the IP world, the Type code is either 0x0800 (2048 decimal) for IP itself, or 0x0806 (2054 decimal) for the Address Resolution Protocol (ARP).

Both RFC 1042 and 802.1H are derivatives of 802.2's *sub-network access protocol* (SNAP). The MAC addresses are copied into the beginning of the encapsulation frame, and then a SNAP header is inserted. SNAP headers begin with a *destination service access point* (DSAP) and a *source service access point* (SSAP). After the addresses, SNAP includes a Control header. Like high-level data link control (HDLC) and its progeny, the Control field is set to 0x03 to denote unnumbered information (UI), a category that maps well to the best-effort delivery of IP datagrams. The last field inserted by SNAP is an organizationally unique identifier (OUI). Initially, the IEEE hoped that the 1-byte service access points would be adequate to handle the number of network protocols, but this proved to be an overly optimistic assessment of the state of the world. As a result, SNAP copies the type code from the original Ethernet frame. The only difference between 802.1H and RFC 1042 is the OUI used.

At one point, many products offered the option to switch between the two encapsulation standards, though this option is much less common. Microsoft operating systems default to using 802.1H for the AppleTalk protocol suite and IPX, and use RFC 1042 for all other protocols. Most access points now conform to the Microsoft behavior, and no longer have an option to switch encapsulation type. In fact, the Microsoft encapsulation selection is so widely supported that it was part of the Wi-Fi Alliance's certification test suite at one point.

Contention-Based Data Service

The additional features incorporated into 802.11 to add reliability lead to a confusing tangle of rules about which types of frames are permitted at any point. They also make it more difficult for network administrators to know which frame exchanges they can expect to see on networks. This section clarifies the atomic exchanges that move data on an 802.11 LAN. (Most management frames are announcements to interested parties in the area and transfer information in only one direction.)

The exchanges presented in this section are atomic, which means that they should be viewed as a single unit. Two distinct sets of atomic exchanges are defined by 802.11. One is used by the DCF for contention-based service; those exchanges are described in this chapter. A second set of exchanges is specified for use with the PCF for contention-free services. Frame exchanges used with contention-free services are intricate and harder to understand. Since very few (if any) commercial products implement contention-free service, these exchanges are not described.

Frame exchanges under the DCF dominate the 802.11 MAC. According to the rules of the DCF, all products are required to provide best-effort delivery. To implement the contention-based MAC, stations process MAC headers for every frame while they are active. Exchanges begin with a station seizing an idle medium after the DIFS.

Broadcast and Multicast Data or Management Frames

Broadcast and multicast frames, which can also be referred to as *group frames* because they are destined for more than one receiving station, have the simplest frame exchanges because there is no acknowledgment. Framing and addressing are somewhat more complex in 802.11, so the types of frames that match this rule are the following:

- Broadcast data frames with a broadcast address in the Address 1 field
- Multicast data frames with a multicast address in the Address 1 field
- Broadcast management frames with a broadcast address in the Address 1 field (Beacon, Probe Request, and IBSS ATIM frames)

Frames destined for group addresses cannot be fragmented and are not acknowledged. The entire atomic sequence is a single frame, sent according to the rules of the contention-based access control. After the previous transmission concludes, all stations wait for the DIFS and begin counting down the random delay intervals in the contention window.

Because the frame exchange is a single-frame sequence, the NAV is set to 0. With no further frames to follow, there is no need to use the virtual carrier-sense mechanism to lock other stations out of using the medium. After the frame is transmitted, all stations wait through the DIFS and begin counting down through the contention window for any deferred frames. See Figure 3-14.




Figure 3-14. Broadcast/multicast data and broadcast management atomic frame exchange

Depending on the environment, frames sent to group addresses may have lower service quality because the frames are not acknowledged. Some stations may therefore miss broadcast or multicast traffic, but there is no facility built into the MAC for retransmitting broadcast or multicast frames.

Unicast Frames

Frames that are destined for a single station are called *directed* data by the 802.11 standard. This book uses the more common term *unicast*. Unicast frames must be acknowledged to ensure reliability, which means that a variety of mechanisms can be used to improve efficiency. All the sequences in this section apply to any unicast frame and thus can apply to management frames and data frames. In practice, these operations are usually observed only with data frames.

Basic positive acknowledgment (final fragment)

Reliable transmission between two stations is based on simple positive acknowledgments. Unicast data frames must be acknowledged, or the frame is assumed to be lost. The most basic case is a single frame and its accompanying acknowledgment, as shown in Figure 3-15.




Figure 3-15. Basic positive acknowledgment of data

The frame uses the NAV to reserve the medium for the frame, its acknowledgment, and the intervening SIFS. By setting a long NAV, the sender locks the virtual carrier for the entire sequence, guaranteeing that the recipient of the frame can send the acknowledgment. Because the sequence concludes with the ACK, no further virtual carrier locking is necessary, and the NAV in the ACK is set to 0.

Fragmentation

Many higher-layer network protocols, including IP, incorporate fragmentation. The disadvantage of network-layer fragmentation is that reassembly is performed by the final destination; if any of the fragments are lost, the entire packet must be retransmitted. Link layers may incorporate fragmentation to boost speed over a single hop with a small MTU.* 802.11 can also use fragmentation to help avoid interference. Radio interference is often in the form of short, high-energy bursts and is frequently synchronized with the AC power line. Breaking a large frame into small frames allows a larger percentage of the frames to arrive undamaged. The basic fragmentation scheme is shown in Figure 3-16.

The last two frames exchanged are the same as in the previous sequence, and the NAV is set identically. However, all frames prior to the penultimate frame use the NAV to lock the medium for the next frame. The first data frame sets the NAV for a long enough period to accommodate its ACK, the next fragment, and the acknowledgment following the next fragment. To indicate that it is a fragment, the MAC sets the More Fragments bit in the frame control field to 1. All nonfinal ACKs continue to extend the lock for the next data fragment and its ACK. Subsequent data frames then

* This is the approach used by Multi-link PPP (RFC 1990).




Figure 3-16. Fragmentation

continue to lengthen the NAV to include successive acknowledgments until the final data frame, which sets the More Fragments bit to 0, and the final ACK, which sets the NAV to 0. No limit is placed on the number of fragments, but the total frame length must be shorter than any constraint placed on the exchange by the PHY.

Fragmentation is controlled by the fragmentation threshold parameter in the MAC. Most network card drivers allow you to configure this parameter. Any frames larger than the fragmentation threshold are fragmented in an implementation-dependent way. Network administrators can change the fragmentation threshold to tune network behavior. Higher fragmentation thresholds mean that frames are delivered with less overhead, but the cost to a lost or damaged frame is much higher because more data must be discarded and retransmitted. Low fragmentation thresholds have much higher overhead, but they offer increased robustness in the face of hostile conditions.

