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

Jim Florwick
Jim Whiteaker
Alan Cuellar Amrod
Jake Woodhams

Design Guide

November, 2013

For further information, questions and comments please contact ccbu-pricing@cisco.com

About the Guide

This design guide provides engineering guidelines and practical techniques for designing, planning, and implementing a wireless LAN (WLAN) within a high-density environment in a university or college campus. High-density is defined as any environment with a large concentration of users, such as a classroom, lecture hall, or auditorium where the users are connected wirelessly, sharing applications and using other network services individually.

This document is intended for wireless network design engineers responsible for designing, deploying, and maintaining today's Wi-Fi networks. Knowledge of Cisco® networking concepts, WLAN technology fundamentals, Cisco Unified Wireless Network (CUWN) features and configurations are prerequisites.

Related Documentation

Cisco Mobility 4.1 Design Guide

Cisco Campus Wireless LAN Controller Configuration Design Guide

Optimize the Cisco Unified Wireless Network to Support Wi-Fi Enabled Phones and Tablets

802.11n: Mission-Critical Wireless

Executive Summary

The demands on WLANs for functionality and scalability are growing due to the rapid proliferation of new network devices and applications. The number of devices and connections per user is steadily increasing. It is common for most users today to not only have a primary computing device but also at least one other smart device. Wireless operators have worked hard to accommodate the increased demand for data services over wireless networks. They have been forced to consider alternative offload strategies, including wirelessly connecting electronic devices (Wi-Fi). Unfortunately, the majority of smartphones being introduced into the marketplace only support Wi-Fi at 2.4 Gigahertz (GHz), which is rapidly increasing pressure on Wi-Fi designers and administrators to design products for the smallest segment of bandwidth available. This trend has driven a dramatic increase in user densities, with many users competing for 2.4 GHz services. According to some projections, this competition for resources has just begun. In addition to this rapid increase in demand for an already congested spectrum, new network devices often are designed for use in the home. This is often not well suited for optimal efficiency in an engineered public wireless space.

Administrators are finding themselves faced with the challenge of providing ever-increasing levels of service in areas where simple pervasive coverage was the singular design goal. Simply adding more access points (APs) often does not enhance service. This design guide focuses on the challenges facing administrators deploying WLANs in higher education and offers practical strategies and design guidance for evaluating and modifying current deployment strategies, improving performance with existing resources, and successfully scaling network accessibility in high-density venues.

The best practices discussed have been gathered from multiple venues and have been used to successfully deploy high-density wireless networks throughout the world. While the guide primarily focuses on requirements for a large, network-connected lecture hall, the principles discussed will provide the reader with the tools necessary to successfully increase density in a wide variety of other shared network environments.

Introduction

While there have been great advances made in the speed and ease of implementation of Wi-Fi networks, the basic nature of radio frequency (RF) is generally unchanged. Increasing the number of users who can access the WLAN in a small physical space remains a challenge. The steps and process for a successful high user density WLAN design that can be proven, implemented, and maintained using Cisco's Unified Wireless Network architecture is detailed. It includes these general steps:

- Plan: Determine application and device requirements such as bandwidth, protocols, frequencies, service level agreement (SLA), etc.
- Design: Determine density, cell sizing, antennas, coverage, site survey, etc.
- Implement: Install, test, tune, establish baseline, etc.
- Optimize: Monitor, report, adjust, review baseline for SLA.
- Operate: Cisco Wireless Control System (WCS) monitoring, troubleshooting tools, capacity monitoring and reporting tools, etc.

The general concepts underlying high-density Wi-Fi design remain true for many environments. But it is important to note that the content and solutions presented here will not fit every WLAN design scenario. Rather, the intent of the guide is to explain the challenges in WLAN design for high-density client environments and to offer successful strategies so that engineers and administrators understand them and are able to articulate the impact design decisions will have.

Target Environmental Characteristics for WLANs in Higher Education Environments

High-density WLAN design refers to any environment where client devices will be positioned in densities greater than coverage expectations of a normal enterprise deployment, in this case a traditional, carpeted office. For reference, a typical office environment has indoor propagation characteristics for signal attenuation. User density is the critical factor in the design. Aggregate available bandwidth is delivered per radio cell, and the number of users and their connection characteristics (such as speed, duty cycle, radio type, band, signal, and SNR) occupying that cell determines the overall bandwidth available per user.

A typical office environment, Figure 1, may have APs deployed for 2,500 to 5,000 square feet with a signal of -67 decibels in milliwatts (dBm) coverage and a maximum of 20 to 30 users per cell. That is a density of one user every 120 square foot (sq. ft.) and yields a minimum signal of -67 dBm.

Figure 1. Typical Office WLAN

In planning and deploying such a WLAN, an AP is typically placed in an area expected to have a higher user density, such as in a conference room, while common areas are left with less coverage. In this way, pre-planning for high-density areas is anticipated. Conference rooms are often placed in clusters, so it is best to design for the maximum capacity of the area. For example, maximum occupancy for the three rooms is 32, so user density would be one user per 28 square feet, Figure 2.

Figure 2. Calculating User Density

In a high-density environment such as a lecture hall or auditorium, the densities of users in the occupied space increase dramatically. User seating is typically clustered very close together to achieve high occupancy. The overall dimensions of the space are really only useful for getting an idea of the free space path loss of the AP signal. User densities are not evenly distributed over the entire space as aisle ways, stages, and podiums represent a percentage of space which is relatively unoccupied. The RF dynamics of the AP are very different from those experienced at the user level. The APs are exposed with an excellent view of the room and the user devices will be packed closely together with attenuating bodies surrounding them.

The single biggest sources of interference in the room are the client devices themselves. For each user sitting in the auditorium who can rest their hand comfortably on the back of the seat in front of them, the distance is approximately three feet, with an average seat width of 24 inches. This yields what is defined as a high-density environment, with less than 1 square meter per device deployed, assuming one or more devices connected per seat.

Figure 3.   Seating and Interference

What is ultimately going to effect the client devices more than any other factor is the degradation of signal-to-noise ratio (SNR) through both co-channel and adjacent channel interference driven by co-located devices. Proper system engineering can minimize the impact by maximizing proper spatial reuse but it cannot be eliminated in highly dense environments entirely. Operating margins become more critical as space is condensed and a bad radio or behavior in the mix can have a large impact within a cell. Client behavior under these conditions will vary widely and trends based on environment and event type have also been reported. There is not much that can be done about the particular client mix or behavior. The design goal is to engineer the network side as robustly as possible and to control and understand all variables.

Within environments that qualify as high-density, there are also submodels built by use case. For example, in a high-density environment such as a public venue or stadium, capacity is planned based on what percentage of users are likely to be active on the network at any one time. In higher education there is a different model, where casual WLAN activity is one use case while activity when a professor is lecturing may increase dramatically, up to 100 percent.

Planning

The WLAN design process can begin in many ways but generally it begins with an expressed desire to provide connections to a specific area where a number of users will participate in a focused activity. To evaluate what is possible, it is first necessary to understand what is required as well as what is possible. There is generally a primary application that is driving the need for connectivity. Understanding the throughput requirements for this application and for other activities that will take place on the network will provide the designer with a per-user bandwidth goal. Multiplying this number by the number of expected connections yields the aggregate bandwidth that will be required.

The required per connection bandwidth will be used to drive subsequent design decisions.

Design Point #1: Establish and Validate a Per-Connection Bandwidth Requirement

How much bandwidth does each user require on average? In Table 1, the nominal throughput requirements for several popular applications and use cases in a higher education setting are shown.

Table 1.   Bandwidth Requirements per Application

Application by Use Case	Nominal Throughput
Web - Casual	500 kilobits per second (Kbps)
Web - Instructional	1 Megabit per second (Mbps)
Audio - Casual	100 Kbps
Audio - Instructional	1 Mbps
On-demand or Streaming Video - Casual	1 Mbps
On-demand or Streaming Video - Instructional	2-4 Mbps
Printing	1 Mbps
File Sharing - Casual	1 Mbps
File Sharing - Instructional	2-8 Mbps
Online Testing	2-4 Mbps
Device Backups	10-50 Mbps

In all cases, it is highly advisable to test the target application and validate its actual bandwidth requirements. Software designers are often required to pick just one average number to represent the application's requirements when there are actually many modes and deployment decisions that can make up a more accurate number. It is also important to validate applications on a representative sample of the devices that are to be supported in the WLAN. Additionally, not all browsers and operating systems enjoy the same efficiencies, and an application that runs fine in 100 kilobits per second (Kbps) on a Windows laptop with Microsoft Internet Explorer or Firefox, may require more bandwidth when being viewed on a smart phone or tablet with an embedded browser and operating system.

Once the required bandwidth throughput per connection and application is known, this number can be used to determine the aggregate bandwidth required in the WLAN coverage area. To arrive at this number, multiply the minimum acceptable bandwidth by the number of connections expected in the WLAN coverage area. This yields the target bandwidth needed for the need series of steps.

Design Point #2: Calculate the Aggregate Throughput Required for the Coverage Area

If this design guide was for a wired rather than wireless network, calculating aggregate throughput requirements would entail dividing the aggregate capacity by the channel bandwidth available. Then, the number of channels would be established and these would be plugged into a switch. But in a WLAN, a channel's speed is effected by multiple factors including protocols, environmental conditions, and operating band of the adapter. Before calculating aggregate throughput, several things must be considered.

In the aggregate throughput calculation, the connections instead of the seats were used as the basis for calculation. The number of connections in a cell is what determines the total throughput that will be realized per connection instead of the number of seats. Most users today carry both a primary computing device (such as a smartphone, tablet computer, or laptop) as well as a second device (such as a smartphone). Each connection operating in the high-density WLAN consumes air time and network resources and will therefore be part of the

aggregate bandwidth calculation. An increase in numbers of device connections is one of the primary reasons older WLAN designs are reaching oversubscription today.

Wi-Fi is a shared medium. Much like an un-switched Ethernet segment, it operates as a half duplex connection. Only one station can use the channel at a time and both the uplink and downlink operate on the same channel. Each channel or cell used in a Wi-Fi deployment represents a potential unit of bandwidth much like an Ethernet connection to a hub. In Ethernet, switching technology was developed to increase the efficiency of the medium by limiting the broadcast and collision domains of a user to a physical port and creating point-to-point connections between ports on an as-needed basis, dramatically increasing the overall capacity.

Users and applications also tend to be bursty (a measure of the unevenness or variations in the traffic flow) in nature and often access layer networks are designed with a 20:1 oversubscription to account for these variances. Application and end user anticipated usage patterns must be determined and also accounted for. Some applications, such as streaming multicast video, will drive this oversubscription ratio down while others may drive this factor even higher to determine an acceptable SLA for each cell's designed capacity.

For 802.11 wireless networks or any radio network in general, air is the medium of propagation. While there have been many advances in efficiency, it is not possible to logically limit the physical broadcast and collision domain of an RF signal or separate it's spectrum footprint from other radios operating in the same spectrum. For that reason, Wi-Fi uses a band plan that breaks up the available spectrums into a group of non-overlapping channels. A channel represents a cell. Using the analogy of Ethernet, a cell represents a single contiguous collision domain.

How many users can access an AP comfortably? Hundreds. But the question should not be how many users can successfully associate to an AP but how many users can be packed into a room and still obtain per-user bandwidth throughput that is acceptable.

802.11 and Scalability: How Much Bandwidth Will a Cell Provide?

