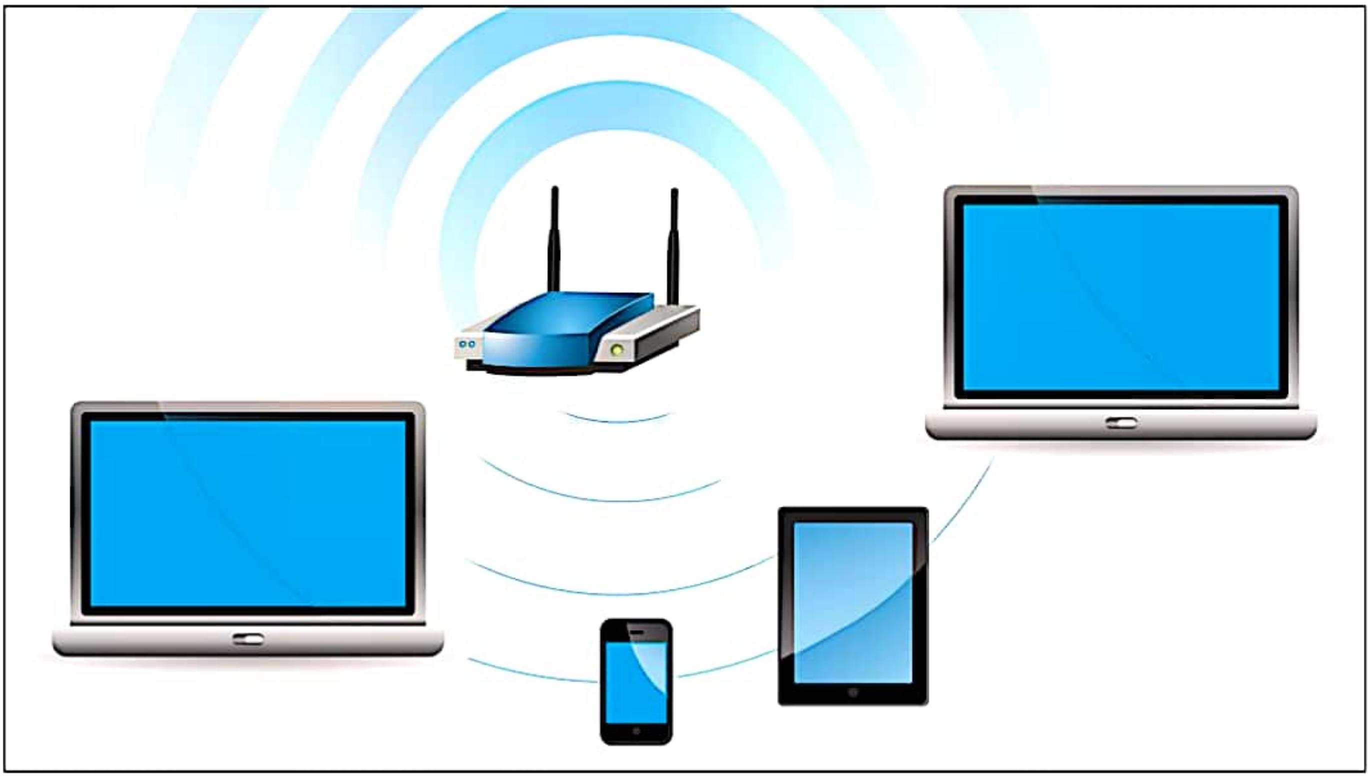

Figure 1. The 802.11 standards have enabled millions of electronic devices to exchange data or connect to the internet wirelessly using radio waves.

Introduction

Wi-Fi is a technology that allows many electronic devices to exchange data or connect to the internet wirelessly using radio waves. The Wi-Fi Alliance defines Wi-Fi devices as any "Wireless Local Area Network (WLAN) products that are based on the Institute of Electrical and Electronics Engineers' (IEEE) 802.11 standards".

The key advantage of IEEE 802.11 devices is that they allow less expensive deployment of Local Area Networks (LANs). For places where running cables to every device is not practical, such as outdoor areas and airports, they can host wireless LANs. Products from every brand name can inter-operate at a basic level of service thanks to their products being designated as "Wi-Fi Certified" by the Wi-Fi Alliance.

Today, millions of IEEE 802.11 devices are in use around the world and they operate in the same frequency bands, this makes the need for their coexistence critical. Even though over time older devices will be retired, some consumers and businesses will still be using the old standards for years. For some businesses the original 802.11b devices meet their needs and the need to change has not occurred. Wider bandwidth 802.11 deployments must therefore be able to "play nicely" with the older standards, both by limiting their impact on nearby legacy WLANs and by enabling communication with legacy stations.

This primer provides a general overview for each of the 802.11 standards, their PHY layer characteristics and their testing requirements. In this document, we use 802.11 and IEEE 802.11 interchangeably.

IEEE 802.11 Standard and Formats

IEEE 802 refers to a family of IEEE standards dealing with Local Area Networks and Metropolitan Area Networks (Table 1). The IEEE 802 family of standards is maintained by the IEEE 802 LAN/MAN Standards Committee (LMSC). An individual Working Group provides the focus for each area.

IEEE 802.11 is a set of medium access control (MAC) and physical layer (PHY) specifications for implementing Wireless Local Area Network (WLAN) communication. The 802.11 family is a series of over-the-air modulation techniques that share the same basic protocol (Table 2). These standards provide the basis for wireless network products using the Wi-Fi brand. The segment of the radio frequency spectrum used by 802.11 varies between countries.

