WLAN RF Design Considerations

This chapter describes the basic information necessary to understand radio frequency (RF) considerations in planning for various wireless local area network (WLAN) environments. The topics of this chapter include:

- Regulatory domains and RF considerations
- IEEE 802.11 standards
- RF spectrum implementations of 802.11b/g/n (2.4-GHz) and 802.11a/n/ac (5-GHz)
- Planning for RF deployment
- Manually fine-tuning WLAN coverage
- Radio resource management (RRM) algorithms

RF Basics

In the United States there are three bands (frequency ranges) allocated for unlicensed industrial, scientific, and medical (ISM) usage (see Figure 3-1).

The ISM bands are designated as the:

- 900 MHz band: 902 to 928 MHz
- 2.4 GHz band (IEEE 802.11b/g/n): 2.4 to 2.4835 GHz
- 5 GHz band (IEEE 802.11a/n/ac):
  - 5.150 to 5.250 GHz (UNII-1)
  - 5.250 to 5.350 GHz (UNII-2)
  - 5.725 to 5.875 GHz (UNII-3/ISM)

Each band has different characteristics. The lower frequency 2.4 GHz band exhibits better range but with limited bandwidth and thus lower data rates. The higher frequency 5 GHz band exhibits less range and is subject to greater attenuation from solid objects but has higher data rates.

The following sections provide a summary of regulatory domains and their operating frequencies.

Regulatory Domains

Devices that operate in unlicensed bands do not require a formal licensing process, but when operating in the ISM bands, the vendor is obligated to follow the government regulations for that region. The regulatory agencies in different parts of the world monitor these bands according to different criteria. WLAN devices must comply with the specifications of the relevant governing regulatory body. Although the regulatory requirements do not affect the interoperability of IEEE 802.11b/g/n- and 802.11a/n/ac-compliant products, the regulatory agencies do set certain criteria in the standard. For example, the emission requirements for WLAN are to minimize the amount of interference a radio can generate or receive from another radio in the same proximity. It is the responsibility of the WLAN vendor to obtain product certification from the relevant regulatory body.

Besides following the requirements of the regulatory agencies, many vendors also ensure compatibility with other vendors through the Wi-Fi Alliance (WFA) certification program (www.wi-fi.org).

Operating Frequencies

The 2.4-GHz band regulations of 802.11b/g/n have been relatively constant, given the length of time they have been in operation. The FCC (U.S) allows for 11 channels, ETSI (most other parts of the world) allows for up to 13 channels, and Japan allows up to 14 channels but requires a special license to operate in channel 14. Countries that adhere to the 5.0-GHz band regulations of 802.11a/n/ac are moving to open additional spectrum for additional channels.

These frequency bands and their associated data rates are described in more detail in the following sections.

2.4 GHz - 802.11b/g/n

The 802.11b standard, which was ratified in 1999, supports data rates of 1, 2, 5.5, and 11 Mbps and enjoys broad user acceptance and vendor support. 802.11b has been deployed by thousands of enterprise organizations as the first standardized method of WLAN communication.

The 802.11g standard, which was ratified in 2003, operates in the same spectrum as and is backwardly compatible with 802.11b. The 802.11g standard supports the additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps and is the most popular method of WLAN communications in the 2.4 GHz band.

The 802.11n standard, which was ratified in 2009, allows for usage in both 2.4 and 5 GHz bands. The 2.4 GHz band supports data rates up to 144 Mbps (assuming 20 MHz bandwidth and single transmitter stream). Note that faster speeds and bonding channels (using spectrum greater than 20 MHz) can permit speeds up to 300 Mbps, but this is typical of a home deployment. Although enterprise networks are limited to a data rate of 20 MHz in the 2.4 GHz band (due to a limited amount of spectrum), faster speeds can be achieved using the 5 GHz band with 802.11n and/or 802.11ac technology.

5 GHz - 802.11a/n/ac

Operating in the unlicensed portion of the 5 GHz radio band, 802.11a/n/ac is immune to interference from devices that operate in the 2.4 GHz band, such as microwave ovens, many cordless phones, and Bluetooth (a short-range, low-speed, point-to-point, personal-area-network wireless standard). Because the 802.11a/n/ac standards operate in a different frequency range, they are not compatible with existing 802.11b or 802.11g-compliant wireless devices. Regardless, 2.4 GHz and 5 GHz band devices can operate in the same physical environment without interference.

Deployment Considerations

Choosing between these two technologies (802.11a/n/ac and 802.11b/g/n) does not involve an equal one-for-one trade-off. They are complementary technologies and will continue to coexist in future enterprise environments. Those responsible for implementing these technologies must be able to make an educated choice between deploying 2.4 GHz-only networks, 5 GHz-only networks, or a combination of them. Organizations with existing 802.11b/g networks cannot simply deploy a new 802.11a network for existing APs and expect to have their 54 Mbps 802.11a/n/ac coverage in the same areas as their 11 Mbps 802.11b/g/n coverage. The technical characteristics of both these bands simply do not allow for this kind of interchangeability in coverage.

802.11a provides data rates of 6, 9, 12, 18, 24, 36, and 48 Mbps with a maximum data rate of 54 Mbps, (although generally at shorter ranges for a given power and antenna gain). 802.11a also has up to 23 non-overlapping channels (depending on the geographic area) compared to only three non-overlapping channels of 802.11b/g. This results in increased network capacity, improved scalability, and the ability to create microcellular deployments without interference from adjacent cells.

The 802.11a/n/ac 5 GHz band is divided into several different sections. Each of the UNII bands presented in Table 3-1 was originally intended for different uses, but all can currently be used by indoor 802.11a/n/ac devices with appropriate power restrictions. Initially, the FCC defined only the UNII-1, UNII-2, and UNII-3 bands, each of which had four channels. The channels were spaced 20 MHz apart with an RF spectrum bandwidth of 20 MHz, thereby providing non-overlapping channels.

There are differing limitations on these three UNII bands with restrictions varying between them for transmit power, antenna gain, antenna styles, and usage. The UNII-1 band is designated for indoor operations and initially had a restriction of permanently attached antennas. The UNII-2 and UNII-3 bands are designated for indoor or outdoor operations and permit external antennas.

The channels in UNII-1 (5.150 to 5.250 GHz) are 36, 40, 44, and 48. The channels in UNII-2 (5.250 to 5.350 GHz) are 52, 56, 60, 64 and require dynamic frequency selection (DFS) and transmitter power control (TPC). The channels in the new frequency range (5.470 to 5.725 GHz) are 100, 104, 108, 112, 116, 120, 124, 128, 132, 136, and 140 also require DFS and TPC. The channels in UNII-3 (5.725 to 5.850) are 149, 153, 157, 161, and 165 and do not require DFS and TPC. Not all channels in a given range can be used in all of the regulatory domains. Figure 3-1 shows the various channels in the UNII-1, 2, and 3 bands, along with the additional 11 new channels.

Figure 3-1 802.11 Channel Capacity

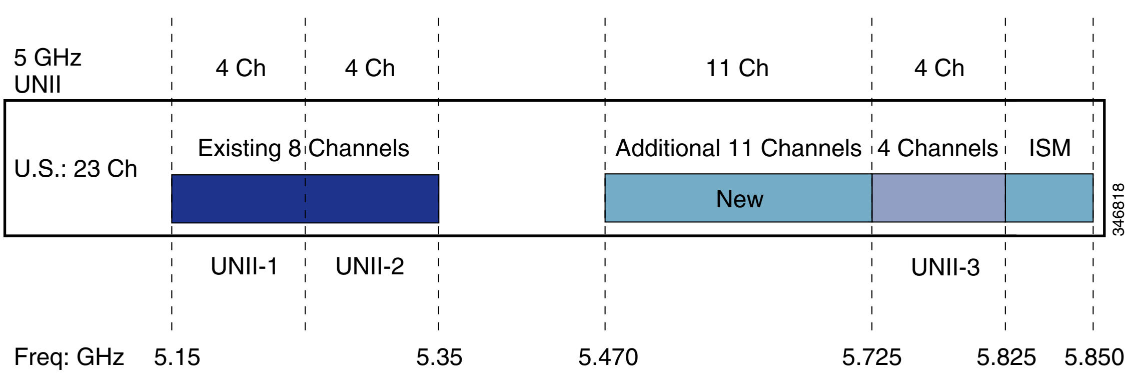

The FCC released a revision to the regulations in 2004 covering the 5 GHz 802.11a channel usage. This revision added 11 additional channels, bringing the available channels capacity to 23 channels, as shown in Figure 3-1. The 11 new channels are for indoor/outdoor use. However, to use the new channels radios must comply with the DFS and TPC features that are part of the 802.11h specification. DFS is required to avoid radar that operates in this frequency range, but it can also be used for other purposes such as dynamic frequency planning. 802.11h has been supported by Cisco Unified Wireless Network since 2010.

DFS dynamically instructs a transmitter to switch to another channel whenever a particular condition such as the presence of a radar signal is met. Before transmitting, the DFS mechanism of a device monitors its available operating spectrum and listens for a radar signal. If a signal is detected the channel associated with the radar signal is vacated or flagged as unavailable for use by the transmitter. Prior to and during operation the transmitting device continuously monitors the environment for the presence of radar. Portions of the 5 GHz band are allocated to radar systems, which allows WLANs to avoid interference with incumbent radar users in instances where they are co-located.

TPC allows the AP to negotiate power levels with a WLAN client during the association process. The AP can inform the WLAN client of the range of allowable transmit power to be used with that AP and might reject clients unable to meet those levels. The WLAN client is able to adjust its transmit power level within the range specified in the TPC negotiations. This ensures that interference from the WLAN is minimized and allows the WLAN client to optimize battery life.

For more information on FCC regulation updates, see the Cisco white paper *FCC Regulations Update* at:

http://www.cisco.com/en/US/products/hw/wireless/ps469/products_white_paper0900aecd801c4a88.shtml

Table 3-1 shows the standard 802.11a/n/ac frequencies.

Table 3-1 Operating Frequency Range for 802.11a/n/ac

Channel ID	36	40	44	48	52	56	60	64	100	104	108	112	116	120	124	128	132	136	140	149	153	157	161	165
Center Freq. in MHz	5180	5200	5220	5240	5260	5280	5300	5320	5500	5520	5540	5560	5580	5600	5620	5640	5660	5680	5700	5745	5765	5785	5805	5825
Band	UNII-1				UNII-2															UNII-3				ISM

Understanding the IEEE 802.11 Standards

IEEE 802.11 is the working group within the Institute for Electrical and Electronics Engineers (IEEE) responsible for wireless LAN standards at the physical and link layer (Layers 1 and 2) of the OSI model, as compared to the Internet Engineering Task Force (IETF), which works on network layer (Layer 3) protocols. Within the 802.11 working group are a number of task groups that are responsible for elements of the 802.11 WLAN standard. Table 3-2 summarizes some of the task group initiatives.

For more information on these working groups see: http://www.ieee802.org/11/

Table 3-2 IEEE 802.11 Task Group Activities

Task Group	Project
MAC	Develop one common MAC for WLANs in conjunction with a physical layer entity (PHY) task group.
PHY	Develop three WLAN PHYs—Infrared, 2.4 GHz FHSS, 2.4 GHz DSSS.
a	Develop PHY for 5 GHz UNII band.
b	Develop higher rate PHY in 2.4 GHz band.
c	Cover bridge operation with 802.11 MACs (spanning tree).
d	Define physical layer requirements for 802.11 operation in other regulatory domains (countries).

Table 3-2 IEEE 802.11 Task Group Activities (continued)

e	Enhance 802.11 MAC for QoS.
f	Develop recommended practices for Inter Access Point Protocol (IAPP) for multi-vendor use.
g	Develop higher speed PHY extension to 802.11b (54 Mbps).
h	Enhance 802.11 MAC and 802.11a/n/ac PHY-Dynamic Frequency selection (DFS), Transmit Power control (TPC).
i	Enhance 802.11 MAC security and authentication mechanisms.
j	Enhance the 802.11 standard and amendments to add channel selection for 4.9 GHz and 5 GHz in Japan.
k	Define RRM enhancements to provide interfaces to higher layers for radio and network measurements.
k	Define Radio Resource Measurement enhancements to provide interfaces to higher layers for radio and network measurements.
m	Perform editorial maintenance, corrections, improvements, clarifications, and interpretations relevant to documentation for 802.11 family specifications.
n	Focus on high throughput extensions (>100 Mbps at MAC SAP) in 2.4 GHz and/or 5 GHz bands.
o	Provide Fast Handoffs in Voice over WLAN (goal is around 50 ms)
p	Focus on vehicular communications protocol aimed at vehicles, such as toll collection, vehicle safety services, and commerce transactions via cars.
r	Develop a standard specifying fast BSS transitions and fast roaming.
s	Define a MAC and PHY for meshed networks that improves coverage with no single point of failure.
t	Provide a set of performance metrics, measurement methodologies, and test conditions to enable manufacturers, test labs, service providers, and users to measure the performance of 802.11 WLAN devices and networks at the component and application level.
u	Provide functionality and interface between an IEEE 802.11 access network (Hotspot) and any external network.
v	Provide extensions to the 802.11 MAC/PHY to provide network management for stations (STAs).
w	Provide mechanisms that enable data integrity, data origin authenticity, replay protection, and data confidentiality for selected IEEE 802.11 management frames including but not limited to: action management frames, deauthentication and disassociation frames.

