

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. Shipments of Wi-Fi certified integrated circuits exceeded a billion units per year in 2011 (ABIresearch, 2012) and are expected to exceed 2.5 billion units per year by 2016 (ABIresearch, 2012), 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.

There has been a proliferation of small scale deployments providing Internet access in coffee shops, airports, hotels, etc., which have come to be known as hotspots. In recent

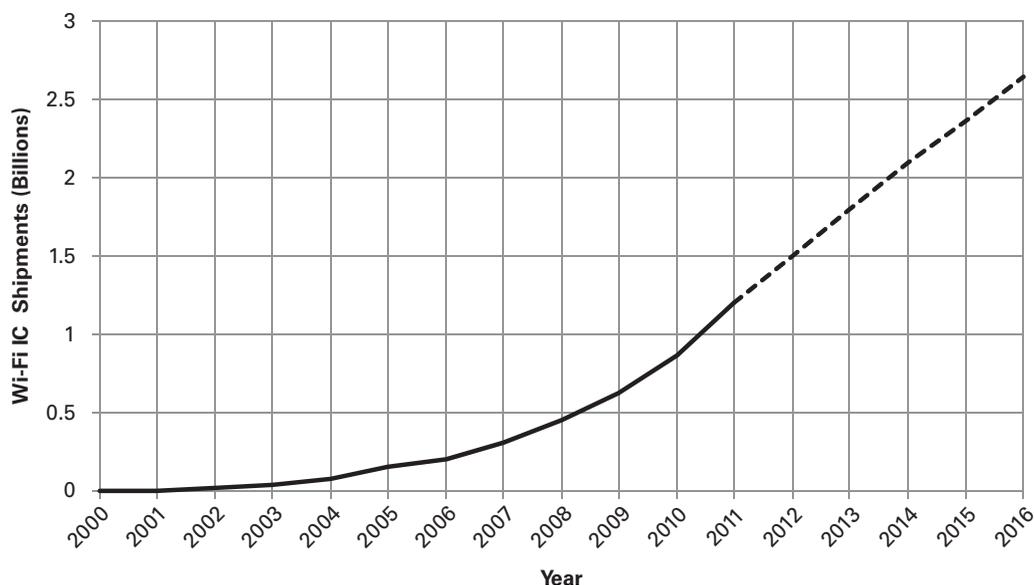


Figure 1.1 Wi-Fi IC shipments. Source: ABIresearch (2007, 2012).

years, carriers with heavily congested cellular networks are deploying hotspots to off-load traffic from their cellular networks. Additionally, when these networks are used in conjunction with virtual private network (VPN) technology, employees can securely access corporate networks from almost anywhere.

WLAN products and systems started with 802.11b, 802.11g, and 802.11a standard amendments, which provided throughput enhancements over the original 802.11 PHYs. Progress in WLAN technology continued with the development of 802.11n. Increased data rates were 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” 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 Wi-Fi Alliance (WFA) to begin interoperability testing and certification of devices based on a subset of the 802.11n D2.0 features in May 2007. Wi-Fi certified 802.11n D2.0 products provide consumers the assurance of interoperability between manufacturers that was not guaranteed by previous “pre-n” products. At the end of 2009, 802.11n was finally approved and the WFA updated the certification program to reflect support for the approved standard. Full interoperability was maintained between 802.11n D2.0 and the approved standard products. These were major steps in speeding up the standardization and certification process of new technology.

As this process was successful for the industry and beneficial for the consumers, 802.11ac will follow a similar path. It is expected that 802.11ac products based on an early draft will be certified and on the market in early 2013. Completion of the 802.11ac is expected by the end of 2013.¹ Certification based on the approved standard will take place in a similar time frame.

1.1 An overview of IEEE 802.11

The IEEE 802.11 standard defines multiple physical layers (PHYs) and a common medium access control (MAC) layer for 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 inherits the 802 reference model and 48-bit universal addressing scheme. The 802 reference model is based on the OSI reference model described in Table 1.1. In this model, the 802.11 MAC and 802.2 logical link control (LLC) sublayers form the data link layer and the 802.11 PHY the physical layer.