IEEE 802.11-1997 or Legacy Mode

The original version of the standard IEEE 802.11 was released in 1997, but is basically obsolete today. It specified bit rates of 1 or 2 megabits per second (Mbit/s). It specified three alternative physical layer technologies:

- Diffuse infrared operating at 1 Mbit/s
- Frequency-hopping spread spectrum operating at 1 Mbit/s or 2 Mbit/s
- Direct-sequence spread spectrum operating at 1 Mbit/s or 2 Mbit/s

The latter two radio technologies used microwave transmission over the Industrial Scientific Medical (ISM) frequency band at 2.4 GHz. Its specified data rate was to be transmitted via infrared (IR) signals or by either frequency hopping or direct-sequence spread spectrum (DSSS) radio signals. IR remains a part of the standard but has no actual implementations.

A weakness of this original specification was that it offered so many choices that interoperability was sometimes challenging. It is really more of a "beta-specification" than a rigid specification, initially allowing individual product vendors the flexibility to differentiate their products but with little to no inter-vendor operability.

IEEE 802 Standards	
802.1	Bridging & Management
802.2	Logical Link Control
802.3	Ethernet - CSMA/CD Access Method
802.4	Token Passing Bus Access Method
802.5	Token Ring Access Method
802.6	Distributed Queue Dual Bus Access Method
802.7	Broadband LAN
802.8	Fiber Optic
802.9	Integrated Services LAN
802.10	Security
802.11	Wireless LAN
802.12	Demand Priority Access
802.14	Medium Access Control
802.15	Wireless Personal Area Networks
802.16	Broadband Wireless Metro Area Networks
802.17	Resilient Packet Ring

Table 1. 802 Family of Standards.

The DSSS version of legacy 802.11 was rapidly supplemented (and popularized) by the 802.11b amendment in 1999, which increased the bit rate to 11 Mbit/s. Wide spread adoption of 802.11 networks only occurred after the release of 802.11b. As a result few networks were implemented using the original 802.11-1997 standard. In this document several sections do not provide further detail about the original legacy mode for this reason.

IEEE 802.11b

802.11b has a maximum raw data rate of 11 Mbit/s and uses the same media access method defined in the original legacy standard. 802.11b products appeared on the market in early 2000 and is a direct extension of the modulation technique defined in the original standard. The dramatic increase in throughput of 802.11b (compared to the original standard) along with substantial price reductions led to the rapid acceptance of 802.11b as the definitive wireless LAN technology.

IEEE 802.11 PHY Standards						
Release Date	Standard	Frequency Band (GHz)	Bandwidth (MHz)	Modulation	Advanced Antenna Technologies	Maximum Data Rate
1997	802.11	2.4 GHz	20 MHz	DSSS, FHSS	N/A	2 Mbits/s
1999	802.11b	2.4 GHz	20 MHz	DSSS	N/A	11 Mbits/s
1999	802.11a	5 GHz	20 MHz	OFDM	N/A	54 Mbits/s
2003	802.11g	2.4 GHz	20 MHz	DSSS, OFDM	N/A	542 Mbits/s
2009	802.11n	2.4 GHz, 5 GHz	20 MHz, 40 MHz	OFDM	MIMO, up to 4 spatial streams	600 Mbits/s
2013	802.11ac	5 GHz	40 MHz, 80 MHz, 160 MHz	OFDM	MIMO, MU-MIMO, up to 8 spatial streams	6.93 Gbits/s

Table 2. IEEE 802.11 PHY Standards.

One disadvantage of the 802.11b devices is that they may have interference issues with other products operating in the 2.4 GHz band. Devices operating in the 2.4 GHz range include microwave ovens, cordless phones, Bluetooth devices, baby monitors and some amateur radio equipment. Interference issues and user density problems within the 2.4 GHz band have become a major issue as the popularity of Wi-Fi has grown.

IEEE 802.11a

The 802.11a standard was added to the original standard and was ratified in 1999. The 802.11a standard uses the same core protocol as the original standard and was the first of the 802.11 family to operate in the 5 GHz band. It uses a 52-subcarrier orthogonal frequency-division multiplexing (OFDM) with a maximum raw data rate of 54 Mbit/s, which typically yields a throughput in the mid-20 Mbit/s. Today, many countries around the world are allowing operation in the 5.47 to 5.725 GHz Band. This will add more channels to the overall 5 GHz band enabling significant overall wireless network capacity. 802.11a is not interoperable with 802.11b since they operate on different frequency bands. However, most enterprise class Access Points have multi-band capability today.

Using the 5 GHz band gives 802.11a a significant advantage, since the 2.4 GHz ISM band is heavily used. Degradation caused by such conflicts can cause frequently dropped connections and degradation of service. However, the higher 5 GHz frequency also brings a slight disadvantage as the effective range of 802.11a is slightly less than that of 802.11b/g. 802.11a signals cannot penetrate as far as those for 802.11b because they are absorbed more readily by walls and other solid objects in their path and because the path loss in signal strength is proportional to the square of the signal frequency. On the other hand, OFDM has fundamental propagation advantages when in a high multipath environment, such as an indoor office, and the higher frequencies enable the building of smaller antennas with higher RF system gain, which counteract the disadvantage of a higher band of operation. The increased number of usable channels and the near absence of other interfering systems (microwave ovens, cordless phones, baby monitors) give 802.11a a significant bandwidth and reliability advantage over 802.11b/g.

Confusion on the release time of 802.11a and 802.11b is common. The 802.11a products started shipping late, lagging 802.11b products due to 5 GHz components being more difficult to manufacture. In addition, first generation product performance was poor and plagued with problems. When second generation products started shipping, 802.11a was not widely adopted in the consumer space primarily because the less-expensive 802.11b was already widely adopted. However, 802.11a later saw significant penetration into enterprise network environments, despite the initial cost disadvantages, particularly for businesses which required increased capacity and reliability over 802.11b/g-only networks. Sections in this document often lead with 802.11b for this reason.