RTS/CTS

To guarantee reservation of the medium and uninterrupted data transmission, a station can use the RTS/CTS exchange. Figure 3-17 shows this process. The RTS/CTS exchange acts exactly like the initial exchange in the fragmentation case, except that the RTS frame does not carry data. The NAV in the RTS allows the CTS to complete, and the CTS is used to reserve access for the data frame.

RTS/CTS can be used for all frame exchanges, none of them, or something in between. Like fragmentation, RTS/CTS behavior is controlled by a threshold set in the driver software. Frames larger than the threshold are preceded by an RTS/CTS exchange to clear the medium, while smaller frames are simply transmitted.




Figure 3-17. RTS/CTS lockout

RTS/CTS with fragmentation

In practice, the RTS/CTS exchange is often combined with fragmentation (Figure 3-18). Fragmented frames are usually quite long and thus benefit from the use of the RTS/CTS procedure to ensure exclusive access to the medium, free from contention from hidden nodes. Some vendors set the default fragmentation threshold to be identical to the default RTS/CTS threshold.




Figure 3-18. RTS/CTS with fragmentation

Powersaving Sequences

The most power-hungry components in RF systems are the amplifiers used to boost a signal immediately prior to transmission and to boost the received signal to an intelligible level immediately after its reception. 802.11 stations can maximize battery life by shutting down the radio transceiver and sleeping periodically. During sleeping periods, access points buffer any unicast frames for sleeping stations. These

frames are announced by subsequent Beacon frames. To retrieve buffered frames, newly awakened stations use PS-Poll frames.

An access point that receives a PS-Poll frame may respond with the requested data immediately, or at its leisure when circumstances permit. In some cases, the PS-Poll response is determined by the 802.11 chipset vendor in the AP. Some chipset vendors support both modes, while others support only one. 802.11 only requires that one method be supported; both approaches are equally standards-compliant.

Immediate response

Access points can respond immediately to the PS-Poll. After a short interframe space, an access point may transmit the frame. Figure 3-19 shows an implied NAV as a result of the PS-Poll frame. The PS-Poll frame contains an Association ID in the Duration/ID field so that the access point can determine which frames were buffered for the mobile station. However, the MAC specification requires all stations receiving a PS-Poll to update the NAV with an implied value equal to a short interframe space and one ACK. Although the NAV is too short for the data frame, the access point acquires that the medium and all stations defer access for the entire data frame. At the conclusion of the data frame, the NAV is updated to reflect the value in the header of the data frame.




Figure 3-19. Immediate PS-Poll response

If the buffered frame is large, it may require fragmentation. Figure 3-20 illustrates an immediate PS-Poll response requiring fragmentation. Like all other stations, access points typically have a configurable fragmentation threshold.

Deferred response

Instead of an immediate response, access points can also respond with a simple acknowledgment. This is called a *deferred response* because the access point acknowledges the request for the buffered frame but does not act on it immediately. One of the advantages of using deferred response is that the software on the access point is somewhat easier to implement because the acknowledgment can be




Figure 3-20. Immediate PS-Poll response with fragmentation

transmitted immediately by the chipset firmware, and the buffered data may be queued for transmission normally.

A station requesting a frame with a PS-Poll must stay awake until it is delivered. Under contention-based service, however, the access point can deliver a frame at any point. A station cannot return to a low-power mode until it receives a Beacon frame in which its bit in the traffic indication map (TIM) is clear.

Figure 3-21 illustrates this process. In this figure, the station has recently changed from a low-power mode to an active mode, and it notes that the access point has buffered frames for it. It transmits a PS-Poll to the access point to retrieve the buffered frames. However, the access point may choose to defer its response by transmitting only an ACK. At this point, the access point has acknowledged the station's request for buffered frames and promised to deliver them at some point in the future. The station must wait in active mode, perhaps through several atomic frame exchanges, before the access point delivers the data. A buffered frame may be subject to fragmentation, although Figure 3-21 does not illustrate this case.




Figure 3-21. Deferred PS-Poll response example

After receiving a data frame, the station must remain awake until the next Beacon is transmitted. Beacon frames only note whether frames are buffered for a station and




Figure 4-1. Generic data frame

The different data frame types can be categorized according to function. One such distinction is between data frames used for contention-based service and those used for contention-free service. Any frames that appear only in the contention-free period can never be used in an IBSS. Another possible division is between frames that carry data and frames that perform management functions. Table 4-1 shows how frames may be divided along these lines. Frames used in contention-free service are discussed in detail in Chapter 9.

Table 4-1. Categorization of data frame types

Frame type	Contention-based service	Contention-free service	Carries data	Does not carry data
Data	✓		✓	
Data+CF-Ack		✓	✓	
Data+CF-Poll		AP only	✓	
Data+CF-Ack+CF-Poll		AP only	✓	
Null	✓	✓		✓
CF-Ack		✓		✓
CF-Poll		AP only		✓
CF-Ack+CF-Poll		AP only		✓

Frame Control

All the bits in the Frame Control field are used according to the rules described in Chapter 3. Frame Control bits may affect the interpretation of other fields in the MAC header, though. Most notable are the address fields, which depend on the value of the ToDS and FromDS bits.

Duration

The Duration field carries the value of the Network Allocation Vector (NAV). Access to the medium is restricted for the time specified by the NAV. Four rules specify the setting for the Duration field in data frames:

1. Any frames transmitted during the contention-free period set the Duration field to 32,768. Naturally, this applies to any data frames transmitted during this period.
2. Frames transmitted to a broadcast or multicast destination (Address 1 has the group bit set) have a duration of 0. Such frames are not part of an atomic




Figure 4-9. Data frames from the AP

Nothing in the 802.11 specification forbids an access point from transmitting Null frames, but there is no reason to transmit them. Access points are forbidden from using the power-saving routines, and they can acknowledge Null frames from stations without using Null frames in response. In practice, access points send Data frames during the contention-based access period, and they send frames incorporating the CF-Poll feature during the contention-free period.

Frames to the AP

Figure 4-10 shows the format of a frame sent from a mobile station in an infrastructure network to the access point currently serving it. The receiver address is the BSSID. In infrastructure networks, the BSSID is taken from the MAC address of the network station in the access point. Frames destined for an access point take their source/transmitter address from the network interface in the wireless station. Access points do not perform filtering, but instead use the third address to forward data to the appropriate location in the distribution system.

Frames from the distribution system have the ToDS bit set, but the FromDS bit is 0. Mobile stations in an infrastructure network cannot become the point coordinator, and thus never send frames that incorporate the contention-free polling (CF-Poll) functions.

Frames in a WDS

When access points are deployed in a wireless bridge (or WDS) topology, all four address fields are used, as shown in Figure 4-11. Like all other data frames, WDS frames use the first address for the receiver of the frame and the second address for the transmitter. The MAC uses these two addresses for acknowledgments and control traffic, such as RTS, CTS, and ACK frames. Two more address fields are necessary to indicate the source and destination of the frame and distinguish them from the addresses used on the wireless link.




