Eldad Perahia and Robert Stacey

Next Generation Wireless LANs

Throughput, Robustness, and Reliability in 802.11n

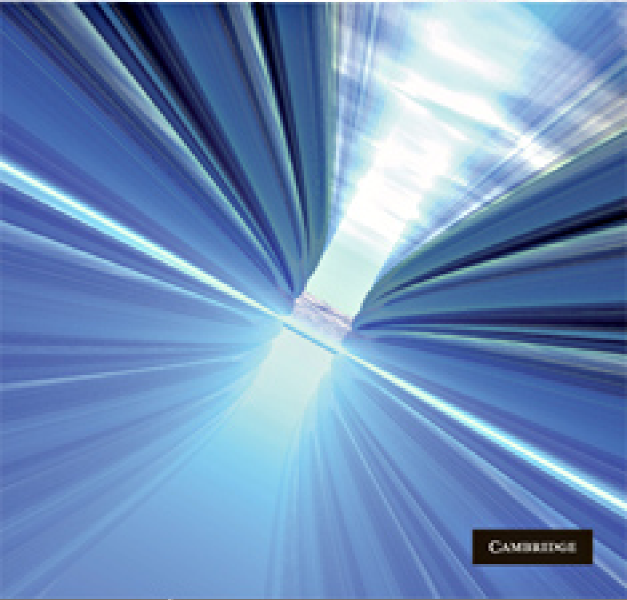

CAMBRIDGE

CAMBRIDGE www.cambridge.org/9780521885843

This page intentionally left blank

Next Generation Wireless LANs

If you've been searching for a way to get up to speed quickly on IEEE 802.11n without having to wade through the entire standard, then look no further. This comprehensive overview describes the underlying principles, implementation details, and key enhancing features of 802.11n. For many of these features, the authors outline the motivation and history behind their adoption into the standard. A detailed discussion of the key throughput, robustness, and reliability enhancing features (such as MIMO, 40 MHz channels, and packet aggregation) is given, in addition to a clear summary of the issues surrounding legacy interoperability and coexistence. Advanced topics such as beamforming and fast link adaption are also covered. With numerous MAC and physical layer examples and simulation results included to highlight the benefits of the new features, this is an ideal reference for designers of WLAN equipment, and network managers whose systems adopt the new standard. It is also a useful distillation of 802.11n technology for graduate students and researchers in the field of wireless communication.

Eldad Perahia is a member of the Wireless Standards and Technology group at Intel Corporation, Chair of the IEEE 802.11 Very High Throughput Study Group, and the IEEE 802.11 liaison to IEEE 802.19. Prior to joining Intel, Dr. Perahia was the 802.11n lead for Cisco Systems. He was awarded his Ph.D. in Electrical Engineering from the University of California, Los Angeles, and has fourteen patents in various areas of wireless communications.

Robert Stacey is a member of the Wireless Standards and Technology group at Intel Corporation. He was a member of the IEEE 802.11 High Throughput Task Group (TGn) and a key contributor to the various proposals culminating in the final joint proposal submission that became the basis for the 802.11n draft standard, and has numerous patents filed in the field of wireless communications.

Next Generation Wireless LANs

Throughput, Robustness, and Reliability in 802.11n

ELDAD PERAHIA AND ROBERT STACEY

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK
Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9780521885843

First published in print format 2008

ISBN-13 978-0-511-43391-7 eBook (Adobe Reader)

ISBN-13 978-0-521-88584-3 hardback

To my wife Sarah and our son Nathan
– Eldad Perahia

To my father, who nurtured and guided an inquiring mind
– Robert Stacey

Brief contents

Contents

Part III Transmit beamforming

Foreword

The first version of the 802.11 standard was ratified in 1997 after seven long years of development. However, initial adoption of this new technology was slow, partly because of the low penetration of devices that needed the "freedom of wireless."

The real opportunity for 802.11 came with the increased popularity of laptop computers just a few years later. This popularity brought a rapidly growing user base wanting network connectivity not only while connected to an Ethernet cable at home or at work, but also in between: in hotels, airports, conference centers, restaurants, parks, etc. 802.11 provided a cheap and easy way to make laptop mobility a reality for anyone who wanted it.

However, technology by itself is rarely sufficient, particularly in the networking space, where interoperability of devices from multiple vendors is almost always the key to market success. Having been formed as WECA in 1999, the Wi-Fi Alliance was ready to provide certification of multi-vendor interoperability.

With the right technology from the IEEE 802.11 Working Group, certified interoperability from the Wi-Fi Alliance, and a real market need based on a growing installed base of laptops, the conditions were ripe for the Wi-Fi market to take off, and indeed it did. By 2007 virtually every new laptop contains Wi-Fi as standard equipment. More importantly, and unlike some other "successful" wireless technologies, many of these devices are used regularly. With this wide use came a growing understanding of the power of cheap, easy-to-deploy, and easy-to-manage interoperable Wi-Fi networks.

The natural next step was for people to ask, "What else can we use Wi-Fi for?" The answer is increasingly becoming "everything, everywhere!" Not just laptops, but now almost anything mobile and even many fixed devices contain Wi-Fi, and they are used in a phenomenal range of applications, including data, voice, games, music, video, location, public safety, vehicular, etc. In 2007, more than 300 million Wi-Fi devices were shipped. By 2012, some analysts are forecasting that more than one billion Wi-Fi devices will be shipped every year.

The 2.4 GHz 802.11b 11 Mb/s DSSS/CCK PHY and the basic 802.11 contention-based MAC provided the basis for a great industry. However, the rapid growth of the Wi-Fi market challenged the capabilities of the technology. It was not long before better security (802.11i certified by the Wi-Fi Alliance as WPA/WPA2^{TM}) and better Quality of Service (802.11e certified by the Wi-Fi Alliance as WMM^{TM} and WMM Power Save) were defined, certified, and deployed.

It was also not long before higher data rates were demanded for greater data density and to support the many new and exciting devices and applications. 802.11a, providing 54 Mbps based on OFDM in the 5 GHz band, failed to garner significant support because two radios were required to maintain backward compatibility with 2.4 GHz 802.11b devices; the cost of two radios was often too high. The real success story was 802.11g, which provided 54 Mbps based on OFDM in the 2.4 GHz band in a way that was backward-compatible with 802.11b.