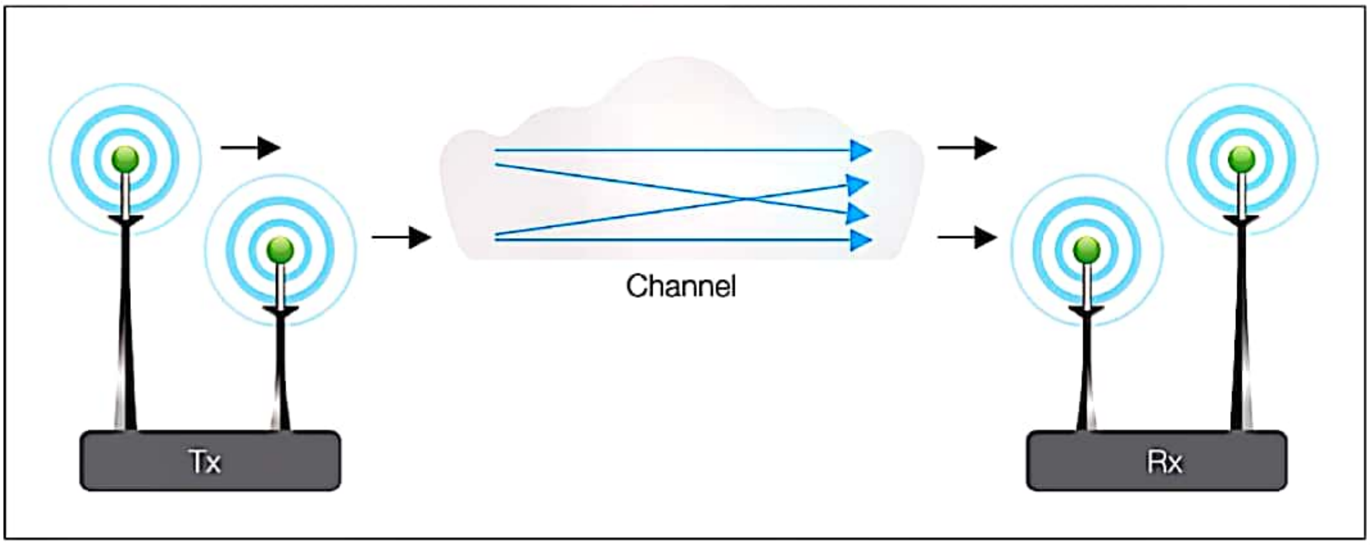

Figure 2. MIMO uses multiple antennas to coherently resolve more information than possible using a single antenna.

IEEE 802.11g

The 802.11g standard was rapidly adopted by consumers starting in January 2003, well before ratification, due to the desire for higher speeds and reductions in manufacturing costs. By summer 2003, most dual-band 802.11a/b products became dual-band/tri-mode, supporting 802.11a and b/g in a single mobile adapter card or access point.

802.11g works in the 2.4 GHz band (like 802.11b), but uses the same OFDM based transmission scheme as 802.11a. It operates at a maximum physical layer bit rate of 54 Mbit/s, exclusive of forward error correction codes. 802.11g hardware is fully backwards compatible with 802.11b hardware. In an 802.11g network, however, the presence of a 802.11b device will significantly reduce the speed of the overall 802.11g network.

Despite its major acceptance, 802.11g suffers from the same interference as 802.11b in the already crowded 2.4 GHz range. Additionally, the success of the standard has caused usage/density problems related to crowding in urban areas. To prevent interference, there are only three non-overlapping usable channels in the U.S. and other countries with similar regulations (channels 1, 6, 11, with 25 MHz separation), and four in Europe (channels 1, 5, 9, 13, with only 20 MHz separation). Even with such separation, some interference due to side lobes exists, though it is considerably weaker.

IEEE 802.11n

The 802.11n amendment includes many enhancements that improve WLAN range, reliability, and throughput. At the physical (PHY) layer, advanced signal processing and modulation techniques have been added to exploit multiple antennas and wider channels. At the Media Access Control (MAC) layer, protocol extensions make more efficient use of available bandwidth. Together, these High Throughput (HT) enhancements can boost data rates up to 600 Mbps – more than a ten-fold improvement over 54 Mbps 802.11a/g.

802.11n operates on both the 2.4 GHz and the 5 GHz bands. Support for 5 GHz bands is optional. IEEE 802.11n builds on previous 802.11 standards by adding multiple-input multiple-output (MIMO) and 40 MHz channels to the PHY layer, and frame aggregation to the MAC layer.

Behind most 802.11n enhancements lies the ability to receive and/or transmit simultaneously through multiple antennas. 802.11n defines many "M x N" antenna configurations, ranging from "1 x 1" to "4 x 4". MIMO uses multiple antennas to coherently resolve more information than possible using a single antenna. One way it provides this is through Spatial Division Multiplexing, which spatially multiplexes multiple independent data streams, transferred simultaneously within one spectral channel of bandwidth. MIMO can significantly increase data throughput as the number of resolved spatial data streams is increased. Each spatial stream requires a discrete antenna at both the transmitter and the receiver.

Figure 3. The broad acceptance and success of 802.11 devices have created the need for new usage models which require higher throughput.

The number of simultaneous data streams is limited by the minimum number of antennas in use on both sides of the link. However, the individual radios often further limit the number of spatial streams that may carry unique data. The M x N = Z notation helps identify the capability of a given radio. The first number M is the maximum number of transmit antennas that can be used by the radio. The second number N is the maximum number of receive antennas that can be used by the radio. The third number Z is the maximum number of data spatial streams the radio can use. For example, a radio that can transmit on two antennas and receive on three, but can only send or receive two data streams would be 2 x 3 : 2.

Another optional 802.11n feature is the 40 MHz channels. Prior 802.11 products use channels that are approximately 20 MHz wide. 802.11n products have the option to use 20 or 40 MHz wide channels, providing the AP has 40 MHz capability as well. Channels operating with a bandwidth of 40 MHz provide twice the PHY data rate available over a single 20 MHz channel. The wider bandwidth can be enabled in either the 2.4 GHz or the 5 GHz mode, but must not interfere with any other 802.11 or non-802.11 (such as Bluetooth) system using the same frequencies.