Direct Sequence Spread Spectrum (DSSS)

Direct sequence spread spectrum (DSSS) encodes redundant information into the RF signal. This provides the 802.11 radio with a greater chance of understanding the reception of a packet over the background noise or interference on the channel. Every data bit is expanded into a string of bits (*chips*) called a chipping sequence or barker sequence. The chipping rate mandated by IEEE 802.11 is 11 chips per bit using binary phase-shift keying (BPSK)/quadrature phase-shift keying (QPSK) at the 1 and 2 Mbps rates and 8 chips (complimentary code keying–CCK) at the 11 and 5.5 Mbps rate. This means that at 11 Mbps, 8 bits are transmitted for every one bit of data. The chipping sequence is transmitted in parallel across the spread spectrum frequency range.

IEEE 802.11b Direct Sequence (DS) Channels

Fourteen channels are defined in the IEEE 802.11b direct sequence (DS) channel set. Each DS channel transmitted is 22 MHz wide but the channel separation is only 5 MHz. This leads to channel overlap so that signals from neighboring channels that are less than are 25 MHz apart can interfere with each other. In the 14-channel DS system (11 usable channels in the US), the only non-overlapping, non-interfering channels possible are channels 1, 6, and 11 (see Figure 3-2).

This channel spacing governs the use and allocation of channels in a multi-AP environment such as an office or campus. APs are usually deployed in a cellular fashion within an enterprise where adjacent APs are allocated non-overlapping channels. Alternatively, APs can be co-located using channels 1, 6, and 11 to deliver 33 Mbps bandwidth to a single area (but only 11 Mbps to a single client). If 802.11g is used in the same manner the aggregate bandwidth is 162 Mbps with a maximum data rate of 54 Mbps. This channel allocation scheme is illustrated in Figure 3-2.

Figure 3-2 IEEE 802.11 DSS Channel Allocations

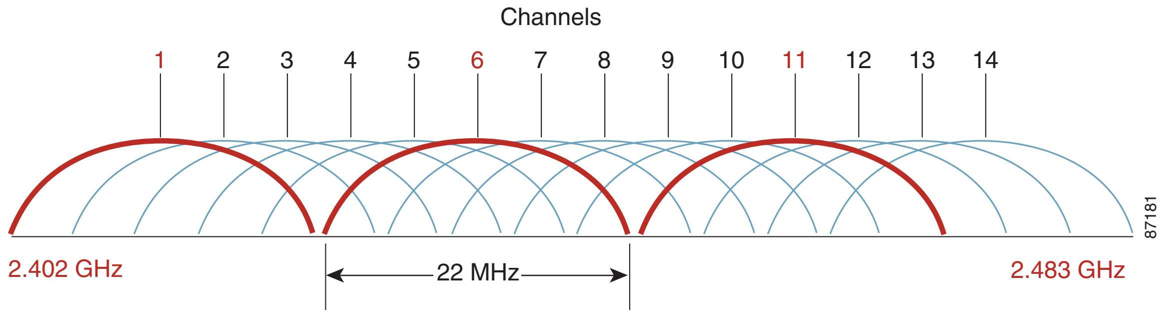

IEEE 802.11g

802.11g provides for a higher data rate (up to 54 Mbps) in the 2.4 GHz band, the same spectrum as 802.11b. 802.11g is backward-compatible with 802.11b and provides additional data rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps. At higher data rates, 802.11g uses the same modulation technique, orthogonal frequency division multiplexing (OFDM), as 802.11a/n/ac (see IEEE 802.11a OFDM Physical Layer, page 3-7).

Table 3-3 lists 802.11g modulation and transmission types for the various data rates.

Table 3-3 802.11g Modulation and Transmission Types

Modulation	Transmission Type	Bits per Subchannel	Data Rate (Mbps)
BPSK	DSSS	NA	1
QPSK	DSSS	NA	2
CCK	DSSS	NA	5.5
BPSK	OFDM	125	6
BPSK	OFDM	187.5	9
CCK	DSSS	NA	11
QPSK	OFDM	250	12
QPSK	OFDM	375	18
16-QAM	OFDM	500	24
16-QAM	OFDM	750	36
64-QAM	OFDM	1000	48
64-QAM	OFDM	1125	54

IEEE 802.11a OFDM Physical Layer

IEEE 802.11a defines requirements for the physical layer of the OSI model, operating in the 5.0 GHz UNII frequency, with data rates ranging from 6 Mbps to 54 Mbps. It uses orthogonal frequency division multiplexing (OFDM), which is a multi-carrier system (compared to single carrier systems). OFDM allows subchannels to overlap, providing a high spectral efficiency. The modulation technique allowed in OFDM is more efficient than spread spectrum techniques used with 802.11b/g/n.

IEEE 802.11a Channels

The 802.11a channel shows the center frequency of the channels. The frequency of the channels is 10 MHz on either side of the dotted line. There is 5 MHz of separation between channels, as shown in Figure 3-3.

Figure 3-3 Channel Set Example

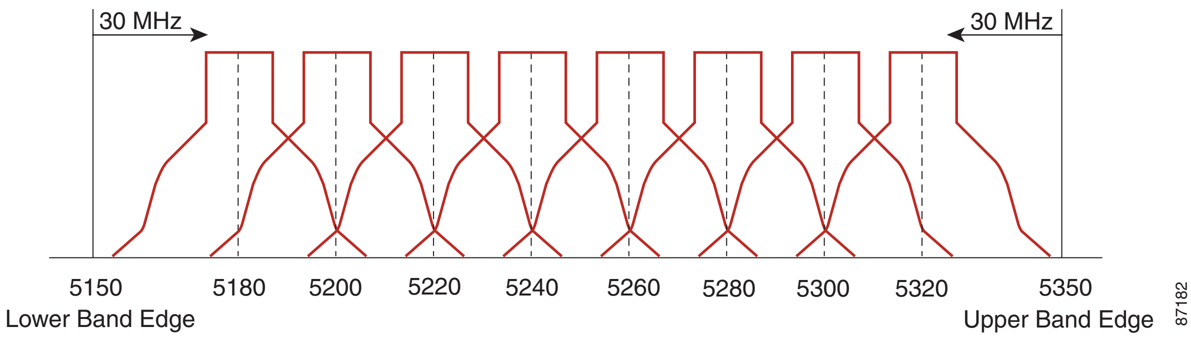

For the US-based 802.11a/n/ac standard, the 5 GHz unlicensed band covers 300 MHz of spectrum and supports 12 channels. As a result, the 5 GHz band is actually a conglomerate of three bands in the United States:

- 5.150 to 5.250 GHz (UNII-1)
- 5.250 to 5.350 GHz (UNII-2)
- 5.725 to 5.875 GHz (UNII-3)

RF Power Terminology

The terms such as *dB*, *dBi*, and *dBm* are used to describe the amount of change in power measured at points in a system, as perceived by the radio or compared to a reference power level. The following sections cover their differences and provide general rules for their use. Effective isotropic radiated power (EIRP) is also described.

dB

The term dB *(decibel)* is mainly used to describe attenuation or amplification of the power level. dB is a logarithmic ratio of a signal to another standardized value. For example, dBm is where the value is being compared to 1 milliwatt, and dBw is where the value is being compared to 1 watt.

The mathematical equation is:

```
power (in dB) = 10 * log10 (signal/reference)
```

Substituting in real numbers (signal 100 mW, reference 1 mW) provides a value in dB of 20 (100 = 10 squared; taking the exponent 2 and multiplying by 10 gives you 20).

Keep in mind that it is logarithmic, meaning that it increases or decreases exponentially and not linearly, and it is a ratio of a given value to a reference. Also keep in mind that it is multiplied by 10.

Given that it is logarithmic, there are general rules to take into consideration. An increase or decrease of 3 dB means that the signal doubled (double the power) or halved, respectively. An increase or decrease of 10 dB means that the signal went up by 10 times or down to $1/10^{th}$ of the original value.

Indoor and outdoor WLAN deployments each offer separate challenges in RF deployments that need to be analyzed separately. However, there are a few general rules for indoor use. For every increase of 9 dB, the indoor coverage area should double. For every decrease of 9 dB, the indoor coverage area should be cut in half.

dBi

The term dBi (dB *isotropic*) describes the forward gain of a real antenna compared with the hypothetical isotropic antenna. An isotropic antenna (a theoretical or imaginary antenna) is one that sends the same power density perfectly in all directions.

Antennas are compared to this ideal measurement, and all FCC calculations use this measurement (dBi). For example, a Cisco omni-directional AIR-ANT4941 antenna has a gain of 2.2 dBi, meaning that the maximum energy density of the antenna is 2.2 dB greater than an isotropic antenna.

dBm

The term dBm (dB *milliwatt*) uses the same calculation as described in the dB section but has a reference value of 1 mW.

Taking into consideration the example given above in the dB section, if the power increased from 1 mW to 100 mW at the radio, the power level would increase from 0 dBm to 20 dBm.

Besides describing transmitter power, dBm can also describe receiver sensitivity. Receiver sensitivity is in represented as minus dBm (-dBm) because the signal reduces in value from its point of transmission. The sensitivity indicates the lowest power the receiver can receive before it considers the signal unintelligible.

Effective Isotropic Radiated Power (EIRP)

Although transmitted power based on the radio setting is rated in either dBm or milliwatts, the maximum energy density coming from an antenna from a complete system is measured as EIRP, which is a summation of the dB values of the various components. EIRP is the value that regulatory agencies such as the FCC or ETSI use to determine and measure power limits, expressed in terms of maximum energy density within the first Fresnel of the radiating antenna. EIRP is calculated by adding the transmitter power (dBm) to antenna gain (dBi) and subtracting any cable losses (dB). For example, if you have a Cisco Aironet bridge connected to a solid dish antenna by a 50 foot length of coaxial cable, substituting in the numbers gives the following:

- Bridge: 20 dBm
- 50 Foot Cable: -3.3 dBm (negative because of cable loss)
- Dish Antenna: 21 dBi
- EIRP: 37.7 dBm

For more information see the Cisco TechNote *RF Power Values* at:
http://www.cisco.com/en/US/tech/tk722/tk809/technologies_tech_note09186a00800e90fe.shtml

Planning for RF Deployment

Many of the RF design considerations are interdependent or implementation-dependent. As a result, there is no *one-size-fits-all* template for the majority of requirements and environments.

The Cisco Prime Infrastructure provides integrated RF prediction tools that can be used to create a detailed wireless LAN design, including CAPWAP AP placement, configuration, and performance/coverage estimates. IT staff can import real floor plans into Cisco Prime Infrastructure and assign RF characteristics to various building components to increase design accuracy.

The Cisco Prime Infrastructure graphical heat maps help IT staff visualize anticipated wireless LAN behavior for easier planning and faster rollout. Cisco Prime Infrastructure also supports irregularly-shaped buildings by offering drawing tools to help organizations easily design and support WLAN deployments in such buildings. Figure 3-4 shows an example of the planning tool.