Figure 4-12. Frame Control field in control frames

Protocol version

The protocol version is shown as 0 in Figure 4-12 because that is currently the only version. Other versions may exist in the future.

Type

Control frames are assigned the Type identifier 01. By definition, all control frames use this identifier.

Subtype

This field indicates the subtype of the control frame that is being transmitted.

ToDS and FromDS bits

Control frames arbitrate access to the wireless medium and thus can only originate from wireless stations. The distribution system does not send or receive control frames, so these bits are always 0.

More Fragments bit

Control frames are not fragmented, so this bit is always 0.

Retry bit

Control frames are not queued for retransmission like management or data frames, so this bit is always 0.

Power Management bit

This bit is set to indicate the power management state of the sender after conclusion of the current frame exchange.

More Data bit

The More Data bit is used only in management and data frames, so this bit is set to 0 in control frames.

Protected Frame bit

Control frames may not be encrypted. Thus, for control frames, the Protected Frame bit is always 0.

Order bit

Control frames are used as components of atomic frame exchange operations and thus cannot be transmitted out of order. Therefore, this bit is set to 0.

Request to Send (RTS)

RTS frames are used to gain control of the medium for the transmission of "large" frames, in which "large" is defined by the RTS threshold in the network card driver.

The diagram illustrates the IEEE 802.11 MAC frame structure, showing the distribution of bytes and bits across various fields.

Bytes:

- 2 bytes: Frame Control
- 2 bytes: Duration
- 6 bytes: Receiver Address
- 6 bytes: Transmitter Address
- 4 bytes: FCS


Bits:

- 2 bits: Protocol
- 2 bits: Type = control
- 4 bits: Sub type = RTS
- 1 bit: To DS
- 1 bit: From DS
- 1 bit: More Frag
- 1 bit: Retry
- 1 bit: Pwr Mgmt
- 1 bit: More Data
- 1 bit: Protected Frame
- 1 bit: Order

Four fields make up the MAC header of an RTS:

There is nothing special about the Frame Control field. The frame subtype is set to 1011 to indicate an RTS frame, but otherwise, it has all the same fields as other control frames. (The most significant bits in the 802.11 specification come at the end of fields, so bit 7 is the most significant bit in the subtype field.)

An RTS frame attempts to reserve the medium for an entire frame exchange, so the sender of an RTS frame calculates the time needed for the frame exchange sequence after the RTS frame ends. The entire exchange, which is depicted in Figure 4-14, requires three SIFS periods, the duration of one CTS, the final ACK, plus the time needed to transmit the frame or first fragment. (Fragmentation bursts use subsequent fragments to update the Duration field.) The number of microseconds required for the transmission is calculated and placed in the Duration field. If the result is fractional, it is rounded up to the next microsecond.



78 | Chapter 4: 802.11 Framing in Detail

Address 1: Receiver Address

The address of the station that is the intended recipient of the large frame.

Address 2: Transmitter Address

The address of the sender of the RTS frame.

Clear to Send (CTS)

The CTS frame, whose format is shown in Figure 4-15, has two purposes. Initially, CTS frames were used only to answer RTS frames, and were never generated without a preceding RTS. CTS frames were later adopted for use by the 802.11g protection mechanism to avoid interfering with older stations. The protection mechanism is described with the rest of 802.11g in Chapter 14.




Figure 4-15. CTS frame

Three fields make up the MAC header of a CTS frame:

Frame Control

The frame subtype is set to 1100 to indicate a CTS frame.

Duration

When used in response to an RTS, the sender of a CTS frame uses the duration from the RTS frame as the basis for its duration calculation. RTS frames reserve the medium for the entire RTS-CTS-frame-ACK exchange. By the time the CTS frame is transmitted, though, only the pending frame or fragment and its acknowledgment remain. The sender of a CTS frame subtracts the time required for the CTS frame and the short interframe space that preceded the CTS from the duration in the RTS frame, and places the result of that calculation in the Duration field. Figure 4-16 illustrates the relationship between the CTS duration and the RTS duration. Rules for CTS frames used in protection exchanges are described with the protection mechanism.

Address 1: Receiver Address

The receiver of a CTS frame is the transmitter of the previous RTS frame, so the MAC copies the transmitter address of the RTS frame into the receiver address of the CTS frame. CTS frames used in 802.11g protection are sent to the sending station, and are used only to set the NAV.




Figure 4-16. CTS duration

Acknowledgment (ACK)

ACK frames are used to send the positive acknowledgments required by the MAC and are used with any data transmission, including plain transmissions, frames preceded by an RTS/CTS handshake, and fragmented frames (see Figure 4-17). Quality-of-service enhancements relax the requirement for a single acknowledgment per Data frame. To assess the impact of acknowledgments on net throughput, see Chapter 25.




Figure 4-17. ACK frame

Three fields make up the MAC header of an ACK frame:

Frame Control

The frame subtype is set to 1101 to indicate an ACK frame.

Duration

The duration may be set in one of two ways, depending on the position of the ACK within the frame exchange. ACKs for complete data frames and final fragments in a fragment burst set the duration to 0. The data sender indicates the end of a data transmission by setting the More Fragments bit in the Frame Control header to 0. If the More Fragments bit is 0, the transmission is complete, and there is no need to extend control over the radio channel for additional transmissions. Thus, the duration is set to 0.

If the More Fragments bit is 1, a fragment burst is in progress. The Duration field is used like the Duration field in the CTS frame. The time required to transmit the ACK and its short interframe space is subtracted from the duration in the most recent fragment (Figure 4-18). The duration calculation in nonfinal ACK frames is similar to the CTS duration calculation. In fact, the 802.11 specification refers to the duration setting in the ACK frames as a *virtual CTS*.




Figure 4-18. Duration in non-final ACK frames

Address 1: Receiver Address

The receiver address is copied from the transmitter of the frame being acknowledged. Technically, it is copied from the Address 2 field of the frame being acknowledged. Acknowledgments are transmitted in response to directed data frames, management frames, and PS-Poll frames.

Power-Save Poll (PS-Poll)

When a mobile station wakes from a power-saving mode, it transmits a PS-Poll frame to the access point to retrieve any frames buffered while it was in power-saving mode. The format of the PS-Poll frame is shown in Figure 4-19. Further details on the operation of power saving modes appears in Chapter 8.




Figure 4-19. PS-Poll frame

Four fields make up the MAC header of a PS-Poll frame:

Frame Control

The frame subtype is set to 1010 to indicate a PS-Poll frame.

Association ID (AID)

Instead of a Duration field, the PS-Poll frame uses the third and fourth bytes in the MAC header for the association ID. This is a numeric value assigned by the access point to identify the association. Including this ID in the frame allows the access point to find any frames buffered for the now-awakened mobile station.

Address 1: BSSID

This field contains the BSSID of the BSS created by the access point that the sender is currently associated with.

Address 2: Transmitter Address

This is the address of the sender of the PS-Poll frame.