The success of 802.11g drove the use of Wi-Fi to new heights and expanded the demands on the technology yet again; everyone wanted more. Fortunately, the technology continued to develop and in 2002 the IEEE 802.11 Working Group started defining the next generation of PHY and MAC features as part of 802.11n. 802.11n will define mechanisms to provide users some combination of greater throughput, longer range and increased reliability, using mandatory and optional features in the PHY (including MIMO technology and 40 MHz channels) and the MAC (including more efficient data aggregation and acknowledgments).

Interestingly, 802.11n operates in both the 2.4 GHz and 5 GHz bands. It is expected that 5 GHz operation will be more popular than when 802.11a was introduced, because 2.4 GHz is now more congested, the number of available channels in the 5 GHz band has been expanded with the introduction of DFS and TPC technology, there is more need for high throughput 40 MHz channels, and the cost of dual-band radios has decreased.

The 802.11n standard is not yet complete, and is unlikely to be ratified by the IEEE until at least mid 2009. Until August 2006, the Wi-Fi Alliance had a policy to not certify 802.11n products until the standard was ratified. However, some vendors decided the market could not wait for ratification of the 802.11n standard and started releasing pre-standard products. These products were often not interoperable at the expected performance levels because they were not based on a common interpretation of the draft 802.11n specification. The problem for the Wi-Fi Alliance was that these products were adversely affecting the reputation of Wi-Fi. The Wi-Fi Alliance decided the only way forward was to certify the basic features of 802.11n from a pre-standard draft. Such a decision is not without precedent. In 2003, certification of WPA started before the 802.11i standard was ratified and in 2004 certification of WMM started before 802.11e was ratified. The Wi-Fi Alliance commenced certification of 802.11n draft 2.0 on 26 June 2007.

The decision has turned out to be the right one for the industry and for users. The Wi-Fi CERTIFIED 802.11n draft 2.0 programme has been remarkably successful, with more than 150 products certified in less than five months. This represents a significantly higher number of certified products than for the 802.11g programme during a similar period after launch. The Wi-Fi Alliance's certification program has helped ensure interoperability for the many products that will be released before the ratification of the 802.11n standard. This is particularly important given that the likely ratification date of the 802.11n standard has been extended by more than a year since the decision to start a certification program was announced by the Wi-Fi Alliance. The next challenge for the Wi-Fi Alliance is to ensure a backward-compatible transition path from the 802.11n draft 2.0 as certified by the Wi-Fi Alliance to the final ratified standard.

Standards are never the most accessible of documents. The 802.11 standard is particularly difficult to understand because it has been amended so many times by different groups and editors over a long period. A draft amendment to the standard, such as 802.11n D2.0, is even harder to interpret because many clauses are still being refined and the refinement process often has technical and political aspects that are only visible to those participating full time in the IEEE 802.11 Working Group.

Books like this one are invaluable because they provide the details and the background that allow readers to answer the questions, "What is likely to be in the final standard and how does it work?" Eldad and Robert should be congratulated on taking up the challenge.

Dr. Andrew Myles
Chairman of the BoD
Wi-Fi Alliance
6 December 2007

Preface

Having worked on the development of the 802.11n standard for some time, we presented a full day tutorial on the 802.11n physical layer (PHY) and medium access control (MAC) layer at the IEEE Globecom conference held in San Francisco in December 2006. Our objective was to provide a high level overview of the draft standard since, at the time, there was very little information on the details of the 802.11n standard available to those not intimately involved in its development. After the tutorial, we were approached by Phil Meyler of Cambridge University Press and asked to consider expanding the tutorial into a book.

Writing a book describing the standard was an intriguing prospect. We felt that a book would provide more opportunity to present the technical details in the standard than was possible with the tutorial. It would fill the gap we saw in the market for a detailed description of what is destined to be one of the most widely implemented wireless technologies. While the standard itself conveys details on what is needed for interoperability, it lacks the background on why particular options should be implemented, where particular aspects came from, the constraints under which they were designed, or the benefit they provide. All this we hoped to capture in the book. The benefits various features provide, particularly in the physical layer, are quantified by simulation results. We wanted to provide enough information to enable the reader to model the physical layer and benchmark their simulation against our results. Finally, with the amended standard now approaching 2500 pages, we hoped to provide an accessible window into the most important aspects, focusing on the throughput and robustness enhancements and the foundations on which these are built.

The book we came up with is divided into three parts. The first part covers the physical layer (PHY), and begins with background information on the 802.11a/g OFDM PHY on which the 802.11n PHY is based and interoperates, and proceeds with an overview of spatial multiplexing, the key throughput enhancing technology in 802.11n. This is followed by details on exactly how high throughput is achieved in 802.11n using spatial multiplexing and wider, 40 MHz channels. This in turn is followed by details on robustness enhancing features such as receive diversity, spatial expansion, space-time block codes, and low density parity check codes.

The second part covers the medium access control (MAC) layer. This part provides background on the original 802.11 MAC as well as the 802.11e quality of service (QoS) enhancements. It gives an overview of why changes were needed in the MAC to achieve higher throughput, followed by details on each of the new features introduced. Given the large installed base of 802.11 devices, coexistence and interoperability are considered crucial to the smooth adoption of the standard. To this end, the book provides a detailed discussion on features supporting coexistence and interoperability.

In the third part we provide details on two of the more advanced aspects of the standard, transmit beamforming and link adaptation. These topics are covered in a section of their own, covering both the PHY and the MAC.

Writing this book would not have been possible without help from our friends and colleagues. We would like to thank Thomas (Tom) Kenney and Brian Hart for reviewing the PHY portion of the book and Solomon Trainin, Tom Kenney, and Michelle Gong for reviewing the MAC portion of the book. They provided insightful comments, suggestions, and corrections that significantly improved the quality of the book.

One of the goals of this book is to provide the reader with a quantitative feel of the benefit of the PHY features in the 802.11n standard. This would have been impossible without the extensive simulation support provided to us by Tom Kenney. He developed an 802.11n PHY simulation platform that includes most of the 802.11n PHY features and is also capable of modeling legacy 802.11a/g. The simulation includes all the 802.11n channel models. Furthermore, Tom modeled receiver functionality such as synchronization, channel estimation, and phase tracking. The simulation also included impairments such as power amplifier non-linearity and phase noise to provide a more realistic measure of performance.