IEEE 802.11ac

The early standards for wireless LAN were designed primarily to connect a laptop PC in the home, office, and to allow connectivity "on the road". The broad acceptance and success of WLAN has created the need for new usage models which would require higher throughput, such as:

- Wireless display
- In-home distribution of HDTV and other content
- Rapid upload/download of large files to/from servers
- Backhaul traffic (mesh, point-to-point, etc.)
- Campus and auditorium deployments
- Manufacturing floor automation

IEEE 802.11ac (aka VHT, Very High Throughput) is a standard under development which will provide throughput in the 5 GHz band. 802.11ac plans to re-use 802.11n (and 802.11a) details where possible. This is advantageous for ensuring backwards compatibility and co-existence and also allows the 802.11ac developers to focus on the new features that are needed to achieve the throughput requirements.

The 802.11ac specification has expected multi-station WLAN throughput of at least 1 Gbps and a single link throughput of at least 500 Mbps. This is accomplished by extending the air interface concepts embraced by 802.11n:

- Wider RF bandwidth (up to 160 MHz)
- More MIMO spatial streams (up to 8)
- Multi-user MIMO
- High-density modulation (up to 256-QAM).

The standard was developed from 2011 through 2013, with final 802.11 Working Group approvals and publication scheduled for early 2014.

All 802.11ac devices are required to support 20, 40, and 80 MHz channels and 1 spatial stream. In addition, several optional features are also defined in 802.11ac:

- Wider channel bandwidths (80+80 MHz and 160 MHz)
- Higher modulation support (optional 256QAM)
- Two or more spatial streams (up to 8)
- Multi-User MIMO (MU-MIMO)
- 400 ns short guard interval
- Space Time Block Coding (STBC)
- Low Density Parity Check (LDPC)

802.11ac devices making use of only the mandatory parameters (80 MHz bandwidth, 1 spatial stream, and 64-QAM 5/6) will be capable of a data rate of approximately 293 Mbps. Devices that take advantage of the optional parameters (8 spatial streams, 160 MHz of bandwidth and 256-QAM 5/6 with a short guard interval) will be able to achieve almost 7 Gbps.

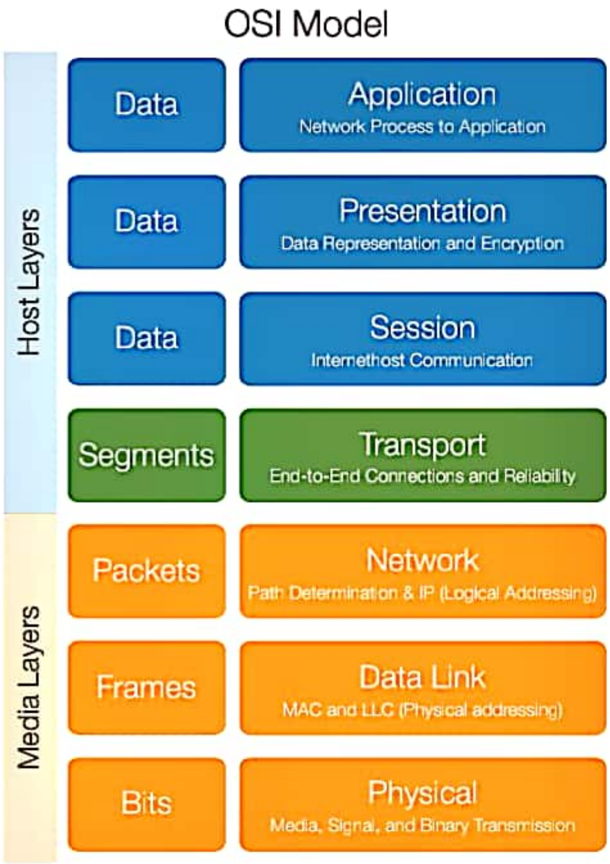

Figure 4. The OSI model describes how information moves from an application program running on one networked computer to an application program running on another networked computer.

Protocol Architecture Overview

The Open Systems Interconnection Reference Model, or the OSI model, was developed by the International Organization for Standardization, which uses the abbreviation of ISO. The OSI model is a layered model that describes how information moves from an application program running on one networked computer to an application program running on another networked computer. In essence, the OSI model prescribes the steps to be used to transfer data over a transmission medium from one networked device to another. The OSI model defines the network communications process into seven separate layers as shown in Figure 4.

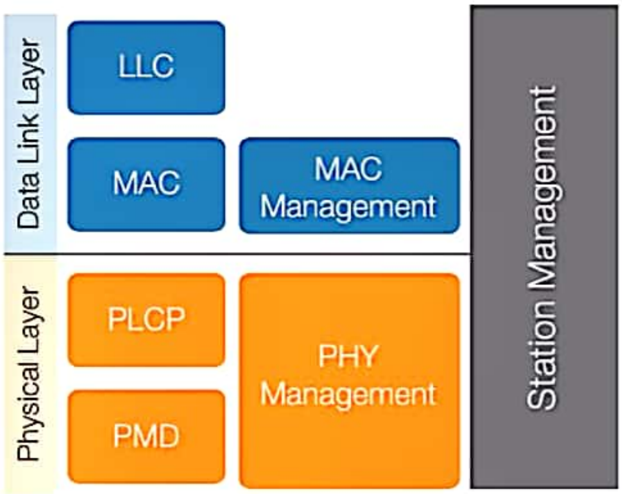

Figure 5. The 802.11 standards focus on the Data Link and Physical Layers of the OSI reference model.

Practical, connectionless LANs began with the pre-IEEE Ethernet specification, which is the ancestor of IEEE 802.3. The Standard 802.11 covers protocols and operation of wireless networks. It only deals with the two lowest layers of the OSI reference model, the physical layer and the Data Link layer (or Media Access Control layer). The goal is for all the 802.11 series of standards to be backward compatible and to be compatible at the Medium Access Control (MAC) or Data Link layer. Each of the 802.11 standards would therefore only differ in physical layer (PHY) characteristics (Figure 5).