To scale 802.11 networks to reliably deliver consistent bandwidth to a large number of users in close proximity, it is important to examine certain WLAN fundamentals under reasonably ideal conditions. Once the rules are understood, the ways to manipulate them to maximum advantage will be presented.

In real WLANs, the actual application throughput is what matters to the end user, and this differs from the signaling speed. Data rates represent the rate at which data packets will be carried over the medium. Packets contain a certain amount of overhead that is required to address and control the packets. The application throughput is carried as payload data within that overhead. Table 2 shows average application throughput by protocol under good RF conditions.

Table 2.   Average Application Throughput by Protocol

Protocol	Throughput (Mbps)
802.11b	7.2
802.11b/g mix	13
802.11g	25
802.11a	25
802.11n - HT20 one spatial stream (1ss) Modulation Coding Scheme 7 (MCS7)	25
802.11n - HT20 2ss Modulation Coding Scheme 15 (MCS15)	70
802.11n - HT40 2ss Modulation Coding Scheme 15 (MCS15)	160

Are 802.11n Data Rates Dependable?

Today, many clients are 802.11n ready and this can provide throughput and efficiency increases in a high-density deployment. Most WLANs, however, will support a mix of client protocols. Evaluating the historic average client mix in a WLAN is possible by either looking at the WLAN controller graphical user interface (GUI) or at Cisco Wireless Control System reports and using this historic mix of information for planning purposes. Unless the WLAN is very unique, most environments will likely be dealing with a diverse mix of clients and protocols for the near future. Consider that other factors, such as the number of connections, can also be expected to vary over time and for these reasons it is often a best practice to build in some buffer to smooth the long term results. The raw speed advantage of 802.11n high throughput (HT) rates is impressive and boosts the overall efficiency and capacity of the design by permitting more users or higher speeds to be realized on the same channel. Figure 4 shows mixed client protocol capacities for a given cell.

Figure 4.   Mixed Wireless Client Protocol Performance in a Cell (802.11a/g/n Data Rates)

The graph above shows throughput rates under varying mixes of HT20 modulation coding scheme-15 (MCS15) 2SS data rates and legacy 802.11a/g (for the purpose of this discussion 802.11a and 802.11g are the same protocol - different bands and are considered equal) data rates within a single isolated cell.

- With either all MCS15 or all 802.11a/g clients, the difference in throughput is 480 percent.
- With a 50/50 mix, there is a 400 percent increase over legacy throughput.
- With a drop to just 25 percent of MCS15 clients, the increase is 300 percent.

In this example using 30 connections, the application throughput to the end user would be 833 Kbps with all legacy connections or 3.9 Mbps with all 802.11n connections. A mix drives throughput down. Other variables, such as user density or environmental noise, can and likely will change over time and will effect the throughput as well.

Using legacy data rates as a nominal value, Table 3 shows the relationship between cell bandwidth and per connection bandwidth.

Table 3.   Data Throughput and User Connections per Wireless Protocol

Protocol	Data Rate (Mbps)	Aggregate Throughput (Mbps)	Example User Count	Average Per User Throughput
802.11b	11	7.2	10	720 Kbps
802.11b	11	7.2	20	360 Kbps
802.11b	11	7.2	30	240 Kbps
802.11b/g	54	13	10	1.3 Mbps
802.11b/g	54	13	20	650 Kbps
802.11b/g	54	13	30	430 Kbps
802.11a	54	25	10	2.5 Mbps
802.11a	54	25	20	1.25 Mbps
802.11a	54	25	30	833 Kbps
802.11n MCS7	72	35	10	3.5 Mbps
802.11n MCS7	72	35	20	1.75 Mbps
802.11n MCS7	72	35	30	1.16 Mbps

A mixed cell containing both 802.11b and 802.11g traffic results in a throughput rate that is less than double that of 802.11b alone and roughly half of 802.11g alone. A similar effect was seen when 802.11n and legacy 802.11a/g rates were compared. Until the inclusion of 802.11n, all advances in Wi-Fi technology have come through incremental increases in encoding technology. 802.11n changed the encoding and streamlined the logistics of bonding 20 MHz channels and increasing the available channel bandwidth. In implementing new technology, it is also necessary to provide a mechanism that allows the old and the new protocols to coexist. It is this mechanism that reduces the overall efficiency of the channel due to additional overhead. An 802.11b modem was not designed to speak 802.11g. In order to avoid collisions, the 802.11b radios must be informed that the channel is needed by 802.11g for a period of time.

In a high-density environment, every available efficiency must be taken advantage of to achieve the desired goal of maximum throughput and access. Figure 5 shows the relationship of per frame air time (channel utilization), frame sizes, and data rates.

Figure 5. Per Frame Channel Utilization, Frame Sizes, and Data Rates in a WLAN

The time scale above is in microseconds (μs). At the top end of the chart, a 2048 byte packet is transmitted at 1 Mbps, taking almost .02 seconds of airtime. Only one packet can be in the air at a time, and the faster that packet gets through, the better use made of the time available. Looking at this from a different perspective, reaching the bandwidth goals while supporting 802.11b and 802.11g will require more radios and cells and more advanced isolation techniques to implement them successfully.

Theoretically, if three radios could be put on the same pole serving all three non-overlapping channels in the same cell, a cell could be created that holds three times the bandwidth in 2.4 GHz and as much as 20 times that in 5 GHz, Figure 6.

Figure 6. Total Capacity of Three 2.4 GHz Radios on One Cell

In 5 GHz, there is more spectrum and the resulting bandwidth for a theoretical single cell increases dramatically, Figure 7.
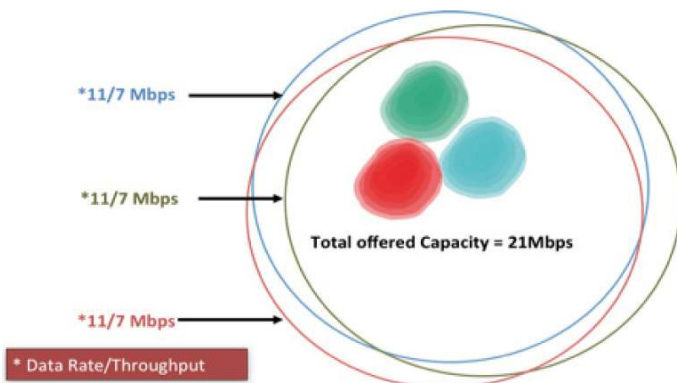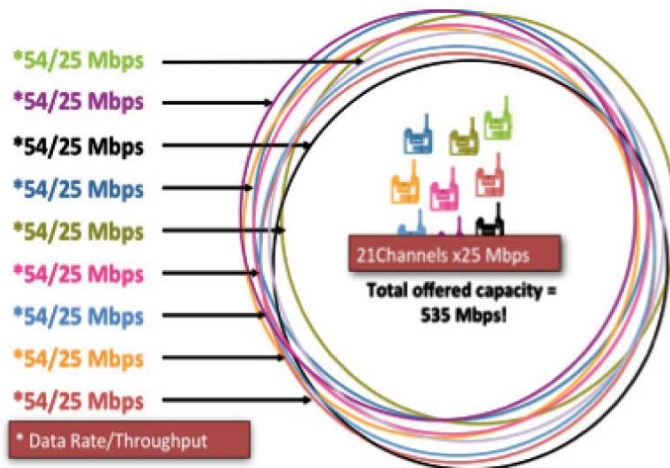
Figure 7. 21 Channels on One 5 GHz Cell, and Resulting Capacity

With today's radio designs, the radio could almost be placed on top of each another, but that would not serve the high-density design well. It would result in the same coverage area of a single cell and likely would not cover the required area, even in a relatively small lecture hall.

Data rates are a function of the received signal strength and the signal to noise ratio (SNR) at the receiver. It is not practical or efficient to lock a radio down to a particular data rate since the radio makes efficiency decisions based on available link conditions. Not every client will respond the same in an otherwise static environment. Variables such as receiver sensitivity, antenna configuration, driver version, and even position within the cell in relation to attenuating or reflecting objects will have a variable effect on the client. An environment that is conducive to good radio efficiency can be helped by appropriate design. The higher the average received signal strength and the better the SNR, the faster the data rate will be.

Cisco's ClientLink technology can increase the signal selectively for legacy 802.11 a/g Orthogonal Frequency-Division Multiplexing (OFDM) clients. Since legacy clients do not support the efficiency gains realized with 802.11n clients, they represent the least efficient clients in our design. Using ClientLink in a high-density user design allows the AP to improve SNR from 3-6 dB on a packet-by-packet basis for a client that is indicating the need to rate shift. This has the overall effect of increasing the range and rate equation for the network and encourages legacy clients to maintain higher data rates under adverse conditions. This is an excellent addition to a high-density design. The document Cisco ClientLink: Optimized Device Performance for 802.11n provides a full discussion of the Cisco ClientLink technology.

In a high-density deployment, channels will be aggregated to increase the total bandwidth. This means moving the APs ever closer together in the design space. A key success factor is Co-Channel Interference (CCI). CCI effects capacity of the cell by reducing the available bandwidth.

What Is Co-Channel Interference and Why Is It Important in High-Density WLANs?

CCI is a critical concept to understand when it comes to understanding the behavior and performance of 802.11 WLANs. It is a phenomenon where transmissions from one 802.11 device bleed into the receive range of other 802.11 devices on the same channel, causing interference and reducing the available spectrum and resulting performance. CCI can cause channel access delays as well as collisions in transmissions that corrupt frames in transit. Figure 8 illustrates how APs on the same channel interfere with each other.

Figure 8. Co-Channel Interference

Basic CCI – AP's on the same Channel interfere with one another

802.11 networks are contention based and rely on Clear Channel Assessment (CCA) mechanisms to judge the medium state (if busy we wait, when free we transmit). In the example above, this client's performance is being impacted because it can hear both APs. To this client, the two AP cells are coupled or acting as one super cell. For the uplink, both APs' transmissions will be seen as a busy channel by the client and the client will simply wait for an opportunity to transmit. Worse yet, on the downlink, transmissions from either AP will potentially collide and retries will increase the contention for the medium and continue to drive the data rates down overall. The effects of CCI are not limited to just the AP cell. In a high-density environment, the clients themselves will have the effect of increasing the overall cell size.

CCA is based on a receive threshold that evaluates the carrier for activity. It is generally a good practice to consider -85 decibels per milliwatt (dBm) as that threshold. Figure 9 shows a coverage model based on data rates. Higher data rates do not propagate as far. If the distances look long in this model, it is because it was calculated using an outdoor open space model rather than an indoor model which assumes attenuating factors in the environment. There are not many walls between the APs and clients in most high-density deployments.

Figure 9.   WLAN Coverage Model Based on Data Rates

In any Wi-Fi design, the effects of CCI can be limited by isolating the individual cells from one another through the use of non-overlapping channels and natural environment attenuation (walls, ceilings, file cabinets and cubes). We would not place two APs on the same channel directly next to one another intentionally. In a normal design, the environment and distances we are covering generally permit adequate coverage without a lot of CCI. But in a high-density network design, the distances are going to be constrained and propagation will be good, as such cell coupling and resulting CCI will become much more likely.

Design Point #3: Choose a High Minimum Data Rate to Support Increased Efficiency, Lower Duty Cycle, and Reduce the Effective Size of the Resulting Cell

CCI is not only an issue that will be faced in aggregating channels within the high-density deployment but something that must be kept in mind regarding existing deployments of surrounding areas. Lecture halls and classrooms tend to be co-located in the same facility, so overall design must be considered.

