

The IEEE 802.11 Standardization Its History, Specifications, Implementations, and Future

Justin Berg
jberg2@gmu.edu

Technical Report GMU-TCOM-TR-8

Abstract

The IEEE 802.11 Standard for Wireless LANs has had a profound impact on the provision of network access and resources to dispersed, and many times varied, network elements. It has not, however, been a static system implemented since its initial ratification. The standard has been under constant amendment and updating, striving to provide new services and capabilities for expanding wireless needs, and address shortcomings in the original standard. From its humble beginnings focusing on interoperability with the broader 802.x standards to provide bridging across various media, to its ongoing search for new RF techniques and spectra for increased support to new applications, IEEE has endeavored to stay ahead of users' requirements for Wireless LAN communications.

1. History of IEEE 802.11

In 1985, the Federal Communications Commission (FCC) deregulated the spectrum from 2.4-2.5 GHz for use by the Industrial,

Scientific, and Medical (ISM) communities. This meant that the spectrum would be available for individual, non-licensed applications [1]. This news was exciting to up-and-coming developers of wireless communications technologies, because they could now develop without spending money on licensing fees. Unfortunately, this led to many developments that were far from the ubiquitous, sprawling networks we see now. At the time, and throughout the development of the 802.11 standard, if wireless network technologies were available, they were usually proprietary, expensive, slow, or simply lacked widespread availability/adaptation – and most suffered from several of these challenges [2].

In the early 1990s, however, the IEEE realized that a wireless communications infrastructure standard was necessary to meet a clearly-desirable market niche. The IEEE established an executive committee, as part of the IEEE 802 standard for Local and Metropolitan Area Networks to focus on developing a wireless LAN standard [2]. The 802.11 committee focused on providing a reliable, fast, inexpensive, robust wireless

solution that could grow into a standard with widespread acceptance, using the deregulated ISM band from 2.4-2.5 GHz.

The original standard, ultimately adopted in 1997, is vastly different from the standard that exists today. The maximum data rate was 2 Mbps. It included forward error correction, and two forms of interference mitigating spread spectrum methods – direct sequence and frequency hopping. It also included a specification for infrared wireless communications, still operating at up to 2 Mbps.

A large part of 802.11's success is its inherent compatibility with current 802 networks, specifically the 802.3 wired Ethernet networks [2]. The independence of physical access (PHY) and media access (MAC) from overlaying communication layers is critical to this compatibility. This compatibility was part of the 802.11 committee's charter [1], but its implementation played a large role in ongoing internetwork growth. The compatibility was built on two pillars – physical layer compatibility and media access layer compatibility. The separation of these layers is critical to, not only the early implementation of the standard, but the ongoing extensibility of the standard.

The physical layer portions of the original standard, and as well as today's standard, focus on allowing the base stations to get wireless broadcasts to one another; transceiving. The broadcast frequencies were in the 2.4 GHz to 2.483 GHz range or in the infrared spectrum (IR) (850-950 nm) [2]. Transmitters used time-division duplex (TDD) radio broadcasts, allowing both uplink and downlink to share the same RF channel, using differential binary phase shift keying (DBPSK) or differential quadrature phase shift keying (DQPSK) signal modulation (Appendix A). Transmitters used either

Direct Sequence Spread Spectrum (DSSS) or Frequency Hopping Spread Spectrum (FHSS) for interference mitigation. Data rates were specified for both 1 Mbps and 2 Mbps operation.

The media access layer (MAC) processes the PHY layer signals into the ubiquitous network layer. Fundamentally, 802.11 uses collision sense media access with collision avoidance (CSMA/CA) for its media access protocol (Appendix B) [2]. The MAC layer also provides several services to assist in the wireless broadcast such as synchronization, power management, frame fragmentation, and frame encryption (WEP - Wired Equivalent Privacy) and authentication, with varying methods of employing these services for both infrastructure-based in distributed (known as ad-hoc) networks. For example, in an infrastructure network, synchronization is performed between all transceivers by using beacons transmitted by the access point. In an ad-hoc network, however, the synchronization responsibility falls to all members of the independent network, creating a sub-network of synchronizers.