The MAC layer provides the functional and procedural means to transfer data between network entities and to detect and possibly correct errors that may occur in the physical layer. It provides access to contention based and contention-free traffic on different kinds of physical layers. In the MAC layer the responsibilities are divided into the MAC sub-layer and the MAC management sub-layer. The MAC sub-layer defines access mechanisms and packet formats. The MAC management sub-layer defines power management, security and roaming services.

The Physical Layer defines the electrical and physical specifications for devices. In particular, it defines the relationship between a device and a transmission medium. The major functions and services performed by the physical layer are the following:

- Establishment and termination of a connection to a communications medium.
- Participation in the process where the communication resources are effectively shared among multiple users. For example, contention resolution and flow control.
- Modulation or conversion between the representation of digital data in user equipment and the corresponding signals transmitted over a communications channel. These are signals operating over the physical cabling (such as copper and optical fiber) or over a radio link.

The Physical layer is divided into three sub layers.

1. The Physical Layer Convergence Procedure (PLCP) acts as an adaption layer.
2. The PLCP is responsible for the Clear Channel Assessment (CCA) mode and building packets for different physical layer technologies.
3. The Physical Medium Dependent (PMD) layer specifies modulation and coding techniques. The PHY management layer takes care of the management issues like channel tuning.

The Station management sub layer is responsible for co-ordination of interactions between the MAC and PHY layers.

This primer focuses on the PHY layer as this is where the design requirements for the device hardware using the different techniques of the 802.11 standards are realized.

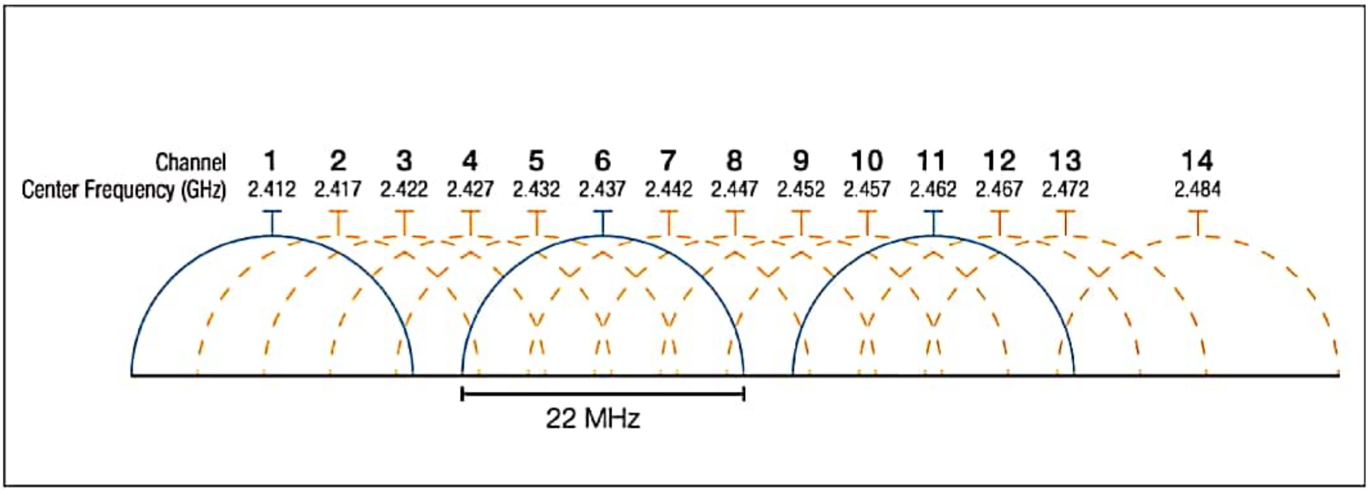

Figure 6. The 2.4 GHz band is divided into 14 overlapping channels.

Channel Allocations and Spectral Masks

The 802.11b, 802.11g, and the low frequency part of the 802.11n standards utilize the 2.400 – 2.500 GHz spectrum located in the ISM band. The 802.11a, 802.11n and 802.11ac standards use the more heavily regulated 4.915 – 5.825 GHz band. These are often referred to as the "2.4 GHz and 5 GHz frequency bands". Each of these spectrums are sub-divided into channels with a center frequency and bandwidth, similar to the way commercial spectrums are sub-divided.

The 2.4 GHz band is divided into 14 channels spaced 5 MHz apart, beginning with channel 1 which is centered on 2.412 GHz (Figure 6). The channel numbering of the 5.725 – 5.875 GHz spectrum is less intuitive due to the differences in regulations between countries.

Channel Bandwidths

Early 802.11 products use channels that are approximately 20 MHz wide. In the US, 802.11b/g radios use one of eleven 20 MHz channels (three non-overlapping: 1, 6, 11) within the 2.4 GHz ISM frequency band. When the OFDM PHY was introduced to 802.11, the channel bandwidth was 20 MHz with later amendments adding support for 5 and 10 MHz bandwidths. 802.11a radios use one of twelve non-overlapping 20 MHz channels in the 5 GHz Unlicensed National Information Infrastructure (UNII) band. 802.11n products can use 20 or 40 MHz wide channels in either the ISM or UNII band.

