

Understanding Wi-Fi Performance

Robin Layland

Lots of factors go into determining the actual capacity of your wireless LAN.

Much has been said about the security of 802.11 wireless LAN (WLAN), or Wi-Fi. The weakness of the original encryption, Wired Equivalent Privacy (WEP), led to the adoption of Wi-Fi Protected Access (WPA), a standard that brings Wi-Fi up to the same level of security as Internet access.

Less has been said about the true performance or capacity of Wi-Fi. Descriptions of the “maximum” throughput, either 11 Mbps (for 802.11b) or 54 Mbps (for 802.11a/g), are generally accompanied by a disclaimer about how your actual throughput will be less. Network managers and architects need to look seriously at Wi-Fi capacity because it will directly affect the user experience and network design, and it is not as easily fixed as security.

Network managers and architects must understand the reality of Wi-Fi capacity for several important reasons:

- So that the network can be properly designed in the first place.
- Because the Wi-Fi network will more than likely have to be expanded in the future to meet growing user demands.
- Probably most importantly, so users have realistic expectations.

Users have become accustomed to high-speed wired networks. Generally, when response times are slow, it is because of applications, not the network. People have forgotten the days when the network itself could be the problem.

Wi-Fi takes us back to those “bad” old days. How bad is a Wi-Fi network? It is not all doom and gloom—wireless LAN throughput is better than dial-up, which many telecommuters rely upon. Wi-Fi can provide adequate performance for most applications, but may not be the best fit for high-volume applications such as CAD/CAM. Users need to understand that the throughput is more like what they receive at home over DSL or cable modems than what they are used to having at the office.

Most vendors in the Wi-Fi arena will readily admit that actual throughput is usually about 40-60 percent of the “full” bandwidth. For 802.11b this means 4.4-6.6 Mbps, and for 802.11a/g, 21-32 Mbps. And that may be the best case.

It is important to understand the factors that affect capacity and performance in a Wi-Fi network to understand why getting 60 percent is something to be happy about, and to understand what actions you can take to reach that level.

Reasons For Limitations

The limitations of Wi-Fi’s capacity are related to the nature of the underlying protocol and the limitation of the radio frequency (RF) technology.

An RF signal loses power as you move farther from the source. (Like all electrical signals, radio waves lose power as they travel through any type of matter—a little loss through air, a big loss through concrete or metal.) Wi-Fi handles this loss of strength by decreasing the throughput available. The throughput does not decrease linearly but instead follows a defined step function. The speeds in Mbps are:

- 802.11b: 11, 5.5, 2 and 1
- 802.11g: 54, 48, 36, 24, 18, 12, 9, 6, 2 and 1
- 802.11a: 54, 48, 36, 24, 18, 12, 9 and 6

The important capacity question is how far out you can count on a particular speed. This is not as clear-cut as you would think, and finding out the actual distance/throughput figure can be complicated. Several factors affect the throughput you can get at a particular distance, including:

- Power output from your Wi-Fi card. Not all Wi-Fi cards are equal; some are “stronger” than others.
- The power output from the access point. Though this is configurable, it’s set at the maximum allowed level—whenever possible. But interference and environmental factors could prevent you from using this setting.
- Type of access point used. The maximum ranges vary from manufacturer to manufacturer; check with your particular vendors.
- The type of antenna used on the access point.
- The density of the building environment.

There are no standards or even generally-agreed speed and distance points. Different vendors and sources will give different values, which is not very helpful to a capacity planner. Table 1

Robin Layland is president of Layland Consulting. He has more than 25 years’ experience in enterprise networking including technical and management positions at American Express and Travelers.



