

ENGINEERING

a WLAN network

White Paper

From idea to implementation— How to successfully deploy a WLAN

Introduction

In a *Wireless LAN (WLAN)* deployment, performance and capacity are impacted by several interacting factors that must all be considered to meet the requirements of a high-availability, high-performance, and mission-critical environment.

Most of the 802.11 technological innovations so far have been aimed at increasing the range and coverage of WLANs, mainly for small deployments not requiring the reuse of radio frequencies. This tutorial addresses the issues and factors that must be considered when planning large WLAN deployments with pre-defined performance and capacity objectives.

This document is aimed primarily at helping network planners and sales engineers to understand the factors affecting the capacity and coverage of Wireless LANs, but professional services staff, system integrators, partners, and some customers may also find it useful.

Section 2 of this document is a brief overview of the state of 802.11 WLAN radio technologies. This section is not intended to be a complete tutorial on 802.11 but rather an introduction to the topics that are relevant to the rest of the paper. It is focused almost entirely on the physical layers and assumes that the reader is familiar with the 802.11 architecture.

Section 3 outlines the issues related to high-capacity, high-performance, and large WLAN deployments capable of supporting mixed voice and data applications. The factors and tradeoffs that must be considered when planning high-performance Wireless LANs are described at a qualitative level.

Section 4 presents the factors affecting the quality of experience for data and voice users and characterizes these applications' typical requirements in terms of delay, loss, and throughput. Finally, Section 5 is an introduction to WLAN planning and site surveys.

Table of Contents

Introduction	3
IEEE 802.11 overview	3
Physical layers	3
Medium Access Control layer	4
– QoS extensions	4
Factors affecting capacity and performance	5
Radio link performance	5
– Rate vs. distance	5
– Differences between 11a, 11b, and 11g	6
– Efficiency vs. packet size	7
Single AP capacity	8
– Effect of data rates on cell capacity	8
– Capacity, coverage, and inter-AP spacing	8
– Contention and collisions	9
Multiple APs, co-channel interference, and capacity	10
– Frequency planning	10
– Co-channel interference and noise rise	11
– Interference in three dimensions	12
Application requirements	13
Quality of experience	13
– Data applications	13
– Voice	14
– Mixed voice and data applications	15
Typical traffic requirements	15
– Data	15
– Voice	15
WLAN planning	16
Site survey	16
– Coverage area	17
– Obstructions	17
– Interferences	17
Coverage vs. capacity	18
Conclusions	18
References	19

IEEE 802.11 overview

The base 802.11 standard was initially ratified by IEEE in 1997. It included an infrared and two spread-spectrum radio physical layers. Their rate was limited to 2 Mbps. Since then, a number of additions to this base standard have been proposed and adopted.

As any 802 protocol, the 802.11 standard covers the physical and *Medium Access Control (MAC)* layers. At the end of 2003, the 802.11 standard consisted of a single MAC and several physical layers as illustrated in Figure 1.

Physical layers

Three faster radio layers have been standardized since 1997: 802.11a and 802.11b in 1999, and 802.11g in 2003. Work on a new high-speed physical layer (*802.11n*) started in late 2003.

The *802.11b* physical layer is a backward-compatible extension of the original *Direct Sequence Spread Spectrum (DSSS)* radio physical layer in the 2.4-GHz band that supports up to 11-Mbps data rates.

The *802.11a* physical layer which was originally defined for the 5-GHz band supports up to 54-Mbps data rates using *Orthogonal Frequency Division Multiplexing (OFDM)* technology. Recently, the *802.11g* amendment extended the use of OFDM to the 2.4-GHz band with some minor modifications required for backwards compatibility with the 802.11b devices operating in this band.

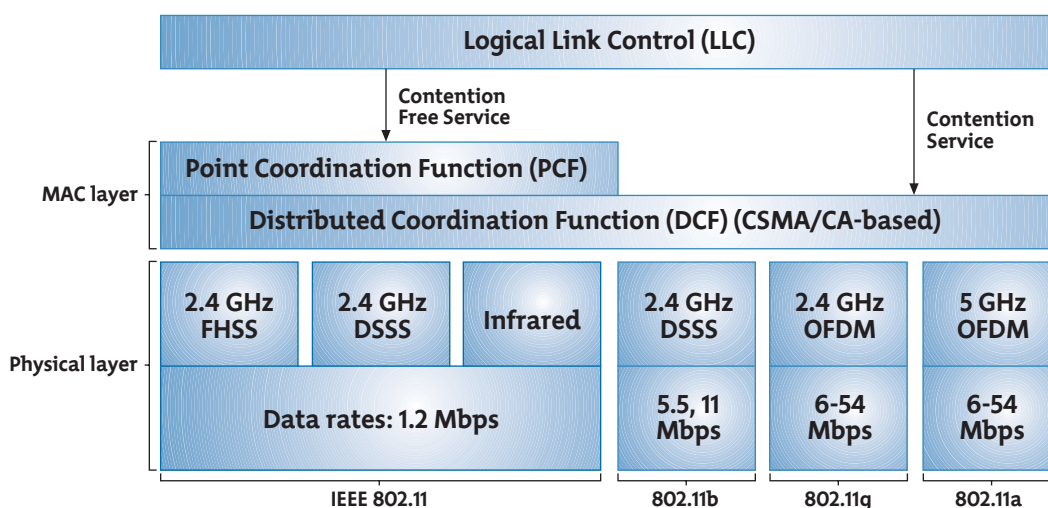


Figure 1. IEEE 802.11 standards

Table 1 below summarizes the essential characteristics of the 802.11 physical layers. The rates that must be supported by any equipment operating in the frequency band, the so-called mandatory rates, are underlined in this table.

	802.11b	802.11g	802.11a
Adopted by IEEE	1999	2003	1999
Technology	DSSS	OFDM and DSSS	OFDM
Frequency band	2.4 GHz	2.4 GHz	5 GHz
Channels (US)	3 non-overlapping	3 non-overlapping	13 increasing to 24 ¹
Physical rates	11, 5.5, 2, and 1 Mbps	All 11a and 11b rates	54, 48, 36, 24, 18, 12, 9, and 6 Mbps

Table 1. 802.11 physical layer characteristics

¹ In the United States, the FCC recently opened up new spectrum in anticipation of increased WLAN usage.

Medium Access Control layer

The 802.11 standard defines a common MAC layer² for all the 802.11 physical layers. The 802.11 frame format is slightly more complicated than the Ethernet frame format because of the need to carry wireless-specific information (e.g., the duration of the transmission), 802.11 security information (e.g., a per-frame message integrity code), and at least one more MAC address which is used to identify the intended radio recipient (i.e., acting as a bridge to the wired network) separately from the MAC layer destination. This extra information is stripped off by the access point as it forwards the frame on a wired 802 LAN.

The 802.11 basic MAC mechanism, called *Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA)*, is based on detecting and avoiding collisions with the transmissions of other devices. This is similar to the well-known Ethernet CSMA/CD protocol except that the Ethernet collision detection which could not be used reliably in a wireless medium has been replaced by a collision avoidance and positive acknowledgment scheme. Stations sense the wireless medium and only attempt to transmit when it appears to be idle. Failure to receive an acknowledgement indicates to the transmitting station that a collision has occurred. If it is configured to retry failed transmissions, the station retransmits the frame with extra precautions: this is the collision avoidance procedure.

In addition to the basic acknowledged frame exchange, an optional *Request-To-Send/Clear-To-Send (RTS/CTS)* mechanism is defined to increase the robustness of the MAC protocol. This mechanism reduces the impact of collisions for long transmissions: a collision is detected at the end of the short RTS/CTS exchange rather than at the end of the transmission of the longer data frame.

