White Paper

Capacity, Coverage, and Deployment Considerations for IEEE 802.11g

Few technologies have received as much anticipation as 802.11g, the IEEE standard for wireless local area networking that operates in the 2.4 GHz band with a top data rate of 54 Mbps. It is easy to understand why—802.11g provides performance that can be comparable to that of the 802.11a WLAN standard that operates in the 5-GHz band, while providing backward compatibility with the legacy 11-Mbps 802.11b standard. This combination of higher performance and backward compatibility is similar in concept to the wildly successful and now-ubiquitous 100-Mbps Fast Ethernet standard from the wired LAN world. Network professionals naturally expect similar ubiquity for products based on the 802.11g standard, sometimes to the exclusion of 802.11b- and 802.11a-based products. This document outlines the capabilities of 802.11g, highlighting both its advantages and its deficiencies. The information provided will enable network professionals to achieve an understanding of the performance and capacity capabilities of 802.11g, range and coverage area expectations, and other information that will help to result in a successful 802.11g deployment.

802.11g in Brief

The IEEE 802.11g WLAN standard can be thought of as an intersection between the 802.11b and 802.11a standards. Like 802.11b, 802.11g operates in the same 2.4-GHz portion of the radio frequency spectrum that allows for license-free operation on a nearly worldwide basis. 802.11g is also limited to the same three nonoverlapping channels as 802.11b. An important mandatory requirement of 802.11g is full backward compatibility with 802.11b, which both provides investment protection for the installed base of 802.11b clients and extracts a substantial performance penalty when operating in this mode.

Like 802.11a, 802.11g uses Orthogonal Frequency Division Multiplexing (OFDM) for transmitting data. OFDM is a more efficient means of transmission than Direct Sequence Spread Spectrum (DSSS) transmission, which is used by 802.11b. When coupled with various modulation types, 802.11g (like 802.11a) is capable of supporting much higher data rates than 802.11b. As noted in Table 1, 802.11g uses a combination of OFDM and DSSS transmission to support a large set of data rates—in fact, all of the data rates supported by both 802.11a and 802.11b.

Table 1 802.11g Data Rates, Transmission Types, and Modulation Schemes

Data Rate (Mbps)	Transmission Type	Modulation Scheme
54	OFDM	64 QAM
48	OFDM	64 QAM
36	OFDM	16 QAM
24	OFDM	16 QAM
18	OFDM	QPSK1
12	OFDM	QPSK
11	DSSS	CCK2
9	OFDM	BPSK3
6	OFDM	BPSK
5.5	DSSS	CCK
2	DSSS	QPSK
1	DSSS	BPSK

QPSK quadrature amplitude modulation

CCK complementary code keying

BPSK biphase shift keying

802.11g Performance and Capacity

Within the context of WLANs, network capacity is roughly the product of throughput multiplied by the number of available channels. As stated above, like 802.11b, 802.11g devices are limited to no more than three nonoverlapping channels. The result is that any increase in network capacity relative to 802.11b will result from increases in throughput alone. The throughput provided by 802.11g networks is dependent upon a number of environmental and application factors, chief amongst them being whether or not the 802.11g network is supporting legacy 802.11b clients.

802.11 networks use Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA), a media access method similar to that of shared Ethernet. Also, 802.11b devices, which share the same 2.4 GHz band as 802.11g, have no means of detecting OFDM transmissions. Although 802.11b devices can sense “noise” in the 2.4-GHz band via their Clear Channel Assessment (CCA) capabilities, they cannot decode any data, management, or control packets sent via OFDM. Given this, the 802.11g standard includes protection mechanisms to provide for coexistence and backward compatibility.