1.1.1 The 802.11 MAC

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

¹ The reader is referred to http://grouper.ieee.org/groups/802/11/Reports/802.11_Timelines.htm for the latest update on the timeline of 802.11ac

Table 1.1 OSI reference model (adapted from ISO/IEC 7498-1, 1994)

OSI reference model layers	Description	Examples	Layer categories
Application	Interacts with the software applications that implement a communicating component	HTTP, FTP, SMTP	Application
Presentation	Establishes context between application-layer entities	MIME, TLS, SSL	
Session	Establishes, manages, and terminates communication sessions	Named pipe, NetBIOS	
Transport	Provides an end-to-end reliable data transfer service, including flow control, segmentation/desegmentation and error control	TCP, UDP	
Network	Provides the means for transferring variable length data sequences from a source device to a destination device. Maintains the quality of service requested by the transport layer	IP (IPv4, IPv6), ICMP, IPsec	
Data link	Provides the means for transferring data between devices	LLC	Data transport
Physical	Provides the electrical and physical specifications for devices	802.11 MAC 802.11 PHY	

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.

1.1.2 The 802.11 PHYs

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.

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 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. Furthermore, in the 5 GHz band, 802.11ac provides a data rate of 1733 Mbps with 80 MHz and four spatial streams, and a maximum data rate of 6933 Mbps with 160 MHz and eight spatial streams. The exponential increase in data rate is illustrated in Figure 1.2.

Table 1.2 Overview of 802.11 PHYs

	802.11	802.11b	802.11a	802.11g	802.11n	802.11ac
PHY technology	DSSS	DSSS/ CCK	OFDM	OFDM DSSS/ CCK	SDM/OFDM	SDM/OFDM MU-MIMO
Data rates (Mbps)	1, 2	5.5, 11	6–54	1–54	6.5–600	6.5–6933.3
Frequency band (GHz)	2.4	2.4	5	2.4	2.4 and 5	5
Channel spacing (MHz)	25	25	20	25 MHz	20 and 40	20, 40, 80, and 160

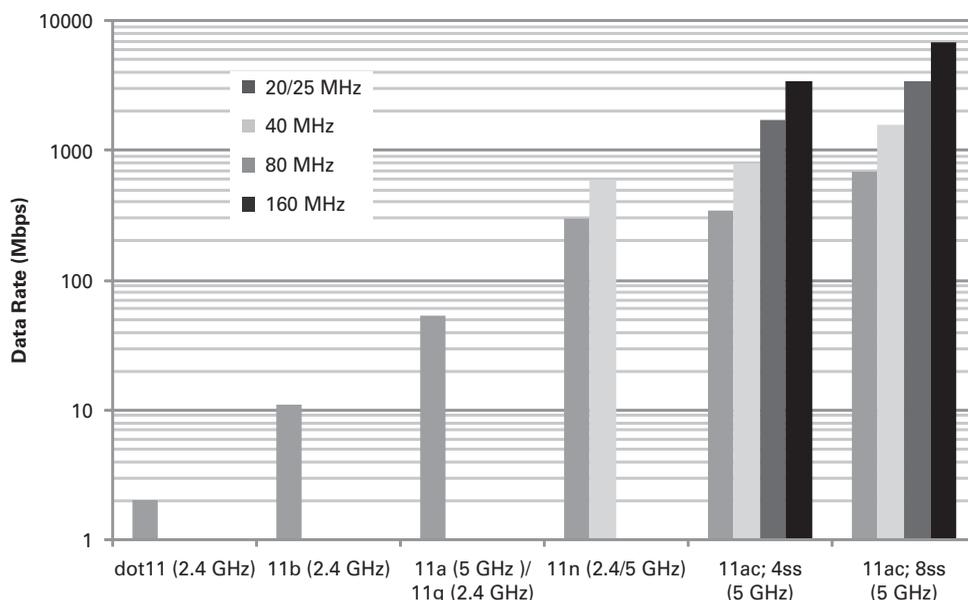


Figure 1.2 Increase in 802.11 PHY data rate.