Slower users hinder faster users

**TABLE 1 WLAN Distance/
Bandwidth Values (ft.)**

| Speed (Mbps) | 802.11b | 801.11a | 802.11g |
|-------------------|---------|---------|---------|
| 1 | 300+ | | 300+ |
| 2 | 250 | | 250 |
| 5.5 (b) / 6 (a/g) | 195 | 200+ | 210 |
| 9 | — | 170 | 185 |
| 11 (b) / 12 (a/g) | 160 | 150 | 170 |
| 18 | — | 125 | 165 |
| 24 | — | 85 | 140 |
| 36 | — | 75 | 115 |
| 48 | — | 50 | 75 |
| 54 | — | 35 | 60 |

shows estimated distance for each speed under good office conditions and with the maximum allowable power level. The distance values are my best estimate derived by combining different sources, and should only be used to give you an idea of the distances. The distances are in feet from the omni-directional access point, so it represents the radius of a circle—the actual cell is twice the size as shown in the table.

The good news on the distance front is that the distance will only get higher as time goes by, because the radio chip manufacturers are improving their products.

A Few Caveats

The bandwidth/distance bands assume transmission in an open room, i.e., nothing that interferes with the RF signal—and that happens only in an ideal world. The reality is that many common objects in a work environment will interfere with the RF signal and reduce its strength, thus reducing the distance for each band.

For example, metal objects such as filing cabinets can degrade the RF signal, and the metal studs used for office construction can interfere with radio waves. Sheetrock and cubicles block the signal; and concrete, cinderblock and metal studs can reduce the signal by 90 percent.

And it's not just objects that can degrade throughput. Microwave ovens are especially bad for 802.11b/g, since microwaves use the same frequency as these radios. A poorly manufactured or inadequately shielded microwave can leak radio waves into the environment, interfering with WiFi. Bluetooth devices and many cordless phones also use the same frequency range as 802.11b/g, creating more interference, and some newer cordless phones interfere with 802.11a's frequency range as well. A general rule is that more devices interfere with 802.11b/g than 802.11a.

Objects can also lower throughput because they cause radio waves to bounce off them. This can create multiple copies of the message—what's called multi-path. Wi-Fi equipment is built to deal with multi-path, but this problem can be bad enough to interfere and cause lower throughput.

Distance also affects throughput in another way. Since users will be at different distances from the access point, their throughput rates will be different. A user who is close to the access point will transmit and receive nearer to the maximum of 11 Mbps or 54 Mbps, while another user at the edge of the range might see only 2 Mbps.

An 802.11b example shows how this works, and why it can be a problem. Say you have two users each receiving 1,000-character messages. One user is close to the access point and receives 11 Mbps while the other is at the outer range and receives only 2 Mbps. The closer user will take only .00073 seconds to receive the message while the farther user will take .04 seconds.

That doesn't sound like much of a difference, but here's why it matters: When the slower user is sending or receiving, the faster users have to wait to send their next messages, until that slower user's message is sent. This constrains how much traffic the faster users can send and receive.

This is why the total throughput for the cell is expressed as an average of all users' throughputs. For example, if your users are equally spaced out over the bandwidth/distance bands (i.e., evenly spread throughout the cell), your total capacity is only 6.2 Mbps for 802.11b and 19.6 Mbps for 802.11a/g. So the average throughput will depend on the average distance of the users from the access point—something that is constantly changing. In fact, the situation is even more complicated, because if the closer users are sending and receiving smaller messages than the farther-out users, capacity is further reduced.