Figure 3-4 Cisco Prime Infrastructure Planning Tool

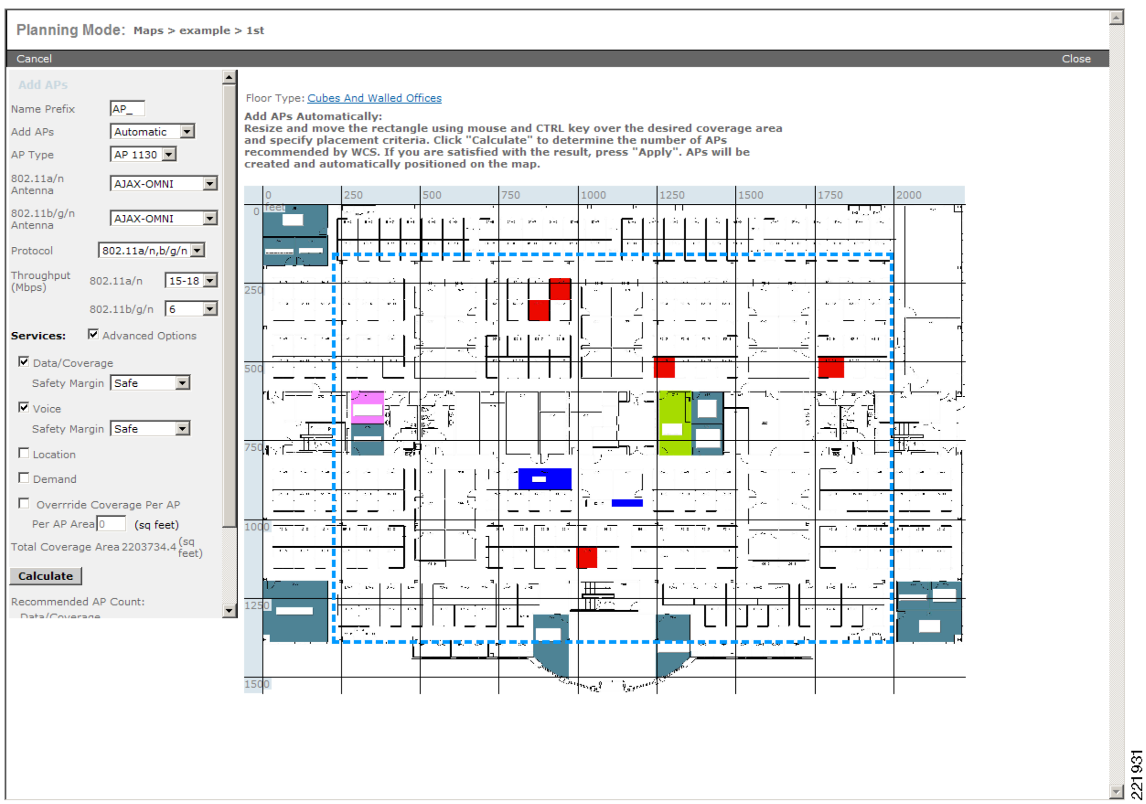

Different Deployment Types of Overlapping WLAN Coverage

How much overlapping WLAN coverage you set in the design of your wireless network depends on the usage, though with limited exceptions, all designs should be deployed to minimize retransmission and data rate shifting. Wireless networks can be deployed for data-only, voice, video, and location-based services (LBS), or a combination of all of these. The difference is in the pattern in which the APs are laid out and the amount of RF overlap in the coverage area.

When planning a WLAN deployment, consideration should be given to future uses of the WLAN deployment. Converting a WLAN deployment to support additional services beyond a data-only deployment is not simply a matter of adding APs; it can require an additional site survey and the possible relocation of existing APs.

For more information on different types of network deployments, see:

- *Cisco Unified Wireless iPhone 792x Deployment Guide*:

  http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7925g/7_0/english/deployment/guide/7925dply.pdf

- *Cisco VoWLAN Troubleshooting Guide: Site Survey and RF Design Validation*:

  http://www.cisco.com/en/US/docs/wireless/technology/vowlan/troubleshooting/8_Site_Survey_RF_Design_Valid.html

- *Cisco Site Survey Guide*: Deploying Cisco 7920 IP Phones:

  http://www.cisco.com/en/US/docs/wireless/technology/7920/site_survey/guide/survovr.html

- *Wireless LAN Design Guide for High Density Client Environments in Higher Education*:

  http://www.cisco.com/web/strategy/docs/education/cisco_wlan_design_guide.pdf

- *Cisco Wireless Mesh Access Points, Design and Deployment Guide*:

  http://www.cisco.com/en/US/docs/wireless/technology/mesh/7.3/design/guide/Mesh_chapter_0100.html

- *Wi-Fi Location-Based Services Design Guide: Best Practices Location-Aware WLAN Design Considerations*:

  http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich5.html

Data-Only Deployments

Data-only deployments do not require a large amount of overlap. This is because 802.11 clients respond to a lower signal from a nearby AP by stepping down their rate and extending the time to transmit. The required overlap is determined by the WLAN data rate requirement described in WLAN Data Rate Requirements, page 3-13. For data-only networks a general rule for the separation of APs is typically 120 to 130 feet. But when making your estimation for AP separation be sure to factor in objects that affect RF coverage such as wall densities, machinery, elevators, or even open spaces with steel cages. Your results can vary depending on the RF environment. Radio resource management (RRM, also known as *Auto-RF* in WLC) has been developed for this type of deployment and is very useful for controlling the RF coverage.

Voice Deployments

Figure 3-5 shows the voice network pattern and overlap.

Figure 3-5 Single Floor Site Survey for Voice

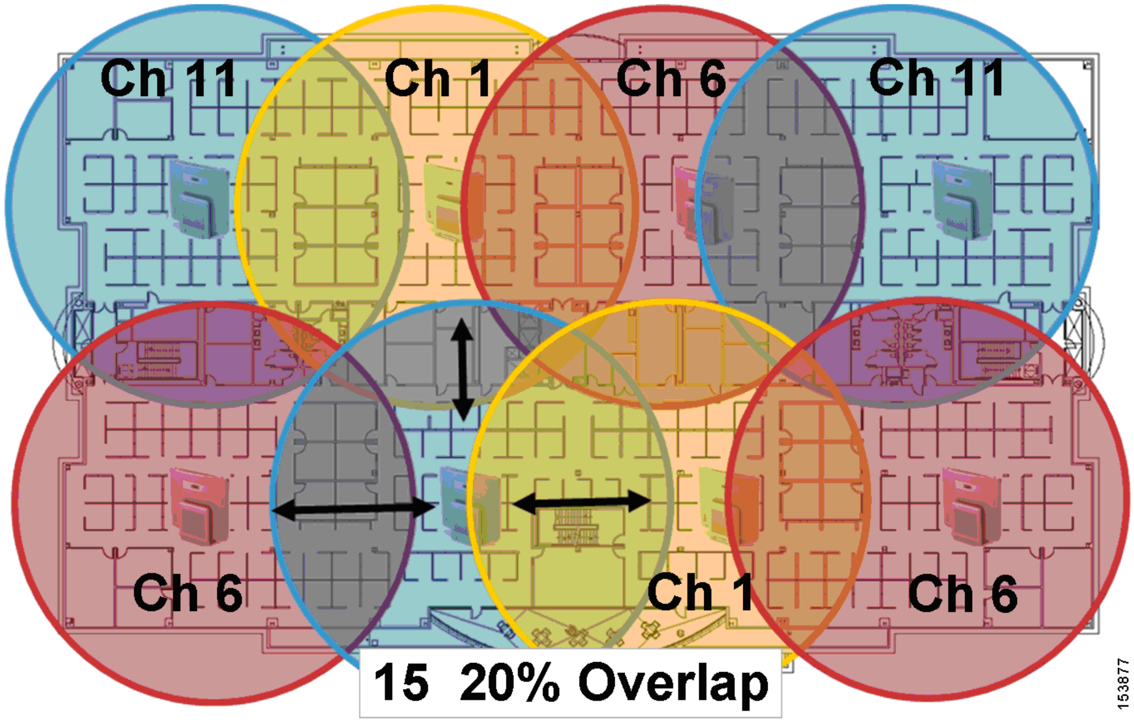

For a voice network the APs are grouped closer together and have more overlap than a data-only installation because voice clients need to roam to a better AP before dropping packets. Generally, you should create smaller cells than for data-only networks and ensure the overlapping cell edges are at or above -67 dBm. This accomplishes a number of things including greater homogeneity across a single cell and reduced processor load in the client device, which increases link stability and reduces latency. Although only one AP might be required for a defined area, Cisco recommends that you have two APs on non-overlapping channels with a received signal strength indication (RSSI) above 35 at all times in your installation for the purposes of redundancy and load balancing. For example, for a Cisco 792x VoIP phone deployment, Cisco recommends that you have an RSSI above 35 at all times in your installation. This is to ensure that the phone has good reception as well as allow some over-subscription. It also enhances roaming choices for the phone.

Keep in mind that designing for low noise background is as important as relatively high energy density within cells. This means that a good baseline power setting for the AP is in the 35 to 50 mW range. This generally requires approximately 15 percent more APs than if you deployed a coverage model at 100 mW.

Pre-site surveys are useful for identifying and characterizing certain challenging areas and potential sources for interference such as existing WLANs, rogue clients, and non-802.11 interference from sources such as microwave ovens and certain cordless telephones. Following a design that is reviewed and approved by all stakeholders, post-site surveys can be considered as an excellent audit mechanism to ensure that the coverage model complies with the intended functional requirements as set forth by the stakeholders. For more information on site surveys, see: *Cisco VoWLAN Troubleshooting Guide: Site Survey and RF Design Validation*:

http://www.cisco.com/en/US/docs/wireless/technology/vowlan/troubleshooting/8_Site_Survey_RF_Design_Valid.html

When making your estimation for separation, remember to factor in objects that affect RF coverage such as wall densities, machinery, elevators, or even open spaces with steel cages. Your results can vary depending on the RF environment. Be sure to include transient dynamics such as forklifts, large groups of people, or large objects moved through the area by crane or similar load bearing devices.

The WLC can be used to create a preliminary site evaluation (*predictive site survey*) that allows for the rapid deployment of a WLAN infrastructure and for the making of RF measurements of the area. A hand-walked site survey is also effective insurance for complex areas such as those commonly found in health care, retail, and manufacturing. For more information on a wireless voice deployments, see Chapter 9, "VoWLAN Design Recommendations."

Location-Based Services Deployments

The last type of deployment to discuss is about location-based services (LBS), which could be the most complex of the network applications because it relies not only on excellent cell coverage but also on the optimal location of APs. LBS deployments can simultaneously track thousands of devices by using the WLAN infrastructure. Examples in LBS include Wi-Fi tag type deployments or asset tracking deployments to locate equipment or devices by way of the wireless network. LBS can also be used simply to indicate where wireless clients are throughout the wireless network in relation to a drawing or diagram. This information can be used to increase wireless infrastructure security by providing the location of rogue clients or APs. One thing it does is greatly improve client troubleshooting capabilities.

APs are laid out in a staggered pattern for location-based services deployments. The staggered pattern allows for more accurate estimation of the location of a device. Figure 3-6 shows a typical pattern.

Figure 3-6 Example of a Single Floor Location Management Deployment

For a discussion of location-based services, see Chapter 11, "Cisco Mobility Services Engine," and the Cisco *Wi-Fi Location Based Services 4.1 Design Guide* at:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Mobility/wifich1.html.

The Cisco 7921G and the Cisco 7920 are Cisco VoWLAN handsets. Their use is one of the most common reasons for deploying QoS on a WLAN.

For more information on the 7920 and 7921G handsets, see:

- *Cisco Unified Wireless IP Phone 7921G*:

  http://www.cisco.com/en/US/products/hw/phones/ps379/products_data_sheet0900aecd805e315d.html

- *Cisco Unified Wireless IP Phone 7920*:

  http://www.cisco.com/en/US/products/hw/phones/ps379/products_data_sheet09186a00801739bb.html

Deploying VoWLAN infrastructure involves more than simply providing QoS on WLAN. When planning a voice WLAN deployment you need to consider site survey coverage requirements, user behavior, roaming requirements, and admission control. These are covered in the following guides:

- *Design Principles for Voice Over WLAN*:

  http://www.cisco.com/en/US/solutions/collateral/ns340/ns394/ns348/net_implementation_white_paper0900aecd804f1a46.html

- *Cisco Unified Wireless IP Phone 7921G Administration Guide*:

  http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7921g/5_0_1/english/administration/guide/21adm501.html

- *Cisco Wireless IP Phone 7920 Design and Deployment Guide*:

  http://www.cisco.com/en/US/docs/voice_ip_comm/cuipph/7920/5_0/english/design/guide/7920ddg.html

WLAN Data Rate Requirements

Data rates affect AP coverage areas. Lower data rates (such as 1 Mbps) can extend the coverage area farther from the AP than higher data rates (such as 54 Mbps) as illustrated in Figure 3-7 (not drawn to scale). Therefore, the data rate (and power level) affects coverage and consequently the number of APs required for the installation for different data rates, as illustrated in Figure 3-8. As part of the planning process, consider the required data rates, the required range, and the required reliability.

Data Rates Compared to Coverage Area

Different data rates are achieved by the AP using different encoding techniques on the wireless link allowing data to be more easily recovered from noise; this can be seen in the different receiver sensitivities for the different data rates. The number of *symbols* (groups of chips) sent out for a packet at the 1 Mbps data rate is greater than the number of symbols used for the same packet at 11 Mbps. This means that sending data at the lower bit rates requires more time than sending the equivalent data at a higher bit rate. And when there is more than one client associated to the radio the lower rate client affects the maximum data throughput of higher rate clients by taking longer to transmit a packet of the same length.

The actual diameter of the coverage, as shown in Figure 3-7, depends on factors such as environment, power level, and antenna gain.