The RTS/CTS exchange also attempts to address the so-called “hidden node” problem which occurs when wireless devices are far enough away from each other (e.g., on opposite sides of an access point) or hidden from each other by some obstruction so that they can’t detect each other’s transmissions. Wireless devices that are hidden from each other incorrectly assess that the channel is clear when the other device in the pair is in fact transmitting. Hidden nodes tend to experience a higher rate of collisions. For all practical purposes, the RTS/CTS exchange is only useful in WLAN environments where there is a high amount of traffic and the packet sizes are large. Otherwise, the time saved in avoiding “hidden-node” collisions is offset by the added delay caused by the RTS/CTS exchange.

Extensions that address some of the shortcomings of the original MAC design in the areas of security (*802.11i*) and Quality of Service (*802.11e*) are expected to be ratified in 2004. The *Wireless Fidelity (Wi-Fi)* Alliance, a nonprofit organization formed in 1999 to certify interoperability of WLAN products based on the IEEE 802.11 specifications, included a subset of the 802.11i security features, which it calls *Wi-Fi Protected Access™ (WPA)*, into its certification tests in the fall of 2003. Similarly, the Wi-Fi Alliance has announced it will begin testing a subset of the 802.11e QoS features, called *Wi-Fi Multimedia Enhancements™ (WME)*, in the fall of 2004.

QoS extensions

The basic 802.11 MAC mechanism has no support for Quality of Service (QoS): all WLAN traffic is treated on a best effort basis. QoS extensions are required in WLAN systems that support mixed voice and data traffic in order to provide a positive end-user experience. Because the different traffic types have different requirements for delay and packet drop rate, QoS extensions allow for traffic types to be treated differently during WLAN medium access, thus allowing specific performance criteria to be met.

A second optional access control mechanism called *Point Coordination Function (PCF)* was included in the original 802.11 standard to provide deterministic access to the wireless medium for time-sensitive traffic. PCF is based on resource reservation and uses polling by the access point to schedule transmissions. Although the PCF mechanism was intended to support QoS, several problems have been identified with its original specification and it has not been widely adopted.

The proposed 802.11e QoS amendment defines a new coordination function called the *Hybrid Coordination Function (HCF)*. HCF has two modes of operation: the *Enhanced Distributed Channel Access (EDCA)* and the optional *HCF Controlled Channel Access (HCCA)*.

EDCA is an extension of the original basic channel access mechanism that provides differentiated access to the wireless medium by varying the amount of time a station senses the channel to be idle, the length of the contention window during a backoff, and the duration a station may transmit after it gets access to the medium.

HCCA fixes the shortcoming of the original PCF polling mechanism and provides deterministic access for predictable, time-sensitive traffic. The access point gains control of the wireless medium as needed to send QoS traffic and to issue QoS polls to stations by waiting a shorter time between transmissions than the stations using the EDCA access procedure.

Proprietary QoS extensions exist in the market particularly on systems targeted towards delivering Voice over WLAN. The need for these proprietary solutions will diminish as deployment and support of 802.11e becomes more widespread and the Wi-Fi Alliance’s WME testing ensures interoperability between different vendor implementations of the standard.

² Some of the parameters controlling the MAC layer behaviour are physical layer dependent.

Factors affecting capacity and performance

The wireless medium is a very dynamic shared environment affected by several interacting factors. Some of these factors can be controlled while others are fundamental limitations of the wireless medium that must be recognized and taken into account when planning a WLAN.

Radio link performance

The wireless medium is not a wire. It is a highly changing environment in which transmission errors are unavoidable and quite common. Wireless signals suffer attenuations as they propagate through space, especially inside buildings where walls, furniture, and other obstacles cause absorptions, reflections, and refractions.

Rate vs. distance

The farther a wireless device is from its *Access Point*, the weaker the signal it receives and the lower the physical rates that it can reliably achieve. The radio link throughput is a function of a number of factors including the physical data rate and the frame error rate (the frame error rate is itself an increasing function of the data rate). A high frame error rate may defeat the speed advantages of a high data rate by causing too many retransmissions. 802.11 devices constantly monitor the quality of the signals received from devices with which they communicate. When their turn comes to transmit, they use this information to select the data rate that is expected to provide the highest throughput (i.e., the best compromise between speed and reliability).

On average, the data rate will fall off in direct relation to the distance from the access point. The red data points in Figure 2 show a plot of the average data rate users might experience as a function of their distance to the AP. In a typical setting where there are walls, cubicles, hallways, etc., the actual achieved data rates will be more random depending on the number and types of obstacles between the mobile unit and the APs at each specific location. These obstacles cause signal absorption and reflections resulting in what is called “shadowing”. The green data points in Figure 2 illustrate what the data rate versus distance might be in an environment where shadowing exists.

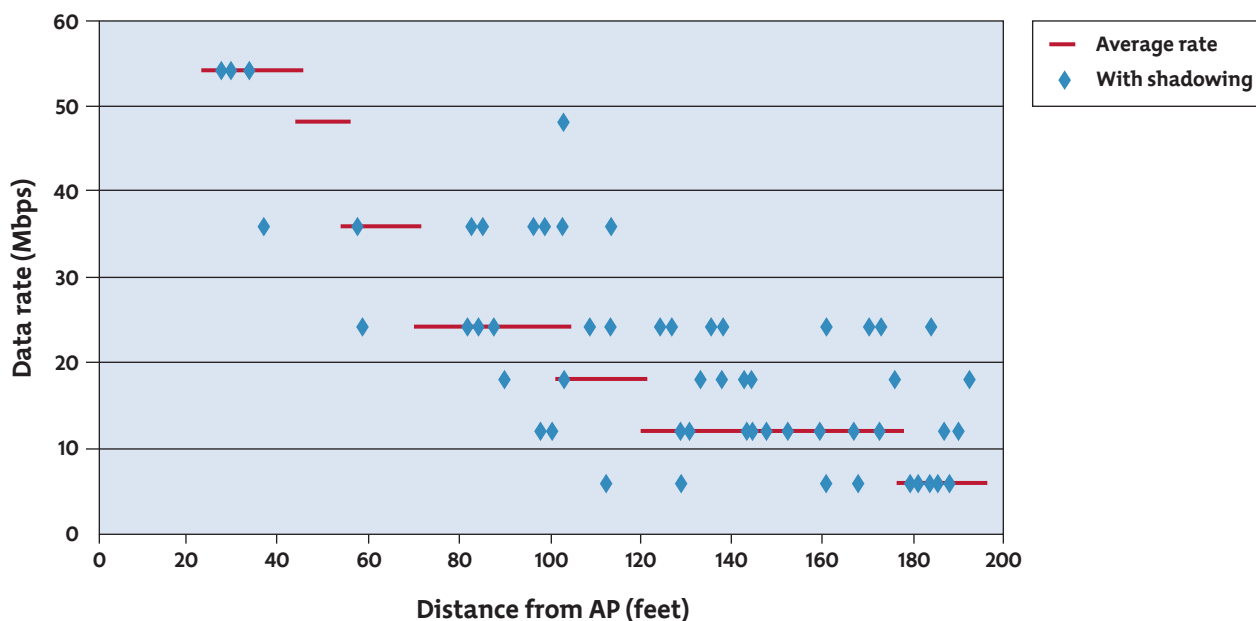


Figure 2. Data rate vs. distance

One example from Figure 2 shows how a user 36 feet from the AP might obtain a data rate of 54-Mbps, and another user the same distance from the AP, but with different obstacles between the AP and its receiver, might only obtain a 36-Mbps data rate.

The coverage of an AP can also be viewed in two dimensions as seen in Figure 3. The left side of Figure 3, labeled “Average AP coverage”, shows the average coverage of an 802.11a AP where the physical rates have been color-coded. It is worth noting that the maximum rate of 54 Mbps is only achieved within the center ring representing only four percent of the total AP coverage area. Note that the figure is a schematic representation and the exact dimensions of the rings depend on a number of factors, including transmit power and receiver sensitivity.