1.1.3 The 802.11 network architecture

The basic service set (BSS) is the basic building block of an 802.11 LAN. Stations that remain within a certain coverage area and form some sort of association form a BSS. The most basic form of association is where stations communicate directly with one another in an ad-hoc network, referred to as an independent BSS or IBSS. This is illustrated as BSS 1 in Figure 1.3.

More typically, however, stations associate with a central station dedicated to managing the BSS and referred to as an access point (AP). A BSS built around an AP is called an infrastructure BSS and is illustrated by BSS 2 and BSS 3 in Figure 1.3. Infrastructure BSSs may be interconnected via their APs through a distribution system (DS).

The BSSs interconnected by a DS form an extended service set (ESS). A key concept of the ESS is that stations within the ESS can address each other directly at the MAC layer. The ESS, being an 802.11 concept, encompasses only the 802.11 devices and does not dictate the nature of the DS. In practice, however, the DS is typically an Ethernet (802.3) LAN and the AP functions as an Ethernet bridge. As such, stations in a BSS can also directly address stations on the LAN at the MAC layer.

1.1.4 Wi-Fi Direct

Recognizing the need for improved peer-to-peer operation, the Wi-Fi Alliance has developed a specification for direct communication between Wi-Fi devices without being associated with an infrastructure BSS. Such communication is possible using an

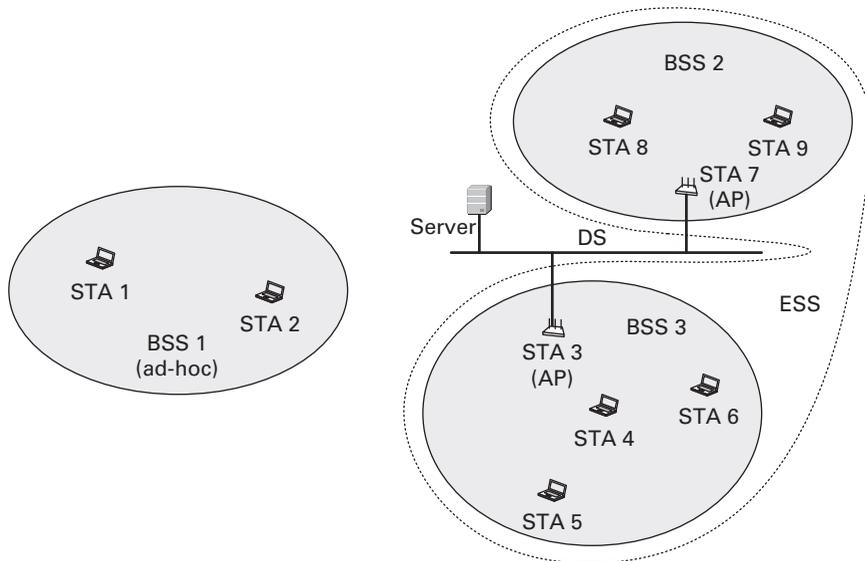


Figure 1.3 BSS, DS, and ESS concepts.

independent BSS, as defined in the 802.11 specification; however, it was preferable to create a mode of operation closer to that of the infrastructure BSS.

In a Wi-Fi Direct network, one device, called the group owner (GO), assumes a role similar to that of an AP while the other devices associate with that device as they would an AP. The Wi-Fi Direct network is thus similar to an infrastructure BSS except that (1) the GO does not provide access to a distribution system, and (2) like its peers, the GO could be a mobile, battery powered device, and thus also need to enter a low power sleep state when idle.

The Wi-Fi Direct standard builds on the 802.11 specification, specifying protocols by which devices can discover each other, how a device assumes the role of group owner and the protocol for absence from the session channel (for power management or to visit an infrastructure BSS channel).

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 Airgo Networks, 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).

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.