When 802.11b clients are associated to an 802.11g access point, the access point will turn on a protection mechanism called Request to Send/Clear to Send (RTS/CTS). Originally a mechanism for addressing the “hidden node problem” (a condition where two clients can maintain a link to an access point but, due to distance cannot hear each other), RTS/CTS adds a degree of determinism to the otherwise multiple access network. When RTS/CTS is invoked, clients must first request access to the medium from the access point with an RTS message. Until the access point replies to the client with a CTS message, the client will refrain from accessing the medium and transmitting its data packets. When received by clients other than the one that sent the original RTS, the CTS command is interpreted as a “do not send” command, causing them to refrain from accessing the medium. One can see that this mechanism will preclude 802.11b clients from transmitting simultaneously with an 802.11g client, thereby avoiding collisions that decrease throughput due to retries. One can see that this additional RTS/CTS process adds a significant amount of protocol overhead that also results in a decrease in network throughput.

In addition to RTS/CTS, the 802.11g standard adds one other significant requirement to allow for 802.11b compatibility. In the event that a collision occurs due to simultaneous transmissions (the likelihood of which is greatly reduced due to RTS/CTS), client devices “back off” the network for a random period of time before attempting to access the medium again. The client arrives at this random period of time by selecting from a number of slots, each of which has a fixed duration. For 802.11b, there are 31 slots, each of which are 20 microseconds long. For 802.11a, there are 15 slots, each of which are nine microseconds long. 802.11a generally provides shorter backoff times than does 802.11b, which provides for better performance than 802.11a, particularly as the number of clients in a cell increases. When operating in mixed mode (operating with 802.11b clients associated) the 802.11g network will adopt 802.11b backoff times. When operating without 802.11b clients associated, the 802.11g network will adopt the higher-performance 802.11a backoff times.

When 802.11g-based networks are operating in the absence of legacy 802.11b clients, network throughput is similar to that of 802.11a. With its OFDM means of transmission and 802.11a backoff scheme, 802.11g can essentially be viewed as 802.11a applied to the 2.4-GHz band. Table 2 illustrates throughput comparisons for 802.11a, 802.11b, and 802.11g. Note that the throughput increase for 802.11g when in mixed-mode operation is relatively modest when compared to 802.11b, and is a fraction of the throughput provided by 802.11g when not supporting legacy clients.

Table 2 Approximate Throughput Comparison for 802.11a, 802.11b, and 802.11g

	Data Rate (Mbps)	Approximate Throughput (Mbps)	Throughput as a Percentage of 802.11b Throughput
802.11b	11	6	100%
802.11g (802.11b clients in cell)	54	8	133%
802.11g (no 802.11b clients in cell)	54	22	367%
802.11a	54	25	417%

The 802.11g standard provides an option called CTS to Self, which is able to provide greater throughput when in mixed-cell mode. As the name suggests, CTS to Self dispenses with the RTS and relies upon the 802.11b client’s clear channel assessment capabilities to check for an open medium. Cisco Systems® intends to provide CTS to Self capabilities, which will provide for increased performance for future products as well as the installed base via a firmware upgrade.

Having established the throughput capabilities of 802.11g in both modes, one can now examine its overall capacity capabilities. Remember, as a result of operating in the same 2.4 GHz band as 802.11b, 802.11g is similarly limited to three nonoverlapping channels. This compares to the 12 unlicensed channels that are available in many parts of the world for 802.11a. With the inclusion of radar detection capabilities that are part of the 802.11h standard, the number of available 5 GHz channels for 802.11a will increase from 12 to 24 on a nearly worldwide basis. Cisco intends to support 802.11h in future 802.11a products that will become available in 2004. Even when not supporting 802.11b clients, 802.11g still provides a fraction of the network capacity provided by 802.11a. Table 3 illustrates a comparison of 802.11b, 802.11g, and 802.11a network capacity.