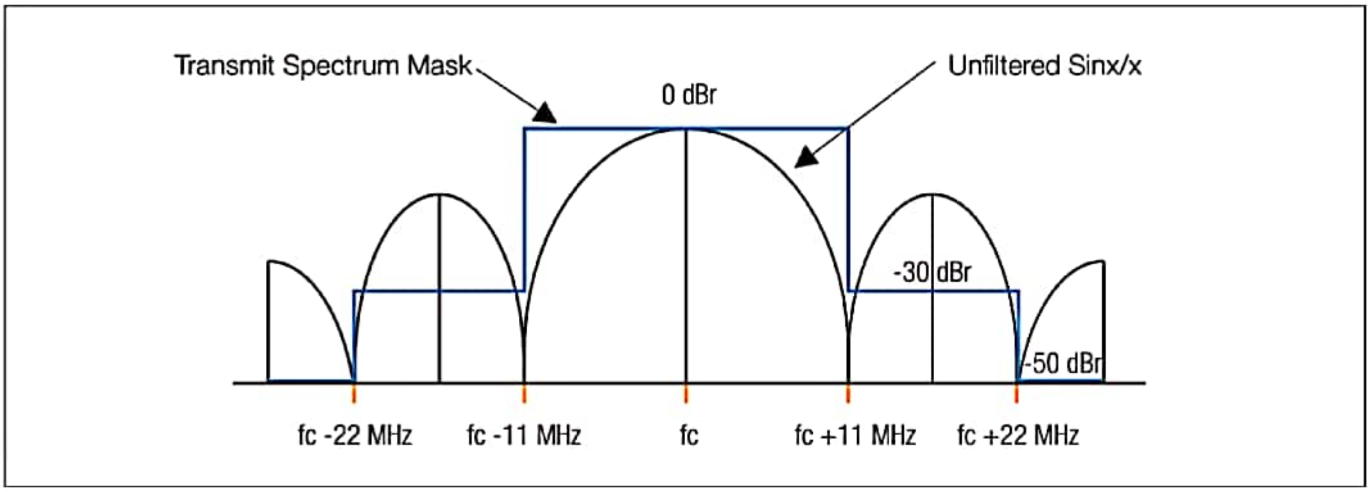

Figure 7. Spectral mask for 802.11b standard.

802.11ac will include support for 80 MHz bandwidth as well as an optional 160 MHz bandwidth. The 802.11ac device is required to support 20, 40, and 80 MHz channel bandwidth reception and transmission. The 80 MHz channel will consist of two adjacent, non-overlapping 40 MHz channels. The 160 MHz channels will be formed by two 80 MHz channels which may be adjacent (contiguous) or non-contiguous. With the newer standards, they are allowed to use a wider bandwidth to boost throughput.

However, it is important to realize that the 2.4 and 5 GHz frequency bands did not get any bigger. Products from all of the 802.11 standards are required to share the same bandwidth. Only if there is spectrum available can the wider bandwidths be used. This is why many 802.11n WLANs may end up using 40 MHz channels only in the 5 GHz band.

Spectral Masks

The 802.11 standard specifies a spectral mask which defines the permitted power distribution across each channel. The spectral mask requires the signal be attenuated to certain levels (from its peak amplitude) at specified frequency offsets. Figure 7 shows the spectral mask used for the 802.11b standard. While the energy falls off very quickly from its peak, it is interesting to note the RF energy that is still being radiated at other channels. This will be discussed further in our next section.

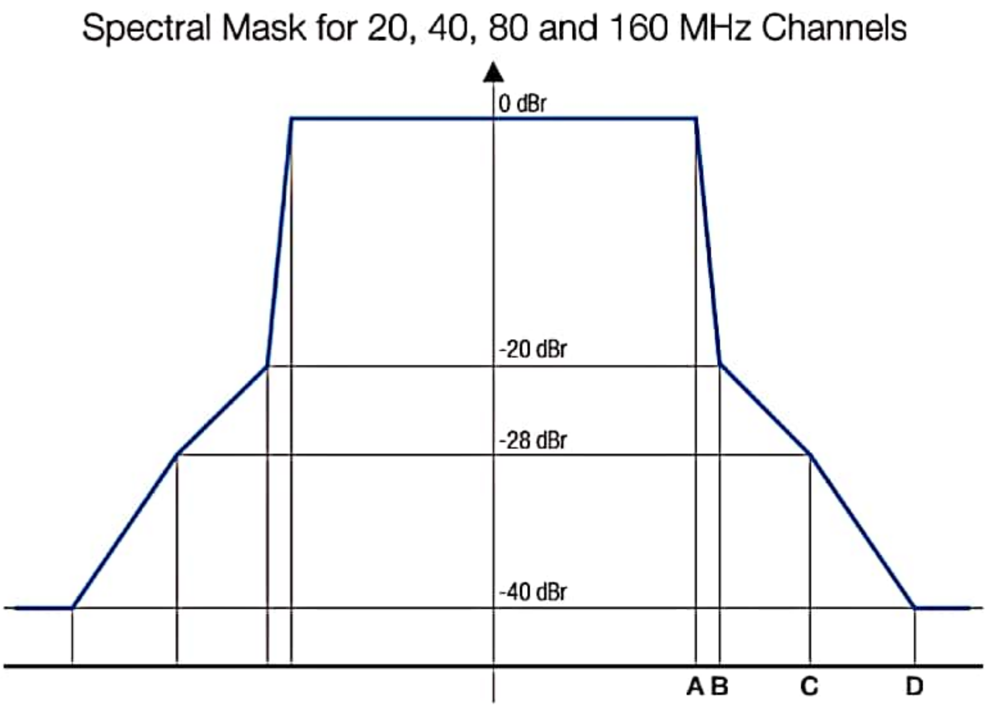

Channel Size	A	B	C	D
20 MHz	9 MHz	11 MHz	20 MHz	30 MHz
40 MHz	19 MHz	21 MHz	40 MHz	60 MHz
80 MHz	39 MHz	41 MHz	80 MHz	120 MHz
160 MHz	79 MHz	81 MHz	160 MHz	240 MHz

Figure 8. OFDM spectral mask used for 802.11a/g/n/ac.

For the 802.11a, 802.11g, 802.11n and 802.11ac standards that are using the OFDM encoding scheme, they have a spectral mask that looks entirely different (Figure 8). OFDM allows for a more dense spectral efficiency, thus it gets higher data throughput than used with the BPSK/QPSK techniques in 802.11b.