Figure 3-7 Data Rate Compared with Coverage

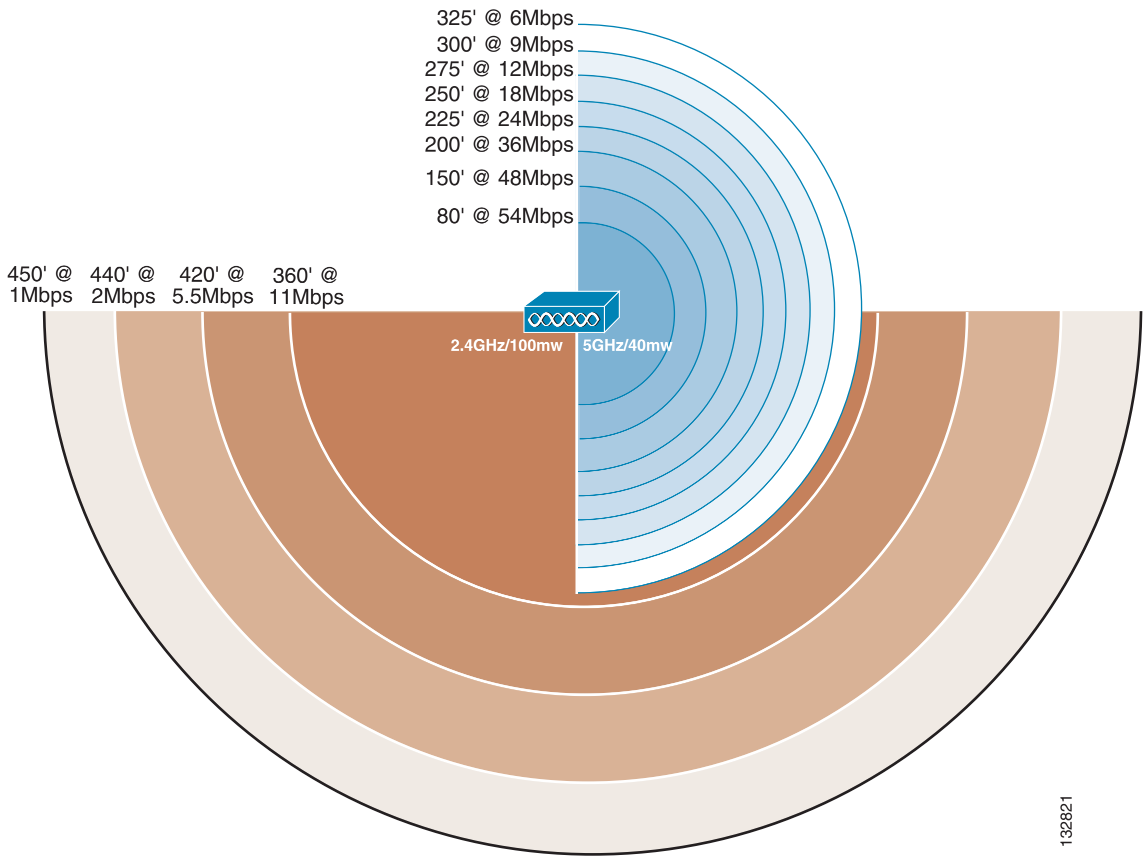

For example, deployed indoors using the standard antennas on the NIC card and APs, the diameter of the 1 Mbps circle is approximately 700 feet (210 meters), and the diameter of the 11 Mbps circle is about 200 feet (60 meters). These coverage diameters depend upon the type of indoor environment. An open office plan building is different from one with offices and solid walls. Increasing the gain of the antenna can increase the distance and change the shape of the radiation pattern to be focused in specific directions rather than being radiated evenly.

AP Density for Different Data Rates

The minimum required reliable data rate has a direct impact upon the number of APs needed in the design, along with power setting, antenna gain, and location. Figure 3-8 shows coverage comparison and AP density for different data rates. Although six APs with a minimum data rate of 11 Mbps might adequately service an area, it might take twice as many APs to support a minimum data rate of 24 Mbps, and more again to support a minimum data rate of 48 Mbps for the same coverage area.

Figure 3-8 Coverage Comparison and AP Density for Different Data Rates

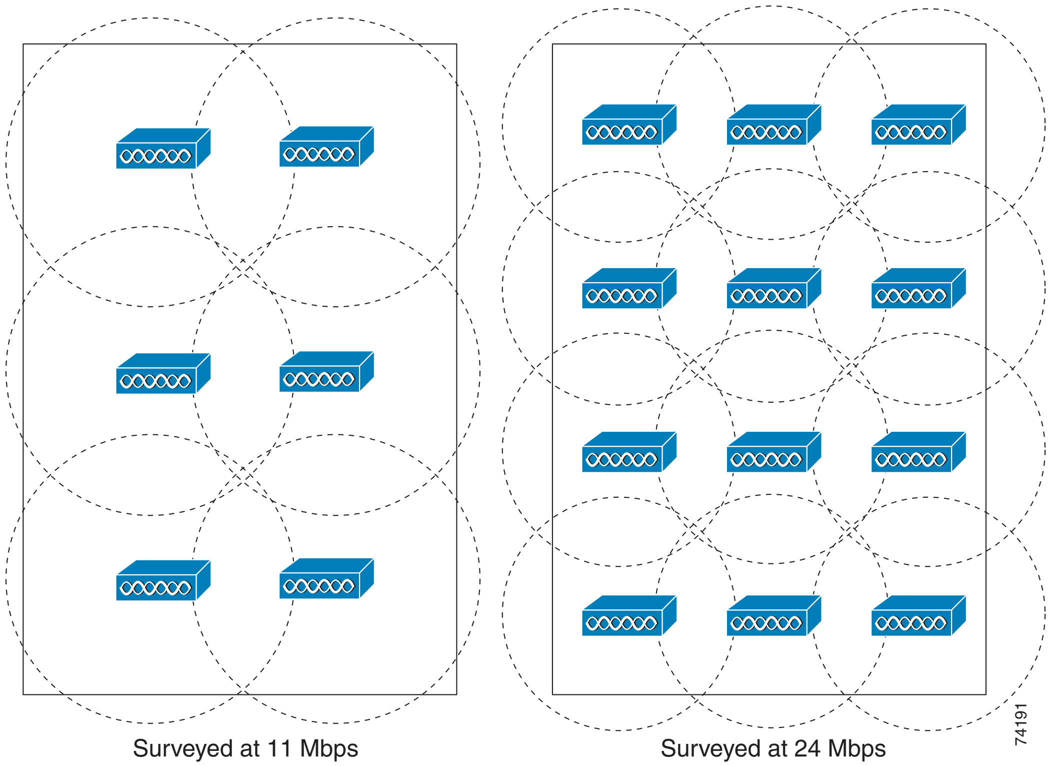

The data rate you choose depends on the type of application to be supported but it should not be greater than the typical requirements because there is trade-off in coverage. In a typical WLAN environment, the higher data rates provide maximum throughput and should minimize performance-related support issues. The physical facility and/or whether the network is client-centric generally dictates range requirements; some clients might not support the higher data rates, longer ranges, or the delay and jitter rates of an infrastructure element such as an AP.

To allow all data rates it might seem logical at first to choose the default configuration of APs and clients. However, there are three key reasons for limiting the data rate to the *highest* rate at which full coverage is obtained:

- Broadcast and multicast (if enabled) are sent at the lowest associated data rate to ensure that all clients can receive the packets. This reduces the throughput of the WLAN because traffic must wait until frames are processed at the slower rate.
- Clients that are farther away, and therefore accessing the network at a lower data rate, decrease the overall throughput by causing delays while the lower bit rates are being serviced. It might be better to force the clients to roam to a closer AP so as not to impact the performance of the rest of the network.
- If a 54 Mbps service is specified and provisioned with APs to support *all* data rates (for example), clients at lower rates can associate with the APs that can create a coverage area greater than planned, thereby increasing the security exposure (by allowing association from outside the building) and potentially interfering with other WLANs.

Client Density and Throughput Requirements

Wireless APs have two characteristics that make actual client data throughput slower than the data rate:

- APs have an aggregate throughput less than the data rate because 802.11 provides a reliable transport mechanism that ACKs all packets, thereby halving the throughput on the channel.
- APs are similar to shared hubs. That is, the channel is shared by all the clients associated to that AP on that channel, which causes collisions to slow data throughput.

With this in mind, you must have an estimate of the maximum number of active associations (active clients). This can be adjusted more or less according to the particular application.

Each cell provides an aggregate amount of throughput that is shared by all the client devices that are within the cell and associated to a given AP. This basically defines a cell as a collision domain. After deciding on the minimum data rate, be sure to consider how much throughput should, on average, be provided to each user of the WLAN.

Take the example of a simple barcode scanning application; 25 Kbps might be more than sufficient bandwidth for such an application because using an 802.11b AP at 11 Mbps of data rate results in an aggregate throughput of 5 to 6 Mbps. A simple division results in a maximum number of 200 users that can theoretically be supported. This number cannot in fact be achieved because of the 802.11 management overhead associated with the large number of clients and packet collisions. For a 1 Mbps system, 20 users can use the same AP for similar bandwidth results.

You can increase the potential per-user throughput by decreasing the number of users contending for the aggregate throughput provided by a single AP. This can be done by decreasing the size of the coverage area or adding a second AP on a non-overlapping channel in the same coverage area. To reduce the coverage area the AP power or antenna gain can be reduced resulting in fewer clients in that coverage area. This means you need more APs for the same overall area, thereby increasing the cost of deployment. An example of this is shown in Figure 3-9.

Figure 3-9 Changing the Output Power to Increase Client Performance

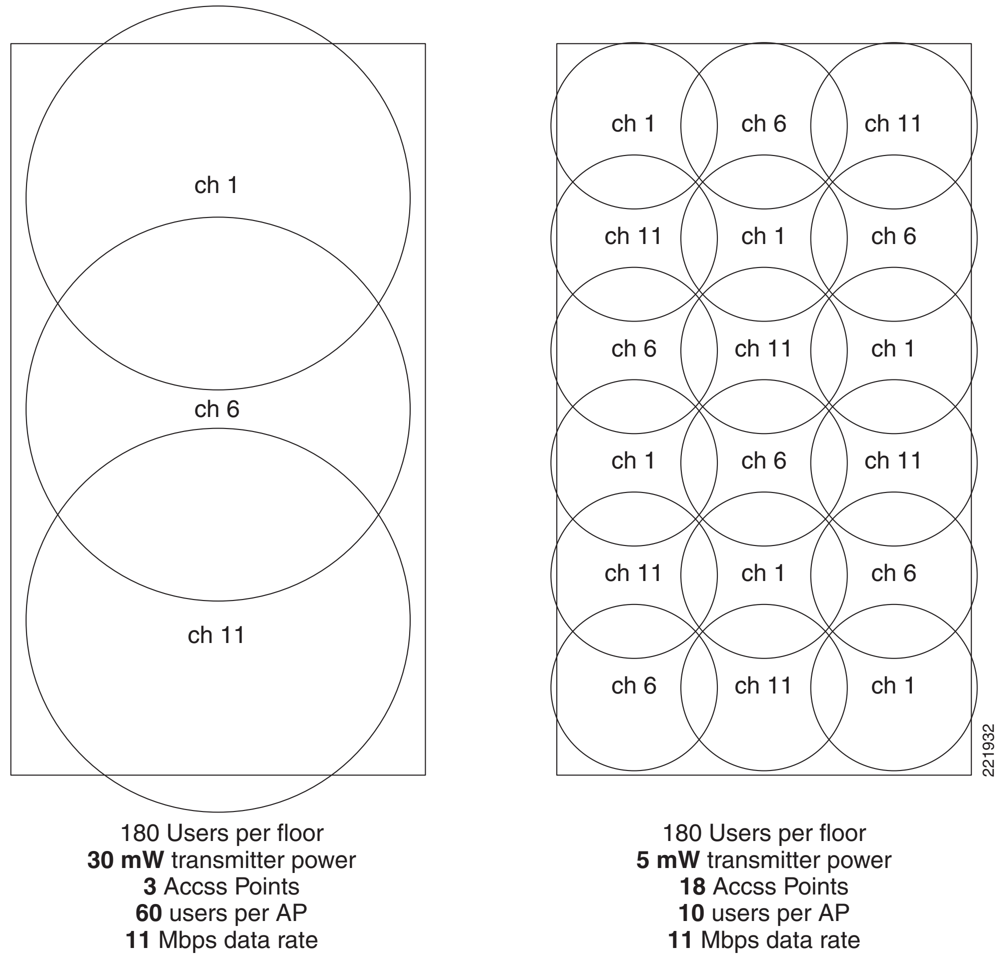

Note Client power should be adjusted to match the AP power settings. Maintaining a higher setting on the client does not result in higher performance and it can cause interference in nearby cells.

WLAN Coverage Requirements

Different enterprises have different coverage requirements. Some need a WLAN to cover specific common areas while others need WLANs to cover each floor of a building. They also might need to cover the entire building including stairwells and elevators or to cover the entire campus including car parks and roads. Apart from impacting the number of APs required, the coverage requirements can introduce other requirements, such as specialized antennas, outdoor enclosures, and lightning protection.

Power Level and Antenna Choice

Power level and antenna choice go hand-in-hand to determine AP placement. Together, these two variables determine where and how powerful the RF is in any given place in the environment. Along with choosing the correct antenna to produce the required coverage area, Cisco recommends that you use RRM to control the power level and provide the optimal channel/power plan. For more information, see Radio Resource Management, page 3-27.