The PS-Poll frame does not include duration information to update the NAV. However, all stations receiving a PS-Poll frame update the NAV by the short interframe space plus the amount of time required to transmit an ACK. The automatic NAV update allows the access point to transmit an ACK with a small probability of collision with a mobile station.

Association ID (AID)

In the PS-Poll frame, the Duration/ID field is an association ID rather than a value used by the virtual carrier-sensing function. When mobile stations associate with an access point, the access point assigns a value called the Association ID (AID) from the range 1–2,007. The AID is used for a variety of purposes that appear throughout this book.

Management Frames

Management is a large component of the 802.11 specification. Several different types of management frames are used to provide services that are simple on a wired network. Establishing the identity of a network station is easy on a wired network because network connections require dragging wires from a central location to the new workstation. In many cases, patch panels in the wiring closet are used to speed up installation, but the essential point remains: new network connections can be authenticated by a personal visit when the new connection is brought up.

Wireless networks must create management features to provide similar functionality. 802.11 breaks the procedure up into three components. Mobile stations in search of connectivity must first locate a compatible wireless network to use for access. With wired networks, this step typically involves finding the appropriate data jack

on the wall. Next, the network must authenticate mobile stations to establish that the authenticated identity is allowed to connect to the network. The wired-network equivalent is provided by the network itself. If signals cannot leave the wire, obtaining physical access is at least something of an authentication process. Finally, mobile stations must associate with an access point to gain access to the wired backbone, a step equivalent to plugging the cable into a wired network.

The Structure of Management Frames

802.11 management frames share the structure shown in Figure 4-20. The MAC header is the same in all management frames; it does not depend on the frame sub-type. Management frames use *information elements*, little chunks of data with a numerical label, to communicate information to other systems.




Figure 4-20. Generic management frame

Address fields

As with all other frames, the first address field is used for the frame's destination address. Some management frames are used to maintain properties within a single BSS. To limit the effect of broadcast and multicast management frames, stations are required to inspect the BSSID after receiving a management frame, though not all implementations perform BSSID filtering. Only broadcast and multicast frames from the BSSID that a station is currently associated with are passed to MAC management layers. The one exception to this rule is Beacon frames, which are used to announce the existence of an 802.11 network.

BSSIDs are assigned in the familiar manner. Access points use the MAC address of the wireless network interface as the BSSID. Mobile stations adopt the BSSID of the access point they are currently associated with. Stations in an IBSS use the randomly generated BSSID from the BSS creation. One exception to the rule: frames sent by the mobile station seeking a specific network may use the BSSID of the network they are seeking, or they may use the broadcast BSSID to find all networks in the vicinity.

FH PHY Convergence Procedure (PLCP)

Before any frames can be modulated onto the RF carrier, the frames from the MAC must be prepared by the Physical Layer Convergence Procedure (PLCP). Different underlying physical layers may have different requirements, so 802.11 allows each physical layer some latitude in preparing MAC frames for transmission over the air.

Framing and Whitening

The PLCP for the FH PHY adds a five-field header to the frame it receives from the MAC. The PLCP is a relay between the MAC and the physical medium dependent (PMD) radio interface. In keeping with ISO reference model terminology, frames passed from the MAC are PLCP service data units (PSDUs). The PLCP framing is shown in Figure 11-9.

Preamble

As in a wired Ethernet, the preamble synchronizes the transmitter and receiver and derives common timing relationships. In the 802.11 FH PHY, the Preamble is composed of the Sync field and the Start Frame Delimiter field.




Figure 11-9. PLCP framing in the FH PHY

Sync

The sync field is 80 bits in length and is composed of an alternating zero-one sequence (010101...01). Stations search for the sync pattern to prepare to receive data. In addition to synchronizing the sender and receiver, the Sync field serves three purposes. First of all, the presence of a sync signal indicates that a frame is imminent. Second, stations that have multiple antennas to combat multipath fading or other environmental reception problems can select the antenna with the strongest signal. Finally, the receiver can measure the frequency of the incoming signal relative to its nominal values and perform any corrections needed to the received signal.

Start Frame Delimiter (SFD)

As in Ethernet, the SFD signals the end of the preamble and marks the beginning of the frame. The FH PHY uses a 16-bit SFD: 0000 1100 1011 1101.

Header

The PLCP header follows the preamble. The header has PHY-specific parameters used by the PLCP. Three fields comprise the header: a length field, a speed field, and a frame check sequence.

PSDU Length Word (PLW)

The first field in the PLCP header is the PLW. The payload of the PLCP frame is a MAC frame that may be up to 4,095 bytes long. The 12-bit length field informs the receiver of the length of the MAC frame that follows the PLCP header.

PLCP Signaling (PSF)

Bit 0, the first bit transmitted, is reserved and set to 0. Bits 1–3 encode the speed at which the payload MAC frame is transmitted. Several speeds are available, so this field allows the receiver to adjust to the appropriate demodulation scheme. Although the standard allows for data rates in increments of 500 kbps from 1.0 Mbps to 4.5 Mbps, the modulation scheme has been defined only for 1.0 Mbps and 2.0 Mbps.* See Table 11-3.

Table 11-3. PSF meaning

Bits (1-2-3)	Data rate
000	1.0 Mbps
001	1.5 Mbps
010	2.0 Mbps
011	2.5 Mbps
100	3.0 Mbps
101	3.5 Mbps
110	4.0 Mbps
111	4.5 Mbps

Header Error Check (HEC)

To protect against errors in the PLCP header, a 16-bit CRC is calculated over the contents of the header and placed in this field. The header does not protect against errors in other parts of the frame.

No restrictions are placed on the content of the Data field. Arbitrary data may contain long strings of consecutive 0s or 1s, which makes the data much less random. To make the transmitted data more like random white noise, the FH PHYs apply a *whitening* algorithm to the MAC frame. This algorithm scrambles the data before radio transmission. Receivers invert the process to recover the data.

Frequency-Hopping PMD Sublayer

Although the PLCP header has a field for the speed at which the MAC frame is transmitted, only two of these rates have corresponding standardized PMD layers. Several

* It is unlikely that significant further work will be done on high-rate, frequency-hopping systems. For high data rates, direct sequence is a more cost-effective choice.

features are shared between both PMDs: antenna diversity support, allowances for the ramp up and ramp down of the power amplifiers in the antennas, and the use of a Gaussian pulse shaper to keep as much RF power as possible in the narrow frequency-hopping band. Figure 11-10 shows the general design of the transceiver used in 802.11 frequency-hopping networks. It is required that the transceiver have a sensitivity of -80 dBm for both 1 Mbps and 2 Mbps transmission.