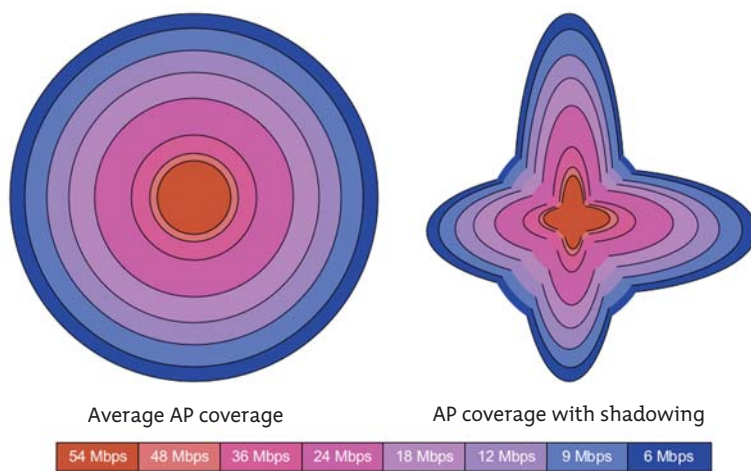


Figure 3. Schematic representation of 802.11a/g AP coverage

Figure 3 on the right, labeled “AP coverage with shadowing,” shows what the AP coverage with shadowing effects might look like. It represents a case where there is an AP positioned at the intersection of two corridors. The walls of a corridor might act as a wave guide carrying the signal further in this direction than expected. Conversely, radio signals might have a hard time penetrating through these same walls.

In order to guarantee that users at any location within the deployed WLAN area can achieve a minimum data rate with a certain degree of probability, it may be necessary to factor in an additional margin that accounts for the effects of shadowing. Depending on the variance of the shadowing in a particular environment, the margin can be determined and factored in as a reduction in the expected range. As an example when doing network dimensioning calculations, a requirement might be to ensure a 95 percent probability that all users achieve; for example, a 36-Mbps minimum data rate. In this case, a reduction in the range based on the typical variations in shadowing effects will have to be factored in. This may result in APs being placed more densely in order to guarantee coverage.

Differences between 11a, 11b, and 11g

Some early 802.11a devices had notoriously bad performance and there is still some confusion in the industry about the range and capacity of WLANs in the 2.4- and 5-GHz bands.

Although some specific building materials exhibit different absorption and reflection characteristics at 2.4 and 5 GHz, differences in the average indoor propagation models at 2.4- and 5-GHz are small and, therefore, it should be possible to achieve roughly equivalent range performance anywhere in the 2.4- to 5-GHz band. However, achieving similar performance at higher frequencies requires more careful radio designs: higher frequencies tend to suffer higher losses inside circuit boards; cheap amplifiers have reduced gain at higher frequencies; and the collecting areas of nearly isotropic antennae are reduced because of the shorter wavelength.

Tests conducted with current generation equipment show that the maximum ranges for 11a, 11b, and 11g are essentially the same. However, the simpler but less efficient DSSS modulation used by 11b puts it at a significant speed disadvantage and it is likely that in the future equipment operating in the 2.4-GHz band will predominantly use OFDM modulation i.e., 802.11g.

There is no significant difference in range between 2.4- and 5-GHz frequencies.
For any distance, the achievable rate is greater with 11a than with 11b.

When 11b and 11g devices operate simultaneously in a WLAN cell, the 11g devices must take some extra precautions to protect their OFDM signals from interference by the 11b devices. Because 11b devices are incapable of decoding OFDM signals, the 11g devices are permanently hidden to them. To avoid interferences from 11b devices, the 11g devices must reveal their presence by prefacing their OFDM transmissions with a CTS-to-Self or an RTS/CTS exchange that is transmitted using DSSS modulation. The overhead associated with this extra transmission is quite substantial and reduces considerably the effective throughput of a cell operating in 11b compatibility mode.

WLANs operating with a mix of 11b and 11g devices have less capacity than WLANs operating with only 11g devices. In the long term, 11b devices will be replaced by 11g devices and thus the network will achieve its full capacity.

Efficiency vs. packet size

The transmission of a packet over the wireless medium involves a fixed overhead (inter-frame gap, preamble, and acknowledgement). The efficiency of the 802.11 MAC and physical layers varies substantially with the size of the packets and the type of exchange involved.




Figure 4. 802.11a vs. 11b maximum effective throughput

Figure 4 above shows the maximum effective throughput as a function of the packet size for the eight 802.11a and the four 802.11b physical rates, assuming a simple two message exchange. It is worth noting that up to about 1000 bytes the lowest 802.11a rate (6 Mbps) outperforms the highest 802.11b rate (11 Mbps). Even for 1500-byte packets and ideal radio conditions (no retransmissions due to bit errors or collisions), the MAC layer throughput is substantially less than the physical layer rate. At the 54-Mbps physical rate, the maximum theoretical throughput of 802.11a is roughly 30 Mbps.

The average packet size depends on the application mix, but it is generally not very high. For example, the average packet size on the Internet is somewhere between 400 and 500 bytes. If voice traffic constitutes a substantial portion of the mix, the average may be even lower (VoIP packets are typically only 200 bytes or shorter).

The achievable WLAN throughput depends on the average packet size. Assuming an average packet size of 500 bytes, the maximum theoretical throughput of a WLAN link is about 16.3 Mbps for 11a, 3.3 Mbps for 11b, and 4.9 Mbps for 11g in 11b-compatibility mode.




Figure 5. 802.11a vs. 11g maximum effective throughput

As mentioned previously, two protection mechanisms could be used in a mixed 11g/11b WLAN: the lower-overhead CTS-to-Self, in which a device transmits a CTS message addressed to itself (the RTS is omitted), or the more robust RTS/CTS³. Figure 5 shows the dramatic impact that even the CTS-to-Self protection mechanism has on the effective throughput of 11g devices.

³ The CTS-to-Self has lower overhead because the RTS is omitted but it is not as robust as a regular RTS/CTS exchange because the device does not know if its CTS transmission was successful.

Backward compatibility with 11b devices is possible but it comes at a high price. In fact, for shorter frames (e.g., TCP acknowledgements) it is often more efficient to transmit the frame directly using DSSS modulation rather than transmitting a CTS-to-Self using DSSS modulation followed by the actual frame using OFDM modulation.

Single AP capacity

With the notable exception of the 11g protection mechanisms, the above analysis of the 802.11 performance has focused on an isolated radio link between an access point and a wireless device. When multiple devices are simultaneously active in a cell, contention, collisions, and other types of interferences further reduce the medium capacity.

The overall reduction in media capacity due to the effects mitigating single AP capacity (described in subsequent sections) can be as much as 60 to 70 percent when compared to the ideal case of all users achieving the maximum data rate and experiencing no collisions nor contention. For example, the single AP capacity for an 802.11a system, in the ideal case where packet sizes are 1500 bytes, is 30 Mbps (see Figure 4). However, for the non-ideal case where wireless devices use all 802.11a rates and experience contention and collisions, the capacity may drop by as much as 60 percent. Another example is an 802.11b system supporting voice traffic (200-byte IP packets). If these WLAN phones are far away from the AP, the WLAN capacity may drop by as much as 70 percent simply because of the low data rates.

Effect of data rates on cell capacity

As explained in section 3.1.1, the rate selected for transmission between a wireless device and its access point limits the throughput that can be achieved between them. There is, however, a more subtle effect. The wireless medium is a shared medium and the 802.11 MAC layer must guarantee that all devices get an equal chance to access the medium. The 802.11 designers chose to measure fairness in terms of number of transmission opportunities and not in terms of number of bytes transmitted or of time spent utilizing the medium.

During periods of heavy usage, devices transmitting short packets are at a significant disadvantage. Furthermore, devices with slow connections seize the medium for a proportionally longer time and drag down the average throughput of the cell. The worst situation occurs when a wireless device at the very edge of a cell transmits or receives large frames at the slowest possible rate and experiences a high error rate due to the combined effects of marginal signal quality and large packet sizes.