Note that there is no 5 GHz spectrum specification in the original 802.11-1997 standard. This frequency allocation was not explored (or at least, published) until shortly after the original standard was adopted. The original standard focused on exploiting the recently-unlicensed 2.4 GHz ISM band, and the practical, and already-in-use infrared spectrum. In fact, the original standard largely overlooks, or at least actively ignores, many compatibility standards that would end up being crucial to widespread acceptance of the standard. For example, the entire standard makes only cursory mention of MAC address space, pointing out that its 48-bit address space is compatible within the broader scope of the IEEE 802 address space,

but is not required to be unique from a global 802 address overlay. This compatible address space, which is still a part of the 802.11 standard today, allows 802.11 networks to interact with the 802.1 LAN specification that provides for bridging between separate physical networks, and is perhaps the cornerstone of the success for the standard. This address compatibility with 802.x networks (and flexibility) played a role in the widespread adoption and interoperability of 802.11 wireless networks [2], even in the face of other, higher-speed competing network standards such as HiperLAN, a competing European standard for wireless network communications, which provided its own convergence to internet protocol (IP) networks, vice relying on 802.1 for internetwork bridging [1].

Despite not having addressed direct compatibility of the 802.11 with 802 networks, the committee left the door open, and in fact immediately fostered the follow on Task Groups to address specific supplemental topics for use within the 802.11 standard framework. The 802.11b task group, TGb, addressed higher speed transmissions within the WLAN environment. The 802.11b Task Group produced the 802.11b amendment, adopted by IEEE in 1999, just two years after the original standard was adopted. It allows for 5.5 Mbps and 11 Mbps data rates, using Direct Sequence Spread Spectrum (DSSS) transmissions [2]. It also prompted the creation of the Wireless Ethernet Compatibility Alliance (WECA); a non-profit association for standardization and promotion of Wi-Fi technologies. From wi-fi.org [5]:

“The Wi-Fi Alliance is a global non-profit industry association of hundreds of leading companies devoted to seamless connectivity. With technology

development, market building, and regulatory programs, the Wi-Fi Alliance has enabled widespread adoption of Wi-Fi worldwide.”

Even today, 802.11b is probably the most widely-recognized, and widely-used 802.11 standard, although 802.11g is quickly surpassing it, with 802.11n up-and-coming in popularity and availability. WECA renamed itself to the Wi-Fi Alliance in October, 2002 [4].

About the same time the 802.11b Task Group was designing the 802.11b amendment, the 802.11a Task Group, TGa, was doing the same for another wireless standard [3]. At the time, many countries had recently opened up some 5 GHz spectrum for unlicensed (but still regulated) use. This spectrum was less “RF dense” than the 2.4 GHz spectrum [2], which includes other interferors such as garage door openers, cordless telephones, microwave ovens, and baby monitors. With less interference high bandwidth available, another, higher capacity standard could be constructed.

The ultimate 802.11a standard included a 54 Mbps data rate using the more-complex orthogonal frequency division multiplexing (OFDM) waveforms (Appendix C), and operated in the 5 GHz range, set aside for the Unlicensed National Information Infrastructure (U-NII) usage [1]. While the standard was completed and adopted in 1999, the more-complex equipment did not begin shipping until 2001.

It is significant to note that while data rates were increased by both 802.11a and 802.11b, that both only increased data bandwidth within RF applications. The IR specification, while still valid, was left behind with 1-2 Mbps maximum throughput, while the RF environment has continued to

increase in data throughput throughout the development of the 802.11 standard.

Not long after 802.11a was adopted, IEEE immediately recognized that the OFDM waveform could benefit the 802.11b standard. Increased data rates would even support bandwidth-hungry multimedia applications as the demand for these applications grew. In July 2000, the 802.11 Task Force G was assigned the task of overlaying the OFDM waveform on the 2.4 GHz spectrum, producing a new standard that was fully backward-compatible with the 802.11b standard. This was no easy feat, but after 3 years the new standard was ratified. The key was in requiring all 802.11g equipment to support complimentary code key (CCK) modulation as a fall-back mechanism to ensure 802.11b compatibility. This fall-back has significant impacts on the total data rate of the network, but allows mixed 802.11b-802.11g network equipment to coexist on the same topology. As 802.11b equipment is phased out and replaced with 802.11g equipment, users can seamlessly upgrade their network without upgrading the entire infrastructure. In June 2003, the amendment was ratified.