Overlapping Channels

The 802.11 use of the term "channel" can often lead to confusion. For radio and TV channels, they are allocated specific frequency spectrums to operate in. As shown in Figure 6 and from the 802.11 spectral masks it is apparent that plenty of RF energy is going into adjacent channels. Since the spectral mask only defines power output restrictions at specific frequency offsets, it is often assumed that the energy of the channel extends no further than these limits. It is more correct to say that, given the separation between channels, the overlapping signal on any channel should be sufficiently attenuated to minimally interfere with a transmitter on any other channel.

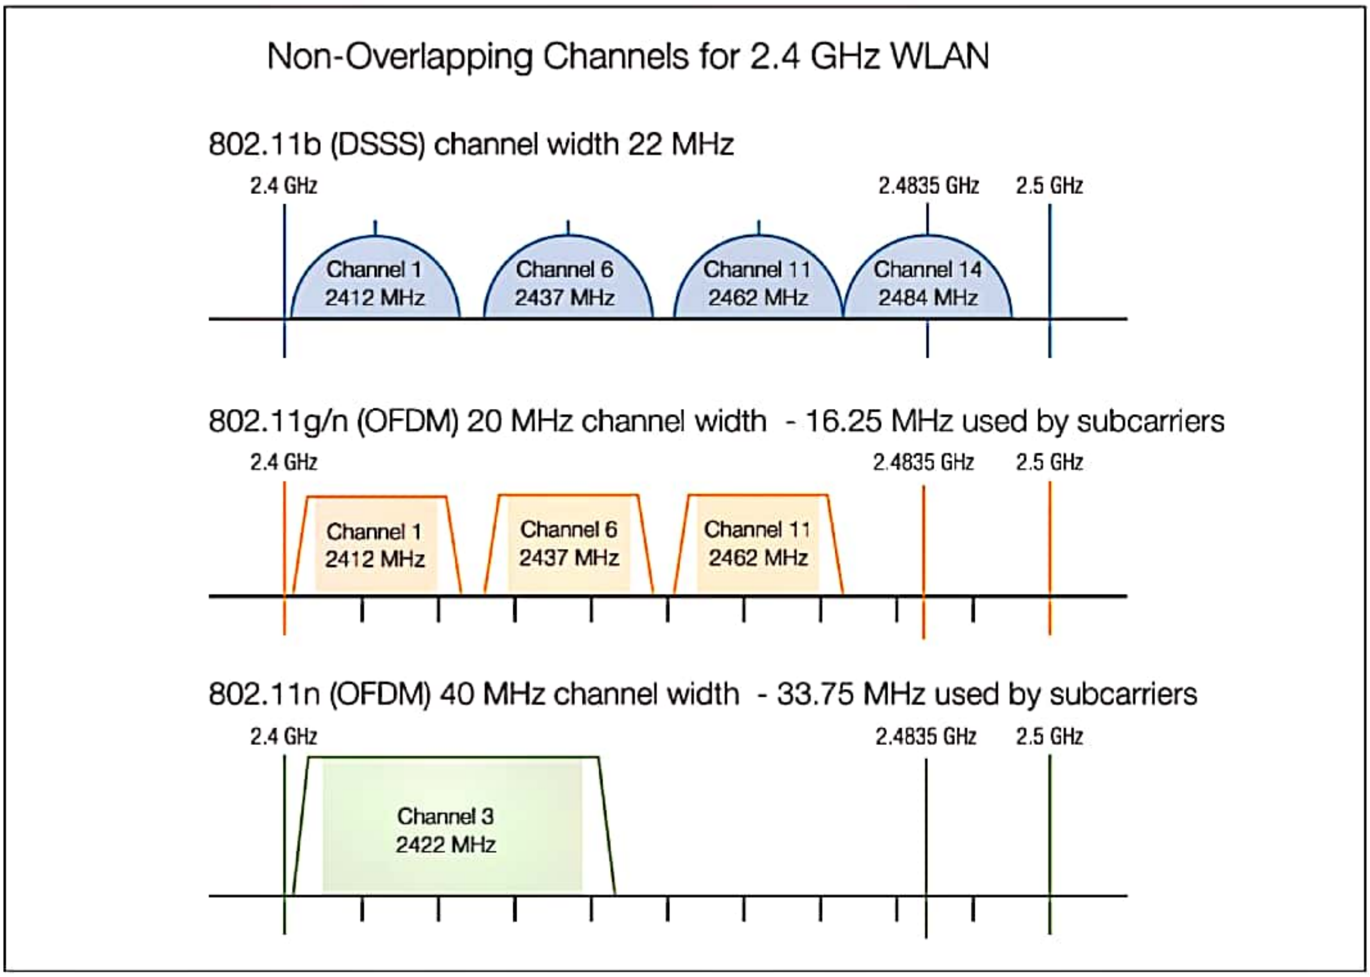

Figure 9. For the 802.11 standards there are only a few channels that are considered non-overlapping.

Confusion often arises over the amount of channel separation required between transmitting devices. The 802.11b standard was based on DSSS modulation and utilized a channel bandwidth of 22 MHz, resulting in three "non-overlapping" channels (1, 6 and 11). 802.11g was based on OFDM modulation and utilized a channel bandwidth of 20 MHz. This occasionally leads to the belief that four "non-overlapping" channels (1, 5, 9 and 13) exist for 802.11g, although this is not the case. Figure 9 highlights the potential non-overlapping channels in the 2.4 GHz bands. Although the "non-overlapping" channels are limited to spacing or product density, the concept has some merit in limited circumstances. Special care must be taken to adequately space AP cells since overlap between the channels may cause unacceptable degradation of signal quality and throughput.