The Cisco Voice over Wireless LAN Design Guide is an excellent resource that presents CCI and best practices for Wi-Fi implementation. As an older document, it does not cover the extreme densities found in a high-density WLAN.

The Cisco WCS and controllers make monitoring co-channel interference and identifying the responsible AP or APs a fairly straightforward exercise. Cisco Radio Resource Management (RRM) algorithms are centralized and are a network-wide resource that continuously evaluates every single AP in the RF network to determine its relationship to every other AP in the system. It does this through the use of over the air (OTA) measurements and observations. Knowing how well other APs can hear a selected AP is a very useful feature when considering or planning a high-density WLAN deployment. Using Cisco WCS, it is possible to evaluate how well APs can hear one another-independent of a channel. This information is shown in a graphic display that shows not only how APs are effecting each other on a particular map, but also how other APs that are not on the map can impact a WLAN as well.

The wireless LAN controller maintains two lists of APs, Figure 10, both transmit and receive (TX and RX) neighbors that indicate how other APs hear a selected AP and how a selected AP hears other APs. This can be viewed using the Wireless LAN Controller (WLC) Configuration Analyzer tool and used to tune the resulting network and identify sources of RF as the APs themselves see it. Since this observation is based on OTA metrics and not based on predictive modeling, these values are independent of the antenna and AP combination.

Figure 10. Dual Lists of APs Maintained by WLAN Controller
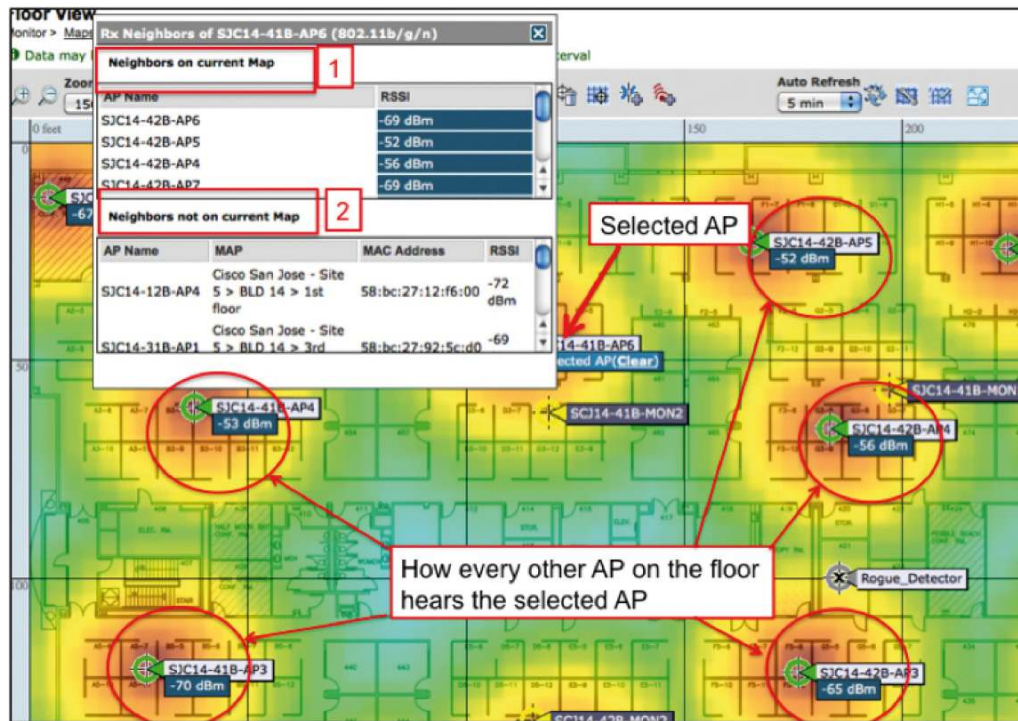
2.4 GHz Channel Reuse in High-Density Wireless Design

In 2.4 GHz there are three non-overlapping channels with which to work in achieving isolation. The RF properties of 2.4 GHz signals give it better range and less attenuation than signals in 5 GHz. In a high-density environment, there is often only one clean channel reuse within a 10,000 square foot area. Channel reuse in such an area is opportunistic at best and it is not possible to estimate without careful advanced survey techniques. Results will vary from no increase in bandwidth to modest gains and will differ from site to site. If faced with such a challenge, consult with a professional with experience in advanced engineering techniques specific to a high-density RF deployment. Adding more APs can reduce the number of users per cell and may appear to give more coverage when the space is empty. But once it fills up, the effect will be that of one large super cell covering the room with limited bandwidth and sporadic connections for all.

Note: Before considering a four-channel plan in 2.4 GHz, see Channel Deployment Issues for 2.4-GHz 802.11 WLANs for an excellent discussion on the issues. The conclusion is that it is better for two APs to share a channel than to have two channels overlapping on the edge. Two APs sharing a channel can demodulate each others' transmissions and share the bandwidth amicably. When two channels overlap at the edge, it is just noise to both and will result in collisions, retransmits, and SNR degradation.

Additionally, if the WLAN is located in a regulatory domain where the bandwidth to deploy four channels is available (e.g., availability of channel 13 and 14) unless the WLAN is sufficiently isolated from every other network it is likely that someone will deploy using the standard 1, 6, 11 model and drastically increase the interference to the WLAN.

If it is necessary to maximize a 2.4 GHz connection, it is possible to increase the bandwidth and efficiency of cells by physically limiting the propagation through the use of antennas and creative placement options. This will require site specific engineering and careful measurement and design. Cisco Advanced Services and experienced Cisco partners can help with this type of design and have achieved amazing results in extremely large and complex environments. This, however, is not always an option for budgetary or aesthetics reasons. We will discuss this in much more detail in the section on AP placement.

5 GHz Channel Reuse in a High-Density Design

In contrast to 2.4 GHz, 5 GHz has many more channels with which to work. As many as 20 channels can be received in the United States and between five and 21 in the rest of the world. Most regions have between 19 and 21 channels. But all 5 GHz channels are not created equally. Limitations on maximum power for parts of the band are not of concern, but Dynamic Frequency Selection (DFS) channels represent some challenges that must be addressed.

Dynamic Frequency Selection and High-Density Design

DFS was implemented so that APs and clients can share the band with radar devices. DFS details how radar is detected and what should be done in the event of detection. APs operating on DFS channels must first listen to a channel for 60 seconds to determine if there is a radar present before transmitting any energy. If an AP is operating on a DFS channel and detects a radar (real or false) it must shut down operations on that channel and abandon it for 30 minutes before that channel can be evaluated again for use.

Cisco APs were some of the first in the industry to support DFS channels. Client support for DFS channels has been inconsistent, however. Client devices do not have the ability to detect radar and rely on the infrastructure established by a DFS certified AP. Most clients today support channels 52-64. Client support for channels 100-140 has been slow in coming. Often it is a matter of not only the hardware but the version of the driver for the client that determines its operating channel range.

Client support has been steadily increasing and to-date Intel 5100 a/g/n, 5300 a/g/n, and 6300 a/g/n all operate on channels 52-64 and 100-140. The Cisco Cius and the Apple iPad and the Cisco 7925 IP phone also support the full range of DFS channels.

The effect of using channels that are not supported by all clients can result in coverage holes for those clients. Channels 100-140 are disabled by default on a Cisco Unified Wireless Network but can be enabled easily in the DCA channel selections by choosing the extended UNII-2 channels. Before doing so, it is highly advisable to inventory the clients and drivers that must be supported.

If DFS channels have been used in a WLAN installation, their suitability within the WLAN will be established. If they have not been enabled previously, it is advisable that the DFS channels are surveyed using Cisco equipment and that monitoring for radar detection is done before enabling the channels. In public and other venues within higher education environments, it is often recommended to avoid using these extended UNII-2 channels due to their current lack of client support. The base UNII-2 channel availability in clients is more pervasive and these are channels that could be considered but ongoing monitoring of client capabilities should not be overlooked.

802.11n - 20 MHz or 40 MHz Channels?

802.11n can operate in a 40 MHz channel by bonding two 20 MHz channels together and this significantly increases throughput. However, this is reserved for burst mode transfers only. It is only practical to do this in 5 GHz because 2.4 GHz is already limited by the number of channels available. If there are enough 5 GHz channels to achieve the WLAN goals using a bonded channel plan (9 in the U.S. if using available DFS channels) to meet throughput goals, consider it. If forced to reuse 5 GHz channels, more consistent results will be delivered using strictly 20 MHz channels and avoiding loss of efficiency due to CCI.

Evaluating Requirements for 2.4 GHz and 5 GHz Connection Support

The essential question for a high-density design is how many channels for each band will be needed to match the client base? This can be a tricky question since even dual band capable clients do not always select the faster 5 GHz band. Since bandwidth in 2.4 GHz is going to be limited, 5 GHz must be relied on to reach the goal.

Dual band adapters have been shipping with most laptops for some time. This does not mean that every laptop is a dual band client, but many are. Simply having a dual band client does not guarantee that it will choose 5 GHz over 2.4 GHz. The Microsoft Windows operating system defaults to a Wi-Fi channel search that starts with the 5 GHz channel 36 and continues searching through all of the 5 GHz channels that the client is capable of. If no 5 GHz AP is found then it will continue the search in 2.4 GHz starting at channel 1. Unless the Windows default is changed or the user has chosen a third party Wi-Fi utility to set spectrum preference to 2.4 GHz, the client radio will first try to associate to a 5 GHz AP. Apple Computer's latest release for Atheros and Broadcom chipsets also searches 5 GHz first.

The Cisco BandSelect feature enables the infrastructure to optimize these types of client connection choices. Where possible, it helps make sure that devices are attaching to the 5 GHz spectrum channels where interference sources tend to be significantly lighter. A much greater channel selection leads to the alleviation of bandwidth challenges.

Tablet computers and smartphones have begun entering the market at a staggering rate. The vast majority of smartphones shipping today operate in 2.4 GHz only. While many of them are 802.11n clients, of these most have implemented a single input single output (SISO) rather than Multiple Input, Multiple Output (MIMO). A SISO device is only capable of supporting MCS7 data rates, or 54 Mbps.

Design Point #4: 5 GHz Support Will Be Critical for High-Density, So Determine the Channel Plan That You Will Support and How It Will Be Administered

Evaluating the particular client mix for the WLAN can be done easily on Cisco wireless networks by utilizing the reporting features in the Cisco Wireless Control System or by reviewing the WLAN controller's connection logs.

Determine the Number of Channels and Cells Needed

A sample high-density WLAN project may include a design that yields 300 Mbps consistently to support 300 concurrent users. Under optimal conditions, 802.11g and 802.11a data rates yield 25 Mbps throughput. However, high-density environment will be less than optimal from a SNR standpoint. A better number to use is 20 Mbps throughput. Table 4 provides a quick reference using 20 Mbps per cell and per channel as the throughput value. Looking strictly at 5 GHz and assuming no channel reuse at this point, it is clear that 1 Mbps per user with 15 channels and 15 cells can be easily supported.