An antenna gives the wireless system three fundamental properties:

- Gain—A measure of increase in power introduced by the antenna over a theoretical (isotropic) antenna that transmits the RF energy equally in all directions.
- Direction—The shape of the antenna transmission pattern. Different antenna types have different radiation patterns that provide various amounts of gain in different directions.
- Polarization—Indicates the direction of the electric field. An RF signal has both an electric and magnetic field. If the electric field is orientated vertically, the wave is said to be vertically polarized.

A good analogy for how an antenna works is the reflector in a flashlight. The reflector concentrates and intensifies the light beam in a particular direction similar to what a parabolic dish antenna does to an RF source in a radio system.

Gain and direction mandate range, speed, and reliability while polarization affects reliability and isolation of noise.

For more information on antenna selection, see the Cisco *Antenna Selection Guide* at:

http://www.cisco.com/en/US/prod/collateral/wireless/ps7183/ps469/product_data_sheet09186a008008883b.html

Omni-Directional Antennas

Omni-directional antennas have a different radiation pattern compared to isotropic antennas; the isotropic antenna is theoretical and therefore all physical antennas are different to the isotropic antenna. The omni-directional antenna features a radiation pattern that is nearly symmetric about a 360 degree axis in the horizontal plane and 75 degrees in the vertical plane (assuming the dipole antenna is standing vertically). The radiation pattern of an omni-directional antenna generally resembles a donut in shape.

Regarding antenna choice, you must consider the RF pattern produced by the antenna because the type of antenna (omni or directional) affects RF coverage by focusing the bulk of the RF energy in a specific direction, pattern, and density.

For example, the omni-directional antenna shown in Figure 3-10 shows an omni-directional antenna RF radiation pattern in the vertical and horizontal direction. This is an actual measurement so it does not follow the donut lines perfectly, but does show from where this shape comes. As described above, other RF-affecting variables (people in the room, amount of devices stored in the facility, leaves on trees for outdoor deployment, interference from different RF sources, and so on) can affect the real RF coverage pattern.

Figure 3-10 Omni-Directional RF Pattern

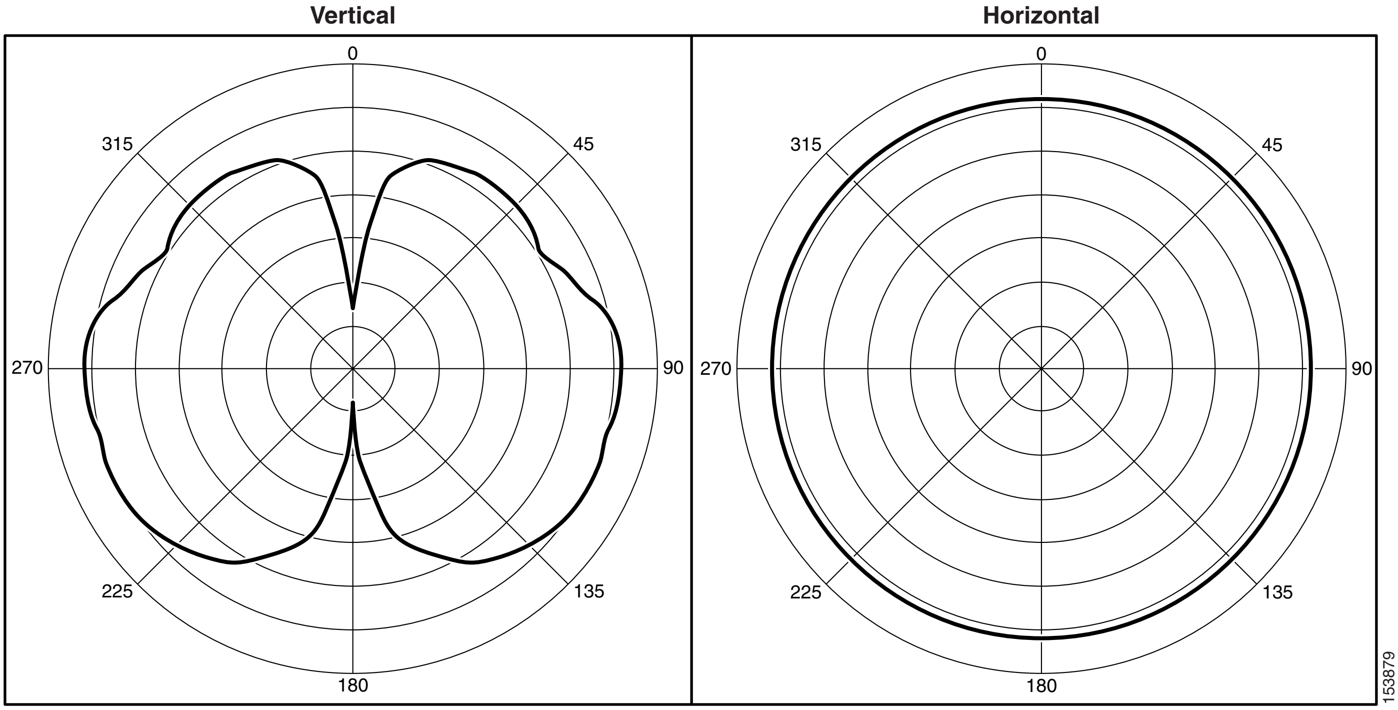

With regards to the pattern in Figure 3-10, this might not be the right antenna to use on a wall especially if it is mounted along an exterior wall where the pattern can radiate outside of the building.

Patch Antennas

The patch antenna is a type of directional antenna. Patch antennas not only radiate away from the wall or place where they are mounted, but also have rear and side lobes that produce a weakened but still potentially useful RF region. Figure 3-11 shows the real horizontal pattern of a patch wall mount antenna. Although most of the coverage area is in front of the patch antenna, notice the back and side RF pattern from the center area. Again, antenna selection is important because it defines the radiation pattern and where wireless connectivity is possible.

Figure 3-11 Patch Wall Mount Antenna Horizontal Plane

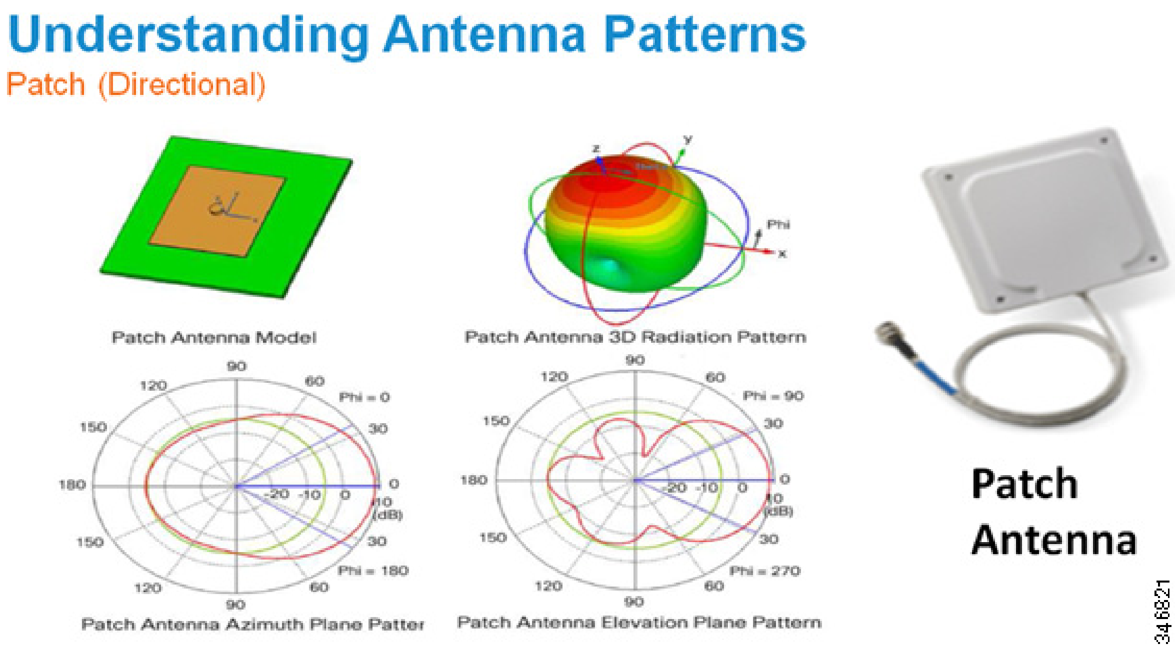

Dipole Antennas

Dipole antennas (see Figure 3-12) are the most basic type of radio antennas. They come in various geometries with different feeding mechanisms and radiating elements. Dipole antennas are the simplest and most practical antennas from a theoretical point of view and as of today, probably the most common antenna type.

Figure 3-12 Dipole Antenna

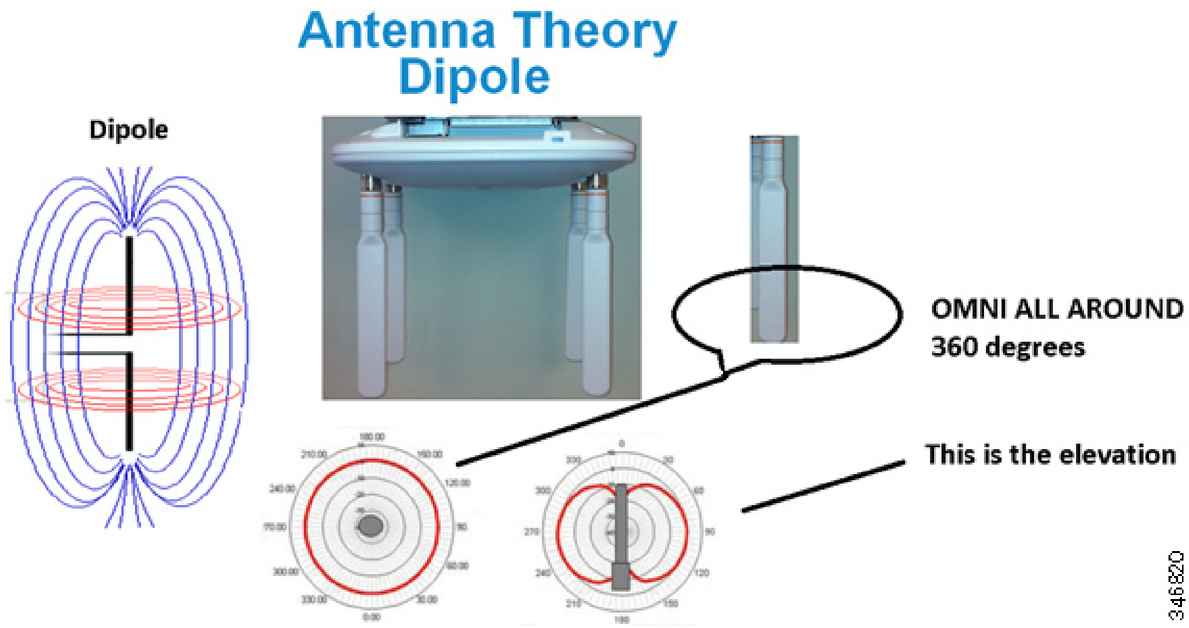

Security Policy Requirements

A well designed RF deployment can effectively minimize unintended RF radiation in areas not requiring coverage. For example, if WLAN coverage is required only in buildings and not outside, then the amount of RF coverage outside of the buildings can be minimized by using the correct power setting, AP placement, and directional antennas pointing inwards towards the center of the building or areas. By tuning RF transmit levels and using the correct antenna for the coverage area, you can reduce the amount of RF that radiates outside the buildings to decrease the security exposure. This can reduce the exposure of wireless network to hackers outside the building or coverage area and avoid a compromise of the wireless network.

RF Environment

The performance of the WLAN and its equipment depends on its RF environment, equipment, selection, coverage design, quality of audits, configurations, and quality of deployment. Adverse environmental variables can disrupt wireless communications by either providing interference on the channel or in some way changing the RF characteristics of the signal. Examples include:

- 2.4 GHz cordless phones and Bluetooth
- Walls fabricated from wire mesh and stucco
- Filing cabinets and metal equipment racks
- Wireless cameras
- Heavy duty electric motors, welders, robots, and things that potentially make sparks
- Fire walls and fire doors as they contain additional metal that reflects radio waves
- Concrete

- Refrigerators
- Air conditioning duct-work
- Other radio equipment such as amateur (ham) radios
- Microwave ovens, especially industrial paint drying equipment
- HVAC ducts near antennas can cause directionality or increased retries
- Large transient elements such as forklifts or metal fabrications
- Other WLAN equipment that is not part of your equipment (for example, a nearby company)

A site survey might be required to ensure that the required data rates are supported in all of the required areas, often driven by the environmental variables mentioned above. A WLC can be an excellent resource for use in site pre-planning and initial identification of RF challenges as well as for identifying channel and power settings.