Figure 11-10. Frequency-hopping transceiver

PMD for 1.0-Mbps FH PHY

The basic frequency-hopping PMD enables data transmission at 1.0 Mbps. Frames from the MAC have the PLCP header appended, and the resulting sequence of bits is transmitted out of the radio interface. In keeping with the common regulatory restriction of a 1-MHz bandwidth, 1 million symbols are transmitted per second. 2GFSK is used as the modulation scheme, so each symbol can be used to encode a single bit. 802.11 specifies a minimum power of 10 milliwatts (mW) and requires the use of a power control function to cap the radiated power at 100 mW, if necessary.

PMD for 2.0-Mbps FH PHY

A second, higher-speed PMD is available for the FH PHY. As with the 1.0-Mbps PMD, the PLCP header is appended and is transmitted at 1.0 Mbps using 2GFSK. In the PLCP header, the PSF field indicates the speed at which the frame body is transmitted. At the higher data rate, the frame body is transmitted using a different encoding method than the physical-layer header. Regulatory requirements restrict all PMDs to a symbol rate of 1 MHz, so 4GFSK must be used for the frame body. Two bits per symbol yields a rate of 2.0 Mbps at 1 million symbols per second. Firmware that supports the 2.0-Mbps PMD can fall back to the 1.0-Mbps PMD if signal quality is too poor to sustain the higher rate.

Carrier sense/clear channel assessment (CS/CCA)

To implement the CSMA/CA foundation of 802.11, the PCLP includes a function to determine whether the wireless medium is currently in use. The MAC uses both a virtual carrier-sense mechanism and a physical carrier-sense mechanism; the physical layer implements the physical carrier sense. 802.11 does not specify how to determine

whether a signal is present; vendors are free to innovate within the required performance constraints of the standard. 802.11 requires that 802.11-compliant signals with certain power levels must be detected with a corresponding minimum probability.

Characteristics of the FH PHY

Table 11-4 shows the values of a number of parameters in the FH PHY. In addition to the parameters in the table, which are standardized, the FH PHY has a number of parameters that can be adjusted to balance delays through various parts of an 802.11 frequency-hopping system. It includes variables for the latency through the MAC, the PLCP, and the transceiver, as well as variables to account for variations in the transceiver electronics. One other item of note is that the total aggregate throughput of all frequency-hopping networks in an area can be quite high. The total aggregate throughput is a function of the hop set size. All sequences in a hop set are orthogonal and noninterfering. In North America and most of Europe, 26 frequency-hopping networks can be deployed in an area at once. If each network is run at the optional 2 Mbps rate and half the airtime is able to carry user payload data, the area can have a total of 26 Mbps throughput provided that the ISM band is relatively free of interference.

Table 11-4. FH PHY parameters

Parameter	Value	Notes
Slot time	50 μ s	
SIFS time	28 μ s	The SIFS is used to derive the value of the other interframe spaces (DIFS, PIFS, and EIFS).
Contention window size	15–1,023 slots	
Preamble duration	96 μ s	Preamble symbols are transmitted at 1 MHz, so a symbol takes 1 μ s to transmit; 96 bits require 96 symbol times.
PLCP header duration	32 μ s	The PLCP header is 32 bits, so it requires 32 symbol times.
Maximum MAC frame	4,095 bytes	802.11 recommends a maximum of 400 symbols (400 bytes at 1 Mbps, 800 bytes at 2 Mbps) to retain performance across different types of environments.
Minimum sensitivity	–80 dBm	

Now encode M in DQPSK (Figure 12-13). In the UTF-8 character set, M is represented by the binary string 01001101 or, as the sequence of four two-bit symbols, 01-00-11-01. In the first symbol period, there is a phase shift of 90 degrees; for clarity, the figure shows the phase shift from a pure sine wave. The second symbol results in no phase shift, so the wave continues without a change. The third symbol causes a phase shift of 180 degrees, as shown by the sharp change from the highest amplitude to the lowest amplitude. The final symbol causes a phase shift of 90 degrees.




Figure 12-13. The letter M encoded in DQPSK

The obvious advantage of DQPSK relative to DBPSK is that the four-level encoding mechanism can have a higher throughput. The cost of using DQPSK is that it cannot be used in some environments because of severe multipath interference. Multipath interference occurs when the signal takes several paths from the transmitter to the receiver. Each path has a different length; therefore, the received signal from each path has a different delay relative to the other paths. This delay is the enemy of an encoding scheme based on phase shifts. Wavefronts are not labeled or painted different colors, so a wavefront could arrive later than expected because of a long path or it could simply have been transmitted late and phase-shifted. In environments where multipath interference is severe, DQPSK will break down much quicker than DBPSK.

The “Original” Direct Sequence PHY

The physical layer itself consists of two components. The Physical Layer Convergence Procedure (PLCP) performs some additional PHY-dependent framing before transmission, while the Physical Medium Dependent (PMD) layer is responsible for the actual transmission of frames.

PLCP Framing and Processing

The PLCP for the DS PHY adds a six-field header to the frames it receives from the MAC. In keeping with ISO reference model terminology, frames passed from the MAC are PLCP service data units (PSDUs). The PLCP framing is shown in Figure 12-14.




Figure 12-14. DS PLCP framing

The FH PHY uses a data whitener to randomize the data before transmission, but the data whitener applies only to the MAC frame trailing the PLCP header. The DS PHY has a similar function called the *scrambler*, but the scrambler is applied to the entirety of the direct-sequence frame, including the PLCP header and preamble.

Preamble

The Preamble synchronizes the transmitter and receiver and allows them to derive common timing relationships. It is composed of the Sync field and the Start Frame Delimiter field. Before transmission, the preamble is scrambled using the direct-sequence scrambling function.

Sync

The Sync field is a 128-bit field composed entirely of 1s. Unlike the FH PHY, the Sync field is scrambled before transmission.

Start Frame Delimiter (SFD)

The SFD allows the receiver to find the start of the frame, even if some of the sync bits were lost in transit. This field is set to 0000 0101 1100 1111, which is different from the SFD used by the FH PHY.

Header

The PLCP header follows the preamble. The header has PHY-specific parameters used by the PLCP. Five fields comprise the header: a signaling field, a service identification field, a Length field, a Signal field used to encode the speed, and a frame-check sequence.

Signal

The Signal field is used by the receiver to identify the transmission rate of the encapsulated MAC frame. It is set to either 0000 1010 (0x0A) for 1-Mbps operation or 0001 0100 (0x14) for 2-Mbps operation.

Service

This field is reserved for future use and must be set to all 0s.

Length

This field is set to the number of microseconds required to transmit the frame as an unsigned 16-bit integer, transmitted least-significant bit to most-significant bit.

CRC

To protect the header against corruption on the radio link, the sender calculates a 16-bit CRC over the contents of the four header fields. Receivers verify the CRC before further frame processing.

No restrictions are placed on the content of the Data field. Arbitrary data may contain long strings of consecutive 0s or 1s, which makes the data much less random. To make the data more like random background noise, the DS PHY uses a polynomial scrambling mechanism to remove long strings of 1s or 0s from the transmitted data stream.