Table 4. Reference Guide for Channels, Connections, and Aggregate Bandwidth in Mbps

# Channels		1	2	3	4	5	6	7	8	9	10	11
Aggregate Bandwidth		20	40	60	80	100	120	140	160	180	200	220
# Connections	600								0.27	0.30	0.33	0.37
	500							0.28	0.32	0.36	0.40	0.44
	400					0.25	0.30	0.35	0.40	0.45	0.50	0.55
	300				0.27	0.33	0.40	0.47	0.53	0.60	0.67	0.73
	200		0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00	1.10
	100		0.40	0.60	0.80	1.00	1.20	1.40	1.60	1.80	2.00	2.20
	90		0.44	0.67	0.89	1.11	1.33	1.56	1.78	2.00	2.22	2.44
	80		0.50	0.75	1.00	1.25	1.50	1.75	2.00	2.25	2.50	2.75
	70		0.57	0.86	1.14	1.43	1.71	2.00	2.29	2.57	2.86	3.14
	60		0.67	1.00	1.33	1.67	2.00	2.33	2.67	3.00	3.33	3.67
	50		0.80	1.20	0.04	2.00	2.40	2.80	3.20	3.60	4.00	4.40
	40	0.02	1.00	1.50	2.00	2.50	3.00	3.50	4.00	4.50	5.00	5.50
	30	0.67	1.33	2.00	2.67	3.33	4.00	4.67	5.33	6.00	6.67	7.33
	20	1.00	2.00	3.00	4.00	5.00	6.00	7.00	8.00	9.00	10.00	11.00
	10	2	4	6	8	10	12	14	16	18.00	20.00	
Aggregate Bandwidth		20	40	60	80	100	120	140	160	180	200	220
# Channels		1	2	3	4	5	6	7	8	9	10	11

# Channels		12	13	14	15	16	17	18	19	20	21
Aggregate Bandwidth		240	260	280	300	320	340	360	380	400	420
# Connections	600	0.40	0.43	0.47	0.50	0.53	0.57	0.60	0.63	0.67	0.70
	500	0.48	0.52	0.56	0.60	0.64	0.68	0.72	0.76	0.80	0.84
	400	0.60	0.65	0.70	0.75	0.80	0.85	0.90	0.95	1.00	1.05
	300	0.80	0.87	0.93	1.00	1.07	1.13	1.20	1.27	1.33	1.40
	200	1.20	1.30	1.40	1.50	1.60	1.70	1.80	1.90	2.00	2.10
	100	2.40	2.60	2.80	3.00	3.20	3.40	3.60	3.80	4.00	4.20
	90	2.67	2.89	3.11	3.33	3.56	3.78	4.00	4.22	4.44	4.67
	80	3.00	3.25	3.50	3.75	4.00	4.25	4.50	4.75	5.00	5.25
	70	3.43	3.71	4.00	4.29	4.57	4.86	5.14	5.43	5.71	6.00
	60	4.00	4.33	4.67	5.00	5.33	5.67	6.00	6.33	6.67	7.00
	50	4.80	5.20	5.60	6.00	6.40	6.80	7.20	7.60	8.00	8.40
	40	6.00	6.50	7.00	7.50	8.00	8.50	9.00	9.50	10.00	10.50
	30	8.00	8.67	9.33	10.00	10.67	11.33	12.00	12.67	13.33	14.00
	20	12.00	13.00	14.00	0.00	16.00	17.00	18.00	19.00	20.00	
	10										
Aggregate Bandwidth		240	260	280	300	320	340	360	380	400	420
# Channels		12	13	14	15	16	17	18	19	20	21

If DFS channels will not be used, with 15 channels needed, use the nine non-DFS channels and reuse six of these. The cells must all be isolated from one another to prevent CCI from robbing available bandwidth.

Non Wi-Fi Interference and the High-Density Network

The important role of non Wi-Fi interference in the high-density network should now be clear. The success of a high-density WLAN will be compromised if any non Wi-Fi interference is operating within the same environment. Non Wi-Fi interference has a much larger impact on throughput in a high-density environment than unmanaged Wi-Fi energy. This is because 802.11 utilizes contention-based access mechanisms to coordinate station access to the channel. All 802.11 devices operate this way. Non Wi-Fi devices operating in the same band do not share these rules and do very well breaking the queuing and back-off mechanisms, forcing all stations within range to wait until the air is free.

When the 802.11 standards were being drafted, there was great concern that proliferation of Wi-Fi networks would create interference for licensed users operating in the same frequency bands. As a result of this concern, Wi-Fi was designed to be very "polite", yielding the band to almost anything else found operating there. Twenty years later, there are many consumer devices sharing the industrial, scientific, and medical (ISM) bands with Wi-Fi. The challenge is that while these devices operate under the same power restrictions of Wi-Fi devices, they are in no way obligated to yield the band for Wi-Fi traffic, and most do not. This creates a problem for normal Wi-Fi operations, since a Wi-Fi modem can only classify energy as:

- Wi-Fi (the energy detected can be demodulated)
- Noise (all remaining energy is considered to be "noise")

The impact of non Wi-Fi interference is logarithmic in its impact on Wi-Fi network operations. The higher the utilization of the Wi-Fi network, the more destructive non Wi-Fi energy will be. This means that if there is interference present and the network is only slightly utilized (e.g., there is ample duty cycle available within the spectrum), the presence of non Wi-Fi energy may not even be noticeable. There is space for both to share the spectrum. However, if the Wi-Fi network is highly utilized, then even a small amount of non Wi-Fi interference can have a large and noticeable effect.

This is why Cisco created CleanAir, a system-level, proactive monitoring, reporting, and mitigation mechanism as a core system resource for the CUWN. A high-density WLAN is designed to take advantage of every bit of the available spectrum. Any spectrum being consumed by even a relatively benign non Wi-Fi interference source will have a large impact in a dense environment. Interference needs to be identified, managed, and eliminated to provide the required bandwidth for a high-density network to work properly.

See [Cisco CleanAir](#) for more details.

Design Point #5: Account for and Manage All Energy Within the Operating Spectrum to Ensure All of It Is Available for Use

Client Duty Cycle

The discussion until now has centered on a use case where every client in the room will be competing for bandwidth simultaneously. This is the case when the users in the room simultaneously access a resource on queue. However there are many instances where the design requirement is to offer access to resources or the Internet for casual use at an event or within a venue such as a sports arena. Planning and sizing for these types of events can be quite different and will be based on expected Client Duty Cycle.

At a sporting event, for example, there are certain areas that will require ubiquitous and instant access during the entire event. Ticketing, vendor sales, staff, and press areas will generally require the highest amount of access. Of these, the press area is the only one that requires a high level of capacity in the arena itself. For the fans attending the event, only a percentage will be active on the WLAN at any one time From experience we see a 20 to 30 percent take rate with some well defined peaks occurring during period breaks. During play, very few fans are accessing the WLAN. However, this is changing as applications such as video replay, instant stats, and concession orders from the seat become more commonplace.

Observation and understanding of the requirements of WLAN users and situational requirements will guide the development of reasonable design goals. 500 users in a room who require simultaneous access to a single resource is a different design challenge than 1,000 or 1,500 users who only occasionally use the wireless network. Also, be aware that user patterns can and do change with time. This has been seen with the increase in the number of network clients per user. Monitoring network access and keeping good statistics will allow wireless engineers to stay on top of user trends on the university campus. Good management platforms such as Cisco WCS or Cisco Network Control System (NCS) are essential for managing the resulting network in real time and monitoring trends in a proactive manner.

Access Point Placement and Coverage Strategies

Often one of the biggest challenges in a high-density environment is access and aesthetics. A large meeting hall is impressive because of its size and a great deal goes into the aesthetics of the environment. The best approach to engineering a specific space is to do a qualified sight survey. Once the APs are mounted, physical adjustments become a lot more complex, so it is best to test while installing and make certain that the coverage that has been defined in the design is what is installed.

APs have evolved rapidly in short period of time. If an AP with external antenna capabilities is to be used, it is essential that an antenna that was designed for that AP also be used. MIMO or 802.11n APs need MIMO antennas to perform properly. Even if HT rates are not being counted on, the antenna and the radios are a system and the system is designed to perform with all of these elements. A good overview of antennas and pattern information can be found in [Antenna Patterns and Their Meaning](#).

Omnidirectional Antennas Versus Directional Antennas for High-Density Coverage

Mounting APs or directional antennas directly overhead in an environment may not be acceptable to the building owners. There are several ways to solve this problem and, depending on the environment and restrictions imposed, several methods may be used together.

The best approach to engineering a specific space for a high-density WLAN is to first do a thorough, active site survey to determine how and where the APs should be optimally installed. This will clarify what is possible in the space and provide a design to work from. Any changes to the optimal placement imposed by restrictions will require another survey because the final throughput for the space will likely change. If the environment and requirements necessitate the use of directional antennas, remember that once the APs are mounted, physical adjustments can become a lot more complex, depending on the mounting location. So it is best to test them while they are being installing to ensure that the anticipated coverage results from what is installed.

Omnidirectional Antennas

Use of an AP with attached low gain omnidirectional MIMO antenna is recommended if mounting is to be done on the ceiling of a modest-sized auditorium (averaging 20 feet or lower) with no channel reuse required in 2.4 GHz or 5 GHz. Omnidirectional antennas provide better ceiling-to-floor coverage, thereby reducing the likelihood that a packet traveling to or from the client has bounced off some object (usually a wall or the ceiling) before reaching the receiving antenna. This reduces the opportunity for multipath interference.

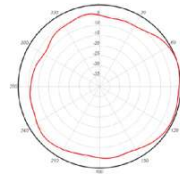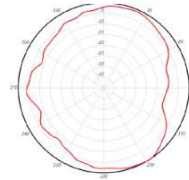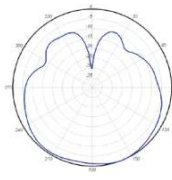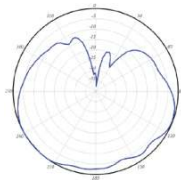**A related consideration is high-gain versus low-gain omnidirectional antennas:**

- **Use of a high gain omnidirectional antenna should be avoided.** This type of antenna will increase the size of the cell and the number of users that will be sharing the bandwidth. Higher gain in an omnidirectional antenna design generally means increased horizontal beamwidth with a decrease in vertical beamwidth. This effect will be more pronounced as the ceiling height increases.
- **The low-gain omnidirectional antenna has less horizontal coverage** and in an auditorium will have less floor coverage than a high-gain antenna. This supports the goal of small channel and small floor size and it will serve to limit the number of users in the coverage area, effectively managing client-based co-channel interference. The low-gain antenna will also provide a better quality signal.

Cisco Indoor Access Points with Internal Antennas

Cisco APs with internal antennas (Table 6) have optimized coverage patterns for consistent RF distribution. The coverage patterns for all currently shipping APs optimized for ceiling mounting, however, will perform well mounted on a wall, under a seat, or beneath the floor as well.

Table 5.   Cisco Aironet 1040, 1140, and 3500i Series Integrated Antennas

Note: The same integrated antennas are used on these devices but AP-1040 only has two elements per band.

Cisco Aironet 1040, 1140, 3500i Series Integrated Antennas			
2.4 GHz, 4 dBi Azimuth Plane Radiation Pattern	5 GHz, 3 dBi Azimuth Plane Radiation Pattern	2.4 GHz, 4 dBi Elevation Plane Radiation Pattern	5 GHz, 3 dBi Elevation Plane Radiation Pattern
Frequency Range	• 2.4-2.5 GHz • 5.15-5.85 GHz		
Gain	• 2.4 GHz: 4 dBi • 5 GHz: 3 dBi		
Polarization	Linear, Vertical		
Azimuth 3 dB Beamwidth	Omnidirectional		

Cisco Aironet 1040, 1140, 3500i Series Integrated Antennas	
Elevations 3 dB Beamwidth	2.4 GHz = 120 Degrees, 5 GHz = 120 Degree
Antenna Connector	Integrated
Mounting	Integrated
Antenna Type	Omnidirectional

MIMO low gain omnidirectional antenna options for the Cisco 802.11n Series APs are shown in Table 7.