RF Deployment Best Practices

Some design considerations can be addressed by general best practice guidelines. The following applies to most situations:

- Cisco recommends that for a given AP the number of users per AP be:
  - 15 to 25 for data-only users
  - 7 to 8 voice users (using Cisco 792x VoIP wireless handsets or similar) when data is present

  This number should be used as a guideline and can vary depending on the handset in use. Check your handset requirements.
- The AP data rates should be limited to those designed and for which the site survey was performed. Enabling lower data rates can cause increases in co-channel interference and greater throughput variations for clients.
- The number of APs depends on coverage and throughput requirements, which can vary. For example, the Cisco Systems internal information systems (IS) group currently uses six APs per 3000 square feet of floor space for data-only operation.

Note Based on the variability in environments, Cisco recommends that a site survey be performed to determine the number of APs required and their optimal placement.

Manually Fine-Tuning WLAN Coverage

A number of factors can affect the WLAN coverage. They include:

- Channel and data rate selection
- Overlapping WLAN coverage for location-based services, voice, or data-only
- Power level
- Antenna choice (directional or omni-directional antenna)

For a given data rate and location, the WLAN designer can alter power levels and/or elect to use a different antenna to effect changes to the coverage area and/or coverage shape. Altering power levels or channel selection can be done manually as described below, or Cisco Prime Infrastructure can do this automatically by way of the RRM algorithms. Cisco recommends that you use RRM to control the power level and channel, keeping in mind that the channel changing algorithm is highly dampened so that only a very disruptive (and persistent) interference source can cause a change to the channel topology. This can then cause clients to reassociate and voice calls to be dropped. Changes in AP power do not impact clients. See Figure 3-16 for more details.

Channel and Data Rate Selection

Channel selection depends on the frequencies that are permitted for a particular region. For example, the North American and ETSI 2.4 GHz band permits allocation of three non-overlapping channels (1, 6, and 11), while the 5 GHz band permits 23 channels.

The channels should be allocated to the coverage cells as follows:

- Overlapping cells should use non-overlapping channels.
- Where channels must be re-used in multiple cells, those cells should have minimal overlap with each other. Figure 3-13 shows this overlap pattern. In 802.11a/n/ac deployments, adjacent channels should be avoided as overlapping cells.

Recommendations for Channel Selection

Figure 3-13 shows an example of a typical 2.4-GHz band channel configuration. While channel selection is typically done automatically, it can also be done manually as described in Manual Channel Selection, page 3-24.

Figure 3-13 Channel Allocated To APs

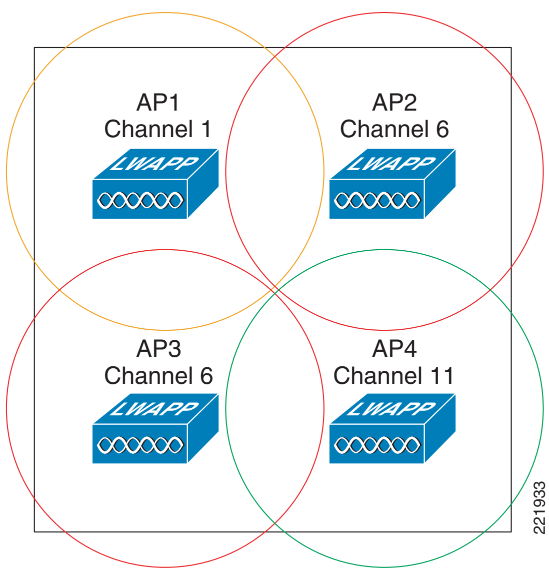

A site survey should be conducted using the same frequency plan as intended for the actual deployment. They are used to test the environment for connectivity and best AP placement. Some sites have high noise backgrounds that could prohibit the use of one or more channels. The survey provides a better estimate of how a particular channel at a particular location will react to the interference and the multipath. Channel selection also helps in planning for co-channel and the adjacent channel interference, and provides information about where you can reuse a frequency (see Figure 3-14).

Note For more information on site surveys, see the publication *802.11 Wireless Network Site Surveying and Installation* (Cisco, 2005), by Bruce E. Alexander.

In multi-story buildings, check the cell overlap between floors, especially where windows might be located, according to the guidelines described here. Careful pre-planning and selection of AP location might be required in approximately 10 percent of the cases. Multi-story structures such as office towers, hospitals, and university classroom buildings introduce a third dimension to coverage planning. The 2.4 GHz waveform of 802.11b and 802.11g can pass through many walls. The 5 GHz waveform of 802.11a/n/ac has approximately half the tendency for a given power to transmit suitable amounts of energy through walls because of its higher frequency. With 2.4 GHz WLANs in particular, you must not only avoid overlapping cells on the same floor, but also on adjacent floors when coverage models include cells that cover windows on both floors. With only three channels, this can be achieved through careful three-dimensional planning.

As a final step, after setting up the WLAN network you should always retest the site using the selected channels and monitor for any interference. Keep in mind that the RRM algorithms are logical and subject to the physical topology of the network. It therefore takes into account the three-dimensional placement of APs and provides the optimal channel/power setting for the sampling interval.

Manual Channel Selection

Figure 3-14 is a screenshot of the Cisco WLC window for configuring one of the 802.11b radios under the wireless selection. On the top-right side, channel 11 has been manually selected and the transmit power is set to 1, the highest level (8 is the lowest).

Note The assignment method should normally be left at the global setting unless there is a desire to manually control these settings. This allows the WLC to dynamically change the channel number and transmit power as determined by RRM. See Radio Resource Management, page 3-27 for more information.

Figure 3-14 Channel Assignment

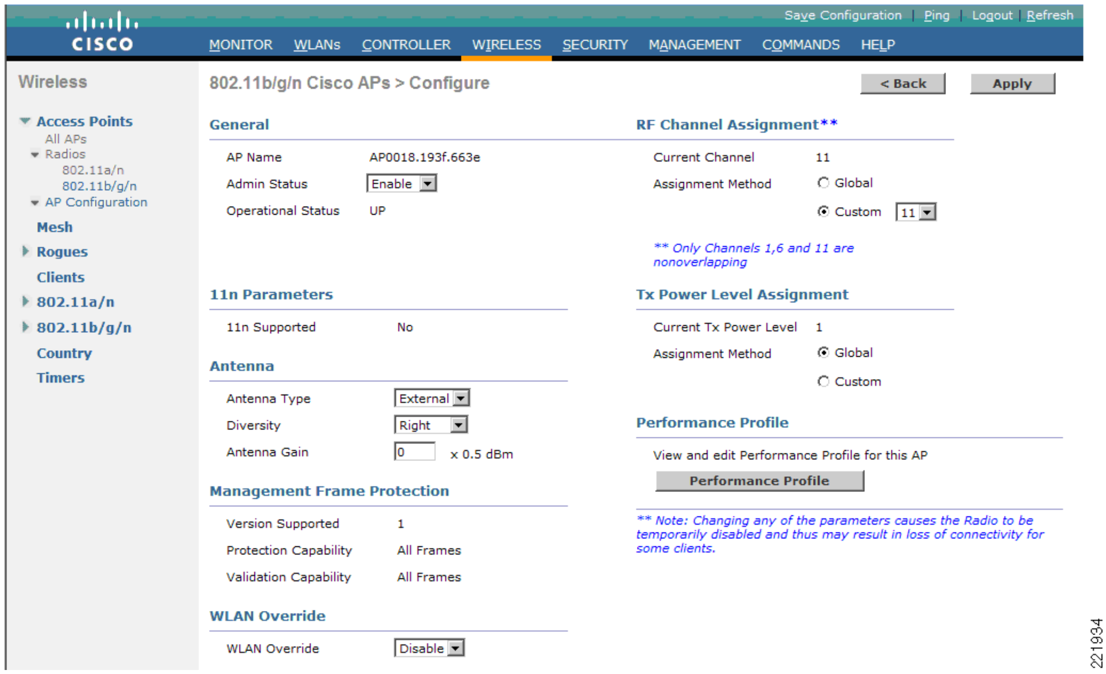

It is also possible to implement a dual band deployment scheme, as shown in Figure 3-15. The top left portion of the diagram shows the 802.11b/g/n-only deployment, which uses the three non-overlapping channels (channels 1, 6, and 11) to map out a pattern that has the least co-channel interference; that is, interference from an AP close by that is on the same channel that is operating at sufficient power levels with its coverage pattern overlapping with that of another access point. It also shows an 802.11a/n/ac deployment, which uses the eight non-overlapping channels. The right side of the diagram illustrates how the channels would be mapped in a dual band deployment.

Figure 3-15 Dual Band Deployment Diagram

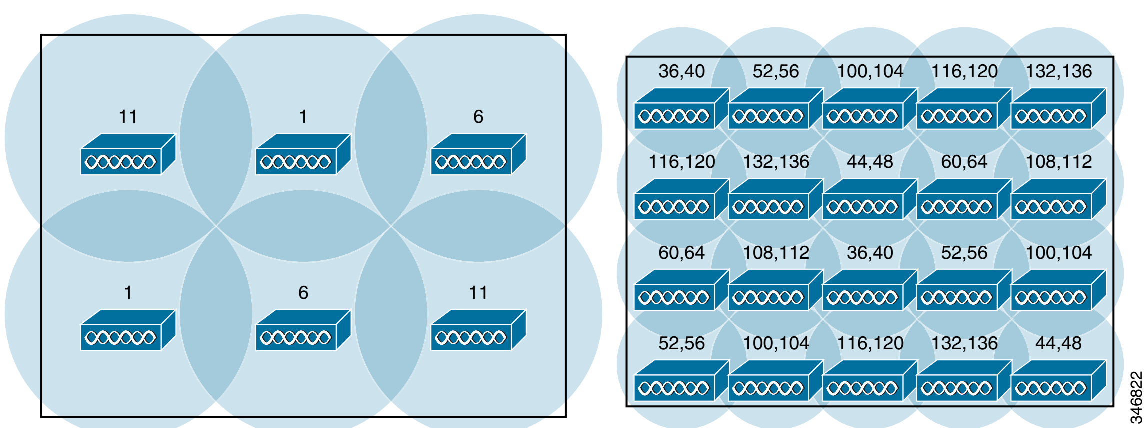

Data Rate Selection

Figure 3-16 is a screenshot of the Cisco WLC window for managing the global 802.11b/g/n parameters. The data rate settings are shown on the right side of the screen.

Figure 3-16 Data Rate Assignment

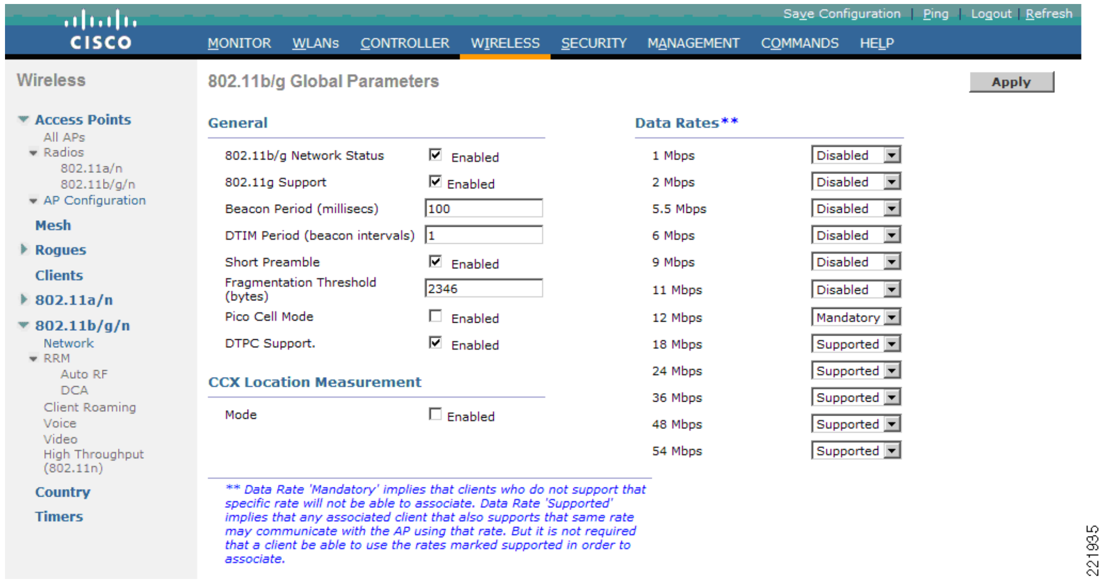

Data Rate Modes

You can use the data rate settings to choose which data rates the wireless device can use for data transmission. There is a direct correlation between data rates, range, and reliability. The lower the data rate the greater the reliability and range for a given power setting. Sites vary for specifics, but a general rule is that for carpeted space you need to increase reliability by an order of magnitude every time you halve the data rate. Range is generally affected by a factor of an approximate 30 percent increase for every halving of data rate. Managing the square footage of the area covered within a -67 dBm edge can be effectively managed using this technique. Setting the data rates to match client, application, or user needs is an effective RF design element that should be considered before deploying APs.