The particular 802.11 MAC media access mechanism causes a rather unexpected averaging effect. If two wireless devices, one connected at a high data rate and another at a low data rate, are generating the same traffic (e.g., same size packets) and saturating the medium, the capacity of the medium will result in a value closer to the lower rate than the higher rate. For example, two users, one connected at 54 Mbps and the other at 6 Mbps, will result in an overall medium capacity of 10.8 Mbps. This is analogous to what happens when trying to compute the average speed of a car. Assume a car travels 108 miles with the first half of the trip covered at high speed, say 54 miles per hour, and the second half at the much slower 6 miles per hour speed (obviously a major traffic jam). It turns out that the average speed for the entire trip is only 10.8 miles per hour and not the 30 miles per hour one may naively have expected.

This averaging effect, sometimes referred to as the “edge user effect”, explains why an increase of the cell coverage may cause a decrease of the throughput even for the devices that are at close range of the access point.

Low data rates caused by large cells or inefficient devices should be avoided in order to maximize the capacity and performance of a WLAN.

Capacity, coverage, and inter-AP spacing

The two main metrics when describing WLAN performance are capacity and coverage. Both of these are directly affected by the inter-AP spacing, or the AP density, in a system. This section describes the interaction between these three parameters, considering that all other parameters remain fixed in the system (transmit powers, receiver sensitivity, etc.). In particular, this section shows how these three parameters interplay when there are enough channels available to avoid reusing frequencies.

Capacity is the amount of throughput that a WLAN system is able to provide to users, taking into account all the mitigating effects on single AP capacity, and multiple AP capacity, as discussed later. The coverage refers to the area over which the wireless signals propagate, and can best be described as the probability that a user will be able to reliably connect to an AP from any location within the WLAN. Capacity and coverage are affected by the inter-AP spacing which is in direct relation to the physical area over which the AP is required to provide connectivity.

When the inter-AP spacing is decreased, the required service area for each AP decreases, and the average distance from a user to the closest AP decreases. Considering that none of the other WLAN parameters have changed, this has two effects. First, it increases the per AP capacity by allowing for the mobile units on the same AP to connect at higher data rates on average. Second, it increases the coverage reliability by increasing margin of safety in the range to achieve the minimum required connect speed even in the case of shadowing.

Coverage reliability is further increased by the fact that, at each location, there are now more likely overlapping RF signals from different APs arriving from different directions and providing more chances for a good signal in the presence of obstacles. Therefore, the coverage, i.e., the probability that a user will be able to associate with at least one AP at the minimum connect speed within the WLAN system, has increased.

Figure 6 depicts an example of a WLAN deployment where in the left hand figure there are 6 cells. This example assumes that the size of the service area yields an average data rate of 12 Mbps and a total system capacity of 6 times 12 or 72 Mbps. If the AP density is doubled, and in turn the AP service area is halved, as in the depiction on the right in Figure 6, users can now connect at a higher average data rate of 24 Mbps (for the cell shown in Figure 3, reducing the service area by 54 percent increases the average rate from 12 to about 24 Mbps). This yields a total system capacity of 12 times 24 or 288 Mbps. So by reducing the area an AP serves, thus allowing for increased average data rates, the available system capacity can be increased by more than the factor of AP density increase. Note that the circles in Figure 6 are depicting the 9 and 18 Mbps service area per AP respectively and not the coverage area. The coverage area per AP for both figures will be the same—a much larger circle than the ones in either side of the figure.

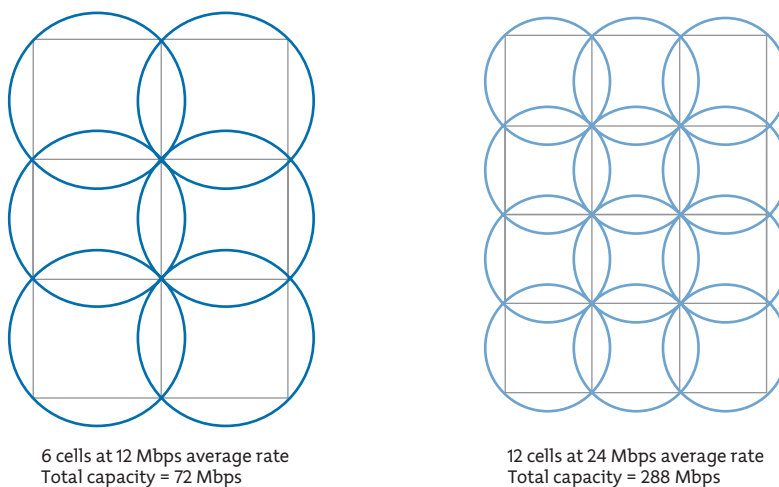


Figure 6. AP service area vs. capacity

Decreasing the service area of an AP increases its capacity. Assuming they can all use different frequencies, increasing the density of APs in an area results in an increase of the overall WLAN capacity due to: 1) the increased number of APs and 2) the increased per-AP capacity. Doubling the number of APs in an area more than doubles its capacity!

This increase in capacity is only assured if enough non-overlapping channel frequencies are available for complete coverage of the denser layout without requiring any frequency reuse. If this is not the case, co-channel interference could affect the overall capacity to the point of dramatically reducing or even completely negating the benefits of the smaller AP service areas. For 802.11b and 802.11g with only three non-overlapping channels, this will often be a limiting factor in medium and large WLAN deployments.

Also, mobile units that have efficient roaming so as to continually connect to the “best” AP, i.e., the AP that gives them the highest data rate, are required in order to achieve a capacity increase. Otherwise, since the transmit power and the coverage per AP have not changed, devices may be able to connect to multiple APs reliably, just at different rates. If a user moves into another AP’s service area but remains associated with the previous AP, it may need to use a lower data rate and may thus create the edge user effect described in the previous section.

Contention and collisions

Instead of the familiar collision detection mechanism used in wired Ethernet networks, 802.11 wireless LANs utilize a collision avoidance mechanism to arbitrate between devices trying to access the shared wireless medium.

A little contention improves the medium efficiency: if a few devices try to simultaneously access the medium, the one that draws the shortest backoff interval gets to transmit while the others wait for the next transmission opportunity. The medium utilization is slightly increased as a result of the shorter deferrals.

Unfortunately, too much contention reduces efficiency: if many devices try to access the medium at the same time, chances are that two or more transmissions will collide. When this happens, the entire duration of the longer transmission is wasted. Furthermore, the devices involved in the collision are forced to wait even longer before attempting to transmit again. Scarce bandwidth is wasted by collisions.

The bandwidth of a wireless LAN cannot be divided up between arbitrarily large numbers of devices. To keep the medium efficiency high, the probability of collisions must be kept low (e.g., a few percent).

The number of active wireless devices that can be efficiently supported in a WLAN cell is limited. If the number of active devices is too high, collisions could destroy the system performance.

Multiple APs, co-channel interference, and capacity

As long as radio frequencies are not reused, the capacity of a WLAN composed of multiple access points is essentially the sum of the single AP capacities of the constituent cells. As soon as the number of access points increases beyond the number of available frequencies, care must be taken to minimize the amount of interference between cells using the same frequencies.

Co-channel interferences are caused by the transmissions of devices in remote cells using the same frequency channel. These signals are usually too faint to be properly detected outside of the originating cell but are still strong enough to disrupt traffic in the nearby cells using the same frequency.

Co-channel interference occurs because of the very nature of the wireless medium: radio signals are not confined to the cell from which they emanate but rather propagate in all directions. Fortunately, even in the absence of obstacles, radio signals get weaker the farther away they travel from their source. Frequency planning involves reusing frequencies across a two- or three-dimensional space while positioning cells that use the same frequencies as far away as possible from each other to minimize interferences.

Wired networks can be scaled by adding additional bandwidth or additional links in order to increase overall user throughput. In large WLAN deployments, however, when frequencies must be reused and co-channel interference is the limiting factor, adding more APs to the system in an attempt to increase capacity may actually have the opposite effect as the amount of interference traffic increases, causing the system performance to drop further.

