

Exhibit 1023



TEAMS ▾

INDIVIDUALS

SIGN IN

TRY NOW

FEATURES ▾

BLOG

CONTENT SPONSORSHIP

SEARCH

802.11n: A Survival Guide by Matthew S. Gast

ChapterA 1.A Introduction to 802.11n-2009

I feel the need...

...the need for speed!
Maverick and Goose, *Top Gun*

802.11n's appeal can be summed up in one word--speed. In the stone ages of 1 Mbps and 2 Mbps wireless LANs, the most charitable thing that one could say about the technology was that it was faster than DSL. With the benefit of a decade of development, wireless LAN technology is mature and widely viewed as "the way you connect to a network," just as Ethernet

was in the mid-1990s. When I started working, visitors would often ask where the nearest Ethernet port could be found. Today, they ask for a password for the wireless LAN. The speed of 802.11n is a large part of the reason why so many devices don't even have the option for Ethernet ports any more. Once 802.11n put wireless LANs on the same performance footing as wired LANs, wireless became the obvious choice.

History

802.11n began life in late 2002, and took exactly seven years to produce from start to finish.^[3] As with all IEEE 802 standards, it began as a study group. The High Throughput Study Group was founded to investigate building an 802.11 physical layer that could provide speeds of 100 Mbps, as measured at the Medium Access Control (MAC) layer. An early rule of thumb for wireless LANs is that the throughput at the MAC layer would be half of the transmission speed. In 802.11b, the 11 Mbps transmission speed topped out at 5-6 Mbps of throughput; in 802.11a/g, the 54 Mbps transmission speed resulted in 25-30 Mbps of throughput. 802.11n's goal was to increase the efficiency of the protocol so that increases in transmission speed would not be held back by protocol overhead.

The first major milestone in the development of 802.11n was the second draft. Although a few products were made based on 802.11n draft 1.0, interoperability was spotty. After resolving some 12,000 comments on the first draft, a much more robust second draft was produced in early 2007. With a great deal of the ambiguity removed, draft 2.0 products had the expected high degree of interoperability. Around the same time, the Wi-Fi Alliance launched a certification program to drive interoperability of the

then-emerging 802.11n standard. That certification program was a wild success. At the member meeting in the summer of 2008, the cake in [FigureA 1-1](#) was tangible (and tasty) evidence of the success of the program.



FigureA 1-1.A 802.11n draft 2.0 celebration cake

I remember when 802.11n was finally approved in September 2009. I was visiting Australia, and woke up to the anticipated news, which was that 802.11n had been approved by the standards board. As a 500-page specification, there was a great deal of work that went into the publication process. One of my favorite bits of memorabilia from my participation in the 802.11 working group is the mouse pad in [FigureA 1-2](#). I keep it around as a reminder that it's more important to get something done right than to get it done as quickly as possible.



FigureA 1-2.A 802.11n meeting schedule mouse pad

The Technology of 802.11n

In the "traditional" 802.11a/b/g world, the link layer and physical layer were separate entities. Improving speed was largely viewed as a task for the physical layer. The way to improve network speed was to make bits go

faster.

Physical Layer (PHY) Enhancements

The basis of the raw speed in 802.11n is Multiple Input/Multiple Output (MIMO) technology, which allows a single radio channel to support multiple data streams. Before 802.11n, the transmitter and receiver were Single Input/Single Output (SISO) devices. From the transmitter's antenna, the same data stream flew out in every direction, bouncing off walls and other obstacles, and then arrived at the receiver. If two paths between the transmitter antenna and the receiver antenna were out of sync, the resulting signal could be quite weak due to the interference between paths. This phenomenon, known as multipath interference, was the bane of network designers because moving an access point slightly could dramatically improve coverage. As client devices moved around, they could move from "hot spots" to "cold spots" due to multipath interference. In a MIMO system, the transmitter and receiver can take advantage of multiple paths. Each path gets a different set of data, and therefore, the resulting transmission is not subject to the same destructive effects of multipath interference. In fact, without multiple paths, the benefits of MIMO are significantly limited.

802.11n came along at an interesting point in the evolution of 802.11. Prior PHYs had been targeted at particular radio bands. The original 802.11 frequency-hopping and direct-sequence PHYs were specified only in the 2.4 GHz ISM band. 802.11a was developed when the 5 GHz band was opened for license-free use in the United States, and subsequently extended when regulations were liberalized in many other countries.

802.11g was an effort to take the technology behind 802.11a and make it available in the 2.4 GHz band. 802.11n, however, was designed when both bands were available, and therefore, it is "frequency agnostic."

Note

802.11n can use both radio bands (2.4 GHz and 5 GHz). Manufacturers use a variety of terms to tell the user what band the radio operates in. One of the more common appellations is "802.11ng" for 802.11n in the 2.4 GHz band, and "802.11na" for 802.11n devices in the 5 GHz band. Dual-band devices may be labeled in a variety of ways, such as "dual-band 802.11n" or "802.11agn."

To increase speed beyond the capabilities offered by MIMO, 802.11n offers the option for wider channels. **By doubling channel width, it is possible to double data rates.** Network administrators must carefully consider a set of trade-offs in using wider channels. In return for higher speed, radio planning becomes more complex due to a higher demand for spectrum, and coexisting with previously installed networks based on 20 MHz channels becomes a concern. Wider 40 MHz channels also have a higher potential to interfere with non-802.11 technologies such as Bluetooth, which is one of the major reasons why 40 MHz channels must be disabled by default in the 2.4 GHz band.

MIMO technology can also be enhanced for beamforming. With an antenna array, it is possible to arrange transmissions such that the energy is "focused" or "directed" towards a particular physical location. By concentrating energy in one direction, it is possible to improve the signal-to-noise

ratio and the transmit speed, though a complex set of trade-offs also limits the raw capability of beamforming.