Data rates are expressed in megabits per second. You can set the mode of each data rate as mandatory, supported, or disabled.

Mandatory Mode

Mandatory mode allows transmission for all packets, both unicast and multicast. The data rate on at least one of the APs must be set to Mandatory, and all clients that associate to the AP must be able to physically support this data rate on their radio to use the network. Additionally, for the wireless clients to associate to the AP they must be able to currently receive packets at the lowest mandatory rate and their radios must physically support the highest mandatory data rate. If more than one data rate is set to mandatory, multicast and broadcast frames are sent at the highest common mandatory transmission rate of all associated clients (the lowest mandatory receive rate of all of the clients). This allows all clients to receive broadcast packets. The lowest mandatory rate is normally set at 1 Mbps.

Supported Mode

Supported mode allows transmission for unicast packets only. The AP transmits only unicast packets at this rate; multicast and broadcast packets are transmitted at one of the data rates set to Mandatory. The wireless clients always attempt to transmit and receive at the highest possible data rate. They negotiate with the AP for the highest data rate set to supported or mandatory to transmit and receive unicast packets. The wireless client devices are able to receive broadcast or multicast packets at any mandatory rate at or below the negotiated rate.

Disabled

The AP does not transmit data in this setting.

Lowest and Highest Mandatory Rate Settings

Multiple clients associated to the AP can have different transmission rates, depending on interference, obstacles, or their distance from the AP. For example, if an 802.11b client is far from the AP and can only transmit and receive at a speed of 1 Mbps because of the distance, it would be able to associate to the AP because the lowest mandatory rate (see Figure 3-16) is set to 1 Mbps. If a second 802.11g client associates to the AP at 54 Mbps, the AP transmits broadcasts and multicasts at 1 Mbps because this is the highest mandatory rate that all clients can receive. If the lowest mandatory rate was set to 5.5 Mbps, the 802.11b client would not be able to associate to the AP because it could not receive broadcast packets at the lowest mandatory rate.

In Figure 3-16, note that the highest mandatory setting is 11 Mbps. The highest mandatory rate tells the AP at what rate the client radio must be able to transmit physically. This does not mean that they are actually transmitting and receiving packets at that rate; it just means that the radio physically supports that rate. The wireless client needs only to be able to receive packets at the lowest mandatory rate. 802.11b devices are able to associate to the AP shown in Figure 3-16 because their radios can physically transmit at 11 Mbps. If a higher data rate (such as 18 Mbps) is set to mandatory, only 802.11g clients are able to associate to the APs.

Setting any of the OFDM rates (rates above 1 Mbps) to mandatory disables 802.11b connectivity. This can, for example, allow the administrator to exclude 802.11b clients from the AP by requiring an 802.11g data rate or setting a minimum transmission rate of all clients by disabling 802.11 rates. The reason this might be done is that the same 1500 byte packet at a lower data rate takes a longer time to transmit and, thus, lowers the effective data rate for all wireless clients associated to the AP.

Radio Resource Management

In the Cisco WLAN *split MAC* architecture (see Chapter 2, "Cisco Unified Wireless Technology and Architecture"), the processing of 802.11 data, management protocols, and AP capabilities is distributed between a CAPWAP AP and a centralized WLC. More specifically, time-sensitive activities such as probe response and MAC layer encryption are handled by the APs while all other functions are sent to the WLCs where system-wide visibility is required.

Real-time RF management of a WLAN requires system-wide visibility and is implemented at the WLC level. The controller learns about the necessary information for an effective RF channel/power plan by way of information forwarded by the APs in the RF network group.

Note An RF network group (or RF group) is not the same as a mobility group. A mobility group defines a mobility domain of 1 to 24 WLCs in which a client would not be required to change IP address during a roaming event. This is accomplished by creating Ethernet over IP tunnels for forwarding client data from an *anchor* controller to the *foreign* controller handling the new AP servicing the client.

Radio resource management (RRM) algorithms adjust the channel and power (using dynamic channel assignment and dynamic transmit power control, respectively) to maintain the RF coverage area. RRM adjusts the power level of the AP to maintain a baseline signal strength with neighboring APs at (configurable) -65 dBm (see Overview of RRM Operation). RRM adjusts the channel of the AP when it notices nearby interference sources on the channel that the AP is currently located. RRM continues to optimize the RF coverage for the best reception and throughput for the wireless network.

Note The transmit power control and dynamic frequency management performed by RRM are not the TPC and DFS required for operation in the UNII-2 bands that are defined in 802.11h.

RRM understands that the RF environment is not static. As different RF affecting variables change (people in the room, amount of devices stored in the facility, leaves on trees for outside deployment, interference from different RF sources, and so on), the RF coverage adjusts to these variables and changes with them. Because these variables change continuously, it is necessary to periodically monitor for the RF coverage and adjust it accordingly.

Overview of RRM Operation

If group mode is enabled in the WLCs, they select a leader in each RF group and form an RF domain. The function of the leader is to collect the network-wide neighbor data packets for the APs from the group of WLCs and do channel/power computation for an optimal system-wide map. If group mode is not enabled, the controllers run computations based only on the neighbor data gathered from the APs connected by way of CAPWAP, trying to optimize the signal to -70 dBm between APs.

The APs transmit RRM neighbor data packets at full power at regular intervals (neighbor data packets include information about the signal strength and radio elements in the environment). These messages contain a field that is a hash of the RF network (group) name, BSSID, and time stamp. The APs accept only RRM neighbor packets sent with the default RF network name (*RF Network Name*).

When neighboring APs receive neighbor data packets they validate data packets before forwarding them to the RF controller. If they can validate the message hash and confirm that it belongs to the same RF group, the packet is sent to the RF group controller; otherwise, the AP drops the neighbor packet. When APs forward the validated messages to the WLC they fill in the CAPWAP packet status field with the SNR and RSSI of the received neighbor packet.

Table 3-4 provides a summary of the various functions of the devices in the system.

Note TPC performs only downward power level adjustments. Coverage hole detection and correction increases power levels on APs.

Table 3-4 Device Function

Device	Functions
RF Group Leader	WLC that collects AP neighbor data from the other WLCs in the RF group and analyzes it for system-wide TPC and DCA. TPC adjusts power levels downward.

Table 3-4 Device Function (continued)

Local WLC	Collects data and runs the Coverage Hole Detection and Correction algorithm. Adjusts power levels upward if necessary for clients
CAPWAP AP	• Sends neighbor messages on all channels at full power at configured interval • Verifies neighbor hash on received neighbor messages • Scans configured channels for noise, interference, IDS/rogue detection, and alerts

RRM must not be confused with *rogue detection* (channel scanning), which is done separately from the RRM algorithm. APs perform rogue detection by periodically scanning all country-specific channels (channel scanning) for rogue APs. The APs goes *off-channel* for a period not greater than 60 ms to listen to the other channels. Packet headers collected during this time are sent to the WLC, where they are analyzed to detect rogue APs, whether service set identifiers (SSIDs) are broadcast or not, rogue clients, ad-hoc clients, and interfering APs.

By default, each AP spends approximately 0.2 percent of its time off-channel. This is statistically distributed across all APs so that no two adjacent APs are scanning at the same time, which can adversely affect WLAN performance. Client packets received by the AP are forwarded to the WLC with the CAPWAP status field filled in. This field provides the WLC with radio information including RSSI and signal-to-noise ratio (SNR) for all packets received by the AP during reception of the packet.

RRM Configuration Settings

Radio resource management (RRM) also called as *Auto-RF* in the WLC, can be turned on or off in the WLC by way of a global setting of Channel Selection (see Figure 3-14). From this window you can also manually set the channel and transmit level for the APs. Additionally, RRM can be turned on or off from the WLC global Auto-RF configuration window. Keep in mind that RRM is performed on a per band basis. RF group computations for the 5 GHz band are separate from the computations for the 2.4 GHz band.

The Auto-RF configuration window is divided into three sections that you move between using the scroll bar on the right. The first section (see Figure 3-17) is for dynamic channel assignment. This allows the WLC to automatically change the channel that the AP is on (for more information, see Dynamic Channel Assignment).

Figure 3-17 Auto-RF (Section 1)

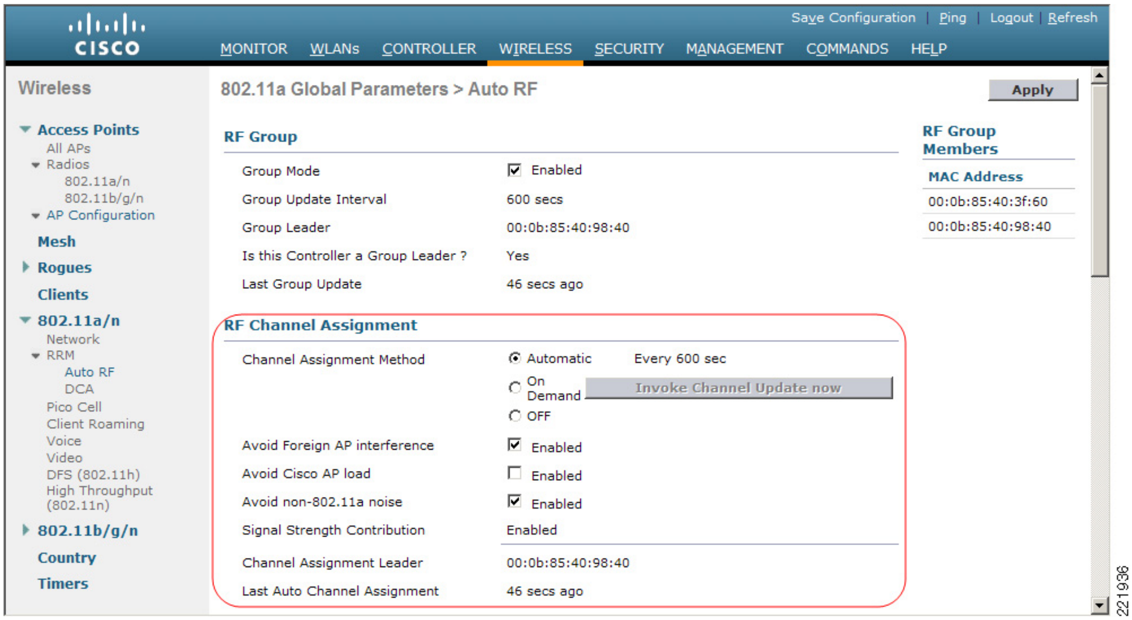

Note For more information, see the *Cisco Wireless LAN Controller Configuration Guide,* http://www.cisco.com/en/US/docs/wireless/controller/7.3/configuration/guide/b_cg73.html

The first group of settings in the Auto-RF configuration window, the RF Group, is used to determine whether the WLC joins the dynamic grouping with the other WLCs in the group. Dynamic grouping helps the WLC find out about APs that are neighbors but might be associated to another WLC in the mobility group. If this is disabled, the WLC only optimizes the parameters of the access points that it knows about (that is, the ones that are associated to it). The group leader indicates the MAC address of the elected leader. You can find the MAC address of the controller in the WLC Inventory window by clicking on the *Controller* menu at the top and then selecting *Inventory.*

The second section is for assigning the transmit (Tx) power level (Figure 3-18). The power level can be fixed for all APs, or it can be automatically adjusted. This window also indicates the number of neighbors the AP has and the power thresholds for which it is adjusting.

Figure 3-18 Auto-RF (Section 2)

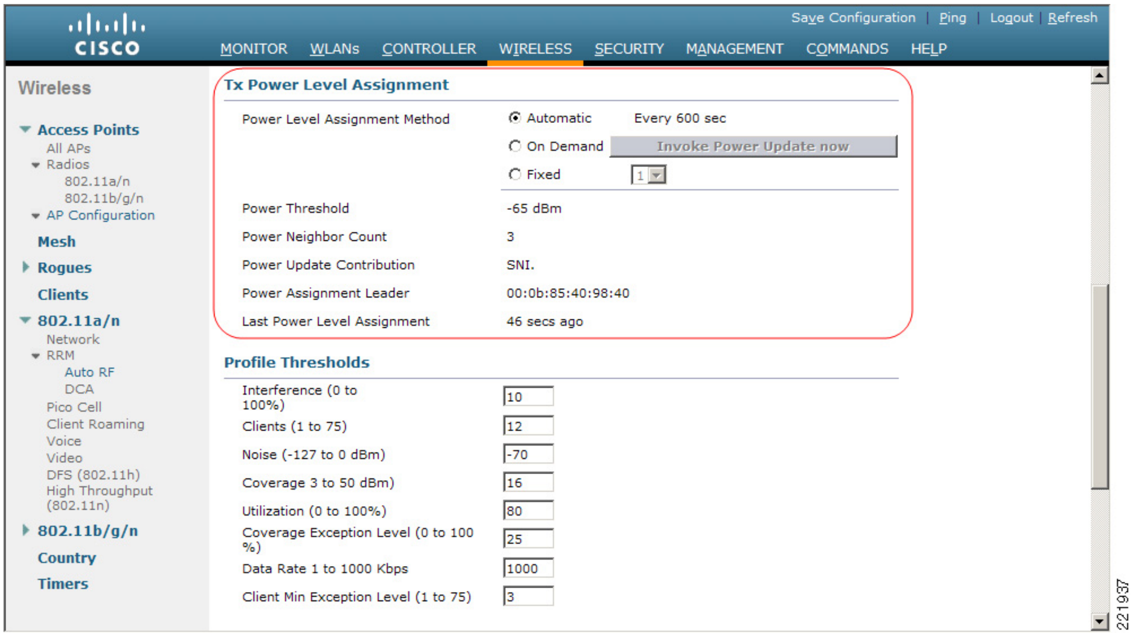

The third section (Figure 3-19) is for profile thresholds.