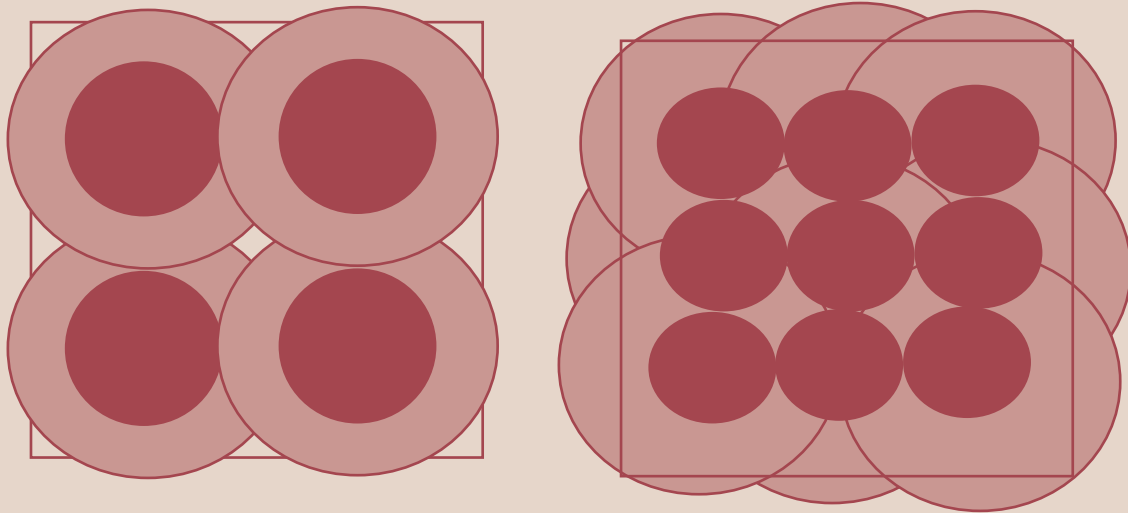
Availability Constraints

The Wi-Fi protocol also affects the capacity available. Wi-Fi uses a contention protocol, called Carrier Sense Multiple Access with Collision Avoidance. Here's how it works: The sender first determines if the Wi-Fi LAN is being used, and if it's not, that device can send its message. It is possible for two senders to both think the network is available and start to send their message at the same time. When this happens, both messages are corrupted and have to be resent, wasting the bandwidth and further reducing the overall throughput. This is the same situation Ethernet faced before switching (remember Ethernet's CSMA/CD?).

The access point can hear everyone, but if two PCs are on opposite sides of the cell, their signals may not reach one another, and both attempt to use the WLAN simultaneously. This is called the "hidden nodes" problem.

When the WiFi LAN is lightly used, collisions are not a big issue, but as utilization increases, it can become a problem. Thus, just when you need throughput the most, the actual capacity is reduced. Additionally, the protocol, by listening before sending, reduces the maximum throughput, since the time it takes to send a message is both the listening period and the transmit time.

FIGURE 1 High Capacity Coverage



Furthermore, while 802.11g is backwards-compatible to support 802.11b devices, such a connection will lower the overall throughput of the 802.11g WLAN, for the same reason that farther-away clients slow down clients nearer to the access point, as described above.

Good News/Bad News

The bad news in all of this discussion is that you should plan on getting only 40 percent of the rated throughput on your Wi-Fi network. If you have 20 users on an 802.11b network, that means you can only plan on each user having 220 kbps; for the same number of users, an 802.11a/g LAN might provide 1 Mbps per user—optimistically.

The good news is twofold. First of all, for most applications, these levels of throughput are adequate. Furthermore, you can take actions to ensure you get the maximum available capacity—perhaps even more than our benchmark 40 percent of the rated capacity.

The first step to ensure you get the maximum capacity is to set up the Wi-Fi network with good RF coverage. That means understanding your environment. The best way to do that is to test it by installing a few access points and then measuring signal strength at various points around the office. Free software is available for Wi-Fi equipped PCs, or test equipment from manufacturers such as Fluke can measure. A reminder for “greenfield” sites: this test must be done *after* you move into the office, since the equipment and walls can greatly affect the measurements.

The signal strength information will then help you to determine how many access points you need and where to put them. Wi-Fi vendors such as Airespace, Trapeze, Extreme and others provide RF planning tools that can automate this process. The tools allow you to input a CAD diagram, generally the same one you have already developed for your wiring design, into their tools.

The next step is to input the measurement data to determine number and location of access points. This will give you a good initial design, but don’t expect it to be perfect. After you have installed the access points, go out and measure the signal strength to make sure the model matches reality. You also need to redo the process any time there are major structural changes in the office.

You may need to consider using special antennas instead of the standard ones that come with most access points. The standard antennas cover a circular pattern. This is good for most open office environments, but doesn’t help with odd-shaped rooms or hallways. Special-purpose antennas are available that allow for different coverage patterns. For example, you can get an antenna that has a narrow pattern, to cover a hallway.