If the Cisco Aironet 1250, 1260, or 3500e Series AP are to be used, there are several external omnidirectional antenna options to allow broader mounting options and to meet differing aesthetic requirements. Consult the Cisco Aironet Antenna Reference Guide for full details.

Table 6.   Omnidirectional Antenna Options for Cisco 802.11n Series APs

Using 802.11n capable radios is a fundamental design consideration for today's high client density environments. The above advantages, especially in a mixed environment, lend themselves well to a high-density deployment.

Product ID	Description	Product	Gain
AIR-ANT2452V-R	2.4 GHz 5.2 dBi Diversity pillar mount ant, RP-TNC connectors		5.2 dBi
AIR-ANT2451NV-R	2.4 GHz 3 dBi/5 GHz 4 dBi 802.11n dual band omni antenna		3 dBi/4 dBi
AIR-ANT2430V-R	2.4 GHz Omni 3 dBi, 3 element ceiling mount		3 dBi
AIR-ANT5140V-R	5 GHz Omni 4 dBi, 3 element ceiling mount		4 dBi
AIR-ANT2422SDW-R	2.4 GHz 2.2 dBi Short white dipole antenna, Qty 1		2.2 dBi
AIR-ANT5135SDW-R	5 GHz 3.5 dBi Short white dipole antenna, Qty. 1		3.5 dBi
AIR-ANT2440NV-R	2.4 GHz 4 dBi 802.11n Omni wall mount antenna		4 dBi
AIR-ANT5140NV-R	5 GHz 4 dBi 802.11n Omni wall mount antenna		4 dBi

Directional Antennas

It is not always going to be possible to solve challenges in a high-density environment using strictly omnidirectional antennas. If a WLAN requires channel reuse within the same floor space or if coverage is required for non-standard areas such as indoor or outdoor arenas, mounting options for a usable design may be limited. Therefore directional antennas come in many coverage patterns that are more suited to challenging environments where an omnidirectional will not be adequate.

When an environment requires the use of directional antennas, the complexity of the design and the implementation both will go up accordingly. It should, however, also be noted that outstanding results can be achieved.
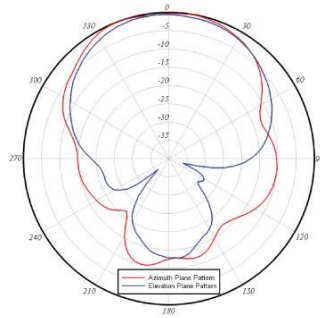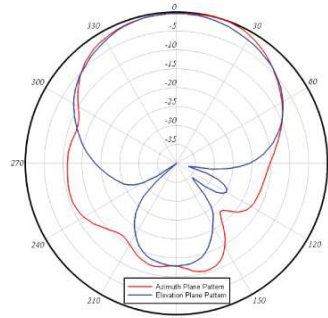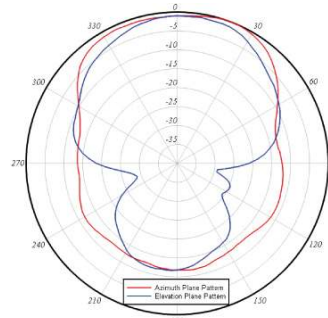
Channel Reuse and Directional Antennas

If mounting on the ceiling in an auditorium with high ceilings greater than 25 feet, it is recommended that high gain directional patch or Yagi style MIMO antennas be used because these antennas provide:

- Better ceiling to floor coverage, if mounted in the ceiling or on catwalks with antennas oriented in a direct downward direction. This creates smaller cells of coverage directly beneath the APs and allows for better channel isolation between adjacent cells while maintaining power levels and sensitivity in the direction of the covered clients.

- Coverage in larger venues with very high ceilings or perhaps where ceiling access is not available. Placing directional antennas at mid level from the sides or from behind the coverage zone and using downtilt can provide controllable coverage zones and better installation options in difficult environments.

- A much narrower coverage lobe or beam, thus allowing smaller cells so the number of cells and channels for both 2.4 GHz and 5 GHz can occupy a common room providing increased cell isolation and reducing CCI. This is critical in any environment where the design solution requires channel reuse for either band.

Directional antennas like omnidirectional antennas are classified by their 3 dB beamwidth in the horizontal and elevation planes. But directional antennas typically have much higher gain than an antenna that is classified as omnidirectional. For example, the AIR-ANT2460NP-R antenna, Table 7, has a 3 dB horizontal beamwidth of 80° and a 75° elevation plane. This is quite useful in isolating the AP and antenna from surrounding energy and providing better coverage in the intended zone.

Table 7.   6 dBi Patch, Three Element Wall Mount

| 6 dBi Patch, Three Element Wall Mount AIR-ANT2460NP-R		
Antenna A Radiation Pattern	Antenna B Radiation Pattern	Antenna C Radiation Pattern
Frequency Range	2402-2484 MHz	
VSWR	2.1	
Gain	6 dBi	

6 dBi Patch, Three Element Wall Mount AIR-ANT2460NP-R	
Polarization	Linear
Azimuth 3 dB Beamwidth	80 degrees
Elevations 3 dB Beamwidth	75 degrees
Antenna Connector	(3) RP-TNC male
Cable Length	36 in (91.4 cm) plenum rated
Dimensions	5.8 in x 11.25 in x 1.13 in (14.7 cm x 28.6 cm x 2.9 cm)
Mounting	Wall mount

Front to back ratio (FTB) is another measurement that is typically provided with directional antenna specifications. Since the gain is increased in one direction, it is reduced elsewhere. The FTB ratio spells out the amount of isolation that can be achieved in the opposite direction of the antenna's intended coverage. For the ANT-2460NP-R, the FTB is nominally 8 dBi. Combined with the density of an average load-bearing wall, that equals minimal leakage to or from the other side.

Here's an example of how this might be used as an advantage in a campus environment: today, most campuses have achieved a level of ubiquitous coverage and it is likely that HD WLAN design requirements will be sharing at least a portion of their airspace with another coverage zone. Typically, when engineering for pure coverage, lower data rates will be enabled to maximize coverage (increase the cell size). This is adequate for the intended coverage zone but will negatively impact the HD coverage zone where the cell size has been carefully engineered to exclude lower data rates. By reengineering the normal coverage zone, the desired coverage can be achieved and the amount of coupling between coverage zones can be reduced, Figure 11.

Figure 11. Example of Border Coverage Zone Management Using Directional Patch Antenna

Border Area between coverage and HD coverage

Table 8 features MIMO directional antenna products for Cisco 802.11n Series APs.

Table 8. MIMO Directional Antenna Options for the Cisco 802.11n Series APs

Product ID	Description H/E plane	Product	Gain
AIR-ANT2460NP-R	2.4 GHz 80°/75° MIMO directional patch		6 dBi
AIR-ANT5160NP-R	5 GHz 65°/65° MIMO directional patch		6 dBi

Product ID	Description H/E plane	Product	Gain
AIR-ANT2410Y-R	2.4 GHz 55°/47° single element yagi (1 Piece, 3 Required)		10 dBi
AIR-ANT25137NP-R	Dual-Band 2.4 GHz 36°/36° 5 GHz 55°/48° MIMO directional patch * Requires AP 3502P		13/7 dBi

Use of Directional Antennas and Downtilt

One challenge often faced in a lecture hall or auditorium is the need to provide more bandwidth than a single use of the channels available in 2.4 GHz will allow. Using a directional antenna can provide cell-to-cell isolation if placed, mounted, and adjusted properly. One aspect of using directional antennas is the concept of mechanical downtilt. Downtilt involves adjusting the antenna down to change the coverage pattern that is created.

The coverage pattern can be adjusted by changing the mounting height or the mechanical downtilt angle, Figure 12.

Figure 12. Adjusting Directional Antenna Using Downtilt

Inner Radius
Distance/FT= H/Tan(A+(BW/2))

Outer Radius
Distance/FT= H/Tan(A+(BW/2))
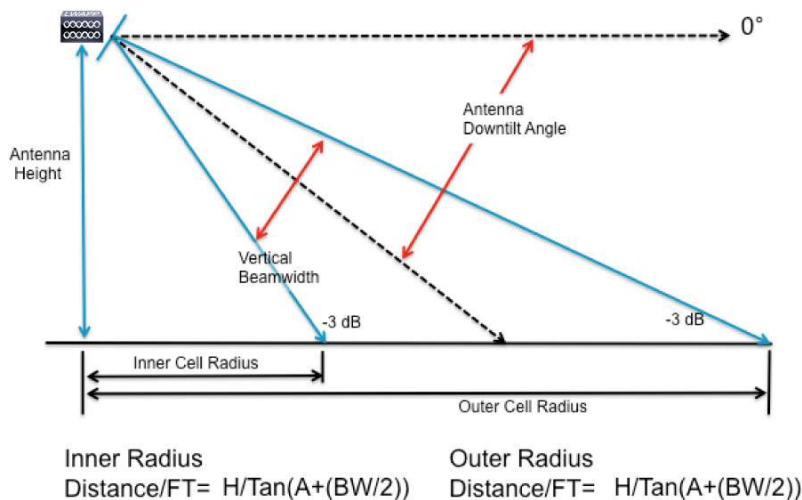
The values for the formula above are:

- H = height of the antenna
- A = downtilt angle
- BW = the 3 dB horizontal beamwidth of the antenna

By adjusting the downtilt of the antenna, it is possible to "dial in" - or add WLAN coverage - to specific areas within the coverage zone. APs and RF energy operate much like light cast by lighting fixtures. It is possible to light an entire warehouse with a bare bulb on the ceiling, but the result is low levels of light in some areas. But if there are multiple fixtures, including some with higher patterns of luminosity to illuminate larger, the result is comprehensive overall lighting. RF is invisible, so measuring the coverage and adjusting it appropriately requires tools to measure the coverage. For each antenna placement, simply walking the area below it and adjusting the antenna to change the pattern based on Received Strength Signal Indication (RSSI) levels to match coverage requirements is generally all that is required at the initial installation. Antennas hear the same as they transmit. If measuring and adjusting are done carefully, using consistent measurements and tools, good results can be achieved.

Any additional tuning can be managed with power threshold adjustments through RRM. In all cases, a full site survey to compare the results to the plan is required once all assets are installed.

Directionality in an antenna increases the overall gain and resulting power that will be delivered. Modest gains (4-8 dBi) can be easily managed by RRM. Some situations, however, require higher gain antennas (10-13- 17 dBi) to achieve a desired coverage area pattern. It is important to pay attention to the effective isotropic radiated power (EIRP) as this will rise with the antenna gain. Physically attenuating the transmission line using good quality RF attenuators rated for the spectrum and power may be required. With TX power set to its minimum (-1 dBm), a 13 dBi antenna will have an EIRP of 12 dBm. If a WLAN design calls for a transmit power of 4-5 dBm, then a 10 dB attenuator will be needed to put the EIRP back in the tunable range of the AP's transmit power. Attenuators will reduce the overall signal level that is on transmit and receive and will thereby reduce the received power of the clients at the AP. This is not a problem since the approach is to compensate for higher antenna gain.