Figure 3-19 Auto-RF (Section 3)

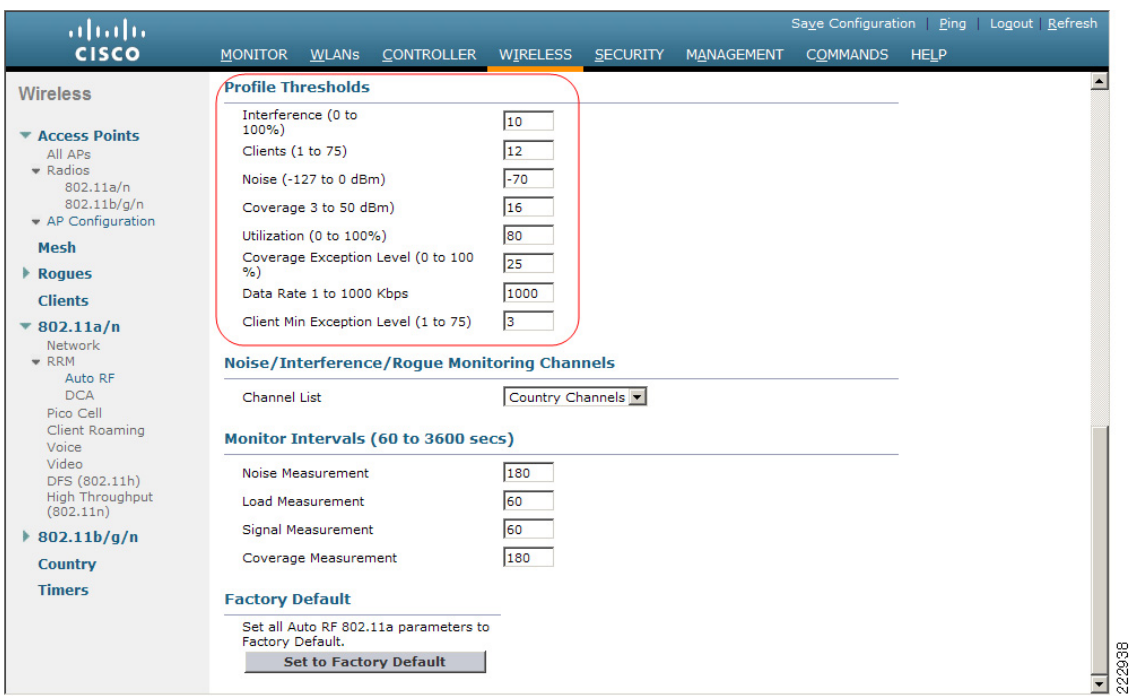

Sample 'show ap auto-rf' Command Output

The WLC analyzes the information passed to it by the APs and determines a pass/fail status for each of these thresholds. These pass/fail profiles are best seen in the output of the following WLC CLI **show ap auto-rf** *radio ap_name* command. This command displays RF statistics from the radio being used.

Note The same information can be seen in graphical form in the WLC **Monitor -> 802.11b/g/n Radios -> Detail** window.

```
show>ap auto-rf 802.11b <access point name>
Number of Slots . . . . . . . . . . . . . . . . . . . . . 2
AP Name . . . . . . . . . . . . . . . . . . . . . . . . . <AP name>
MAC Address . . . . . . . . . . . . . . . . . . . . . . . 00:0b:85:1b:df:c0
Radio Type . . . . . . . . . . . . . . . . . . . . . . . . RADIO_TYPE_80211b/g
Noise Information
    Noise Profile . . . . . . . . . . . . . . . . . . . . . . PASSED
    Channel 1 . . . . . . . . . . . . . . . . . . . . . . . . -93 dBm
    Channel 2 . . . . . . . . . . . . . . . . . . . . . . . . -90 dBm
 .
 .
 .
    Channel 11 . . . . . . . . . . . . . . . . . . . . . . . . -95 dBm
Interference Information
    Interference Profile . . . . . . . . . . . . . . . . . . . FAILED
    Channel 1 . . . . . . . . . . . . . . . . . . . . . . . . -69 dBm @ 31 % busy
    Channel 2 . . . . . . . . . . . . . . . . . . . . . . . . -58 dBm @ 26 % busy
 .
 .
 .
    Channel 11. . . . . . . . . . . . . . . . . . . . . . . -68 dBm @ 26 % busy
Load Information
    Load Profile . . . . . . . . . . . . . . . . . . . . . . . PASSED
    Receive Utilization . . . . . . . . . . . . . . . . . . . . 0 %
    Transmit Utilization . . . . . . . . . . . . . . . . . . . . 0 %
    Channel Utilization . . . . . . . . . . . . . . . . . . . 26 %
    Attached Clients . . . . . . . . . . . . . . . . . . . .  2 clients
Coverage Information
    Coverage Profile . . . . . . . . . . . . . . . . . . . .. PASSED
    Failed Clients . . . . . . . . . . . . . . . . . . . . . . 0 clients
Client Signal Strengths
    RSSI -100 dBm. . . . . . . . . . . . . . . . . . . . . . . 0 clients
    RSSI -92 dBm . . . . . . . . . . . . . . . . . . . . . . . 0 clients
 .
 .
 .
    RSSI -52 dBm . . . . . . . . . . . . . . . . . . . . . . . 1 clients
Client Signal To Noise Ratios
    SNR 0 dBm . . . . . . . . . . . . . . . . . . . . . . . .  0 clients
    SNR 5 dBm . . . . . . . . . . . . . . . . . . . . . . . .  0 clients
    SNR 10 dBm . . . . . . . . . . . . . . . . . . . . . . . . 0 clients
 .
 .
 .
    SNR 45 dBm . . . . . . . . . . . . . . . . . . . . . . . . 1 clients
Nearby APs
Radar Information
Channel Assignment Information
    Current Channel Average Energy . . . . . . . . . . . . . . -68 dBm
    Previous Channel Average Energy . . . . . . . . . . . . . -51 dBm
    Channel Change Count . . . . . . . . . . . . . . . . . . . 21
```

```
    Last Channel Change Time . . . . . . . . . . . . . . . . Thu Mar 9 12:18:03 2006
    Recommend Best Channel . . . . . . . . . . . . . . . . . 11
RF Parameter Recommendations
    Power Level . . . . . . . . . . . . . . . . . . . . . . . 1
    RTS/CTS Threshold . . . . . . . . . . . . . . . . . . . . 2347
    Fragmentation Threshold . . . . . . . . . . . . . . . . . 2346
    Antenna Pattern . . . . . . . . . . . . . . . . . . . . . 0
```

The following sections describe some of the WLC RRM settings.

Dynamic Channel Assignment

The 802.11 MAC layer uses Carrier-Sense Multiple Access/Collision Avoidance (CSMA/CA). Because of CSMA/CA-managed channel sharing, two APs on the same channel in the same vicinity get approximately half the capacity of what two APs would get on different channels in the same vicinity. This is caused by the 802.11 MAC sensing that the channel is busy, and deferring sending frames until the channel has become free. If the 802.11 MAC defers traffic that is not part of its own AP cell, this is considered interference. Interference from another AP on the same channel is commonly called *co-channel interference*, and is to be expected in most 2.4 GHz band 802.11 deployments. This is because there are insufficient non-overlapping channels to prevent some channel overlap from occurring. One of the goals of design, planning, and dynamic radio management is to minimize the amount of co-channel overlap, which minimizes co-channel interference and therefore maximizes AP traffic capacity. The Cisco Unified Wireless Network addresses this problem and other co-channel interference issues by dynamically allocating AP channel assignments to avoid conflict. Because the WLC, or a designated WLC (RF Group Leader), has system-wide visibility it can control how channels are *reused* to minimize co-channel interference.

The WLC examines a variety of real-time RF characteristics to efficiently handle channel assignments, including the following:

- Noise—Noise limits signal quality at the client and AP and can vary in range and periodicity. There are numerous potential types and effects of interference. For one, an increase in noise reduces the effective cell size. At regular intervals the WLC reassesses the RF environment of an AP and then optimizes channel selection to avoid noise sources while still maintaining overall system capacity. Channels that become unusable because of excessive noise can be avoided. If other wireless networks are present, the WLC shifts its channel usage to complement the other networks. For example, if one network is on Channel 6, the adjacent WLAN is assigned Channel 1 or 11. This increases the capacity of the network by limiting the sharing of frequencies. If a channel is used so much that no capacity is available, the WLC might choose to avoid this channel.
- Client load—Client load is taken into account when changing the channel structure to minimize the impact on the clients currently on the WLAN. The WLC periodically monitors the channel assignment in search of the best assignments. Change occurs only if it significantly improves the performance of the network or corrects the performance of a poorly performing AP.

The WLC combines the RF characteristic information to make system-wide decisions. The end result is an optimal channel configuration in a three-dimensional space, where APs on the floor above and below are factored into an overall WLAN configuration.

Interference Detection and Avoidance

Interference (as it pertains to a Cisco Unified Wireless Network deployment) is defined as unwanted RF signals in the same frequency band that can lead to a degradation or loss of service. These signals can either be from 802.11 or non-802.11 sources such as certain microwave ovens or certain cordless phones. It can, in certain instances, also include various sources of electromagnetic interference (EMI) such as arc welders or radar facilities. APs constantly scan all channels for major sources of interference and regularly reports this and additional information to the WLC by way of a management link interface (CAPWAP tunnel).

If the amount of 802.11 interference hits a predefined threshold, the WLC attempts to rearrange channel assignments to optimize system performance in the presence of the interference. This might result in adjacent APs being on the same channel. Logically this represents a better scenario than staying on a channel that is otherwise totally unusable because of interference. For example, the WLC can respond to a rogue AP on channel 11 by shifting nearby APs to channel 1 or channel 6.

Dynamic Transmit Power Control

The appropriate AP transmit power settings are essential to maintaining a coverage area, not only to ensure correct (not maximum) amount of power covering an area, but also to ensure that excessive power is not used, which would add unnecessary interference to the radiating area. AP power settings are also used to control network redundancy by helping to ensure real-time failover in the event of the loss of an AP. The WLC is used to dynamically control the AP transmit power-level based on real-time WLAN conditions. In normal instances, power can be minimized to gain extra capacity and reduce interference among the APs. RRM attempts to balance APs so that they see their neighbors at -65 dBm. If an AP outage is detected, power can be automatically increased on surrounding APs to fill the coverage gap created by the loss.

RRM algorithms are designed to create the optimal user experience. For example, if the power of an AP is turned down to Level 4 (where Level 1 = highest and Level 8 = lowest) and the RSSI value of a user drops below an acceptable threshold, the AP power is increased to provide a better experience to that client. When dynamic transmit power control (DTPC) is enabled, the AP adds channel and transmit power information to *beacons* (information elements that contain information such as channel, RF power, network name, and so forth). Client devices using DTPC receive the information and adjust their settings automatically.

Coverage Hole Detection and Correction

The coverage hole detection and correction algorithm is aimed at determining coverage holes based on the quality of client signal levels and then increasing the transmit power of the APs to which those clients are associated.

The algorithm determines whether a coverage hole exists when client signal-to-noise ratio (SNR) levels pass below a given SNR threshold. The SNR threshold is considered on an individual AP basis and based primarily on the transmit power of each AP.

When the average SNR of a single client dips below the SNR threshold for at least 60 seconds, this is seen as an indication that the WLAN client does not have a viable location to roam to. The AP transmit power of that client is increased, correcting the coverage hole.

Client and Network Load Balancing

The IEEE 802.11 standard does not define the process or reasons for client roaming; therefore, it cannot be easily predicted what clients will do in any given situation. For example, all users in a conference room can associate with a single AP because of its close proximity, ignoring other APs that are farther away but with greater free capacity.

The WLC has a centralized view of client distribution across all APs. This can be used to influence where new clients attach to the network if there are multiple *good* APs available. If configured, the WLC can proactively use AP probe responses to guide clients to the most appropriate APs to improve WLAN performance. This results in a smooth distribution of capacity across an entire wireless network. Keep in mind that this load balancing is done at the time of client association, not after a client is connected.