The overall reduction in media capacity due to co-channel interference effects can be as much as 85 percent when compared to the ideal single AP case when all users achieve the highest data rate. For example, using a packet size of 1500 bytes, in an 11a, two-dimensional, multiple-AP system the capacity can drop by as much as 85, 75, or 60 percent for a reuse plan of 4, 7, and 12 frequencies respectively. In an 11b system supporting voice traffic (200-byte IP packets), the capacity can drop by as much as 80 percent when the three non-overlapping channels are used.

Frequency planning

Frequency planning involves assigning frequencies to access points in a way that maximizes the distance between cells using the same frequency (the reuse distance). For a given cell layout, a larger number of frequencies implies a larger reuse distance. The number of frequencies used in the planning is called the reuse plan. Figure 7 below illustrates the concept of frequency planning on an ideal hexagonal grid. The layout on the left uses seven different frequencies and each frequency is used exactly seven times. The access point at the center of the network is surrounded by a ring of six cells using the same frequency (denoted by the blue color).

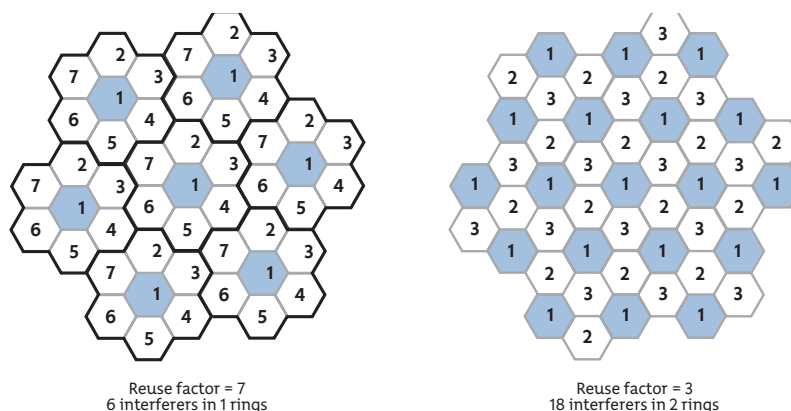


Figure 7. Frequency planning in two dimensions

The layout on the right of Figure 7 uses only three frequencies. The center access point is surrounded by 18 cells using the same frequency. These cells are likely to generate noticeable interferences when their traffic load is high. Furthermore, in this case the first ring of interferers is 35 percent closer to the central access point than it was when seven frequencies were used.

In two dimensions, for a given cell size, the maximum reuse distance that can be achieved on a regular grid (square or hexagonal) grows as the square root of the number of available frequencies: for example, the distance that can be achieved with 12 frequencies is double the distance that can be achieved with only three. This doubling of the reuse distance translates into a significant reduction of the level of interferences (about 9 dB).

In a real world deployment, frequency planning may involve exploiting barriers to the propagation of radio signals that are present in the environment. For example, concrete walls or floors can be used very effectively in conjunction with distance to prevent access points using the same frequencies from interfering with each other.

In order to minimize interferences, cells using the same frequency should be kept as far away from each other as possible by using all the available non-overlapping channels. Simulations and practical experience show that the number of frequencies must be greater than three to achieve suitable reuse distances in large WLANs.

Co-channel interference and noise rise

In North America, the number of non-overlapping channels available for WLANs in the 2.4-GHz band is three. In the 5-GHz band this number is presently 13 and is expected to increase soon to 24. For deployments requiring fewer frequencies than are available, the range and capacity of a single access point deployment is possible in each of the cells. There is no interference between cells and the total capacity of the network grows linearly with the number of cells (Figure 8a).

Unless the number of available frequencies is very large, a large WLAN deployment will require some amount of frequency reuse. If frequencies are reused and the reuse plan is sufficiently large, traffic in other cells using the same frequency band will result in an increase of the background noise of a given cell. The amount of noise rise depends on the traffic load in the interfering cells and the reuse plan.

A key parameter related to the frequency reuse factor and governing the co-channel interference phenomenon is the ratio of the distance between interfering cells to the cell radius (i.e., not the absolute cell size). Increasing the cell size does not alleviate the co-channel interference problem: with a larger cell radius the interference is reduced because the interferers are further away but the cell edge is equally further away from the access point and the signal strength is reduced in the same proportion as the interferences.




Figure 8. Impact of co-channel interference

In a lightly loaded WLAN with a sufficiently large reuse factor, the noise rise due to co-channel interferences is negligible: the range and capacity of a cell are essentially the same as that of a single access point deployment (Figure 8b).

In a heavily loaded WLAN, interferences from neighboring cells using the same frequency band significantly increase the level of background noise, thereby reducing the **Signal to Noise Ratio (SNR)**, the capacity, and the range of the cell (Figure 8c). Increasing the signal strength does not solve this problem since the level of noise due to co-channel interference is proportional to the signal strength.

If the WLAN was not designed with sufficient overlap between cells, coverage gaps may result from this reduction of the cell range. Connect rates and overall cell capacity will also drop as the background noise rises (Figure 8d). This effect will only manifest itself in a heavily loaded wireless system and will typically⁴ not be detected during the initial site survey.

For a given cell layout, the noise rise due to co-channel interferences is essentially independent of the transmit power and the cell size. The frequency reuse factor, the traffic load, and the way radio signals propagate in the environment are the key parameters affecting the noise rise.

Interference in three dimensions

The above analysis assumed a two-dimensional cell layout. WLANs are often used in multi-storey buildings where possible interferences from the floors above and below must also be taken into account. Figure 9 shows a frequency plan for a large building. Eight different frequencies are used in this plan: the four odd-numbered frequencies are used on the odd-numbered floors and the even-numbered frequencies are used on the even-numbered floors. The frequency plans for the floors and the corresponding plans for vertical sections are shown on the left and right sides of the picture respectively. The cells using frequency 1 are colored in blue. As can be seen by the large number of cells in close proximity of each other in the perspective view at the center of the figure, the interference problem is potentially much more serious in three dimensions.




Figure 9. Frequency planning in three dimensions

First of all, in three dimensions, a greater reuse factor is required to achieve the same reuse distance. In two dimensions, the reuse distance grows as the square root of the reuse factor; in three dimensions, it only grows as the cube root. In other words, if a reuse factor of four was necessary to achieve a given reuse distance in two dimensions, eight frequencies would be required in three dimensions. More worryingly, nine frequencies in two dimensions would imply 27 frequencies in three dimensions.

Secondly, because in three dimensions a cell also has neighbors above and below, the number of close interferers is doubled in the three-dimensional equivalent of the hexagonal cell packing: there were only six interferers on the first ring in two dimensions, but there are 12 close interferers in three dimensions (e.g., four on the same floor, four above, and four below). For the same reuse distance, this doubling of the number of interferers could potentially result in a 3 dB noise rise. Alternatively, an even greater reuse factor would be required to compensate for this effect.

Finally, in three dimensions, the number of interferers that are further away grows as the square of distance (vs. linearly in 2D): there are $10n^2+2$ interferers in ring n in three dimensions whereas there are only $6n$ in the corresponding ring in two dimensions. For example, there are 42 interferers in the second ring of interferers in three dimensions but the corresponding number in two dimensions is only 12.

⁴ A notable exception is when the interference is generated by an already established nearby Wireless LAN.

Fortunately, techniques are available to mitigate these effects. For example, concrete floors attenuate radio signals in the 2.4- to 5-GHz range by several decibels. This increases the effective distance between cells on different floors. In addition, antenna technology such as azimuth control can be applied to confine most of the radio energy emitted by access points to the floors where they are deployed. This, however, does not work for the mobile wireless devices that, since they have no predetermined spatial orientation, must use omnidirectional antennae.

Co-channel interference issues are more challenging in multi-storey buildings requiring even more distinct channels to ensure adequate separation of cells using the same frequency.