The simulation supports both 20 MHz and 40 MHz channel widths. With the 40 MHz simulation capability, Tom generated the results given in Figure 5.8 in Section 5.1.5 modeling MCS 32 and Figure 5.9 in Section 5.1.7 which illustrates the range and throughput improvement of 40 MHz modes. With the MIMO/SDM capability of the simulation in both AWGN channel and 802.11n channel models, Tom produced the results for Figures 5.12–5.15 in Section 5.3. By designing the simulation with the flexibility to set the transmitter and receiver to different modes, he also produced the results given in Figure 5.18 in Section 5.4 modeling the behavior of a legacy 802.11a/g device receiving a GF transmission. Tom also incorporated short guard interval into the simulation with which the results for sensitivity to time synchronization error in Figures 5.20–5.22 in Section 5.5 were generated.

Tom designed the simulation with the ability to select an arbitrary number of transmitter and receiver antennas independent from the number of spatial streams. Using this capability he produced the results for receive diversity gain in Figures 6.2–6.4 in Section 6.1 and spatial expansion performance in Figures 6.5 and 6.6 in Section 6.2. Tom also incorporated space-time block coding and low density parity check coding into the simulation and generated the results given in Figures 6.8, 6.9, 6.14, 6.15, and 6.16 in Section 6.3 and Figure 6.24 in Section 6.4.

To accurately model the performance of a transmit beamforming system, it is important to include aspects like measurement of the channel state information, beamforming weight computation, and link adaptation. Tom incorporated all of these functions into the simulation to generate the waterfall curves in Figures 12.11–12.16 and the throughput curves in Figures 12.17 and 12.18 in Section 12.18.

We sincerely hope our book provides you with insight and a deeper understanding of the 802.11n standard and the technology upon which it is built.

Abbreviations

μs	microseconds
2G	second generation (cellular)
3G	third generation (cellular)
AC	access category
ACK	acknowledgement
ADC	analog-to-digital converter
ADDBA	add block acknowledgement
ADDTS	add traffic stream
AGC	automatic gain control
AID	association identifier
AIFS	arbitration inter-frame space
A-MPDU	aggregate MAC protocol data unit
A-MSDU	aggregate MAC service data unit
AoA	angle of arrival
AoD	angle of departure
AP	access point
APSD	automatic power save delivery
A-PSDU	aggregate PHY service data unit
AS	angular spectrum
ASEL	antenna selection
AWGN	additive white Gaussian noise
BA	block acknowledgement
BAR	block acknowledgement request
BCC	binary convolution code
BF	beamforming
BICM	bit interleaved coded modulation
bps	bits-per-second
BPSCS	coded bits per single carrier for each spatial stream
BPSK	binary phase shift keying
BSS	basic service set
BSSID	BSS identifier
BW	bandwidth
CBPS	coded bits per symbol
CBPSS	coded bits per spatial stream

CBW	channel bandwidth
CCA	clear channel assessment
CCDF	complementary cumulative distribution function
CCK	complementary code keying
CFP	contention free period
CP	contention period
CRC	cyclic redundancy code
CS	carrier sense
CSD	cyclic shift diversity
CSI	channel state information
CSMA	carrier sense multiple access
CSMA/CA	carrier sense multiple access with collision avoidance
CSMA/CD	carrier sense multiple access with collision detection
CTS	clear to send
CW	contention window
DA	destination address
DAC	digital-to-analog converter
dB	decibels
dBc	decibels relative to carrier
dBi	decibels isotropic relative to an antenna
dBm	decibel of measured power referenced to one milliwatt
DBPS	data bits per OFDM symbol
dBr	dB (relative)
DC	direct current
DCF	distributed coordination function
DELBA	delete block acknowledgement
DIFS	DCF inter-frame space
DLS	direct link session
DS	distribution system
DSL	digital subscriber line
DSSS	direct sequence spread spectrum
DTIM	delivery traffic indication message
DVD	digital versatile disc
EDCA	enhanced distributed channel access
EIFS	extended inter-frame space
ERP	enhanced rate PHY
ESS	extended service set
ETSI	European Telecommunications Standards Institute
EVM	error vector magnitude
EWC	Enhanced Wireless Consortium
FCC	Federal Communications Commission
FCS	frame check sequence
FEC	forward error correction
FFT	fast Fourier transform

FHSS	frequency hopped spread spectrum
FS	free space
FTP	file transfer protocol
GF	Greenfield
GF-HT-STF	Greenfield High Throughput Short Training field
GHz	gigahertz
GI	guard interval
GIF	graphics interchange format
GPS	global positioning system
HC	hybrid coordinator
HCCA	HCF controlled channel access
HCF	hybrid coordination function
HEMM	HCCA, EDCA mixed mode
HT	high throughput
HTC	high throughput control
HT-DATA	High Throughput Data field
HT-LTF	High Throughput Long Training field
HTSG	High Throughput Study Group
HT-SIG	High Throughput Signal field
HT-STF	High Throughput Short Training field
HTTP	hypertext transfer protocol
Hz	Hertz
IBSS	independent basic service set
IC	integrated circuit
IDFT	inverse discrete Fourier transform
IEEE	Institute of Electrical and Electronic Engineers
IFFT	inverse fast Fourier transform
IFS	inter-frame space
IP	Internet Protocol
IPv6	Internet Protocol version 6
IR	infrared
ISI	inter-symbol interference
ISM	industrial, scientific, and medical
JPEG	Joint Photographic Experts Group
kHz	kilohertz
km/h	kilometers per hour
LAN	local area networking
LDPC	low density parity check
LLC	logical link control
L-LTF	Non-HT (Legacy) Long Training field
LNA	low noise amplifier
LOS	line-of-sight
LSB	least significant bit
L-SIG	Non-HT (Legacy) Signal field