Link Layer (MAC) Enhancements

A decade ago, it was routine to estimate the flat-out top speed of a wireless LAN by taking the "headline" rate promised by the PHY and cutting it in half. Various forms of protocol overhead, such as medium contention and positive acknowledgment, meant that the best-case scenario for 802.11 was that the 802.11 MAC had approximately 50% efficiency. The network may advertise a 54 Mbps data rate, but the best you can hope for is generally 25-27 Mbps. In many cases, the estimate of 50% efficiency was overly generous. Just as with many other markets, the wireless LAN market has vendors that produce implementations with many optional performance-enhancing features (both standards-based options and vendor-specific options), in addition to vendors that produce only the bare minimum required to ship a product.

802.11 manages contention by allocating access to the network medium for transmission, but not all transmitters are created equal. As speeds increase, the required interframe spaces have a higher cost. In an 802.11b network, the opportunity cost of a 10 μ s short interframe space (SIFS) is about 14 bytes. When the data rate increases to a moderate 802.11n rate, the cost of the same SIFS jumps to well over a kilobyte.

A significant portion of 802.11n is devoted to improving the efficiency of the MAC. While many users of 802.11 equipment focus on the high data rates, the improved efficiency is a major source of the speed improve-

ments. The main technique in 802.11n for enhancing efficiency is frame aggregation. Each transmitter needs to pay a cost to access the medium, but frame aggregation spreads that cost over several smaller frames. Depending on the type of data being transmitted, aggregation can improve efficiency from 50% to about 75%.

Although not directly related to transmission efficiency, the 802.11n MAC extends the power-saving capabilities of 802.11. 802.11n radio cards have high power consumption. While this is not a huge drawback in a device with extensive reserves of battery power such as a laptop, reducing power consumption in smaller battery-operated devices such as phones and tablets is necessary.

802.11n: First We Take the LAN, Then We Take the World

802.11n offers great power, but has significant complexity. That complexity may be expressed in a variety of ways. When the 2012 revision of 802.11 first incorporated the 802.11n amendment, the text grew in size by about 50%. Another favorite factoid of mine is that there is a clause that lays out the possible data rates for each potential value of the array of parameters selected for transmission. The resulting tables are a notable fraction of the overall specification.

Note

802.11n has over 300 different data rates! If you feel confused, you're not alone.

Bob O'Hara, the original technical editor of the 802.11 standard, once admonished me to "never confuse standardization with building a product."^[4] The complexity of 802.11n is an excellent illustration of his point. With such a complex standard, it is simply not possible to overcome the hardware engineering challenges and have the maximum capabilities available when the standard is published. In the case of 802.11n, product support of capabilities comes in waves. The first wave of 802.11n that hit in 2007 was composed of products that supported sustained 100 Mbps throughput, and, with the configuration of protocol options for maximum speed, up to a 300 Mbps link rate. The second wave hit in 2011, with products that featured data rates up to 450 Mbps. Meanwhile, many low-power devices upgraded to 802.11n single-stream technology with small improvements in speed.

The interplay between standardization and building product is, in effect, managed by the Wi-Fi Alliance, an industry trade organization. All major wireless LAN vendors are members of the Wi-Fi Alliance. With a complex standard such as 802.11, trade organizations serve a valuable purpose in focusing development effort on the most valuable features for end users.

[5]

The Wi-Fi Alliance works to create certification programs for emerging standards. As a general rule, the Wi-Fi Alliance will have an initial certification program based on a solid subset of the standard that is available as the technology is developed, and then they launch a second program

based on the approved standard. While the technology is being developed, the Wi-Fi Alliance certification programs often serve as a key target for product manufacturers, who will build products capable of passing the certification test.

[3] As with all 802.11 standards, the best place to check on status is the [official 802.11 timeline](#). In addition to estimating the milestone dates for drafts that are in progress, it offers a look back at the dates of milestones for existing standards.

[4] Or, as one of the reviewers (who you will note I have carefully not named) put it, "there are also parts of the standard that are there for political reasons to get support from particular groups, but which will never be implemented. They are put there to ensure perfection of the standard, because a standard is perfect when everybody is equally unhappy with the outcome."

[5] In the past, various industry groups such as the Frame Relay Forum, the ATM Forum, the MPLS Forum, and the DSL Forum have served similar roles in their respective technologies.

Get *802.11n: A Survival Guide* now with the O'Reilly learning platform.

O'Reilly members experience live online training, plus books, videos, and digital content from nearly 200 publishers.

START YOUR FREE TRIAL

DOWNLOAD THE O'REILLY APP

ABOUT O'REILLY[Teach/write/train](#)[Careers](#)[Community](#)[partners](#)[Affiliate program](#)[Submit an RFP](#)[Diversity](#)[O'Reilly for](#)[marketers](#)**SUPPORT**[Contact us](#)[Newsletters](#)[Privacy policy](#)**INTERNATIONAL**[Australia & New](#)[Zealand](#)[Hong Kong & Taiwan](#)[India](#)[Indonesia](#)[Japan](#)

Take O'Reilly with you and learn anywhere, anytime on your phone and tablet.

**WATCH ON YOUR BIG SCREEN**

View all O'Reilly videos, Superstream events, and Meet the Expert sessions on your home TV.

**DO NOT SELL MY PERSONAL INFORMATION****O'REILLY®**

© 2023, O'Reilly Media, Inc. All trademarks and registered trademarks appearing on oreilly.com are the property of their respective owners.

[Terms of service](#) • [Privacy policy](#) • [Editorial independence](#)