



Global Knowledge™  
Experts Teaching Experts

Expert Reference Series

# Which WLAN Technology is Best?

## Which WLAN Technology is Best? (Your Throughput May Vary!)

By Rick Murphy

### Introduction

In June of 2003 the IEEE 802.11 committee ratified the “g” version of their wireless LAN family of technologies. This new standard builds onto the existing 802.11 WLAN protocol, which was originally released in 1997. In addition, it takes advantage of the physical layer techniques designed for 802.11a to provide much higher data rates. The excitement over 802.11g is because it offers the greatly improved speeds of 802.11a but is backwards compatible with existing 802.11b networks. In other words, you can phase-in the upgrade to “g” on an existing 802.11b network.

Over the past months, the industry press—and even some manufacturers—have made claims concerning the superiority of 802.11g over the other 802.11 family members. In some cases, oversimplification and misconception have skewed these assertions. The concern is that these “assertions” are being taken for granted as truth and are becoming “common knowledge.” Some of these assumptions are:

- (1) 802.11g has replaced 802.11a and has made 802.11a obsolete.
- (2) 802.11g is fully backwards compatible with 802.11b.
- (3) 802.11g provides better coverage than 802.11a due to the frequency characteristics of 2.4 GHz and 5.0 GHz.
- (4) 802.11g provides the same high speed as 802.11a.
- (5) 802.11g has better security mechanisms than the other 802.11 family members.
- (6) 802.11g is less expensive than 802.11a.

In order to test these assumptions, the vLogic engineering team recently conducted a series of tests in our lab environment. We tested 802.11b, 802.11a, and 802.11g side-by-side under several scenarios. The results were surprising at first, but after further consideration and some additional research, they made perfect sense.

### Baselining the vLogic Networking Lab

“Baselining” is the process of establishing a benchmark to use as a reference point during comparative studies. In our case, the vLogic 100BT 802.3 network lab was used as the reference network for comparison during the various 802.11 WLAN studies. In order to have consistent test results, it is necessary to use the same application throughout all of the studies. Our choice for performance testing software for all tests was Qcheck from Ixia (formerly NetIQ). Qcheck is an application that allows for several types of performance tests (response time, throughput, streaming, and traceroute) using up to four different protocols (TCP, UDP, SPX, or IPX). The software is configured and managed from a console that issues the test setup instructions to a pair of endpoints. These endpoints execute the tests and return the results to the console to be viewed by the administrator. Qcheck was used rather than “ping” or simple file transfers in order to set a reliable baseline that could later be used on many different types of network access devices and scenarios. Qcheck uses the same protocols and APIs that real-world applications use, as opposed to ping, which uses only ICMP. In addition, Qcheck can simulate streaming video and audio traffic and show the number of lost packets over the link at various stream rates.

Qcheck works by communicating between at least two network-attached devices using the pre-selected utilities and protocols mentioned above. In order to prepare to use Qcheck, you first install a main console utility on one of the devices and then install an endpoint utility on the other devices that you wish to check. NetIQ currently has endpoint software for 20 different operating systems, including Windows CE (Pocket PC 2002). Ixia offers all of these utilities for free at their download site:

<http://www.ixiacom.com/enterprise/Qcheck.php>



To begin the study, Qcheck was installed on a Windows 2003 server connected to the 100BT Ethernet segment by a 3Com 3C905 full duplex 100BaseTx NIC. Then a test machine (in this case a Dell Inspiron 8200 with integrated 100BaseT Ethernet adapter) was set up with the Qcheck endpoint software. The endpoint software is very small and runs as a background system service. While idle, the endpoint consumes very few resources as it waits for test instructions from the console.

To begin the first test, we set the field “From Endpoint 1” to “localhost” and the “To Endpoint 2” field to the IP address of the remote station (in this case 192.168.10.10). You can find the address information easily on the remote by typing “ipconfig” from a system prompt.

The next step was to select a test protocol. We clicked on the “TCP” button and the “Response Time” button to set the test choices. The default parameters of “3 Iterations” and “data Size 100 Bytes” were left as they were; we then clicked the “Run” button. Down in the “Results” window, the message “testing” displayed for few seconds, and then that message was replaced by the actual results of the test. In this case the “Response Time Results” were “Minimum = 1 ms,” “Average = 1 ms,” and “Maximum = 1 ms.” I interpreted this as 1 millisecond for the response time.