L-STF	Non-HT (Legacy) Short Training field
LTF	Long Training field
m	meters
MAC	medium access control
MAI	MRQ or ASEL indication
MAN	metropolitan area networking
Mbps	megabit per second
MCS	modulation and coding scheme
MF	mixed format
MFB	MCS feedback
MFSI	MCS feedback sequence indication
MHz	megahertz
MIB	management information base
MIMO	multiple-input multiple-output
ML	maximum likelihood
MMPDU	MAC management protocol data unit
MMSE	minimum mean-square-error
MPDU	MAC protocol data unit
MPEG	Moving Picture Experts Group
MRC	maximal-ratio combining
MRQ	MCS request
Msample/s	mega-samples per second
MSB	most significant bit
MSDU	MAC service data unit
MSE	mean-square-error
MSFI	MCS feedback sequence identifier
MSI	MCS request sequence identifier
NAV	network allocation vector
NDP	null data packet
NF	noise figure
NLOS	non-line-of-sight
nsec	nanosecond
OBO	output back-off
OBSS	overlapping BSS
OFDM	orthogonal frequency division multiplexing
OSI	open systems interconnection
PA	power amplifier
PAR	project authorization request
PAS	power angular spectrum
PC	point coordinator
PCF	point coordination function
PCO	phased coexistence operation
PDU	protocol data unit
PER	packet error rate

PHY	physical layer
PIFS	PCF inter-frame space
PLCP	physical layer convergence procedure
PPDU	PLCP protocol data unit
ppm	parts per million
PSD	power spectral density
PSDU	PLCP service data unit
PSMP	power-save multi-poll
PSMP-DTT	PSMP downlink transmission time
PSMP-UTT	PSMP uplink transmission time
QAM	quadrature amplitude modulation
QoS	quality of service
QPSK	quadrature phase shift keying
R	code rate
RA	receiver address
RD	reverse direction
RDG	reverse direction grant
RF	radio frequency
RIFS	reduced inter-frame space
RMS	root-mean-square
RSSI	received signal strength indication
RTS	request to send
Rx	receive
SA	source address
SAP	service access point
SCP	secure copy protocol
SDM	spatial division multiplexing
SDU	service data unit
SE	spatial expansion
SIFS	short inter-frame space
SIG	Signal field
SIMO	single-input, multiple-output
SISO	single-input, single-output
SMTP	simple mail transfer protocol
SNR	signal-to-noise ratio
SOHO	small-office, home-office
SS	spatial stream
SSC	starting sequence control
SSID	service set identifier
SSN	starting sequence number
STA	station
STBC	space-time block coding
STF	Short Training field
STS	space-time stream

SVD	singular value decomposition
SYM	symbol
TA	transmitter address
TBTT	target beacon transmission time
TC	traffic category
TCLAS	traffic classification
TCM	trellis coded modulation
TCP	transmission control protocol
TDD	time division duplexing
TGn	Task Group n
TGy	Task Group y
TID	traffic identifier
TIFF	tagged image file format
TRQ	training request
TS	traffic stream
TSID	traffic stream identifier
TSPEC	traffic specification
TV	television
Tx	transmit
TxBF	transmit beamforming
TXOP	transmit opportunity
TXTIME	transmit time
UDP	user datagram protocol
USA	United States of America
VoIP	voice over IP
VPN	virtual private network
WEP	wired equivalent privacy
WFA	Wi-Fi Alliance
WLAN	wireless local area network
WM	wireless medium
WNG SC	Wireless Next Generation Standing Committee
WWiSE	world wide spectral efficiency
XOR	exclusive-or
ZF	zero-forcing
ZIP	ZIP file format

1 Introduction

Wireless local area networking has experienced tremendous growth in the last ten years with the proliferation of IEEE 802.11 devices. Its beginnings date back to Hertz's discovery of radio waves in 1888, followed by Marconi's initial experimentation with transmission and reception of radio waves over long distances in 1894. In the following century, radio communication and radar proved to be invaluable to the military, which included the development of spread spectrum technology. The first packet-based wireless network, ALOHANET, was created by researchers at the University of Hawaii in 1971. Seven computers were deployed over four islands communicating with a central computer in a bi-directional star topology.

A milestone event for commercial wireless local area networks (WLANs) came about in 1985 when the United States Federal Communications Commission (FCC) allowed the use of the experimental industrial, scientific, and medical (ISM) radio bands for the commercial application of spread spectrum technology. Several generations of proprietary WLAN devices were developed to use these bands, including WaveLAN by Bell Labs. These initial systems were expensive and deployment was only feasible when running cable was difficult.

Advances in semiconductor technology and WLAN standardization with IEEE 802.11 led to a dramatic reduction in cost and the increased adoption of WLAN technology. With the increasing commercial interest, the Wi-Fi Alliance (WFA) was formed in 1999 to certify interoperability between IEEE 802.11 devices from different manufacturers through rigorous testing. Since 2000, shipments of Wi-Fi certified integrated circuits (IC) reached 200 million per year in 2006 (ABIresearch, 2007). Shipments are expected to exceed a billion units per year by 2012 (ABIresearch, 2007), as illustrated in Figure 1.1.

Such large and sustained growth is due to the benefits WLANs offer over wired networking. In existing homes or enterprises, deploying cables for network access may involve tearing up walls, floors, or ceilings, which is both inconvenient and costly. In contrast, providing wireless network connectivity in these environments is often as simple as installing a single wireless access point. Perhaps more importantly though, the proliferation of laptops and handheld devices has meant that people desire connectivity wherever they are located, not just where the network connection is located. Network connectivity in a conference room or while seated on the sofa in the living room are just two examples of the flexibility afforded by WLANs.