AP Placement Options

Overhead

The most common method of achieving even coverage is to evenly space the APs directly over the clients they will serve. There are multiple options to accommodate overhead mounting of the APs in an unobtrusive manner. Although many people do not consider any AP to be a welcome stylistic addition to a room, APs with internal antennas can be flush mounted to a variety of surfaces and offer an option with less impact on a room's aesthetics. In these cases, a flush mount antenna can be much less obtrusive. External antennas increase the cost and complexity of the installation slightly, but can be justified if the end result is the ability to cover the room at a sufficient density and meet aesthetic requirements. Once the decision is made to incorporate external antennas, numerous options are opened for shaping the RF cell through the use of directional antennas. Channel reuse in 2.4 GHz can be achieved in smaller spaces by using directional antennas overhead. Ceiling height and antenna choice will determine cell boundaries and taking measurements is required.

In Figure 13, assuming the room is 9,000 square feet, using the internal antenna enables AP nine channels of 5 GHz, and three channels of 2.4 GHz to be provided comfortably. Using an external omnidirectional antenna, the results would be much the same. Using omnidirectional antennas on 5 GHz and directional antennas on 2.4 GHz, one, two, or three additional 2.4 GHz channels could be added within this space. Throughput improvements would largely be gained by more even client distribution and less resulting CCI at the client. Some additional capacity will be gained, but only to the extent that CCI can be eliminated between the cells and this will depend on ceiling height, antenna pattern, and power levels in 2.4 GHz.

Figure 13. Nine AP WLAN Deployment

Ceiling heights can be much higher in a college lecture hall than they would be in a normal classroom environment. A normal ceiling height may be 8-12 feet but in a lecture hall it could be 20 feet or more. This will impact the resulting RF levels seen at the client's position if it is not taken into consideration in the design. Because Cisco APs centralize their RF management, AP to AP neighbor relations will not be effected between the APs, but the APs will be closer to one another than they are to the client and this will require some threshold tuning. Typically, an adjustment of 3-8 dB on the transmit power control (TPC) threshold will be sufficient to increase the power to a comfortable level at the floor. This is a one-time adjustment that biases the entire installation. Very good results can be achieved in this way. If directional antennas are used in this way, the additional gain associated with the antenna will generally offset the path loss associated with a high ceiling, but with very high ceilings (30+ feet) adjustments to TPC may be required, depending on the antenna chosen.

If overhead mounting is not an option, (e.g., there is no access above the ceiling level for cables or the entire room was designed so that nothing is on the ceiling or if there is a large skylight) there are many other mounting options.

Side Mounting

Depending on the dimensions of the room, it may be possible to cover the entire room from the sides. If the room is wider than two cells will accommodate (assuming each cell is covering one half of the room), it will be necessary to use directional antennas and mechanical downtilt to cover individual sections of the room for each radio. Large rooms will have aisles separating seating sections and this space may be used to design cell overlap areas. If mechanical downtilt and directional antennas are used to achieve this, the higher the antenna may be mounted and the larger the resulting cell will be. Good results can be achieved with antennas mounted as low as eight feet up on the wall with 30-60 degrees of downtilt.

Front and Rear Mounting

The fronts and rear of rooms are other areas where there will generally be open space between the edge of the room and the users. In a theater style room, look for mounting positions under or in front of the raised platform or stage and the overhead facing out. It is possible to place APs in both locations, using the lower position on or near the floor and antennas such as the AIR-ANT2460NP-R at 80 degrees horizontal and 75 degrees vertical to cover the first rows. (Note that two of these provide almost a full 180 degrees of coverage and two channels of 2.4 GHz). For 5 GHz, use the AIR-ANT5160NP-R at 65 degrees horizontal and vertical. Additional APs can be mounted near the ceiling facing back into the room using the AIR-ANT25137NP-R at 55 degrees vertical and horizontal to cover the section behind the first rows (with 55 degrees, 3-4 channels). This pattern can be repeated from the rear of the room and it will provide a lot of channels from the perimeter.

In Figure 14, APs have been mounted low to the floor, providing a cell that will use the users to attenuate the propagation distance. APs have also been mounted near the ceiling using mechanical downtilt to manage the resulting cell size.

Figure 14.   APs Mounted Low to the Floor
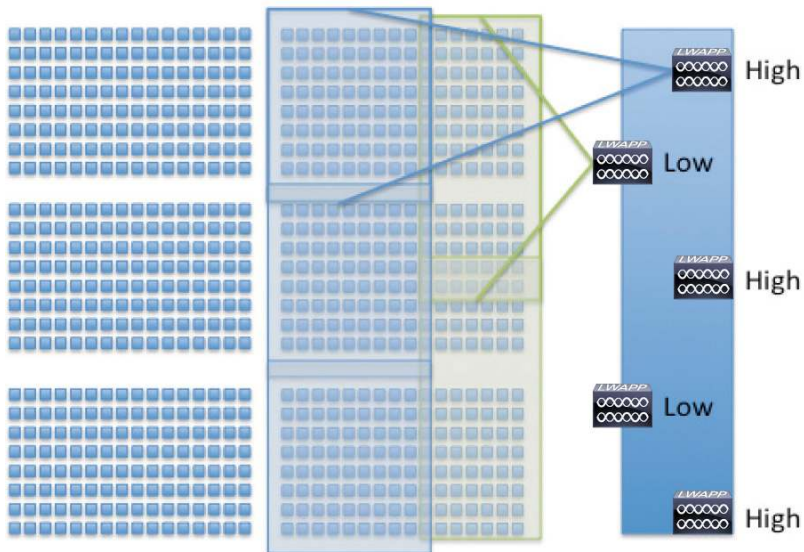
Shadows

Features in a room such as support columns or balconies can represent natural shadows within a room. Under normal densities, such features would require an additional AP to ensure coverage for users placed in such a shadow. In a high-density installation, these features can be used to advantage, increasing the channel reuse by designing to maximize this effect and then filling with another channel.

Under Seat Mounting

One of the optimal ways to cover a large, dense area is from underneath the users. This provides two advantages. First, the users themselves attenuate the signals (so more APs can be accommodated with channel reuse than other more open methods would allow). Second, this can generally provide a great way to hide APs. To do this successfully, a few experiments must be performed to evaluate the propagation characteristics of the mounting location. Under seat or under desk mounting can provide from 6 dB to 20 dB of attenuation to the cell, depending on the openness of the mounting options.

It is important to note that metal chair legs and desk components will interact with the antenna of the AP and change the pattern of the radiation. Surveying the results of placement decisions with a good tool is necessary before making permanent mounting decisions. This can be done by selecting a representative section of the room and placing a minimum of 4 APs (even more is better) temporarily in place. Analyzing more than one mounting option will yield comparative data that will be helpful in making a final decision.

The solution chosen should provide decent attenuation and a usable pattern. Having shadows in this type of mounting is not a problem and may be useful if the shadows are consistent and filled in by other cells. This may sound like considerable work but the results can be well worth the effort. Generally, use of APs with internal antennas is recommended for this type of installation. Experiment with the AP antennas pointed both up and down. Do not be overly concerned about the power levels; in an empty room they will generally look very high when compared to a normal installation. As a final test, select a portion of the room-perhaps 25 percent of the seating area-and measure the site with users present. This can be done using a non production service set identifier (SSID) to prevent attracting live users. What we are evaluating at this stage is placement, not user throughput. Look at CCI and the coverage pattern to evaluate the effectiveness of the design.

Under Floor Mounting

Mounting APs under the floor with the intent of radiating through to the users can provide a great deal of attenuation. This in theory will work and will allow the creation of more pico cells within a given environment. However, based on experience, the number of facilities where this has proven practical remains small. Mounting beneath the floor should not be ignored as an option and it should be investigated. The occasional corner case where this is a viable option can be quite useful to the design. The RF density of the floor and access available will be the limiting factors. Pre-stressed concrete floors or floors containing a lot of reinforcing materials are generally not suitable. It is also important to consider the reflective potential of the floor material and where this will reflect to in the rest of the infrastructure.

Bringing It All Together

Consider a high-density WLAN deployment in two rooms that together are 26,000 square feet, Figure 15. Room A has seating for 1,450 and room B has seating for 1,500. Both rooms are separated by a double air-gaped temporary wall, which provides good RF attenuation. The design goal is to provide coverage for the participants, assuming a 25 percent duty cycle. The focus is on 2.4 GHz coverage, since this is the critical resource for most all PDAs and several tablet models. 5 GHz exhibits similar advantages in the denser room and room B had far more happy 5 GHz users simply because propagation distances were less overall between AP and clients.

Figure 15. Two-Room High-Density WLAN
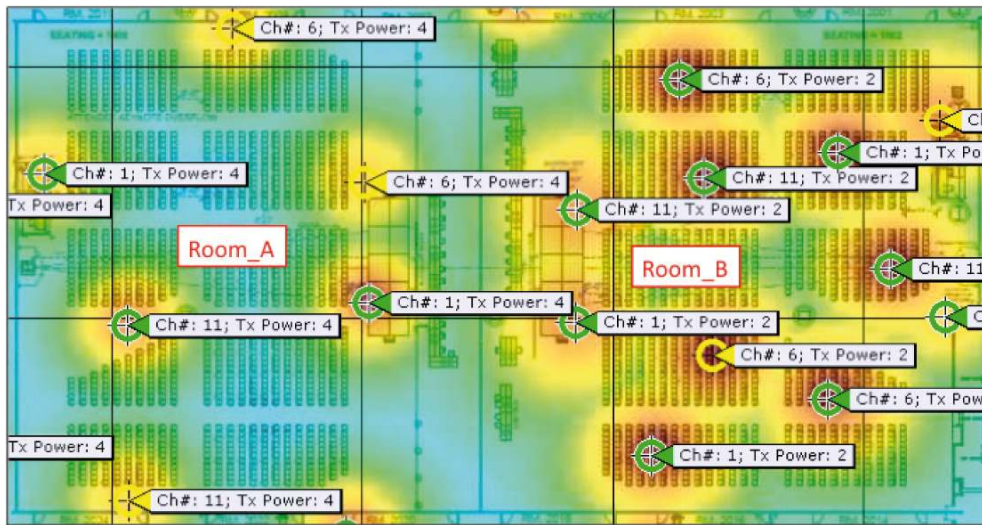
Room A has a total of six APs deployed. Two APs have directional antennas used from the sides, two APs with internal antennas are positioned at the front, one internal antenna AP is positioned at the rear of the room, and one internal AP is floor mounted to fill in for a shadow behind a pillar in the room. Power levels are set to 11 dBm and result in good coverage levels for the room. Minimum data rates are set to 9 Mbps. BandSelect and Cisco ClientLink are enabled.

Room B has a total of 12 APs deployed. Eight are positioned under the seating in the room. Two APs are at the rear and two APs are in the front of the room. Power levels were reduced to 5 dBm. Minimum data rates were held at 12 Mbps, BandSelect and Cisco ClientLink were enabled.

These deployments resulted in significantly different user experiences in each room:

- Room A did not have an even balance of clients on all six of the APs, with a pronounced number staying on the AP that was near the entrance to the room indicating that they never roamed. Performance was moderate, but there were no user complaints.
- Room B had a much more balanced load distributed across all the APs in the room and performance was significantly better, as indicated by the throughput observed in the room and on the network supplying it.