To begin the second test, “Throughput,” all station address settings were left the same as the previous test. Then with the TCP button selected, we clicked the “throughput” button. The default settings for “Throughput” are a bit low for testing on a 100BaseT network, so we changed the “data Size” from its default of 100 bytes to 1000 bytes. Then we clicked “Run” and waited for the “testing” message to be replaced by the “Results.” In this case, our wired connection was delivering 94.118 Mbps of actual data throughput. This test was performed on a separate lab network that had only ten workstations connected to the segment. Of these ten, only two were inserting user data traffic onto the LAN. All other stations were unattended and would have only been inserting signaling information onto the LAN. In addition, all privacy and encryption protocols were disabled on the LAN segment. (NOTE: Be sure all internal firewalls are disabled before running Qcheck tests.)

After performing the initial tests, our results established a baseline maximum for throughput at 94.118 Mbps and a best case for response time of 1 millisecond.

**Wireless Baseline:**




Figure 1: vLogic WLAN Test Results Chart

The focus for the wireless tests was on actual data throughput using several different WiFi Access Points and station adapters. In all cases, response time and streaming were found to be similar to the wired network results. The effects of using different 2.4 GHz channels, as well as two levels of encryption (40-bit and 104-bit WEP), were also tested separately and found to have no significant effect on the test network.

For the main wireless throughput comparisons we chose six test points throughout the office premises. Three indoor test points and three outdoor test points were selected to represent typical impediments to WLAN throughput, such as interior and exterior walls, floors, and various distances from the access points. The construction materials used within the vLogic labs and offices are standard sheet rock walls with wooden studs on 16” centers and plywood subfloors with carpeting. Following is a description of the six test points.

**Indoor Test Points:**

Test Point #1:  
 Clear line of sight to AP, 15 ft. from AP.  
 Signal: -45 dBm.

Test Point #2:  
 From one floor below, 1 wooden floor between AP, 20 ft. from AP.  
 Signal: -57 dBm,

Test Point #3:  
 Elevated office, 1 interior wall between AP, 35 ft. from AP.  
 Signal: -64 dBm.

**Outdoor Test Points:**

Test Point #4:  
 Seating area, 1 exterior wall between AP, 40 ft. from AP.  
 Signal: -68 dBm.

Test Point #5:  
 Farthest back corner of premises, 1 exterior wall between AP, 110 ft. from AP.  
 Signal: -73 dBm.

Test Point #6:  
 Farthest front corner of premises, 2 interior and 2 exterior walls between AP, 82 ft. from AP.  
 Signal: -80 dBm.

### Test Results

(The following results were chosen from a “best of five” set.)

AP Type	Station Type	Test Point #1	Test Point #2	Test Point #3	Test Point #4	Test Point #5	Test Point #6
Cisco 802.11b	Dell 802.11b	5.483 Mbps	5.513 Mbps	5.413 Mbps	4.360 Mbps	5.319 Mbps	4.651 Mbps
Cisco 802.11b	Linksys 802.11g	6.390 Mbps	6.410 Mbps	6.405 Mbps	6.436 Mbps	5.857 Mbps	5.517 Mbps
Linksys 802.11g	Linksys 802.11g	16.601 Mbps	16.375 Mbps	10.192 Mbps	10.687 Mbps	4.273 Mbps	1.376 Mbps
Linksys 802.11a	Netgear 802.11a	19.746 Mbps	20.531 Mbps	17.662 Mbps	4.908 Mbps	4.386 Mbps	6.478 Mbps

Figure 2: vLogic WLAN Test “Best of Five” Results

### Reviewing the results

The most surprising thing about these results was that the 802.11a results at the farthest test points were comparable with the results shown by the 802.11g tests. One of the most often cited weaknesses of 802.11a is that it does not propagate as far as 802.11b or 802.11g, which both operate in the 2.4 GHz ISM band. This is due to the fact that 802.11a operates in the higher 5 GHz frequency. It is a well-known property of radio frequency physics that higher frequencies attenuate more quickly than lower frequencies. This is true if all other components of the signal are identical, such as output power, cable and connector losses, and antenna gain.

There are two items of interest that must be taken into consideration when comparing 802.11g and 802.11a coverage areas. First, the difference in the wavelength between the 2.4 GHz carrier waves and the 5.0 GHz carrier waves is not sufficient to show a significant difference in free space path loss within the tested coverage areas. For example, the FSL (Free Space Loss) of a carrier wave propagating on 802.11b/802.11g channel 6 (centered at 2.437 GHz) over a 100 ft. distance is approximately -98.79 dB. This compares to a carrier wave propagating on 802.11a channel 52 (centered at 5.260 GHz), which shows -104.47 dB of FSL over the same 100 ft. distance from the AP. This difference of approximately -6 dB is noticeable but is not extreme.