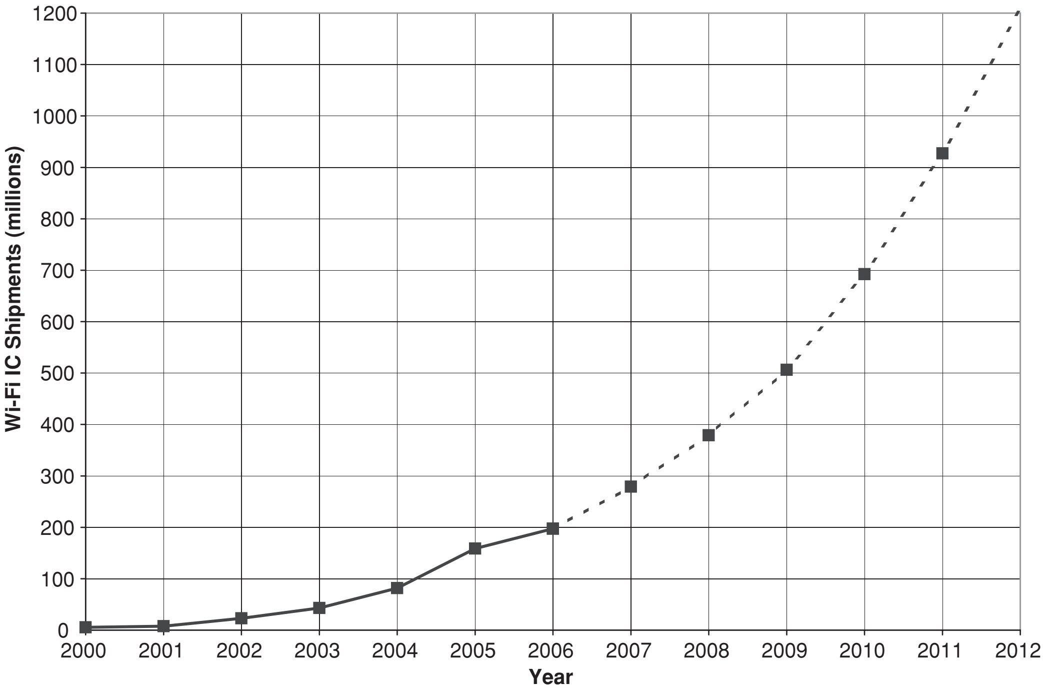

Figure 1.1 Wi-Fi IC shipments. Source: ABIresearch (2007).

Building on the convenience of mobility from the cellular world, WLANs are now enabling Internet access at little or no cost in public wireless networks. In 2005, Google offered to deploy a free Wi-Fi service covering San Francisco at no cost to the city. There has also been a proliferation of small scale deployments providing Internet access in coffee shops, airports, hotels, etc., which have come to be known as hotspots. Additionally, when these networks are used in conjunction with virtual private network (VPN) technology, employees can securely access corporate networks from almost anywhere.

The vast majority of WLAN products and systems today are based on the 802.11b, 802.11g, and 802.11a standard amendments, which provide throughput enhancements over the original 802.11 PHYs. Progress in WLAN technology continues with the development of 802.11n. Increased data rates are achieved with the multiple-input multiple-output (MIMO) concept, with its origins by Foschini (1996) at Bell Labs. In 2004, Atheros demonstrated that 40 MHz devices could be produced at almost the same cost as 20 MHz devices. During a similar time frame, the FCC and ETSI adopted new regulations in the 5 GHz band that added an additional 400 MHz of unlicensed spectrum for use by commercial WLANs. These events paved the way for the broad acceptance of 40 MHz operating modes in 802.11n. When spectrum is free, increasing the channel bandwidth is the most cost effective way to increase the data rate.

Typically product development lags standardization efforts and products are released after the publication of the standard. An interesting event occurred in 2003 when Broadcom released a chipset based on a draft version of the 802.11g amendment, prior to final publication. This set a precedent for the flurry of "pre-n" or "draft-n" products in 2005 and 2006, as industry players rushed to be first to market. Most of these products were either proprietary implementations of MIMO, or based on draft 1.0 of 802.11n, and thus unlikely to be compliant with the final standard.

Through early 2007, major improvements and clarifications were made to the 802.11n draft resulting in IEEE 802.11n draft 2.0. To continue the market momentum and forestall interoperability issues, the IEEE took the unusual step of releasing 802.11n D2.0 to the public while work continued toward the final standard. This allowed the WFA to begin interoperability testing and certification of devices based on a subset of the 802.11n D2.0 features in May 2007. WFA certified 802.11n D2.0 products provide consumers the assurance of interoperability between manufacturers that was not guaranteed by previous "pre-n" or "draft-n" products. These were major steps in speeding up the standardization and certification process of new technology.

1.1 History of IEEE 802.11

The IEEE 802.11 working group began development of a common medium access control (MAC) layer for multiple physical layers (PHY) to standardize wireless local area networking. As a member of the IEEE 802 family of local area networking (LAN) and metropolitan area networking (MAN) standards, 802.11 interfaces with 802.1 architecture, management, and interworking, and 802.2 logical link control (LLC). The combination of 802.2 LLC and 802.11 MAC and PHY make up the data link and physical layers of the Open Systems Interconnection (OSI) reference model, as described in Table 1.1.

Table 1.1 OSI reference model (Zimmerman, 1980; Teare, 1999)

<table>
<tr><th>OSI Reference Model layers</th><th>Description</th><th>Examples</th><th>Layer categories</th></tr>
<tr><td>Application</td><td>Interacts with software applications that implement a communicating component</td><td>Telnet, FTP, SMTP</td><td rowspan="3">Application</td></tr>
<tr><td>Presentation</td><td>Coding and conversion functions that are applied to application layer data</td><td>QuickTime, MPEG, GIF, JPEG, TIFF</td></tr>
<tr><td>Session</td><td>Establishes, manages, and terminates communication sessions</td><td>ZIP, AppleTalk, SCP, DECnet Phase IV</td></tr>
<tr><td>Transport</td><td>Accepts data from the session layer and segments the data for transport across the network</td><td>TCP, UDP</td><td rowspan="4">Data transport</td></tr>
<tr><td>Network</td><td>Defines the network address</td><td>IP, IPv6</td></tr>
<tr><td>Data link</td><td>Transit of data across a physical network link</td><td>802.2 LLC
</td></tr>
<tr><td>Physical</td><td>Electrical, mechanical, procedural, and functional specifications</td><td>802.11 PHY</td></tr>
</table>

The initial version of the 802.11 standard was completed in 1997. Influenced by the huge market success of Ethernet (standardized as IEEE 802.3), the 802.11 MAC adopted the same simple distributed access protocol, carrier sense multiple access (CSMA). With CSMA, a station wishing to transmit first listens to the medium for a predetermined period. If the medium is sensed to be “idle” during this period then the station is permitted to transmit. If the medium is sensed to be “busy,” the station has to defer its transmission. The original (shared medium) Ethernet used a variation called CSMA/CD or carrier sense multiple access with collision detection. After determining that the medium is “idle” and transmitting, the station is able to receive its own transmission and detect collisions. If a collision is detected, the two colliding stations backoff for a random period before transmitting again. The random backoff period reduces the probability of a second collision.