DS Physical Medium Dependent Sublayer

The PMD is a complex and lengthy specification that incorporates provisions for two data rates (1.0 and 2.0 Mbps). Figure 12-15 shows the general design of a transceiver for 802.11 direct-sequence networks.




Figure 12-15. Direct-sequence transceiver

Transmission at 1.0 Mbps

At the low data rate, the direct-sequence PMD enables data transmission at 1.0 Mbps. The PLCP header is appended to frames arriving from the MAC, and the entire unit is scrambled. The resulting sequence of bits is transmitted from the physical interface using DBPSK at a rate of 1 million symbols per second. The resulting throughput is 1.0 Mbps because one bit is encoded per symbol. Like the FH PMD, the DS PMD has a minimum power requirement and can cap the power at 100 mW if necessary to meet regulatory requirements.

Transmission at 2.0 Mbps

Like the FH PHY, transmission at 2.0 Mbps uses two encoding schemes. The PLCP preamble and header are transmitted at 1.0 Mbps using DBPSK. Although using a slower method for the header transmission reduces the effective throughput, DBPSK is far more tolerant of noise and multipath interference. After the preamble and header are finished, the PMD switches to DQPSK modulation to provide 2.0-Mbps service. As with the FH PHY, most products that implement the 2.0-Mbps rate can detect interference and fall back to lower-speed 1.0-Mbps service.

CS/CCA for the DS PHY

802.11 allows the carrier sense/clear channel assessment function to operate in one of three modes:

Mode 1

When the energy exceeds the energy detection (ED) threshold, it reports that the medium is busy. The ED threshold depends on the transmit power.

Mode 2

Implementations using Mode 2 must look for an actual DSSS signal and report the channel busy when one is detected, even if the signal is below the ED threshold.

Mode 3

Mode 3 combines Mode 1 and Mode 2. A signal must be detected with sufficient energy before the channel is reported busy to higher layers.

Once a channel is reported busy, it stays busy for the duration of the intended transmission, even if the signal is lost. The transmission's duration is taken from the time interval in the Length field. Busy medium reports must be very fast. When a signal is detected at the beginning of a contention window slot, the CCA mechanism must report a busy medium by the time the slot has ended. This relatively high performance requirement must be set because once a station has begun transmission at the end of its contention delay, it should seize the medium, and all other stations should defer access until its frame has concluded.

Characteristics of the DS PHY

Table 12-4 shows the values of a number of parameters in the DS PHY. In addition to the parameters in the table, which are standardized, the DS PHY has a number of parameters that can be adjusted to balance delays through various parts of an 802.11 direct-sequence system. It includes variables for the latency through the MAC, the PLCP, and the transceiver, as well as variables to account for variations in the transceiver electronics. One other item of note is that the total aggregate throughput of all direct-sequence networks in an area is much lower than the total aggregate throughput of all nonoverlapping frequency-hopping networks in an area. The total aggregate throughput is a function of the number of nonoverlapping channels. In North America and most of Europe, three direct-sequence networks can be deployed in an area at once. If each network is run at the optional 2 Mbps rate and the efficiency of the protocol allows 50% of the headline rate to become user data throughput, the total throughput is 3 Mbps, which is dramatically less than the frequency-hopping total aggregate throughput.

Table 12-4. DS PHY parameters

Parameter	Value	Notes
Slot time	20 μ s	
SIFS time	10 μ s	The SIFS is used to derive the value of the other interframe spaces (DIFS, PIFS, and EIFS).
Contention window size	31 to 1,023 slots	

Table 12-4. DS PHY parameters (continued)

Parameter	Value	Notes
Preamble duration	144 μ s	Preamble symbols are transmitted at 1 MHz, so a symbol takes 1 μ s to transmit; 144 bits require 144 symbol times.
PLCP header duration	48 μ s	The PLCP header is 48 bits, so it requires 48 symbol times.
Maximum MAC frame	4–8,191 bytes	
Minimum receiver sensitivity	–80 dBm	
Adjacent channel rejection	35 dB	See text for measurement details.

Like the FH PHY, the DS PHY has a number of attributes that can be adjusted by a vendor to balance delays in various parts of the system. It includes variables for the latency through the MAC, the PLCP, and the transceiver, as well as variables to account for variations in the transceiver electronics.

Complementary Code Keying

802.11 direct-sequence systems use a rate of 11 million chips per second. The original DS PHYs divided the chip stream up into a series of 11-bit Barker words and transmitted 1 million Barker words per second. Each word encoded either one bit or two bits for a corresponding data rate of 1.0 Mbps or 2.0 Mbps, respectively. Achieving higher data rates and commercial utility requires that each code symbol carry more information than a bit or two.

Straight phase shift encoding cannot hope to carry more than a few bits per code word. DQPSK requires that receivers distinguish quarter-cycle phase differences. Further increasing the number of bits per symbol would require processing even finer phase shifts, such as an eighth-cycle or sixteenth-cycle shift. Detecting smaller phase shifts is more difficult in the presence of multipath interference and requires more sophisticated (and thus expensive) electronics.

Instead of continuing with straight phase-shift keying, the IEEE 802.11 working group turned to an alternate encoding method. Complementary code keying (CCK) divides the chip stream into a series of 8-bit code symbols, so the underlying transmission is based on a series of 1.375 million code symbols per second. CCK is based on sophisticated mathematical transforms that allow the use of a few 8-bit sequences to encode 4 or even 8 bits per code word, for a data throughput of 5.5 Mbps or 11 Mbps. In addition, the mathematics underlying CCK transforms allow receivers to distinguish between different codes easily, even in the presence of interference and multipath fading. Figure 12-16 illustrates the use of code symbols in CCK. It is quite similar to the chipping process used by the slower direct-sequence layers; the difference is that the code words are derived partially from the data. A static repeating code word such as the Barker word is not used.




Figure 12-16. Code symbols in CCK

Barker spreading, as used in the lower-rate, direct-sequence layers, uses a static code to spread the signal over the available frequency band. CCK uses the code word to carry information, as well as simply to spread the signal. Several phase angles are used to prepare a complex code word of eight bits.

High Rate Direct Sequence PHY

To distinguish it from the original direct sequence PHY, the high rate PHY that runs at 11 Mbps is abbreviated as HR/DSSS. Like its predecessor, it is split into a convergence procedure that prepares frames for radio transmission, and a medium-dependent layer that turns the bits into radio waves in the air.

PLCP Framing and Scrambling

The long headers required by the original PHY greatly reduce performance. The 802.11 MAC requires an acknowledgment for every data frame, and the 192 microsecond preamble is much, much longer than the MAC acknowledgment. At the 11 Mbps data rate, the preamble and PLCP framing header sucks up 25% of the time used to transmit a 1,500 byte frame and its corresponding MAC acknowledgment. As long as a new PHY was being developed, the designers of 802.11b came up with a new “short” framing format that improves protocol efficiency and, hence, throughput. Using the short headers cuts the preamble and PLCP framing overhead cuts the preamble and framing overhead to 14%. While still substantial, it is a dramatic improvement. Figure 12-17 shows the PLCP framing specified in 802.11b. When 802.11b was first released, short headers were not supported by all devices because of the large installed base of 2 Mbps direct sequence equipment. At this point, almost every card can safely support the short preamble, and most access point vendors use it by default.