The second item of interest in this comparison is that the output power of the APs is not necessarily equal. The manufacturers of the 802.11b and 802.11g access points that were used in the tests have limited the maximum output signal power of the equipment to 100 mW of output signal power. But with the 802.11a access point, the maximum output power at the AP is relative to the three bands of frequencies allowed for use in the US. The 802.11a Linksys AP that was used in the above tests defaults to channel 52 upon initialization, and that is the channel that was used during the tests. Channel 52 resides in the middle UNII (Unlicensed National Information Infrastructure) band, which is regulated by the FCC to have a maximum output power of 200 mW. By default, the 802.11a Linksys AP is set to the “Full” power setting on channel 52, which equates to twice the output power as the 2.4 GHz APs.

Another surprising result of the vLogic tests was the major discrepancy in advertised data rate (signaling rate) and the actual user throughput. 802.11a and 802.11g products have an advertised data rate of 54 Mbps, but the actual data throughput when using these products is much less than the advertised rate. During the vLogic tests we found that if the user terminal was located within a few feet (1-3 ft.) of the access point being tested, with no obstructions and no other wireless devices in the vicinity providing

interference, then we could get approximately 50% of the advertised data throughput. However, because it was felt that this scenario was not realistic or representative of a typical network, the original test point data is not included in the result set.

### Understanding the results

Generally speaking, the reason for such a difference between the advertised data rate and the actual user throughput is due to overhead. Because of the hostile nature of the air interface (compared to guided media such as fiber and metallic), it was necessary for the 802.11 design committee to build several methods into the 802.11 protocol that are used to insure the fair use of the channel and to provide a guarantee of delivery of user data across the air medium.

### CSMA/CA

When there are many stations on the same shared network medium (whether that medium is wired, cable or air), there must be a way to assure fair access to that medium by all. In all IEEE 802 defined networks, this is the responsibility of the rules defined by the MAC (Media Access Control) sub layer of the Data Link layer (Layer 2 of the OSI reference model). The method used to provide this fair access in 802.11 WLANs is known as CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance).

The CSMA/CA method is similar to the traditional MAC layer technique used in 802.3 (Ethernet) LANs known as CSMA/CD (Carrier Sense Multiple Access with Collision Detection), except that in wireless there is no way for the participating stations to know if a collision has occurred (Collision Detection) somewhere out in the air medium, so the stations have to assume that a collision *always* occurs unless they are specifically notified that it did not. This notification takes the form of a positive acknowledgement message (ACK), which is always sent by the receiving station in response to each reception of unicast wireless data frames. If the transmitting station does not receive an ACK within a specific amount of time after sending the data frame, it assumes the frame was lost due to a collision, and it retransmits the data. This overhead, plus additional overhead caused by other factors such as station synchronization and pre-approval of transmissions for enhanced collision avoidance, reduces the actual throughput of an 802.11 WLAN channel to approximately 50% of its advertised data rate. This rate is consistent with and similar to wired, non-switched Ethernet networks, which offer actual throughputs of between 30% and 70% of their actual channel bandwidth, depending on the amount of congestion on the network.

In fact, a recent clinical study performed by Atheros, a manufacturer of 802.11 chipsets, and published under the title of “802.11 Wireless LAN Performance” shows the advertised data rates and the maximum “theoretical” data rates to be very different (see Figure 3).

WLAN Type	Modulation	Maximum Link Rate	Maximum Theoretical TCP Rate	Maximum Theoretical UDP Rate
802.11b	CCK	11 Mbps	5.9 Mbps	7.1 Mbps
802.11g (with 802.11b)	OFDM/CCK	54 Mbps	14.4 Mbps	19.5 Mbps
802.11g (11g-only mode)	OFDM/CCK	54 Mbps	24.4 Mbps	30.5 Mbps
802.11a	OFDM	54 Mbps	24.4 Mbps	30.5 Mbps

Figure 3: Source: Atheros Communications

These figures were arrived at by assuming “1500 byte frames, encryption enabled, default 802.11 MAC configurations, zero packet errors,” and operating at a close range to the access point. Notice that the chart shows two different throughput rates for 802.11g, because when 802.11b operates in the same area as 802.11g, then 802.11g suffers additional overhead. This overhead is the result of protection mechanisms in the protocol that are designed to allow the two different networks to operate together.

## **802.11b and 802.11g together**

If 802.11b and 802.11g networks are to be used together in the same operating space (this is defined as the effective radio reception area generated around an access point), then it is important to understand the three different ways these networks can be set up. 802.11b/802.11g network areas can be setup as “overlapping networks,” as “mixed-mode networks,” or as “exclusive networks.”