Application requirements

Wireless capacity is a finite, somewhat scarce resource and, therefore, engineering with reckless abandon, as is customary for wired LANs, is usually not an option for WLANs. A good handle on the applications and the load they will generate is essential to avoid overestimating the cost of the WLAN. Furthermore, because of the shared nature of the wireless medium, deploying too many access points might simply contribute to polluting the air waves without any tangible increase in capacity.

Key application parameters to consider when planning a WLAN deployment include the application mix (e.g., the relative proportion of voice and data traffic), the expected traffic demands (e.g., WLAN voice usage will be a lot heavier if wireless phones replace rather than supplement desktop phones), and the application performance requirements.

These application requirements must eventually translate into constraints for the WLAN deployment such as minimum rate and maximum cell coverage.

Quality of experience

The quality of experience is the overall performance of a system from the point of view of its end-users. It is a measure of how the system enables the users to do what they need to do when they need to do it. The performance metrics and performance targets used to assess the user's quality of experience depend upon the type of application.

Data applications

For data applications, the time elapsed between the moment the user issues a command and the time when the output of the command is displayed at the user's device is a key measure of the perceived quality of the system. Different applications have varying response time requirements.

Interactive applications such as, for example, terminal emulation, in which users issue commands and expect 'immediate' results, typically require end-to-end, round-trip delays of less than 400ms. This may seem like a lot but includes not only the network propagation, serialization and queuing delays, but also the TCP timeouts and retransmissions, and the processing delays in the end systems. The wireless link rate has only a limited effect on the response time when small amounts of data are transmitted. The delay and quality of the connection (i.e., the probability of losses caused by transmission errors or buffer overflows) are the most important factor in providing good quality of experience for applications that use small transactions.

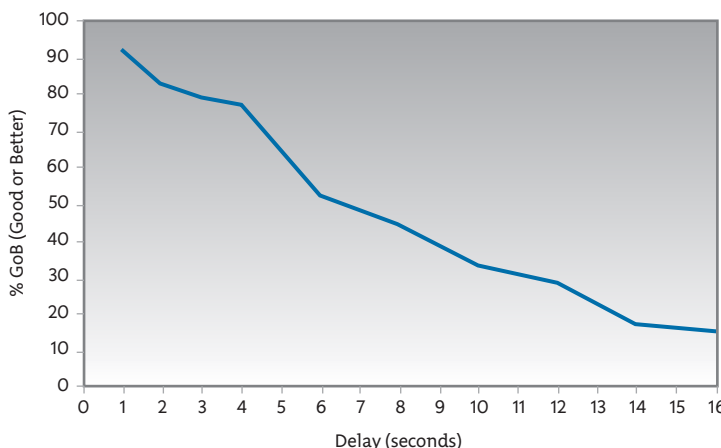


Figure 10. Impact of delay on browser-based applications

For applications in which larger amounts of data are transferred (i.e., a few kilobytes) such as e-mail and browser-based applications, users are usually prepared to tolerate longer delays of the order of a few seconds. As shown in Figure 10 above, user satisfaction decreases dramatically when this delay is more than four seconds. Because of the larger amount of data that must be transferred, the rate of the connection is more important in this case. Link quality remains essential though because TCP, the transport protocol used by a majority of data applications, interprets packet losses as an indication that it should slow down its transmission rate. Even if the link rate is high, the TCP throughput will not be very high unless the packet loss rate remains quite low.

The throughput and quality of a connection are related to the probability of losses caused by transmission errors or buffer overflows. A loss rate of less than 1 percent (preferably less than 0.1 percent) is recommended to achieve satisfactory quality of experience for TCP-based data applications.

Voice

Voice applications are very sensitive to delay and distortions. Conversational voice has much more stringent delay requirements than streaming audio (e.g., voice mail).

For excellent conversational voice quality of experience, the end-to-end one-way delay should be less than 150 milliseconds (see Figure 11 below). Beyond this point some users may notice the excessive delay. When the one-way delay exceeds 300 milliseconds many users will most probably be dissatisfied with the voice quality.

A number of factors contribute to the one-way delay of VoIP connections and, therefore, the WLAN can only use a small fraction of the total 150 milliseconds delay budget. Significant contributors to the delay include the packetization delay (typically 20 or 30ms but could be more if larger packets or codecs requiring a look-ahead are employed), the propagation delay (five milliseconds per 1000 kilometers in optical fiber could result in a 30-millisecond delay for a trans-continental call), and the de-jitter and playout delay (in some implementations this may be as high as the packetization delay—20 or 30 milliseconds even if the actual jitter is low). Other delay contributors include processing delays in the end systems and queuing delays in the network routers and switches⁵. In the case of a long distance call, the fraction of the delay budget that can be spent in the WLAN is typically less than half⁶ of the 150 milliseconds total.

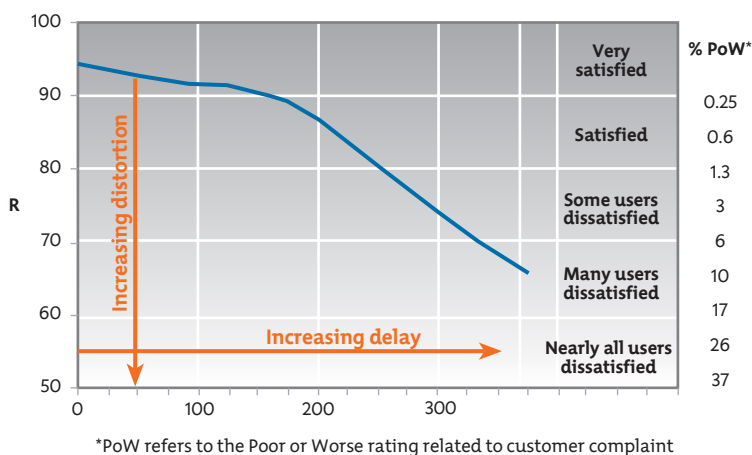


Figure 11. Impact of delay on perceived voice quality

Distortions include artifacts caused by inadequate sound level, echo level, speech interruptions, packet loss, and codec artifacts. Frequent interruptions of the speech path due to network failures (e.g., during handoffs) should not exceed 80 milliseconds. Long interruptions of three seconds or more are perceived as a call drop. These can only be tolerated if they are very infrequent.

In the absence of sophisticated packet loss concealment, the packet loss rate should be kept below 0.1 percent for conversational voice applications. Interruptions of the speech path during handoffs should generally not exceed 80 milliseconds.

⁵ Besides the WLAN deployment, an assessment of the wired LAN is required prior to any IP Telephony deployment to ensure that it does not become the bottleneck.

⁶ In fact the WLAN delay would need to be a lot less if both ends of the call are on a WLAN or if packetization is happening twice in the call (packet-TDM-packet call).

Mixed voice and data applications

The traffic streams generated by voice and data applications have very different characteristics and it is even more challenging to meet the requirements of both types of traffic with one network. Voice traffic is made up of short packets fairly evenly distributed in time. As long as they do not exhaust the medium capacity, several voice streams can coexist on the wireless LAN without any noticeable impact on voice quality.

Data applications, on the other hand, tend to generate bursts of rather large packets. These bursts often involve several kilobytes of data. Bursty streams can coexist on the same LAN without any significant degradation of data application performance. But in a wireless network, if no precautions are taken, even a single bursty data stream can temporarily saturate the medium and cause voice quality impacting delays and losses on an otherwise lightly loaded WLAN.

A mixed voice and data environment deployed with the current 802.11 standard is unlikely to result in satisfactory quality of experience, especially for the voice users. In the absence of any QoS mechanisms, a very effective workaround involves separating voice and data traffic users by frequency bands, such as, for instance, using 802.11b/g for voice users and 802.11a for data users. This approach is relatively easy to implement with multi-mode APs but requires two channels per cell, one in the 2.4-GHz band and another one in the 5-GHz band. More importantly, the number of frequencies available for planning each of the voice and data WLANs is smaller and the co-channel interference problems are therefore more serious in a large deployment.