Reviewing the duty cycle in the two rooms from the same period, there was a significant difference in the RF conditions between room A and B. Over a 10 minute period, the duty cycle for channels 6,11 in Room A averaged near 100%, and channel 1 exhibited spikes of 100%, as shown in Figure 16.

Figure 16.   Duty Cycle Results in Room A
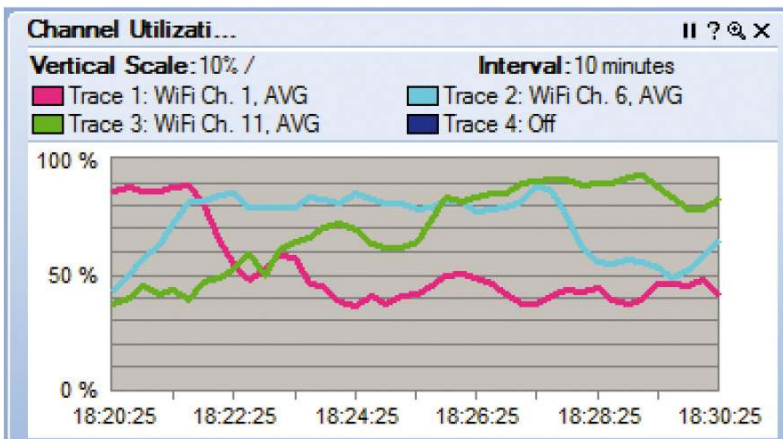
Over a 10 minute period, the duty cycle for channels 1,6,11 in Room B averaged 50 percent with brief spikes of 70%, Figure 17.

Figure 17.   Two-Room High-Density WLAN

The duty cycle measures the actual RF utilization; it is the percentage of time Tx is active for a given frequency. Channel utilization is different, and normally a higher value than RF duty cycle as it adds protocol timers into the metric to show overall Wi-Fi channel availability. By adding the contention window minimum (CW_min), network allocation vector (NAV) and CCA high times together, it shows the percentage of time a Wi-Fi station cannot access the channel. As mentioned previously, Wi-Fi is contention based and while signals are present above the CCA threshold no station can access the channel. 100 percent RF duty cycle is bad. In this case, it is attributable to inefficient spectrum use.

From the physical layer, these spectrum views, Figures 18-19, were taken using Cisco Spectrum expert and CleanAir from a monitoring AP located in the center of each room.

Figure 18. Spectrum View of Room A
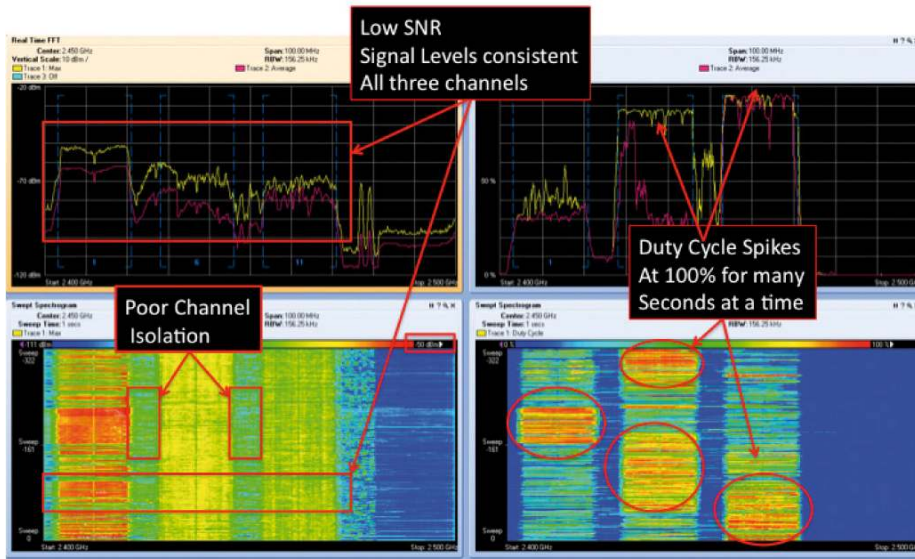
Figure 19. Spectrum View of Room B
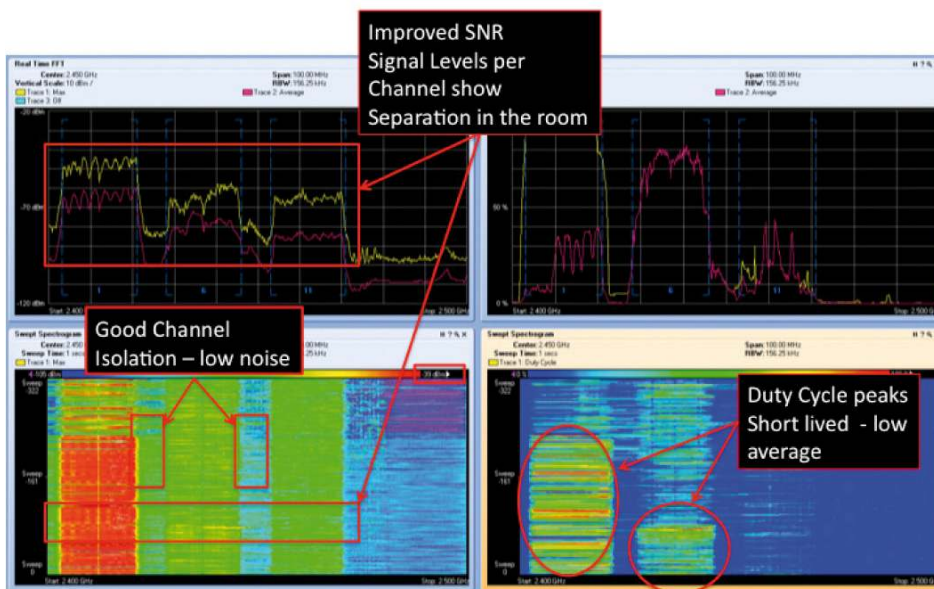
Note that for Room A the scale of the swept spectrogram in the lower left corner of the display had the red level set to -39 dBm and in Room B it was adjusted to -50 dBm. Signals were quite a bit hotter in Room A. Room A had fewer channels and the APs were further from the client; this in turn caused the clients to increase their power. The combination of the increased power and fewer channels in roughly the same space caused more noise and faster deterioration of the overall available spectrum.

In high-density WLAN design, more APs and less power mean higher spectral efficiency. Higher spectral efficiency encourages higher speeds and requires less airtime for the same relative load and less airtime dedicated to recovery and retransmissions or overhead.

The closer that the APs and clients can be to one another, the higher the spectral efficiency will be. If the APs and clients must be operated at a distance due to architectural requirements, use of directional antennas will improve the cell isolation, and reduce the power required for Tx in both directions.

The foremost failure in most high-density designs is not enough APs or channels. The secondary reason for failure in most high-density designs is poor channel isolation; too many APs and channels in the space.

Cisco Unified Wireless Network Best Practices

The following sections cover some of the basic design decisions that should be considered and why they are important. It is understood that not all of these recommendations will be possible in a legacy design. These points are highly recommended for the support of a high-density WLAN design. The market and the technology continue to evolve rapidly, and these recommendations will also continue to evolve.

Information in this section is not intended to replace the recommended design guides for campus and WLAN deployments referenced elsewhere within this document. This is a partial list of considerations that pertain specifically to high-density design considerations. Local Cisco sales teams or Cisco partners are the best resource for questions concerning the goals and requirements of a particular WLAN design.

Pre-Deployment Site Inspection and Validation

Prior to deploying a WLAN, the site should be evaluated based on the following:

RF Interference

- The site may have pre-existing WLANs as part of the same RF network.
- There may be nearby WLANs that are not part of the existing RF network.
- There may be nearby sources of microwaves and/or surveillance cameras or other forms of non Wi-Fi interference.
- There will most likely be Bluetooth.
- The WCS and the WLAN controller provide accurate reporting on CCI conditions.
- Cisco CleanAir provides in-depth interference analysis and monitoring 7x24.

Site Survey

- Visually inspect the site for multipath potential and placement options for the APs and antennas.
- Walking the site using a good survey tool is highly recommended. Using a Live RF tool such as AirMagnet Surveyor will enable an evaluation of RF propagation and the ability to actively transmit data and evaluate range and data rate coverage.

WLAN Design Tools

Providing a high-density WLAN is clearly a challenging project but these networks are increasingly being deployed. Knowing what is in place and using a spectrum today are essential steps. Knowing where a spectrum is going at installation is critical. Maintaining the environment also requires regular observation and occasional troubleshooting, as WLAN dynamics have a tendency to change over time. Most of the management concerns can be eliminated by implementing Cisco WCS, but the importance of the site survey during the planning and deployment cannot be emphasized enough. The need for good tools and accurate assessments increase with the density of the deployment. There is a fine line between interference and usable spectrum in a high-density WLAN.

Cisco WCS version 7.0 MR1 enjoys integration with both Ekahau Site Survey 5.1 and AirMagnet/Fluke Survey Pro and Planner v 8.1. These integrations will reduce the effort involved in processing and sharing information between the management platform and survey tools. This reduces both the time and effort required while maximizing accuracy for planning and lifecycle management of the resulting network.

- For more information on Cisco WCS and Ekahau Site Survey: Recording: Integration between Ekahau Site Survey and Cisco WCS.
- For more information on Cisco WCS and AirMagnet Surveyor and Planner pro integration see: Webinar: Cisco WCS Integration with AirMagnet Survey and Planner.

Calibration

When the word calibration is used in the context of tools, it generally refers to the need to test and adjust the accuracy of certain measurement tools against a reference signal to certify that the results produced are accurate. In wireless networking, most tools rely on a network adapter used for site surveys or to conduct specific measurements. Network adapters rarely if ever require calibration. Network adapters can and do vary widely on their assessment of the physical layer. Here are some best practices that can be used to ensure that data produced from multiple sources is useful and accurate when compared.

For laptop-based tools, it is important to be certain that the network adapter used is consistent between different tools and data sets that will be compared. This means that all laptops used for surveying should have the same adapter and driver software. Different results can be derived even when using the same adapter with the same driver software on different platforms. Antennas and the placement of the antennas are generally driven by the physical space available in a platform and design compromises will vary platform to platform. This will effect how the radio perceives the world. Compare different platforms in a static environment before committing to a survey.

Best practices include:

- Use external USB adapters for survey equipment.
- The software can usually be licensed to the adapter and moved from platform to platform, keeping the measurement tool consistent.
- Always use the same driver software whether from the device manufacturer or from the tool manufacturer on all tools being used to quantify results in a given environment.
- When updating a driver, compare the results between the new and old driver. This will help make sense of older data sets.
- Record what adapters, platforms, and driver software were used when collecting data sets. This will make remembering the details later much easier if the need arises.

Infrastructure Readiness

The wired LAN should remove all protocol traffic not required by the applications on the WLAN endpoints. The wired LAN should be designed to maintain the quality of service (QoS) configurations of the WLAN. The LAN should be designed to support a burst of authentication traffic.

SSID Assignment

Users connect via distinct SSIDs for each user segment, with each SSID being segregated to its own respective VLAN. This wireless connection is secured by respective wireless authentication protocols. Table 9 provides an example of SSIDs that are used in universities.

Table 9.   SSIDs Used in University Campuses

SSID	User Group
IT	Internal technical management
Faculty	University employees
Student	Student
Guest	Public internet users and guests

Note: The number of SSIDs should be kept to a minimum to avoid a negative performance impact because of excessive management traffic. Each SSID requires a separate beacon message that will be broadcast at the lowest mandatory data rate and can significantly impact the performance in a high-density design.