### **Overlapping networks**

In the overlapping network design there are two separate networks: an 802.11b access point and its associated stations are operating within range and on the same channel as an 802.11g access point and its associated wireless stations. Under ideal circumstances the 802.11b and 802.11g devices would not transmit at the same time but would defer their transmissions until the medium was unused to avoid excessive collisions. 802.11b devices take longer to transmit the same amount of data than do 802.11g devices. As an example (which does not include normal protocol overhead), an 802.11g device operating at 54 Mbps could transmit 1500 bytes of data in 222 uSec. This compares with an 802.11b device operating at 11 Mbps, which would require 1091 uSec to transmit the same 1500 byte packet. This would require all 802.11g devices to defer for the full 802.11b transmission time before attempting to use the medium themselves. Because of this, the 802.11g users could expect to see a reduction in their throughput of up to 83% due to the presence of nearby 802.11b devices.

### **Mixed-mode networks**

In a case where a single access point has been configured to support both 802.11g and 802.11b stations, the situation becomes even more complicated. When operating in “compatibility-mode,” the 802.11g access point and all 802.11g stations must downshift to using the same slot times and preamble lengths that the 802.11b stations are capable of. This is in addition to the deferment time outlined above. This effect can account for another 20% reduction in actual throughput for the 802.11g users. Another detrimental effect of using 802.11g and 802.11b within the same radio area is that, since 802.11g uses OFDM and 802.11b uses CCK, the transmitting devices may not be able to understand each other’s control messages properly. This could result in completely uncoordinated transmissions within the same coverage area, which might result in an unacceptable amount of collisions, resulting in a non-performing WLAN. To counteract this problem, a protection mechanism known as RTS/CTS, is used for mixed-mode 802.11g/802.11b networks. In fact, there are two protection methods that may be used to insure the simultaneous operation of 802.11b and 802.11g; these are referred to as RTS/CTS-mode and CTS-only mode. The CTS-only mode is a bit more efficient and is the mode shown in Figure 3 above. If the network must resort to RTS/CTS-mode instead, then the maximum theoretical TCP rate is reduced even further to about 11.8 Mbps.

### **Exclusive-mode networks**

The reason that 802.11g is affected so greatly by 802.11b is because they both operate in the same 2.4 GHz ISM band. In order to get the most throughput from an 802.11g network, the use of exclusive-mode is required. In this mode, the network does not allow 802.11b devices to use the resources of the 802.11g network. In addition to setting the 802.11g access points for 802.11g devices only, it would be necessary to set the 802.11g access point channel to one of the three non-overlapping channels not in use by a nearby 802.11b access point. In this scenario, the 802.11g users would be able to experience actual data throughput equal to users of 802.11a networks. Unfortunately, the limitation of having only three non-overlapping channels to use for all 2.4 GHz networks severely restricts the size and user density of both 802.11b and 802.11g networks.

### **Assertions put to the test**

We can now approach the assertions mentioned at the beginning of this article one by one to verify their accuracy.

*Assertion #1: 802.11g has replaced 802.11a and has made 802.11a obsolete.*

A case could be made that 802.11g is a better choice than 802.11b, and since 802.11g can be phased into 802.11b infrastructures (although at a high cost to performance), it can be considered as a viable replacement to 802.11b. However, 802.11g is not a replacement to 802.11a. 802.11a allows unaffected usage within the same areas where 802.11b and 802.11g are operating, with no degradation in throughput caused by the other network devices. This is because 802.11a operates in a completely separate radio band (5 GHz-UNII band), which is unaffected by other current WLAN technologies. In addition, 802.11a's 12 non-overlapping channels allow for the creation of extremely large WLANs (using roaming) and extremely dense (in number of users) WLANs when compared with 802.11g or 802.11b.

*Assertion #2: 802.11g is fully backwards compatible with 802.11b.*

While this assertion is true on the surface, in actuality, the usage of overlapping 802.11b/802.11g networks or the operation of mixed-mode networks severely degrades the throughput of the 802.11g users. 802.11g can provide either backwards compatibility with 802.11b or high-speed performance, but not both at the same time.

*Assertion #3: 802.11g provides better coverage than 802.11a due to the frequency characteristics of 2.4 GHz and 5.0 GHz.*

Again, this may be true on the surface, but in practice the actual difference in free space path loss due to the difference in frequency is not appreciable. In addition, 802.11a devices can be set to use the additional UNII band channels, which allow for more powerful signal output than what is allowed to be used with 802.11g/802.11b devices.