Table 3 Network Capacity Approximations for 802.11b, 802.11g, and 802.11a

	Throughput (Mbps)	Channels	Capacity (Mbps)
802.11b	6	3	18
802.11g (mixed mode operation)	8	3	24
802.11g (no legacy support)	22	3	66
802.11a	25	12	300
802.11a (with 802.11h support)	25	24	600
802.11g Range and Coverage			

As a matter of physics, there is an inverse relationship between wavelength and range. All other things being held equal, a signal transmitted in a lower portion of the frequency spectrum will carry further than a signal transmitted in a higher band. Additionally, a longer waveform (from lower in the spectrum) will tend to propagate better through solids (like walls and trees) than a shorter waveform. Because 802.11g operates in the same 2.4 GHz portion of the radio frequency spectrum as does 802.11b, it will share its fundamental advantage over the 5 GHz-based 802.11a. With 802.11b and 802.11g all things are not, however, held equal. Another fundamental rule is that as data rates increase, range decreases. 802.11b uses DSSS to support data rates of 11, 5.5, 2, and 1 Mbps each, with correspondingly longer ranges as the data rates decrease. 802.11g uses OFDM to support data rates of 54, 48, 36, 24, 18, 12, 9, and 6 Mbps each, with correspondingly longer ranges as the data rates decrease. The higher data rates supported by 802.11g result in shorter range than the range supported by the maximum 802.11b data rate. Still, OFDM is a more efficient means of transmission than is DSSS, meaning that at a given range, higher OFDM-based data rates will be supported than DSSS-based data rates (all other things being held constant).

Other factors to consider are transmit power and receive sensitivity. The selection of either DSSS or OFDM transmission type has an effect on the maximum power the transmitter can use, as well as the capability of the receiver, particularly at higher data rates. The reason for this is that higher-order modulation schemes such as the 64 quadrature amplitude modulation (QAM) used to support 54 and 48 Mbps data rates requires a high degree of acuity on the part of the receiver. High power coming from the radio's transmitter tends to desensitize the receiver, a phenomenon known as Error Vector Magnitude (EVM). This leads to a counterintuitive effect, whereby increasing transmit power tends to decrease the range of the device at these higher data rates. Radios operating in 802.11g mode therefore use lower transmit power than when operating in 802.11b mode.

The result of differing waveforms, data rates, transmission types, and transmission characteristics are numerous ranges for the three 802.11 technologies. Of course, any stated ranges are no more than approximations, as several environmental factors can have a dramatic impact on range and resulting coverage area. Still, one can discern the relative ranges the different technologies provide from Table 4 below. Note the significant range differences for 802.11a relative to 802.11g. Also note the superior range of OFDM data rates relative to DSSS data rates in 802.11g, as illustrated by the greater range at 18 Mbps than at 11 Mbps.

Table 4 Comparative Ranges in an Open Indoor Office Environment Through Cubicle Walls

Data Rate (Mbps)	802.11a (40 mW with 6 dBi gain diversity patch antenna) Range	802.11g (30 mW with 2.2 dBi gain diversity dipole antenna)	802.11b (100 mW with 2.2 dBi gain diversity dipole antenna)
54	45 ft (13 m)	90 ft (27 m)	-
48	50 ft (15 m)	95 ft (29 m)	-
36	65 ft (19 m)	100 ft (30 m)	-
24	85 ft (26 m)	140 ft (42 m)	-
18	110 ft (33 m)	180 ft (54 m)	-
12	130 ft (39 m)	210 ft (64 m)	-
11	-	160 ft (48 m)	160 ft (48 m)
9	150 ft (45 m)	250 ft (76 m)	-
6	165 ft (50 m)	300 ft (91 m)	-
5.5	-	220 ft (67 m)	220 ft (67 m)
2	-	270 ft (82m)	270 ft (82m)
1	-	410 ft (124 m)	410 ft (124 m)

Because 802.11g operates in the same frequency band as 802.11b, it may share a common set of antennas. Cisco currently offers an extensive selection of 2.4 GHz antennas to achieve a variety of coverage patterns. Because WLAN operation in the 5-GHz range is a relatively new phenomenon and because 802.11a volumes are currently far lower than those of 802.11b, there is not nearly as wide a selection of antennas available. As a result, today one has a greater ability to shape a coverage area through the use of specialized antennas with 802.11g than 802.11a. Additionally, the more limited propagation characteristics of the 5 GHz waveform can further restrict the shape of a coverage area due to obstructions that would represent far less an impediment to a 2.4 GHz-based technology.