With wireless it is not possible to detect a collision with one’s own transmission directly in this way: thus 802.11 uses a variation called CSMA/CA or carrier sense multiple access with collision avoidance. With CSMA/CA, if the station detects that the medium is busy, it defers its transmission for a random period following the medium going “idle” again. This approach of always backing off for a random period following another station’s transmission improves performance since the penalty for a collision is much higher on a wireless LAN than on a wired LAN. On a wired LAN collisions are detected electrically and thus almost immediately, while on wireless LAN collisions are inferred through the lack of an acknowledgement or other response from the remote station once the complete frame has been transmitted.

There is no doubt that the simplicity of this distributed access protocol, which enables consistent implementation across all nodes, significantly contributed to Ethernet’s rapid adoption as the industry LAN standard. Likewise, the adoption by the industry of 802.11 as the wireless LAN standard is in no small part due to the simplicity of this access protocol, its similarity to Ethernet, and again the consistent implementation across all nodes that has allowed 802.11 to beat out the more complex, centrally coordinated access protocols of competing WLAN technologies such as HyperLAN.

The original (1997) 802.11 standard included three PHYs: infrared (IR), 2.4 GHz frequency hopped spread spectrum (FHSS), and 2.4 GHz direct sequence spread spectrum (DSSS). This was followed by two standard amendments in 1999: 802.11b built upon DSSS to increase the data rate in 2.4 GHz and 802.11a to create a new PHY in 5 GHz. 802.11b enhanced DSSS with complementary code keying (CCK), increasing the data rate to 11 Mbps. With higher data rates, IEEE 802.11b devices achieved significant market success, and markets for IR and FHSS PHYs did not materialize.

The development of 802.11a introduced orthogonal frequency division multiplexing (OFDM) to 802.11. Even though 802.11a introduced data rates of up to 54 Mbps, it is confined to the 5 GHz band and, as a result, adoption has been slow. New devices wishing to take advantage of the higher rates provided by 802.11a but retain backward compatibility with the huge installed base of 802.11b devices would need to implement two radios, one to operate using 802.11b in the 2.4 GHz band and one to operate using 802.11a in the 5 GHz band. Furthermore, international frequency regulations in the 2.4 GHz band uniformly allowed commercial use, whereas in 1999 and 2000 the non-military use of the 5 GHz band was limited to select channels in the United States.

Table 1.2 Overview of 802.11 PHYs

	802.11	802.11b	802.11a	802.11g	802.11n
PHY technology	DSSS	DSSS/CCK	OFDM	OFDM DSSS/CCK	SDM/OFDM
Data rates	1, 2 Mbps	5.5, 11 Mbps	6–54 Mbps	1–54 Mbps	6–600 Mbps
Frequency band	2.4 GHz	2.4 GHz	5 GHz	2.4 GHz	2.4 and 5 GHz
Channel spacing	25 MHz	25 MHz	20 MHz	25 MHz	20 and 40 MHz

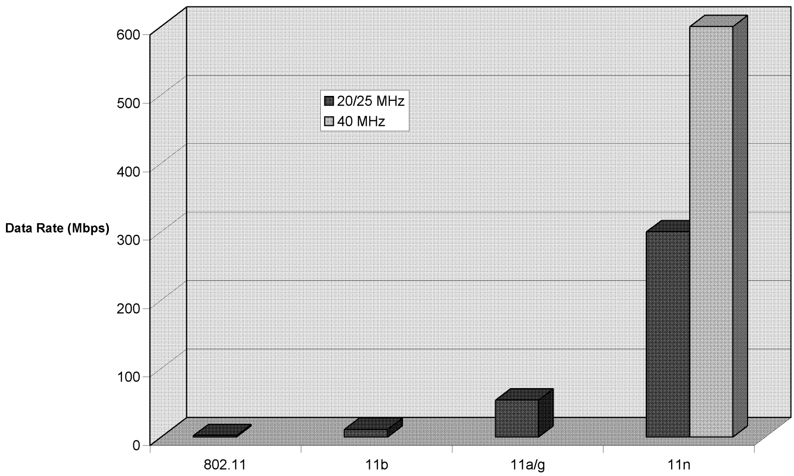

Figure 1.2 Increase in 802.11 PHY data rate.

In 2001, the FCC permitted the use of OFDM in the 2.4 GHz band. Subsequently, the 802.11 working group developed the 802.11g amendment, which incorporates the 802.11a OFDM PHY in the 2.4 GHz band, and adopted it as part of the standard in 2003. In addition, backward compatibility and interoperability is maintained between 802.11g and the older 802.11b devices. This allows for new 802.11g client cards to work in existing 802.11b hotspots, or older 802.11b embedded client devices to connect with a new 802.11g access point (AP). Because of this and new data rates of up to 54 Mbps, 802.11g has experienced large market success. A summary of the high level features of each PHY is given in Table 1.2.

With the adoption of each new PHY, 802.11 has experienced a five-fold increase in data rate. This rate of increase continues with 802.11n with a data rate of 300 Mbps in 20 MHz and 600 Mbps in 40 MHz. The exponential rate of increase in data rate is illustrated in Figure 1.2.

1.2 History of high throughput and 802.11n

1.2.1 The High Throughput Study Group

Interest in a high data rate extension to 802.11a began with a presentation to the Wireless Next Generation Standing Committee (WNG SC) of IEEE 802.11 in January 2002. Market drivers were outlined, such as increasing data rates of wired Ethernet, more data rate intensive applications, non-standard 100+ Mbps products entering the market, and the need for higher capacity WLAN networks (Jones, 2002). The presentation mentioned techniques such as spatial multiplexing and doubling the bandwidth as potential approaches to study for increasing data rate.

After many additional presentations, the High Throughput Study Group (HTSG) was formed with its first meeting in September 2002. The primary objective of HTSG was to complete two documents necessary for the creation of the High Throughput Task Group (TGn). These are the project authorization request (PAR) form and five criteria form. The PAR defined the scope and purpose of the task group as follows:

The scope of this project is to define an amendment that shall define standardized modifications to both the 802.11 physical layers (PHY) and the 802.11 medium access control layer (MAC) so that modes of operation can be enabled that are capable of much higher throughputs, with a maximum throughput of at least 100 Mbps, as measured at the MAC data service access point (SAP). IEEE (2006)

By this statement, the standard amendment developed by TGn must contain modes of operation that are capable of achieving at least 100 Mbps *throughput*. Throughput is the measure of "useful" information delivered by the system and by using throughput as the metric, both MAC and PHY overhead must be considered. 802.11a/g systems typically achieve a maximum throughput of around 25 Mbps; thus this statement required at least a four fold increase in throughput. Meeting this requirement would in essence mandate PHY data rates well in excess of 100 Mbps as well as significant enhancements to MAC efficiency.