As 802.11 enjoyed widespread adoption by home and business users alike, more scrutiny was placed on security. The initial standard included a MAC-level security protocol called WEP, Wired Equivalent Privacy [6]. WEP was intended to provide confidentiality and authentication for connecting users. By using a very small subset (up to four) of pre-shared keys, a user could identify itself as a valid user to an access point, and encrypt every packet of the session [7]. The intent of WEP was not to be a bulletproof security protocol for wireless networks, but to provide reasonable session privacy, like that which could be expected from a direct-connection (wired) connection.

Unfortunately, WEP was rife with vulnerabilities (Appendix D), and continued bad press caused 802.11 users to demand better security [7]. Another task group, Task Group I, was set up to address MAC-level security in an effort to address security problems with WEP [6].

The Task Group model, however, took too long to address the concerns of equipment manufacturers. The Wi-Fi Alliance began implementing additional security enhancements to provide customers with additional security features. Many members of the Wi-Fi Alliance were part of Task Group I, and these enhancements would be seen as part of the final 802.11i amendment. These original security implementations, labeled Wireless Protected Access included many enhancements to address the weaknesses of WEP, including the use of extended initialization vectors (IV) (56-bits), rotating initialization vectors, more robust integrity checks, and protection against replay/redirection attacks [6].

In June 2004, the 802.11i amendment was ratified. The security enhancements in it became known as WPA2, Wireless Protected Access v2. It was largely a mirror of the WPA enhancements from the Wi-Fi Alliance, with some small, but significant, improvements. First, it incorporated the use of the Advanced Encryption Standard for encrypting and protecting data [8]. The AES was selected/adopted by the National Institute of Standards and Technology (NIST) in November 2001, and was not available when WEP was being designed nearly 10 years earlier. Next, enhanced integrity checks leveraging the AES CCMP (counter mode with cipher block chaining with message authentication code protocol, a recursive acronym) provides additional authentication. 802.11i also supports several implementations of using external

authentication mechanisms, including 802.1X authentications and/or RADIUS [8].

Meanwhile, the IEEE was going through another exercise to increase wireless data rates. Recognizing the seemingly unquenchable bandwidth thirst of users, the IEEE set out to exceed 54 Mbps as an upper data rate limit by creating Task Group n (TGn) in September 2003 [3]. By using multiple-input multiple-output (MIMO) transmitting methods, 802.11n would allow multiple data streams, separated spatially, to increase the overall data rate [9]. This access method, as with 802.11g, is backward-compatible with previous 2.4 GHz implementations of 802.11, as well as 802.11a in the 5 GHz and 3.7 GHz spectra (802.11a was extended to 3.7 GHz by the 802.11y amendment in Nov 2008) [9].

While 2.4 GHz implementations include the largest number of users worldwide, unfortunately the 2.4 GHz spectrum is heavy on interference. While MIMO can provide additional and higher data rates, and protection against some interferences (Appendix E), there is a limit as to how much data can be transferred in the congested spectrum. The 802.11n amendment, ratified in September 2009, can support data rates up to 600 Mbps, but in its current implementation, with the congestion in the 2.4 GHz spectrum, the maximum supported transmission rate is 104 Mbps. This is still a significant increase over the 802.11g amendment, but leaves significant room for growth, should 802.11n be deployed in other RF environments. Indeed, as the amendment does not specify the exact spectrum, the largest performance gains will be realized in the 5 GHz and 3.7 GHz ranges, where significantly less interference is found.

Also, in 2003, the IEEE began consolidating the standard amendments from the Task Groups, and rolling them into a consolidated baseline standard. The standard

up to this point had been known as IEEE 802.11-1999. After many amendments, 802.11a, 802.11b, 802.11d, 802.11e, 802.11g, 802.11h, 802.11i, 802.11j, were all rolled into one consolidated standard – IEEE 802.11-2007. While technically these amendments no longer exist, as they are now part of the baseline standard, most still refer to them by their parent amendment designations to easily identify the specific capability or function of the 802.11 wireless LAN standard [9]. The following illustration is a timeline of when specific amendments were initiated (i.e., the Task Group was formed), and when the amendment was ratified, terminated, or rolled into the core standard.

2. The 802.11 Standard

The 802.11 standard, while a single standard, has many manifestations that allow wireless network access. It covers everything from how synchronization should be performed, to how infrared (IR) wireless networks should be configured, to spread spectrum chip rates for different applications. This paper cannot touch on all portions of the standard. Indeed, the 1200+ page standard (not including its many several-hundred page amendments/enhancements) will require this paper leave many topics unexplored, and many, many more topics completely undiscovered.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.