The most important capacity variable in the RF design process is cell size. The larger the coverage area, the fewer access points and thus the lower cost—but this comes at the loss of capacity. Figure 1 shows the trade-off. The diagram shows a simple square office building. The first diagram shows coverage of the floor using only four access points; the high-capacity areas are shown in dark red. While the entire floor is covered, a significant part of the floor receives lesser capacity. If the goal is coverage and not high capacity, this design works.

If the goal is to provide high bandwidth to most users, four access points do not suffice. More access points are needed so that the “red” area of high capacity covers most of the floor. Thus, the right part of the diagram shows a design with nine access points, in which the red high capacity area covers most of the floor.

Avoiding Interference

Unfortunately it is not as simple as just adding more access points until the entire floor is covered by the high-capacity areas. Each access point

generates RF signal. The RF signals from one access point can interfere with those from another access point, reducing capacity. In Figure 1, for example, there was some overlap on the left, but a lot of overlap on the right. The example also assumed that this was a one-story building and you did not have to worry about interference from the floors above and below. In real environments, the radio wave travels through the floors, making the problem even more complicated.

The Wi-Fi protocols attempt to overcome the problem of interference by having the different access points operate on different channels. For example, 802.11b/g uses three channels and 802.11a uses 8–12 channels, depending on whether the WLAN is within a building or outside, and where in the world you are using it (government licensing and regulations vary). The channels divide the available radio frequencies into different bands such that if the access points next to each other are operating on different channels, they don't interfere with each other.


In the example with four access points, it is easy to design a scheme that provides coverage with minimum interference with just the three 802.11b/g channels, as shown in Figure 2. However, the task becomes much harder, if not impossible, with just three channels in the example that uses nine access points. Thus, 802.11a, with its eight channels, is a better choice for high-capacity systems, since it is easier to assign different channels to adjacent access points.

When channels overlap, the protocol has to choose which access point the device should associate with. This has been solved by adding intelligence to the access point or its controller. The access points communicate among themselves and determine which has the best connection.

Note that “best” does not just mean which access point has the strongest signal, but can include other factors, such as which access point is currently supporting the fewest clients.

The 802.11 implementation makes these calculations automatically, but you can affect this. It is possible to tell the client to associate only with an access point that provides a minimum capacity. For example, with 802.11b, you can set the client to only associate with access points if the 5.5-GHz band is available.

Just as important as understanding your potential throughput is knowing what is actually happening on your Wi-Fi network. This can be accomplished by monitoring and measuring via the SMNP MIB supported by every enterprise-level Wi-Fi device. Additional understanding can be gained by adding probes to the network from vendors like AirMagnet and ReefEdge or from specialized Wi-Fi management software from companies such as Wavelink. Probes are passive access points that monitor the Wi-Fi network—they can be an important addition to a Wi-Fi network, but they do add extra cost.



You may need to add access points—without adding interference

Two of the most important variables to manage are the number of collisions and level of interference. As traffic increases, the number of collisions indicates how efficiently the network is operating. There will always be some collisions due to randomness. However, as the number of collisions increases, there will be a point at which throughput decreases due to the time the retransmissions take. The exact number that is acceptable depends on several variables, including number of users and average message size. With experience it is possible to determine the unacceptable number.

Interference tells you if other devices are causing problems, decreasing the maximum throughput. The solution is generally to decrease the cell size—fewer people per access point—or eliminate any non-Wi-Fi devices that are interfering.

Conclusion

What is the best strategy for a high-capacity WiFi network? Use 802.11a in a dense access point design. Next would be to use pure 802.11g in a less dense design. Note that 802.11b can have a role in both of these designs. Access points are increasingly supporting all three standards.

Overall, Wi-Fi can provide more than adequate capacity for most situations, but you have to stay on top of it. It is more like the very early LAN that requires more time in planning and managing □

Companies Mentioned In This Article

Airespace (www.airespace.com)
 AirMagnet (www.airmagnet.com)
 Extreme Networks
 (www.extremenetworks.com)
 Fluke (www.fluke.com)
 ReefEdge (www.reefedge.com)
 Trapeze Networks
 (www.trapezenetworks.com)
 Wavelink (www.wavelink.com)