Additional explanatory notes were included with the PAR outlining many evaluation metrics. These include throughput at the MAC SAP, range, aggregate network capacity, power consumption, spectral flexibility, cost complexity flexibility, backward compatibility, and coexistence (IEEE, 2006).

The five criteria form requires that the study group demonstrate the necessity of creating an amendment to the standard. The criteria include (1) broad market potential, (2) compatibility with existing IEEE 802.1 architecture, (3) distinct identity from other IEEE 802 standards, (4) technical feasibility, and (5) economic feasibility (Rosdahl, 2003). The goal is to create a standard amendment which results in marketable products, but that will also be differentiated from other potentially similar products.

In addition to completing the PAR and five criteria forms, HTSG also began development of new multipath fading MIMO channel models (Erceg *et al*., 2004) and usage models (Stephens *et al*., 2004). The channel models and usage models were used to create a common framework for simulations by different participants in the standard development process.

1.2.2 Formation of the High Throughput Task Group (TGn)

The PAR was accepted and approved by the 802 working group, creating Task Group n (TGn) with the first meeting of the task group held in September 2003. The standard amendment developed by the task group would be proposal driven, meaning that members of the task group would make partial or complete technical proposals, with the complete proposals proceeding through a down-selection process culminating in a single proposal upon which the standard amendment would be based. Partial proposals would be informative and could be incorporated in a complete proposal along the way. To that end, the task group began development of the functional requirements (Stephens, 2005) and comparison criteria (Stephens, 2004) documents. These two documents would provide, respectively, the technical requirements complete proposals must meet and criteria by which complete proposals would be compared.

The task group began with nine functional requirements. One of the functional requirements was a catch-all, requiring that proposals meet the PAR and five criteria. A second requirement was a reiteration of the PAR requirement to achieve 100 Mbps throughput at the top of the MAC. Furthermore, since it was expected that not all regulatory domains would allow a single device to use multiple 20 MHz channels (an easy way to achieve the throughput objective), the second requirement added a restriction that 100 Mbps throughput be achieved in a single 20 MHz channel. To enforce efficient use of spectrum, another requirement was added for a mode of operation with a spectral efficiency of at least 3 bps/Hz.

Four functional requirements addressed operational bands and backward compatibility. One of these requirements was that the protocol should support operation in the 5 GHz band due to the large availability of spectrum there. Another requirement was that at least some modes of operation be backward compatible with 802.11a systems. Noteworthy was the fact that there was no requirement to support operation in the 2.4 GHz band. However, if a proposal did support 2.4 GHz band operation, it was required that there be modes of operation that were backward compatible with 802.11g systems. In this context, some flexibility was given, allowing an 802.11n AP to be configured to accept or reject associations from legacy stations.

The 802.11e amendment to the standard, nearing completion at the time, added many features for improving the quality of service (QoS) in 802.11 systems. Many of the perceived applications for 802.11n involved real time voice and video which necessitate QoS. Therefore a functional requirement was included which mandated that a proposal allow for the implementation of 802.11e features within an 802.11n station.

The comparison criteria in Stephens (2004) outlined metrics and required disclosure of results which would allow for comparison between proposals under the same simulation setup and assumptions. The comparison criteria incorporated the simulation scenarios and usage models defined in Stephens *et al.* (2004). During the development of the comparison criteria, the task group realized that members of the task group did not always share the same definitions for common terms. Therefore definitions for goodput, backward compatibility, and signal-to-noise ratio (SNR) were provided. The comparison criteria covered four main categories: marketability, backward compatibility and coexistence with legacy devices, MAC related criteria, and PHY related criteria.

Under marketability, the proposal must provide goodput results for residential, enterprise, and hotspot simulation scenarios. Goodput is defined by totaling the number of bits in the MAC service data units (MSDU) indicated at the MAC service access point

(SAP), and dividing by the simulation duration (Stephens, 2004). Two optional criteria included describing the PHY and MAC complexity. The PHY complexity was to be given relative to 802.11a.

To ensure backward compatibility and coexistence with legacy devices, a proposal was required to provide a summary of the means used to achieve backward compatibility with 802.11a and, if operating in 2.4 GHz, 802.11g. Simulation results demonstrating interoperability were also required. The goodput of a legacy device in an 802.11n network and the impact of a legacy device on the goodput of 802.11n devices were also to be reported.

The MAC related criteria included performance measurements and changes that were made to the MAC. In the residential, enterprise, and hotspot simulation scenarios a number of different metrics were to be captured and reported. These included the ability to support the service requirements of various applications, including QoS requirements. Measurements of aggregate goodput of the entire simulation scenario were required to indicate network capacity. MAC efficiency was to be provided, which is defined as the aggregate goodput divided by the average PHY data rate. To ensure reasonable range for the new modes of operation, throughput versus range curves were also to be provided.

The PHY related criteria included PHY rates and preambles, channelization, spectral efficiency, PHY performance, and PHY changes. In addition, the comparison criteria also defined PHY impairments to be used in combination with channel models for PHY simulations. Each proposal was required to generate simulation results for both additive white Gaussian noise (AWGN) and non-AWGN channels. Furthermore, simulation conditions to analyze the impact on packet error rate (PER) of carrier frequency offset and symbol clock offset were also defined.

1.2.3 Call for proposals

The TGn call for proposals was issued on May 17, 2004, with the first proposals presented in September 2004. Over the course of the process two main proposal teams emerged, TGn Sync and WWiSE (world wide spectral efficiency). The TGn Sync proposal team was founded by Intel, Cisco, Agere, and Sony with the objective of covering the broad range of markets these companies were involved in, including the personal computer (PC), enterprise, and consumer electronics markets. The WWiSE proposal team was formed by Broadcom, Conexant, and Texas Instruments. These semiconductor companies were interested in a simple upgrade to 802.11a for fast time to market. Many other companies were involved in the proposal process and most ended up joining one of these two proposal teams.