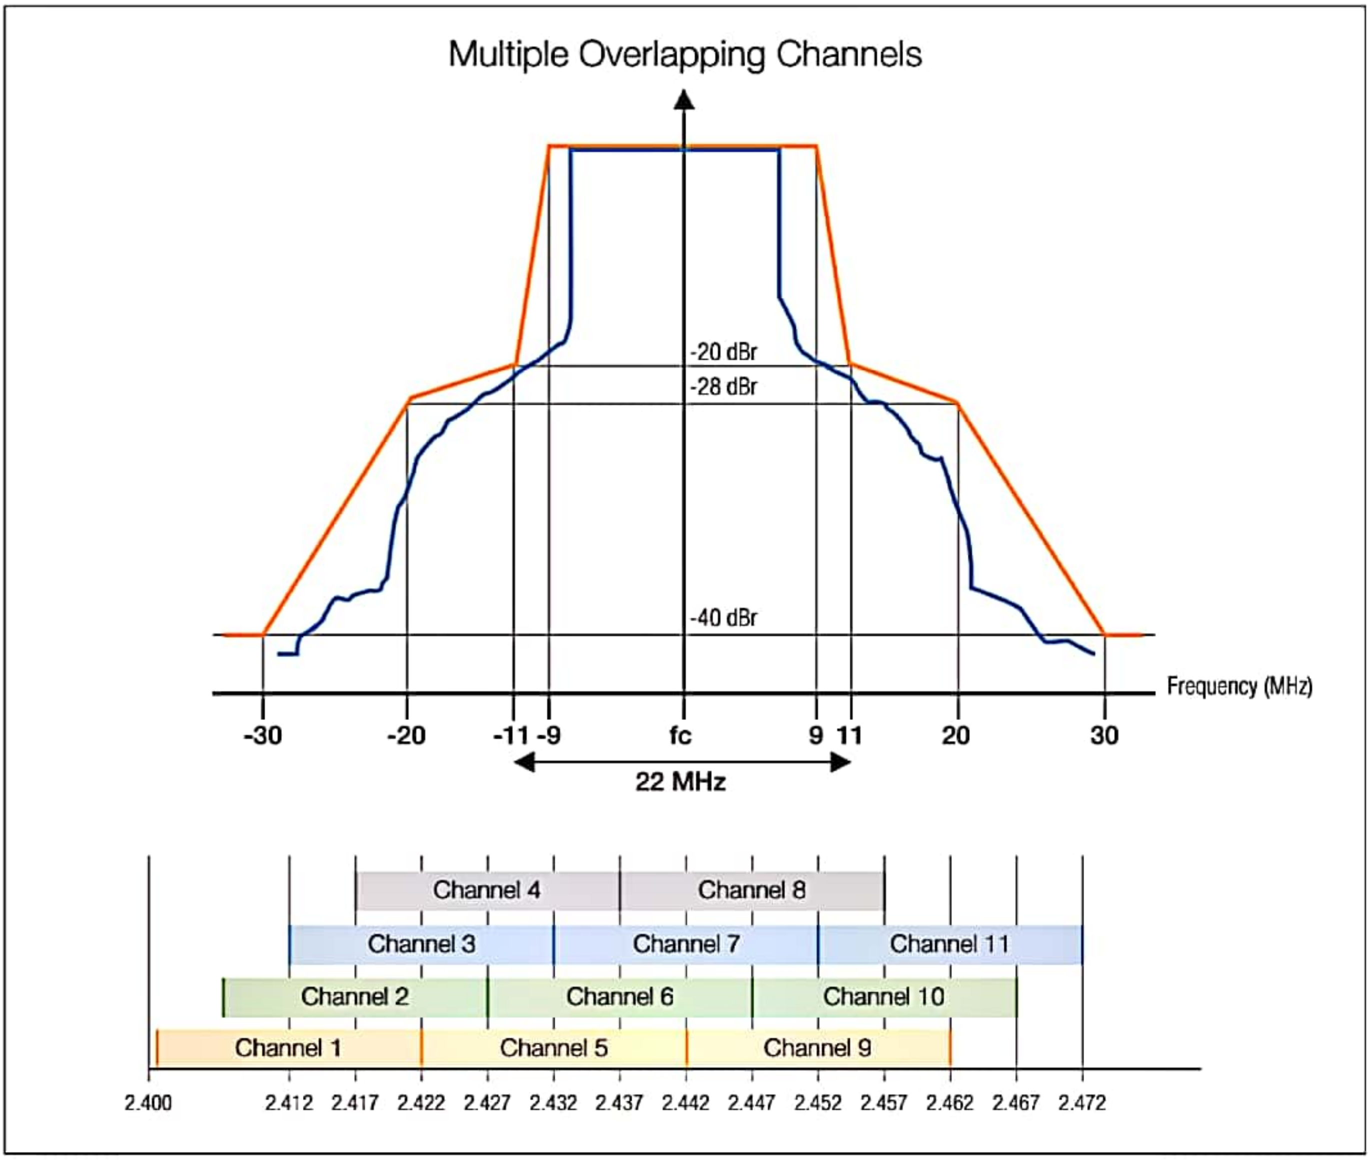

Figure 10. RF energy "bleeds" into frequencies for several adjacent channels, resulting in access points that may actually consume multiple overlapping channels.

ISM use is further complicated by 802.11b/g/n channel overlap. When an 802.11b/g/n radio transmits, the modulated signal is designed to fall within its bandwidth from the channel center frequency. However, RF energy ends up "bleeding" into frequencies for several adjacent channels. As a result, each 802.11b/g/n access point actually consumes multiple overlapping channels (see Figure 10). Transmitting on a 40 MHz 802.11n channel in the ISM band would exacerbate this scarcity by consuming 9 channels: the center frequency plus four channels on the left and four on the right. Finding adjacent unused channels in the congested ISM band is rare; thus, 40 MHz 802.11n operation would very likely interfere with existing 802.11b/g APs. To mitigate this, 802.11n APs using 40 MHz channels are required to listen for legacy (or other non-40 MHz HT) devices and provide coexistence mechanisms.