Naturally, the optional short format may be used only if all stations support it. To prevent networks configured for the short format from disappearing, 802.11b requires that stations answering Probe Requests from an active scan return a response using the same PLCP header that was received. If a station that supports




Figure 12-17. HR/DSSS PLCP framing

only the long PLCP header sends a Probe Response, an access point returns a response using the long header, even if the BSS is configured for the short header.

Preamble

Frames begin with the preamble, which is composed of the Sync field and the SFD field. The preamble is transmitted at 1.0 Mbps using DBPSK.

Long Sync

The Long Sync field is composed of 128 one bits. It is processed by the scrambler before transmission, though, so the data content varies. High-rate systems use a specified seed for the scrambling function but support backwards compatibility with older systems that do not specify a seed.

Short Sync

The Short Sync field is composed of 56 zero bits. Like the Long Sync, it is also processed by the scrambler.

Long SFD

To indicate the end of the Sync field, the long preamble concludes with a Start of Frame Delimiter (SFD). In the long PLCP, the SFD is the sequence 1111 0011 1010 0000. As with all IEEE specifications, the order of transmission from the physical interface is least-significant bit first, so the string is transmitted right to left.

Short SFD

To avoid confusion with the Long SFD, the Short SFD is the reverse value, 0000 0101 1100 1111.

The PLCP header follows the preamble. It is composed of the Signal, Service, Length, and CRC fields. The long header is transmitted at 1.0 Mbps using DBPSK. However, the short header's purpose is to reduce the time required for overhead transmission so it is transmitted at 2.0 Mbps using DQPSK.

Long Signal

The Long Signal field indicates the speed and transmission method of the enclosed MAC frame. Four values for the 8-bit code are currently defined and are shown in Table 12-5.

Table 12-5. Signal field values

Speed	Value (msb to lsb)	Hex value
1 Mbps	0000 1010	0x0A
2 Mbps	0001 0100	0x14
5.5 Mbps	0011 0111	0x37
11 Mbps	0110 1110	0x6E

Short Signal

The Short Signal field indicates the speed and transmission method of the enclosed frame, but only three values are defined. Short preambles can be used only with 2 Mbps, 5.5 Mbps, and 11 Mbps networks.

Service

The Service field, which is shown in Figure 12-18, was reserved for future use by the first version of 802.11, and bits were promptly used for the high-rate extensions in 802.11b. First of all, the Length field describes the amount of time used for the enclosed frame in microseconds. Above 8 Mbps, the value becomes ambiguous. Therefore, the eighth bit of the service field is used to extend the Length field to 17 bits. The third bit indicates whether the 802.11b implementation uses locked clocks; clock locking means that transmit frequency and symbol clock use the same oscillator. The fourth bit indicates the type of coding used for the packet, which is either 0 for CCK or 1 for PBCC. All reserved bits must be set to 0. The Service field is transmitted from left to right (b0 to b7), which is the same in both the short and long PLCP frame formats. (Further changes are made by 802.11g, which will be discussed in Chapter 14.)




Figure 12-18. Service field in the HR/DSSS PLCP header

Length

The Length field is the same in both the short and long PLCP frame formats and is the number of microseconds required to transmit the enclosed MAC frame. Approximately two pages of the 802.11b standard are devoted to calculating the value of the Length field, but the details are beyond the scope of this book.

CRC

The CRC field is the same in both the short and the long PLCP frames. Senders calculate a CRC checksum using the Signal, Service, and Length fields. Receivers can use the CRC value to ensure that the header was received intact and was not damaged during transmission. CRC calculations take place before data scrambling.

The data scrambling procedure for the HR/DSSS PHY is nearly identical to the data scrambling procedure used with the original DS PHY. The only difference is that the scrambling function is seeded to specified values in the HR/DSSS PHY. Different seeds are used for short and long PLCP frames.

HR/DSSS PMD

Like the DS PHY, the 802.11b PHY uses a single PMD specification. The general transceiver design is shown in Figure 12-19.




Figure 12-19. HR/DSSS transceiver

Transmission at 1.0 Mbps or 2.0 Mbps

To ensure backwards compatibility with the installed base of 802.11-based, direct-sequence hardware, the HR/DSSS PHY can transmit and receive at 1.0 Mbps or 2.0

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, inter-frame 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 chip-makers. 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:

$$\text{Data rate (Mbps)} = 0.0675 \times \text{channel bandwidth} \times \text{number of spatial streams} \times \text{coded bits per subcarrier} \times \text{code rate}$$

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 MHz channels	40 MHz channels
Greenfield	2TX20GF	1TX40GF
	3TX20GF	2TX40GF
	4TX20GF	3TX40GF
		4TX40GF
Mixed mode	2TX20MM	1TX40MM
	3TX20MM	2TX40MM
	4TX20MM	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 be 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 μ s	
SIFS time	16 μ s	The SIFS is used to derive the value of the other interframe spaces (DIFS, PIFS, and EIFS).
RIFS time	2 μ s	
Contention window size	15 to 1,023 slots	
Preamble duration	16 μ s	
PLCP header duration	4 μ s	
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.

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:

$$\text{Data rate (Mbps)} = 12 \times \text{channel bandwidth factor} \times \text{number of spatial streams} \\ \times \text{coded bits per subcarrier} \times \text{code rate} \times \text{guard interval factor}$$

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 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 transmitter

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 queuing, 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.

Comparison and Conclusion

Michael

Michael

The message integrity check algorithm specified as part of TKIP.

MIMO

Multiple-Input/Multiple-Output. An antenna configuration that uses more than one transmission antenna and more than one receiver antenna to transmit multiple data streams. 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, respectively.

MPDU

MAC Protocol Data Unit. A fancy name for frame. The MPDU does not, however, include PLCP headers.

MSDU

MAC Service Data Unit. The data accepted by the MAC for delivery to another MAC on the network. MSDUs are composed of higher-level data only. For example, an 802.11 management frame does not contain an MSDU.

NAV

Network Allocation Vector. The NAV is used to implement the virtual carrier sensing function. Stations will defer access to the medium if it is busy. For robustness, 802.11 includes two carrier-sensing functions. One is a *physical* function, which is based on energy thresholds, whether a station is decoding a legal 802.11 signal, and similar things that require a physical measurement. The second function is a *virtual* carrier sense, which is based on the NAV. Most frames include a nonzero number in the NAV field, which is used to ask all stations to politely defer from accessing the medium for a certain number of microseconds after the current frame is transmitted. Any receiving stations will process the NAV and defer access, which prevents collisions. For more detail on how the NAV is used, see “Contention-Based Data Service” in Chapter 3.