The key features of all the proposals were similar, including spatial division multiplexing and 40 MHz channels for increased data rate, and frame aggregation for improved MAC efficiency. The proposals differed in scope (TGn Sync proposed numerous minor improvements to the MAC while WWiSE proposed limiting changes) and support for advanced features such as transmit beamforming (initially absent from the WWiSE proposal).

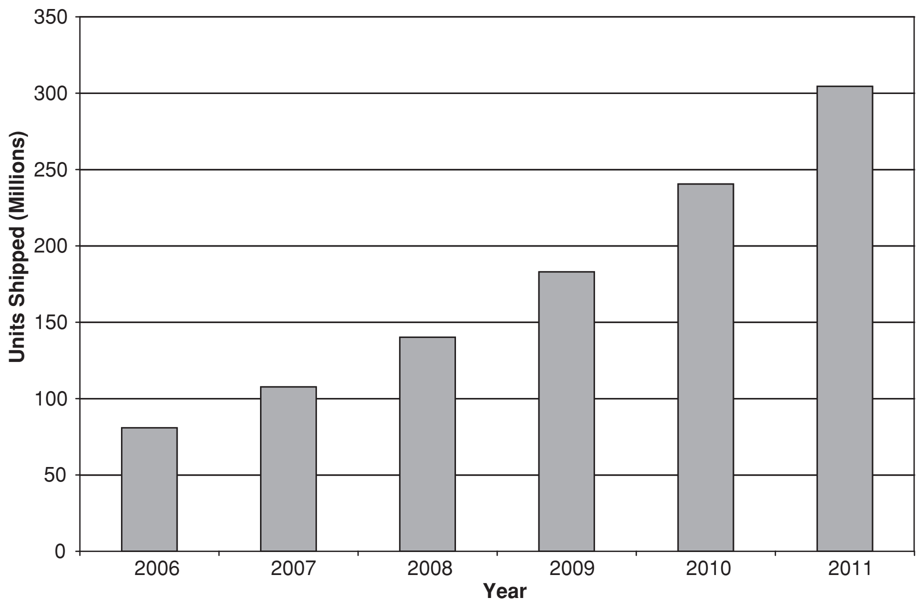

Figure 1.3 Worldwide converged mobile device shipments. Source: IDC (2007).

A series of proposal down-selection and confirmation votes took place between September 2004 and May 2005. During that time, mergers between proposals and enhancements to proposals took place. The TGn Sync proposal won the final down-selection vote between it and WWiSE, but failed the confirmation vote in May 2005.

1.2.4 Handheld devices

During this period interest arose in a new emerging market of converged Wi-Fi and mobile handsets. The shipment of dual mode Wi-Fi/cellular handsets had grown significantly from 2005 to 2006. Some participants in the proposal process believed that handsets would be the dominant Wi-Fi platform within a few years (de Courville *et al*., 2005). A projected world wide growth of converged mobile devices was given in IDC (2007) and is illustrated in Figure 1.3.

A contentious issue for handheld proponents was the high throughput requirement for 100 Mbps throughput. This, in essence, would force all 802.11n devices to have multiple antennas. This is a difficult requirement for converged mobile devices, since they already contain radios and antennas for cellular 2G, 3G, Bluetooth, and in some cases GPS. Concern was raised that mandating 802.11n devices to have multiple antennas would force handset manufacturers to continue to incorporate single antenna 802.11a/g into handsets and not upgrade to 802.11n. Not only does this diminish the capabilities of the handset device, it burdens all future 802.11n deployments with continued coexistence with 802.11a/g embedded in these new handset devices.

For this reason an ad hoc group was formed to create functional requirements supporting single antenna devices. Two new requirements were added to the functional requirements document in July 2005. The first requirement mandated that a proposal define single antenna modes of operation supporting at least 50 Mbps throughput in a 20 MHz channel. The second requirement dictated that an 802.11n AP or station interoperate with client devices that comply with 802.11n requirements but incorporate only a single antenna. This requirement resulted in 802.11n making mandatory at least two antennas in an AP, but only one antenna in a non-AP device.

1.2.5 Merging of proposals

After the failed confirmation vote, a joint proposal effort was started within the task group to merge the two competing proposals. Due to entrenched positions and the large membership of the group, the joint proposal effort proceeded very slowly. As a result, Intel and Broadcom formed the Enhanced Wireless Consortium (EWC) in October 2005 to produce a specification outside the IEEE that would bring products to market faster. With much of the task group membership ultimately joining the EWC, this effort had the effect of breaking the deadlock within the IEEE, and the EWC specification, which was essentially a merger of the TGn Sync and WWiSE proposals, was adopted as the joint proposal and submitted for confirmation to TGn where it was unanimously passed in January 2006.

1.2.6 802.11n amendment drafts

The joint proposal was converted to a draft 802.11 standard amendment for higher throughput (TGn Draft 1.0), and entered letter ballot. In letter ballot, IEEE 802.11 working group members (not just task group members) vote to either adopt the draft as is or reject it with comments detailing changes needed. The draft requires at least a 75% affirmative vote within the IEEE 802.11 working group in order to proceed to sponsor ballot where it is considered for adoption by the broader IEEE standards association. TGn Draft 1.0 entered letter ballot in March 2006 and, not unusually, failed to achieve the 75% threshold for adoption. Comment resolution began May 2006 on the roughly 6000 unique technical and editorial comments submitted along with the votes.

With resolution of the TGn Draft 1.0 comments, TGn Draft 2.0 went out for letter ballot vote in February 2007 and this time passed with 83% of the votes. However, there were still 3000 unique technical and editorial comments accompanying the letter ballot votes. It is typical for the task group to continue comment resolution until a minimum number of negative votes are received; thus comment resolution for TGn Draft 2.0 continued between March 2007 and September 2007, resulting in TGn Draft 3.0. Since TGn Draft 2.0 passed, TGn Draft 3.0 and possible later drafts only require a recirculation ballot in which comments may only address clauses that changed between the drafts.

At the time this book went to press, the standard amendment was in recirculation ballot and would continue there until a minimum number of negative votes and comments were received. It will then proceed to sponsor ballot. Whereas letter ballot includes only