Deployment Considerations for 802.11g

802.11g can be regarded as a superset of 802.11b, providing all the functionality of, and backward compatibility with, 802.11b, plus the higher performance associated with OFDM transmission. With Cisco 802.11g-based products, there are other reasons to consider these devices outside of the advantages specific to 802.11g itself:

- Cisco 802.11b-based products provide an integrated RC4 encryption engine to provide Wired Equivalent Privacy (WEP) and WPA security without significant performance degradation. Cisco 802.11g-based products provide an encryption engine that supports both the RC4 and the Rijndael algorithm for WEP, WPA, and Advanced Encryption Standard (AES) encryption to provide high-speed support for the upcoming 802.11i standard and FIPS-140 compliance.

- The Cisco 802.11g radio represents a fifth generation in 2.4 GHz technology building upon earlier 802.11 and 802.11b radios. Even when operating in 802.11b DSSS mode, the Cisco 802.11g radio provides superior range performance relative to the previous generation 802.11b-only radio. For applications where 802.11g is not a requirement, one should still consider the 802.11g radio for its 802.11b performance.
- Because 802.11g heavily relies on earlier 2.4-GHz technology, it takes advantage of the significant cost reduction engineering and economies of scale that have resulted from the increasing volumes of 802.11b products. As such, the build cost of an 802.11g is similar to that of an 802.11b radio. Cisco has chosen to price 802.11g-based products identical to its 802.11b products. Customers get increased performance, more robust security, and improved range for exactly the same price.

With the above facts, there is little reason not to purchase 802.11g-based devices over 802.11b. It is equally apparent, however, that 802.11g alone does not provide anything approaching the sort of capacity provided by 802.11a. An answer to this dilemma can be summed up in two words—dual band.

Dual-Band Deployments

In the same way that 802.11g leverages 802.11b technology, 802.11a leverages 802.11g technology. At a high level, an 802.11 radio can be considered to have three major parts—a MAC, a physical layer controller (PHY), and a radio frequency (RF) front end. The MAC, where functions such as the encryption engine discussed above reside, is common across 802.11a, 802.11b, and 802.11g. Support for OFDM and DSSS transmission types resides in the PHY. As the name implies, the RF Front end is the portion of the radio that actually transmits and receives signals, and as such is specific to the frequency band in which it operates. As a result, an 802.11a radio can re-use the same MAC and PHY as an 802.11g radio. The incremental difference between an 802.11g radio and a dual-band radio supporting the 802.11a or 802.11g standard is the addition of a 5 GHz RF Front end, which adds a relatively small and decreasing amount of incremental build cost. It is expected that for mainstream client devices like notebooks and desktop PCs, neither 802.11a nor 802.11g will exist as a standalone client-side radio to any great degree. Rather, by the end of 2004, the great majority of client radios in form factors like Mini-PCI (for embedded use), CardBus, PCI, and USB will provide dual-band 802.11a/g support.

Vendors participating in the 802.11 market recognize that this incremental cost buys a substantial amount of additional capacity for a relatively small price. Most industry experts agree that the predominant client-side radio shipping in 2004 will be a dual-band 802.11a/g radio. When associating in a WLAN infrastructure, these dual-band client radios will typically look first for an 802.11a access point, followed by 802.11g, and finally 802.11b. In this way, they take the fullest advantage of whatever capacity is available, while providing for backward compatibility.

Given the absence of incremental cost and the upcoming availability of supporting clients, Cisco believes that at a minimum, networking professionals should begin the process of migrating their infrastructures from 802.11b to 802.11g. Customers need not “forklift upgrade” their current 802.11b infrastructures—they can choose to either upgrade only the radios within the access points (this feature is available on Cisco Aironet® 1100 and 1200 series access points) or simply mix in new 802.11g access points with existing 802.11b access points for a heterogeneous, yet interoperable, 2.4-GHz infrastructure.