NIST

National Institute of Standards and Technology. The United States government agency responsible for setting technology standards for the federal government. NIST standards are used by most non-classified agencies and have been adopted by many other organizations throughout all types of government.

OFDM

Orthogonal Frequency Division Multiplexing. A technique that splits a wide frequency band into a number of narrow frequency bands and inverse multiplexes data across the subchannels. 802.11a and 802.11g are based on OFDM. 802.11n uses MIMO to transmit multiple OFDM data streams.

OSI

Open Systems Interconnection. A baroque compendium of networking standards that was never implemented because IP networks actually existed.

PBCC

Packet Binary Convolution Coding. An alternative method of encoding data in 802.11b networks that has not been widely implemented. PBCC was also proposed for consideration for 20+ Mbps networks, but was rejected.

PC

Point Coordinator. A function in the access point responsible for central coordination of access to the radio medium during contention-free service.

PCF

Point Coordination Function. The set of rules that provides for centrally coordinated access to the medium by the access point.

PCMCIA

Personal Computer Memory Card International Association. An industry group that standardized the ubiquitous “PCMCIA card” form factor and made it possible to connect a wide variety of peripherals to notebook computers. 802.11 interfaces are available almost exclusively in the PCMCIA form factor.

Also expanded humorously as People Who Can't Manage Computer Industry Acronyms because of its unwieldy length and pronunciation.

PDU

See protocol data unit.

PER

Packet Error Rate. Like the bit error rate, but measured as a fraction of packets with errors.

PHY

Common IEEE abbreviation for the physical layer.

physical-layer convergence procedure

The upper component of the PHY in 802.11 networks. Each PHY has its own PLCP, which provides auxiliary framing to the MAC.

PIFS

PCF Inter-Frame space. During contention-free service, any station is free to transmit if the medium is idle for the duration of one PCF inter-frame space.

PLCP

See physical-layer convergence procedure.

PMD

Physical Medium Dependent. The lower component of the PHY, responsible for transmitting RF signals to other 802.11 stations.

PMK

Pairwise Master Key. The root of all keying data between a supplicant and an authenticator. It may be derived from an EAP method during authentication, or supplied directly as a pre-shared key.

PPDU

PLCP Protocol Data Unit. The complete PLCP frame, including PLCP headers, MAC headers, the MAC data field, and the MAC and PLCP trailers.

PRF

Pseudo-Random Function. An 802.11i function that expands a small amount of data into a larger amount of data. The PRF is used most notably to expand master keys into key hierarchies.

protocol data unit

Layers communicate with each other using protocol data units. For example, the IP protocol data unit is the familiar IP packet. IP implementations communicate with each other using IP packets. See also service data unit.

PS

Power Save. Used as a generic prefix for power-saving operations in 802.11.

PSDU

PLCP Service Data Unit. The data the PLCP is responsible for delivering, i.e., one MAC frame with headers.

PSK - Phase Shift Keying

A method of transmitting data based on phase shifts in the transmitted carrier wave.

PSK - Pre-shared Key

In 802.11i, refers to the authentication method that depends on a statically configured authentication key that must be distributed manually. Also called WPA-PSK.

PTK

Pairwise Transient Key. Key derived from pairwise master key that includes keys used by encryption and integrity protocols, but also includes keys to distribute dynamic keys.

QPSK

Quadrature Phase Shift Keying. A modulation method that encodes bits as phase shifts. One of four phase shifts can be selected to encode two bits.

RA

Receiver Address. MAC address of the station that will receive the frame. The RA may also be the destination address of a frame, but not always. In infrastructure networks, for example, a frame destined for the distribution system is received by an access point.

RADIUS

Remote Authenticated Dial-In User Service. A protocol used to authenticate dial-in users that has become more widely

- physical medium, dynamics, 6
- PIFS (PCF Inter-Frame space), 40, 603
- planning (see project planning)
- PLCP (Physical Layer Convergence Procedure), 14
 - framing, 249
 - HR/DSSS, 268
 - OFDM PLCP, 290
 - processing and, 263
 - sublayer, 225
 - whitening, 249
 - WWiSE PLCP, 321
- PLME (physical-layer management entity), 170
- PLW (PSDU Length Word), PLCP framing and, 250
- PMD (Physical Medium Dependent), 14, 249
 - definition, 603
 - ERP PMD, 307
 - HR/DSSS, 271
 - WWiSE PMD, 324
- PMD (Physical Medium Dependent) sublayer, 225, 265
 - frequency hopping and, 250
- PMK (Pairwise Master Key), 164, 603
 - caching, 168
- PN codes (pseudorandom noise code), 254
- PoE (power over Ethernet), 422
 - types, 423
- points coordinators, 37
- polling list, PCF, 216
- port authentication, 141
- Power Capability information element, 99
- power conservation, 188
- Power Constraint information element, 99
- power management
 - IBSS, 193
 - infrastructure networks, 188
 - PCF and, 223
 - performance tuning, 581
 - TGnSync, 332
- power management bit, control frames, 77
- power savers, 59
- PPDU (PLCP Protocol Data Unit), 603
- preamble, OFDM PLCP, 291
- Preamble, PLCP framing and, 249
- preauthentication, 180
 - Beacon frames, 181
 - dynamic keys, 182
 - key caching and, 182
- PRF (Pseudo-Random Function), 603
- Prism-based cards, network analysis, 557
- Privacy bit, management frames, 86
- PRNG (pseudorandom number generator), 115
- Probe Request management frames, 107, 108
- ProbeDelay parameter, scanning, 172
- processing, PLCP framing and, 263
- project planning, 519
 - access point placement, 537–545
 - channel layout, 543
 - final report, 545
 - preliminary plan, 540
 - radio resource management, 543
 - coverage, antennas, 546–552
 - network requirements, 521
 - coverage, 522–525
 - physical layer, 532–537
 - requirements
 - aesthetic considerations, 520
 - application characteristics, 520
 - coverage area, 520
 - logical network architecture and, 520
 - mobility, 520, 531
 - network integration, 531
 - performance, 525
 - physical considerations, 520
 - project management, 521
 - security, 520
 - site environmental considerations, 521
 - throughput, 519
 - user density, 520
 - user population, 520
 - site surveys, 519
- Protected Frame bit, control frames, 77
- protocol architecture, future, 594
- protocol data unit, 603
- protocol stack, security and, 500–503
- protocols, 497
 - 802.1X-based, 483
 - Frame Control subfield, 47
 - higher-layer, encapsulation, 54
 - MAC address filtering, 483
 - network-layer authentication, 483
 - security
 - layered, 128
 - protocol stack, 500–503
 - selecting, 500
 - SNAP (sub-network access protocol), 55
 - version, control frames, 77
 - WEP shared key authentication, 483
- PS (Power Save), 603
- PSDU (PLCP Service Data Unit), 263, 603
- PSF (PLCP Signaling), 250