*Assertion #4: 802.11g provides the same high speed as 802.11a.*

This assertion is true, but only when 802.11g access points are set to "802.11g-only" modes and then only when they are operating with no interference from neighboring access points using the same frequency ranges.

*Assertion #5: 802.11g has better security mechanisms than the other 802.11 family members.*

This assertion is untrue. Security mechanisms in 802.11g are identical to security mechanisms built into all members of the 802.11 WLAN family of protocols. However, an additional 802.11 task group known as 802.11i has been hard at work creating a better security mechanism for all 802.11 protocols. This new security addendum is divided into three parts. One of these parts ("WPA" or WiFi Protected Access) was made available at about the same time as the 802.11g protocol was standardized and was included for many of the 802.11g devices that were released after June 1, 2003. WPA should also become available for all other WiFi branded devices, but this is at the discretion of the manufacturers.

*Assertion #6: 802.11g is less expensive than 802.11a.*

This assertion is true. However, the difference in price is not significant. In many, cases it is possible to purchase 802.11g access points with an additional 802.11a access point incorporated within one chassis for under \$100 more than purchasing the 802.11g access point alone. 802.11a station adapters currently sell for well under \$100 each.

### **So which one is best?**

Unfortunately, there is no "one size fits all." The answer to this question will depend on what the intention of the network designer is.

### *Advantages of 802.11b*

Although there are still a lot of the original 802.11 WLAN products (2 Mbps) in use today, the majority of deployed 802.11-based networks are of the 802.11b variety. The price of these devices has dropped so much that many terminal device manufacturers, such as laptop, PDA, barcode reader, and cell phone

vendors, are integrating 802.11b chipsets directly into their devices. For this reason, support of 802.11b is required in all public wireless locations, such as WiFi hotspots and conference areas.

#### *Advantages of 802.11g*

Because 802.11g is backwards compatible with 802.11b, it is a logical upgrade for public venues that are now using 802.11b. If user expectations are kept reasonable, 802.11g can be seen as providing a slight to moderate speed increase, while allowing a mixed user base of 802.11b and 802.11g services. In an existing network where the user density and the location of the access points is to remain the same, then 802.11g can be seen as a drop-in replacement for 802.11b access points that will give exactly the same coverage patterns as the 802.11b access point coverage areas. In addition, the 2.4 GHz band and the power outputs specified for 802.11b/802.11g have more worldwide acceptability, which may make these networks a better choice for global deployments. Another advantage of 802.11g over 802.11a is that the current state of PDAs cannot make use of the 802.11a station adapters due to bus limitations on the PDAs. If the use of PDAs with WLAN capability is a requirement, then for now 802.11b/802.11g technology is a better choice.

#### *Advantages of 802.11a*

802.11a moves the WLAN out of the congested (due to popularity) 2.4 GHz ISM band in favor of the clear 5.0 GHz UNII band. This unlicensed band offers three sub bands called the lower, middle, and upper UNII bands. Each of these bands has 4 non-overlapping channels with unique power level settings that make it applicable for selected applications. The lower and middle bands can be used indoors, while the upper and middle bands can be used outdoors. New, additional, non-overlapping channels in the 5.0 to 6.0 GHz range are expected to be offered soon by the FCC, as well. 802.11a can operate within the same areas as 802.11b and 802.11g with no degradation from the simultaneous operations. Because 802.11a is less common, there is some additional privacy in using this technology over the more popular 802.11b and 802.11g systems.

### **Summary**

There is no “silver bullet” solution in WLAN technology today. The wireless network engineer should choose a WLAN technology based on the strengths or weaknesses of the standard and the design requirements of the network.

© 2003 vLogic, Inc. - All Rights Reserved

### **About the Author**

Rick Murphy is President of vLogic, Inc., and CEO of vTown WiFi, as well as a noted author and speaker on wireless networking.

### **Sources:**

“802.11 Wireless Networks: The Definitive Guide.” Matthew S. Gast, O’Reilly.

“Certified Wireless Network Administrator (CWNA) Study Guide.” Planet3 Wireless, Osborne/McGraw-Hill.

“A Detailed Examination of the Environmental and Protocol Parameters That Affect 802.11G Network Performance.” Proxim white paper,

[http://www.proxim.com/learn/library/whitepapers/parameters\\_802.11g\\_performance.pdf](http://www.proxim.com/learn/library/whitepapers/parameters_802.11g_performance.pdf).

“802.11 Wireless LAN Performance.” Atheros white paper,

[http://www.atheros.com/pt/atheros\\_range\\_whitepaper.pdf](http://www.atheros.com/pt/atheros_range_whitepaper.pdf).