Customers need not perform a new site survey when deploying 802.11g access points. As stated above, the standard provides for several data rates with inversely proportional supported ranges. To be specific, at its 802.11b, 11-Mbps range, the same radio will provide for an approximately 18 Mbps, 802.11g data rate. Also, this fifth-generation

2.4 GHz radio will outperform the previous generation 802.11b-only radio, meaning that the 802.11g radio will provide for an approximate 24 Mbps data rate at the 11 Mbps range of installed 802.11b access points (which have previous-generation radios).

While the initial performance increase provided by 802.11g will be modest, this is likely to improve over time. First, a firmware upgrade scheduled for the first half of 2004 is planned to include CTS to Self capabilities, which are understood to improve performance when operating in mixed-mode environments. Moreover, new 802.11b devices will become increasingly less common as vendors migrate to 802.11g and 802.11a. This will, over time, decrease the incidence of 802.11b devices in the installed base and decrease the incidence of mixed-mode operation. By selecting OFDM as the transmission type for access point management beacons, customers can configure 802.11g access points such that they specifically exclude 802.11b client associations. This will, of course, dramatically improve the performance of associated 802.11g clients. It does, however, complicate channel reuse patterns and 802.11b coverage, given the small number of 2.4 GHz channels that one has available.

In the same way that one can install 802.11g access points without necessarily performing a new site survey, one can similarly provide 802.11a support. Cisco Aironet 1200 Series Access Point simultaneously support both a 2.4 GHz 802.11b or 802.11g radio and a 5 GHz 802.11a radio. By adding the 802.11a radio to access points that are deployed based upon an 802.11g (or for that matter, 802.11b) site survey, some amount of 802.11a coverage will be provided. 802.11a radios can be added to Cisco Aironet 1200 Series access points at time of purchase or as a field upgrade. As discussed previously, this 802.11a coverage will not provide nearly the same data-rate support as the 802.11g “side” of the access point, and it is likely that there will even be areas known as “nulls” that are completely without 802.11a coverage. Still, upcoming dual-band 802.11a/g clients have the intelligence to switch between 5 GHz and 2.4 GHz, taking advantage of 802.11a capacity in areas where it is available and switching to 802.11g (or even 802.11b) when it is not. Client devices will endeavor to maintain data rate over frequency band. For example, a client associated to an 802.11a radio at 54 Mbps will seek out an 802.11g radio supporting 54 Mbps before it “downshifts” from 54 Mbps to a lower data rate on the 802.11a radio.

As the number of WLAN-enabled client devices on the network increases and the network evolves to more commonly support 802.11a through dual-band clients, customers can begin to add single-band 802.11a-only access points to their infrastructures. In so doing, they can cover any null areas and generally bring 802.11a coverage up to data-rate parity with 802.11g coverage. The very large number of 802.11a channels (particularly when coupled with 802.11h) will tend to make the process of site surveying for this follow-on deployment far more quick, easy, and less expensive than an initial 802.11b or 802.11g site survey.

In Summary

802.11g is an exciting new technology that offers additional performance, while providing investment protection for 802.11b clients through backward compatibility. By using previous technologies and economies of scale, 802.11g devices are available at little or no additional cost relative to 802.11b. As such, there are many reasons to begin migrating from 802.11b to 802.11g.

802.11g is not, however, a panacea for WLAN capacity issues. When operating in mixed-mode environments, the throughput and capacity improvements are very modest. Even with advances like CTS to Self and when operating in 802.11g-only mode, capacity will be limited by the small number of available 2.4 GHz channels. Recognizing this, vendors will soon begin providing dual-band 802.11a/g clients at attractive prices. Vendors like Cisco today provide dual radio access points that simultaneous support 5 GHz 802.11a and 2.4 GHz 802.11g for association with almost any type of 802.11 client. As WLAN capacity needs increase, network professionals are well-advised to begin deploying a dual-band infrastructure to access the far greater capacity available with 802.11a. As such, 802.11g should be viewed as a portion of an overall WLAN architecture, not a substitute for 802.11a. 802.11g is a “bridge” technology and an ideal means for migrating from low-capacity 802.11b networks to the high-capacity, dual-band WLANs of the very near future.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

ETMG 203179—JS 10/03