QoS mechanisms are the most efficient way to ensure that the voice traffic gets higher priority access to the wireless medium than the competing data traffic, and will thus allow for a more satisfactory quality of experience when voice and data traffic share the same channels. Once standard QoS mechanisms are supported by multi-mode devices, the combined channels available in the 2.4- and 5-GHz bands can be used to engineer large WLANs.

The advent of multi-mode 802.11a/g wireless devices supporting WME QoS mechanisms enables a positive wireless multimedia experience.

Typical traffic requirements

Applications requirements must be characterized in terms of a typical traffic load and associated performance metrics such as delay and packet drop rate.

Data

Unlike conversational voice traffic which consists of a steady, uninterrupted stream of packets⁷, data traffic tends to be bursty with short spurts of intense activity separated by long idle periods. The total amount of traffic generated by interactive data applications tends to be very small on average. For planning purposes, this traffic should be characterized either by collecting representative traffic traces or, if this is not an option, by estimating it based on similar deployments.

For example, for good Web browsing performance in a typical office environment, the total number of users connected to an 802.11b access point should not exceed 30. This figure assumes that most users are connected at 11 Mbps and the rest at 5.5 Mbps. The number of data users that could be simultaneously supported on an 802.11a access point is roughly three to four times higher than that of 802.11b, assuming most users are sufficiently close to the access point for connecting at high rates (24 Mbps or higher).

In a mixed voice and data environment, the total number of active users should be kept substantially lower to guarantee good performance for the low priority data applications.

Voice

The load generated by a single VoIP call on a WLAN is a function of three factors: the packetization interval, whether voice activity detection is enabled, and the type of codec employed.

The packetization interval dictates how many packets per second are generated by a voice call and is by far the most important factor in controlling the WLAN load. Since the number of voice packets is inversely proportional to the packetization interval, it is tempting to use long packetization intervals⁸. Unfortunately, because of its impact on the one-way delay, the packetization interval cannot be too large. In order to mitigate some of the 802.11 overheads, VoIP deployments on 802.11 WLANs often use 30 milliseconds packetization intervals that are longer than the 20 milliseconds intervals habitually used in wired networks.

⁷ Even when silence suppression is enabled, voice traffic is pretty much constant with at least one side of the conversation active at any time.

⁸ The reduced number of packets also increases the battery life of small devices, an important consideration for 802.11 wireless phones.

Because it also affects the number of packets, silence suppression could successfully be used to reduce the load generated by voice traffic. The more aggressive the silence suppression is, the higher the savings (and the more severe the voice clipping). With very aggressive settings, it is possible to reduce by as much as 40 percent the total number of voice packets, provided the resulting degradation in voice quality can be tolerated.

Finally, the choice of codec affects the size of the VoIP packets but not their number and in fact has only a very small influence on the WLAN load. Because of the high encapsulation and 802.11 overheads, a codec that compresses the voice signal has a very small impact on the WLAN capacity. For example, the time required to transmit at 11 Mbps a 30 milliseconds G.711 voice packet on an 802.11b WLAN is 798 microseconds on average. The time required to transmit the equivalent G.729 packet is 645 microseconds. The eight-fold voice compression afforded by the G.729 codec translates into a measly 19 percent reduction in the WLAN load. Such a small saving usually does not warrant the degradation in perceived voice quality caused by the compression artifacts⁹.

In order to plan the WLAN, the anticipated number of simultaneous calls must be estimated. This will depend on a number of factors such as the expected user density and the amount of time users are expected to spend on the phone. For instance, if wireless phones completely replace the wired phones, voice traffic on the WLAN will be much higher than if they merely supplement them.

WLAN planning

Planning a WLAN deployment requires a fair amount of site specific information. Some of this information, such as floor plans, target coverage area, location of Ethernet switches and hubs, and density distribution of the user population (e.g., cubicle square footage in a cube farm) could be obtained in advance from the customer. When it is available, this information can provide a rough estimate of the number of access points required to achieve the coverage and capacity targets of the wireless LAN. Nonetheless a thorough site survey may be required to determine the optimum placement of access points, especially in hostile environments.

Site survey

The physical environment of the planned WLAN must be well characterized. If possible, blueprints or floor plans should be obtained prior to conducting the site survey. If these are not available, floors plans will have to be drawn during the site survey.

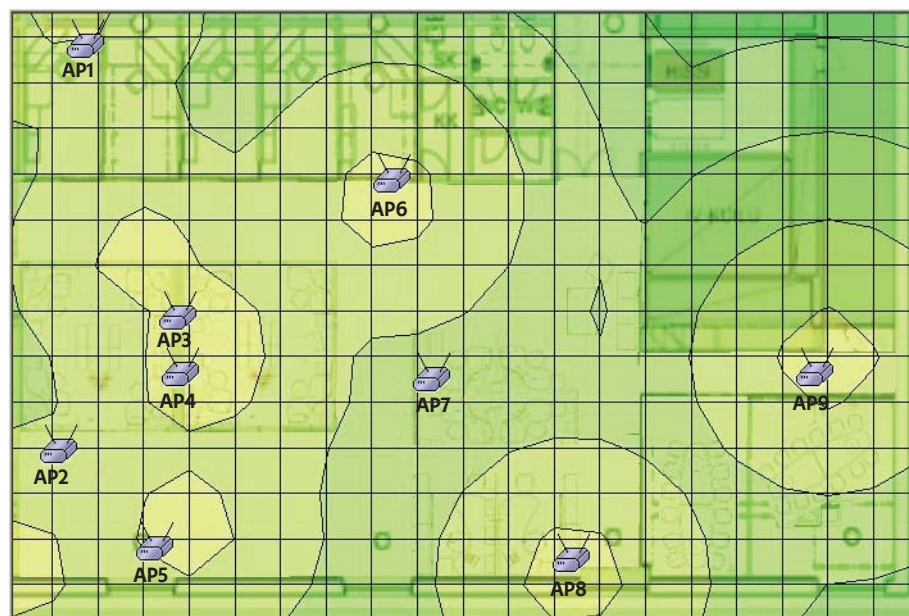


Figure 12. Site survey view

Data about the physical environment (coverage area shape and obstructions) and radio conditions (signal strength and interferences) that is collected during the site survey is typically represented graphically. In the figure above, the color-coded signal strength is overlaid on top of a floor plan showing the access points' locations.

⁹ Some VoIP implementations do not support voice activity detection for the G.711 codec. A motivation for using the G.729 codec in a WLAN deployment might be that VAD is an integral part of the standard.

Coverage area

The size of an area fails to take into consideration its shape. A narrow, elongated area such as, for instance, a hospital wing may require more access points than its surface area would indicate. This is simply because some of the roughly circular coverage of the access points will necessarily fall outside the area of interest. Generally, irregular areas will require more access points than regular ones or external antennae with specific radio propagation patterns (see Figure 13). For example, semi-directional antennae could be used to provide coverage from the side of an atrium and highly directional antennae would be useful down hallways.

Conversely, if the coverage area includes multiple adjacent floors, depending on the type of floor building materials, it may be possible to take advantage of the fact that radio signals penetrate through ceilings to provide coverage between floors. For example, coverage of a three storey building might be achieved by deploying access points only on the first and third floors.

Obstructions

In addition to the shape of the coverage area, other physical characteristics such as the types of building materials used for the floors and supporting walls are required for proper placement of the access points. Details about building materials may be gathered from a blueprint. Care should be taken to mark on the floor plan obstructions such as large concrete or marble columns or metallic objects that might hinder the propagation of radio signals.

Metallic filing cabinets, fire doors, elevator shafts, air ducts, and suspended fluorescent light casings cause reflections of the radio signals resulting in severe disruptions. Water absorbs the radio frequencies used by Wireless LANs. Fish tanks or waterfalls that are often used to decorate atriums or commercial spaces constitute impenetrable barriers to WLAN radio signals.