Wireless LAN Controller and Feature Specific Configuration Recommendations

Here are some specific configuration recommendations and their function for a high-density environment. This is not a list of specific requirements and individual configuration details will likely be different.

WLANs=>WLANs #n Advanced Tab

- MFP Client Protection - optional or disabled, disabled will lower overhead slightly
- Client load balancing - enabled
- Client BandSelect - enabled

Wireless=>Access Points=>802.11a/b=>Configure

- 11n Parameters - ClientLink enabled, Clientlink enables beamforming to the legacy client and is activated when the client SNR indicates a rate shift is needed. This will add 3-6 dB for the downstream to the client and likely prevent a rate shift
- RF Channel Assignment - Global
- TX Power Level Assignment - Global

Wireless=>Advanced=>

- Load Balancing: Adjust client window size for your environment. Default is 5, will need to be more in line with expected load levels of 30-50 clients. Maximum denial count can be adjusted, but the default is generally sufficient. Must be enabled or disabled per SSID and WLAN
- BandSelect: Default values should be good except for Acceptable Client RSSI. Set this to reflect the expected power levels in the room. This will also prevent clients outside of the high-density area from being monitored in the queue

Wireless=>802.11a/802.11b=>Network

Network - Data Rates

- Mandatory = Client must support in order to associate
- Supported = Optional: client may rate shift to this rate if desired
- Disabled = Not supported on the radio

Recommendations: Set the minimum acceptable data rate for the cell as supported. Choose a higher data rate as the target client speed and set as mandatory. The first or lowest mandatory rate is the speed at which beacons

will be sent. The highest mandatory rate is the speed at which multicast will be broadcast. The design for a high-density WLAN should use a minimum mandatory data rate of 18 or 24 Mbps.

Wireless=>802.11a/802.11b=>RRM

RRM should be used to manage both power and channel within a high-density WLAN. This will ensure that power levels and channel selection are optimized. RRM will monitor the OTA propagation and utilization metrics and adjust all APs based on their view of one another for optimal bandwidth/performance. Refer to the WLC Configuration Guide for a complete discussion regarding configuration.

Specific recommendations for high-density WLAN deployments include:

Transmit Power Control Algorithm (TPC)

- **Min/Max power** sets a limit on the minimum and maximum power that a radio can use. It overrides TPC recommendations and effects only the local controller and associated APs.
- **TPC power threshold** is used to adjust the power at the edge of the cell. The default is -70 dBm. Tune for specific installation requirements. It may be necessary to adjust the TPC power threshold in an environment with a ceiling over 25 feet high.

Dynamic Channel Assignment (DCA) Algorithm

- Defaults should suffice with the following exceptions:
  - Avoid foreign AP interference. This tells the DCA to accommodate channels in use by a foreign AP and to attempt to work around that assignment. If on the border with another coverage zone under a separate RF group, enable this (default). Disable this if there are other networks operating that are outside of your control.
  - DCA channel sensitivity (medium is the default) should work fine in most applications. Set to low to make a channel change highly unlikely. Set to high to encourage channel changes in response to environmental changes.
  - In 802.11a only environments with a channel width of 20 MHz/40 MHz, most high-density environments will benefit more from more available channels. But certain applications, such as video, will utilize 40 MHz effectively.
  - DCA channels 100-140 and 165 disabled by default enable extended UNII-2 channels to select 100-140 automatically.

Coverage Hole Detection Algorithm

- Coverage threshold: Data RSSI should be set for the minimum acceptable RSSI level to support the minimum data rate required in the high-density coverage area. This will allow Coverage Hole to report instances where a client exceeded that threshold.

Note: Default values should be sufficient for client exception levels. Enabling a realistic threshold will allow the system to track and report clients by MAC address that fall below the threshold. If location services are available, this can be plotted on the map display. This can be very useful in troubleshooting the final deployment or to identify a new client type that has different RF capabilities.

General - Profile Threshold for Traps

- Adjust these for expected levels in the high-density coverage zone. It will make the traps and alerts meaningful

Wireless=>802.11a/802.11b=>High Throughput (802.11n)

- Should be enabled for both bands, with all MCS rates supported

Wireless=>Media Stream

- Enable if supporting video streaming applications
- For more information: Cisco Unified Wireless Network Solution: VideoStream Deployment Guide

For all of the above commands, it is strongly suggested that WCS and the AP and Controller Templates be used to make changes and assignments. Using the AP groups feature is also highly recommended for organizing APs by coverage zone and WLAN SSID configuration.

Conclusion

The performance of a high-density WLAN within a higher education environment is highly dependent on how well the network requirements are understood before the network is deployed. A good understanding of these concepts will enable the designer to modify the design to accommodate the unforeseen.

It is important to have alternate options available, based on changing factors, and to remain flexible in approach as new requirements and challenges are presented. For example, in most venues, aesthetics will be of paramount importance, and anticipating a design response to an aesthetic critique in advance is useful. Understanding the performance differences between an optimal solution and the less intrusive (or more aesthetically acceptable, with hidden APs) solution will leave the designer prepared to have that discussion and to re-set expectations if necessary.

This paper has presented several design areas as concepts. Recommended values have been offered as examples, based on experience with past solutions. The recommended values should be treated as a starting point and will provide a robust performance window. Performance may vary based on conditions that are beyond the designer's control, but understanding these concepts and controls will allow the designer to design and communicate a realistic expectation of performance.

Appendix A: 5 GHz Channels Available Worldwide by Regulatory Domain

These are channels available by the Cisco regulatory domain worldwide to date. Note the change in status for channels 120 to 132 for the -A region. Interference with terminal Doppler Weather Radar caused a change in regulations in October 2009. APs introduced before October 2009 will have these channels available. APs introduced after this date will not.

Channel	Frequency	-A	-E	-P	-S	-C	-I	-K	-N
		20/40 MHz	20/40 MHz	20/40 MHz	20 MHz	20 MHz	20 MHz	20 MHz	20/40 MHz
36	5180	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
40	5200	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
44	5220	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
48	5240	Yes	Yes	Yes	No	No	Yes	Yes	Yes
52	5260	Yes	Yes	Yes	No	No	Yes	Yes	Yes
56	5280	Yes	Yes	Yes	No	No	Yes	Yes	Yes

Channel	Frequency	-A	-E	-P	-S	-C	-I	-K	-N
		20/40 MHz	20/40 MHz	20/40 MHz	20 MHz	20 MHz	20 MHz	20 MHz	20/40 MHz
60	5300	Yes	Yes	Yes	No	No	Yes	Yes	Yes
64	5320	Yes	Yes	Yes	No	No	Yes	Yes	Yes
100	5500	Yes	Yes	Yes	No	No	No	Yes	Yes
104	5520	Yes	Yes	Yes	No	No	No	Yes	Yes
108	5540	Yes	Yes	Yes	No	No	No	Yes	Yes
112	5560	Yes	Yes	Yes	No	No	No	Yes	Yes
116	5580	Yes	Yes	Yes	No	No	No	Yes	Yes
120	5600	No	Yes	Yes	No	No	No	Yes	No
124	5620	No	Yes	Yes	No	No	No	Yes	No
128	5640	No	Yes	Yes	No	No	No	Yes	No
132	5660	No	Yes	Yes	No	No	No	No	Yes
136	5680	Yes	Yes	Yes	No	No	No	No	Yes
140	5700	Yes	Yes	Yes	No	No	No	No	Yes
149	5745	Yes	No	No	Yes	Yes	Yes	Yes	Yes
153	5765	Yes	No	No	Yes	Yes	Yes	Yes	Yes
157	5785	Yes	No	No	Yes	Yes	Yes	Yes	Yes
161	5805	Yes	No	No	Yes	Yes	Yes	Yes	Yes
165	5825	Yes	No	No	Yes	Yes	Yes	Yes	Yes
Total with DFS		20	19	19	8	5	13	21	21
Total without DFS		9	4	4	8	5	9	9	9

This is a country reference of Cisco Regulatory Domain settings on all APs and WLAN controllers.

Country Settings

Country	Reg Domain	Country	Reg Domain
Algeria	NA	Mexico	-N
Argentina	NA	Montenegro	-E
Australia	-N	New Zealand	-N or -A
Bahrain	-E	Nigeria	-E
Brazil	-t	Oman	NA
Bulgaria	-E	Pakistan	NA
Canada	-A	Panama	-N
Chile	-s	Paraguay	-A
China	-C	Peru	-A
Colombia	-A	Philippines	-A
Costa Rica	-A	Puerto Rico	-A
Croatia	-E	Romania	-E
Dominican Republic	-N	Russian Federation	-R
Ecuador	-A	Saudi Arabia	-E
Egypt	-I	Serbia	-E
EU Countries[*]	-E	Singapore	-s
Hong Kong	-N	South Africa	-E
India	-N	Taiwan	-A

Country	Reg Domain	Country	Reg Domain
Indonesia	NA	Thailand	NA
Israel	-I	Turkey	-I
Japan	-P	Ukraine	-E
Kazakhstan	NA	United Arab Emirates	-C
Kuwait	NA	United States	-A
South Korea	-K	Uruguay	-A
Malaysia	-C	Venezuela	NA
Macedonia	-I	Vietnam	-E

* EU Countries include: Austria, Belgium, Cyprus, Czech Republic, Denmark, Estonia, Finland, France, Germany, Gibraltar, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Liechtenstein, Lithuania, Luxembourg, Malta, Monaco, Netherlands, Norway, Poland, Portugal, Slovakia, Slovenia, Spain, Sweden, Switzerland, United Kingdom.

Summary and Resources

Cisco is uniquely positioned to provide the end-to-end architecture and component solutions that will deliver critical benefits to students, teachers and superintendents, and district CIOs and IT staff. Each will benefit differently from the Borderless Classroom, but the impacts will be an agent for change in the education community, while allowing schools and districts to adhere to compliance requirements to keep their students, parents, teachers, and staff safe.

Cisco is committed to the development and optimization of the Borderless Classroom and to further enhancing the depth and reach of the solutions, applications, and infrastructure presented here.

Resources & Contacts

Public Sector Strategic Programs & Solutions: Leo Kahng

Public Sector Solutions Management: Lori Bush

Herndon, VA Center of Excellence: Harman Cheema

Cisco Borderless Classroom mailer: connectedclass@cisco.com

The Cisco Borderless Networks Architecture

The Cisco Enterprise Medianet Architecture

Legal Disclaimer

The information contained in this Architecture document is proprietary and confidential to Cisco Systems, Inc. (Cisco). The Architecture document is furnished in confidence to the party requesting this Architecture document with the understanding that it will not, without the express written permission of Cisco, be used or disclosed for other than evaluation purposes, by only the recipient of this document.

This Architecture document is not and should not be construed as an offer to contract. Some information contained in this Architecture document may reference future technology under development. All such information is subject to change.

It should be noted that, in preparing this Architecture document, Cisco has made certain assumptions. Except as expressly stated in this Architecture document or as otherwise expressly agreed upon in writing by the involved parties, any descriptions, documentation, or references to Third-Party Products, to the extent they are provided in this Architecture document, shall be for informational purposes only.

Trademarks

Every effort has been made to identify trademark information in the accompanying text. However, this information may unintentionally have been omitted in referencing particular products. Product names that are not so noted may also be trademarks of their respective manufacturers.

CISCO

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**