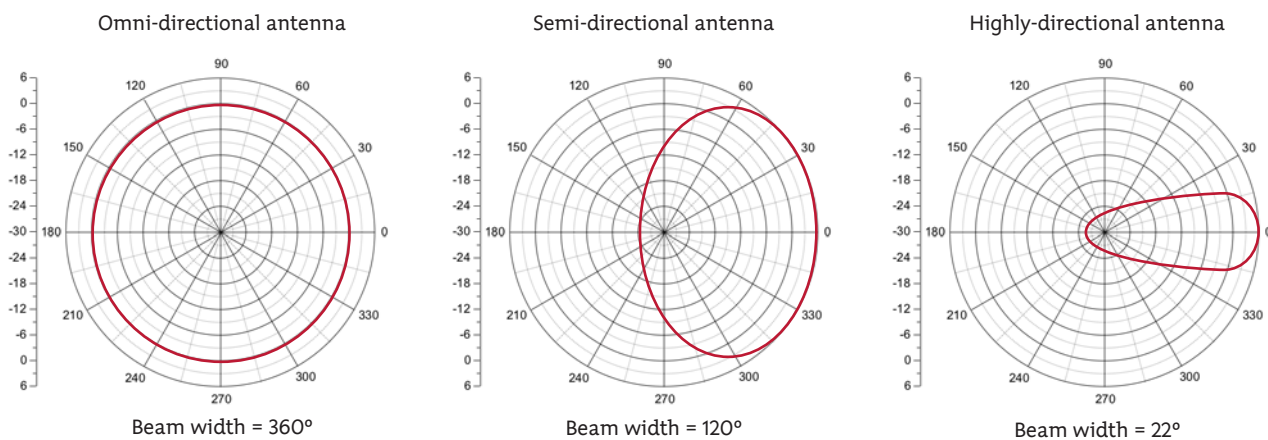


Figure 13. Antenna types and radio patterns

Other impediments to the propagation of WLAN radio signals can be easily overlooked. Energy-efficient glass windows use microscopically thin, virtually invisible, metallic layers to suppress radiative heat flows. To reduce convective heat transfers, argon gas is used instead of air to fill the gap between the sheets of glass. As it turns out, argon gas absorbs radio energy around 2.45 GHz. These energy-efficient windows, although perfectly transparent to visible light, are practically impenetrable to the 802.11b radio signals. Wireless coverage of a patio area might not be possible from an access point inside the building if the two are separated by energy-efficient glass walls.

Interferences

Sources of interferences such as Bluetooth equipment (using the entire 2.4-GHz band), microwave ovens (using frequencies near 802.11b channel 9) and other RF devices (e.g., cordless phones) operating in close proximity to the WLAN must be identified. Traditionally, most of the interference problems have been in the 2.4-GHz band, but, for this very reason, newer devices tend to simultaneously use the 5-GHz band which is less crowded and has many more non-overlapping channels.

In a shared environment such as a multi-tenant building or a mall, care should be taken to uncover other WLANs operating in the vicinity. These can easily be overlooked during a site survey unless a systematic scan of all the frequencies is performed.

Coverage vs. capacity

Simple site surveys, while guaranteeing coverage do not by themselves guarantee capacity or performance targets will be met. Because of the very nature of the shared medium and the dependence of effective throughput on packet sizes, the WLAN traffic characteristics also need to be taken into account in order to guarantee a satisfactory performance for all users and applications.

In larger deployments where channels are reused, the WLAN performance can be degraded by co-channel interference and a simple site survey, while verifying a specific data rate with no interfering traffic, may not take into account the data rate reduction due to the noise rise¹⁰.

Due to the interactions and interferences that are only manifest in a full deployment, further analysis is required in order to guarantee both coverage and capacity. This is where automated real-time planning tools can help refine the number, placement, and configuration of access points.

The WLAN user population, the usage patterns, and the application mix will probably change over time, especially during the early phases of the WLAN deployment.

Detailed activity reports and intelligent management systems are required to monitor the health of the Wireless LAN (e.g., adjust the power levels to minimize interferences and maximize capacity and performance), and automatically identify problem areas before they impact the users' quality of experience.

Conclusions

In a WLAN deployment, performance and capacity are impacted by many factors ranging from the quality of the physical wireless link to the amount and type of traffic generated at the application layer. All these factors must be properly characterized and taken into account in order to meet the requirements of a high-availability, high-performance, and mission-critical environment.

Although a WLAN deployment remains a complex task and requires careful planning, quality tools are now available to facilitate the process. These WLAN planning tools can automatically determine the number, location, and configuration of access points required to meet predetermined capacity and coverage targets.

Based on environment parameters such as coverage area size and shape, and application requirements, engineering guidelines can provide a rough estimate of the number of access points and of the target cell size required to meet the Wireless LAN objectives. Site surveys help refine this initial estimate and identify good placements for the access points especially in challenging environments.

Once deployed, the access points themselves should dynamically assign channels and perform power adjustments to minimize interferences and prevent coverage holes. They should also monitor the use of the radio resources, dynamically balance the traffic load between APs, and identify problem areas before they impact the users' quality of experience.

¹⁰ Some site survey tools are incorporating co-channel interference into the overall performance grid; however, the noise rise is completely dependent on the type and amount of traffic being generated in neighboring areas and may be difficult to predict.

References

- [802.11] IEEE Std 802.11, *Wireless LAN Medium Access Control and Physical Layer Specifications*
- [802.11a] IEEE Std 802.11, *Wireless LAN Medium Access Control and Physical Layer Specifications: High-Speed Physical Layer in the 5 GHz Band*
- [802.11b] IEEE Std 802.11, *Wireless LAN Medium Access Control and Physical Layer Specifications: Higher-Speed Physical Layer Extensions in the 2.4 GHz Band*
- [802.11g] IEEE Std 802.11, *Wireless LAN Medium Access Control and Physical Layer Specifications: Further Higher Data Rate Extension in the 2.4 GHz Band*
- [VoWLAN] *Straight Talk on Converged Wireless LANs*, Nortel Networks white paper.
- [QoE] *Designing QoS-Enabled Networks for Voice & Data User Quality-of-Experience (QoE)*, Nortel Networks white paper.

In the United States:

Nortel Networks
35 Davis Drive, Research Triangle Park, NC 27709 USA

In Canada:

Nortel Networks
8200 Dixie Road, Suite 100, Brampton, Ontario L6T 5P6 Canada

In Caribbean and Latin America:

Nortel Networks
1500 Concorde Terrace, Sunrise, FL 33323 USA

In Europe:

Nortel Networks
Maidenhead Office Park, Westacott Way, Maidenhead Berkshire SL6 3QH UK

In Asia Pacific:

Nortel Networks
Level 5, 495 Victoria Avenue, Chatswood, NSW, 2067, Australia, Phone: (61) 2 8870 5200

In Greater China:

Nortel Networks
Sun Dong An Plaza, 138 Wang Fu Jing Street, Beijing 100006, China, Phone: (86) 10 6528 8877

Nortel Networks is an industry leader and innovator focused on transforming how the world communicates and exchanges information. The company is supplying its service provider and enterprise customers with communications technology and infrastructure to enable value-added IP data, voice and multimedia services spanning Wireless Networks, Wireline Networks, Enterprise Networks, and Optical Networks. As a global company, Nortel Networks does business in more than 150 countries. More information about Nortel Networks can be found on the Web at:

www.nortelnetworks.com

For more information, contact your Nortel Networks representative, or call 1-800-4 NORTEL or 1-800-466-7835 from anywhere in North America.

*Nortel Networks, the Nortel Networks logo, and the globemark design are trademarks of Nortel Networks. All other trademarks are the property of their owners.

Copyright © 2004 Nortel Networks. All rights reserved. Information in this document is subject to change without notice. Nortel Networks assumes no responsibility for any errors that may appear in this document.

