

# White Paper

## A DETAILED EXAMINATION OF THE ENVIRONMENTAL AND PROTOCOL PARAMETERS THAT AFFECT 802.11G NETWORK PERFORMANCE

---

Issues Surrounding the Deployment of 802.11g Networks .....	1
Introduction.....	1
Introduction to 802.11g.....	1
Simplest Situation: 802.11g Only Devices.....	3
Effect of Propagation.....	3
Effect of the Frequency Band and the Number of Channels .....	5
Adding Complexity: Co-existence with 802.11b Devices .....	8
Technical Terms.....	8
Framing and the Preamble.....	8
Channel Access and the Slot Time .....	9
Overlapping BSS Case.....	11
Mixed-Mode Networks .....	12
Protection Mechanisms; The Final Complication.....	13
Proxim Solutions – Taming the Complexity.....	16
802.11b Only Mode.....	17
802.11g Only Mode.....	17
802.11b/g Mixed Mode.....	17
Mode Performance.....	18
Two Slot Architecture Migration.....	18
Summary .....	19

### Introduction

Boasting higher data rate schemes matching those available from 802.11a in the 5 GHz band, and backward compatibility with 2.4 GHz 802.11b, 802.11g is gaining acceptance in the Wireless Local Area Network (WLAN) market. Despite growing interest, early adopters face numerous challenges that may impact radio throughput performance or reduce the stability that currently exists within legacy networks. These include longer slot-time duration than used in 802.11a, Request-to-Send/Clear-to-Send (RTS/CTS) or CTS-Only protection mechanisms for simultaneous b/g user coordination, and fewer available channels for frequency reuse than are available in the 5 GHz band. This white paper examines the challenges involved in deploying an 802.11g-based network, and offers concrete solutions to help users maximize the effectiveness of the 802.11g network depending on their particular infrastructure and deployment requirements.

### Introduction to 802.11g

With low cost and high-speed data rate capabilities, the popularity of IEEE 802.11-based WLANs is growing exponentially. Currently, the two most important IEEE standards are: 802.11b in the 2.4 GHz Industrial-Scientific-Medical (ISM) band, which uses Complementary Code Keying (CCK) at the higher

data rates, and 802.11a in the 5 GHz Unlicensed National Information Infrastructure (U-NII)/ISM bands in the US, and license-free 5 GHz bands elsewhere, which uses Orthogonal Frequency Division Multiplexing (OFDM). The former offers data rates of 1, 2, 5.5, and 11 Mbps, while the latter is capable of supporting higher data rates: 6, 9, 12, 18, 24, 36, 48 and 54 Mbps. IEEE 802.11g, expected to be formally ratified in July of 2003, is drawing attention recently based on the use of the same data rates and OFDM technology as employed in 802.11a, plus backward compatibility with 802.11b devices. In theory, both 802.11g and 802.11a use almost the same PHY specification and, therefore, should have similar throughput performance. In reality, 802.11g throughput performance will be significantly different from 802.11a for the following reasons:

1. 802.11g mandates the use of a 20  $\mu$ s slot time in order to be compatible with current 802.11b devices. The use of a 9  $\mu$ s slot time as is used in 802.11a is optional.
2. 802.11g shares the same 2.4 GHz spectrum as 802.11b. When both 802.11g and 802.11b devices are present, the performance impact may be significant if no coordination is employed between .11b and .11g users.
3. Frequency-dependent propagation loss favors 802.11g, that is, free space path loss is greater at 5 GHz than at 2.4 GHz. However, the prevalence of non-WLAN devices in the 2.4 GHz ISM band, e.g., Bluetooth devices, cordless phones, microwave ovens, etc. raises the probability of .11g devices encountering interference harmful to WLANs.
4. There are fewer available channels in the 2.4 GHz band than in the 5 GHz bands. For example, only three non-overlapping channels exist in the US 2.4 GHz ISM band compared with 13 available channels in 5 GHz U-NII band. (The number of channels is even larger in other regulatory domains, and is likely to increase to more than 20 channels in the US.) Unlike in the home/SOHO application, frequency reuse is necessary for enterprise/public space use to support coverage and capacity requirements. Co-channel interference due to frequency reuse is more likely when fewer channels are available.

Table 1 summarizes the key difference between the three WLAN systems.

	<b>802.11a</b>	<b>802.11b</b>	<b>802.11g</b>
<b>Operating frequencies</b>	5 GHz U-NII/ISM Bands	2.4 GHz ISM Band	2.4 GHz ISM Band
<b>FCC regulation</b>	Part 15.407/15.247	Part 15.247	Part 15.247
<b>Modulation techniques</b>	OFDM	Barker Code/CCK	Barker Code/CCK/OFDM
<b>Data rates (Mbps)</b>	6,9,12,18,24,36,48,54	1,2,5.5,11	1,2,5.5,11 6,9,12,18,24,36,48,54
<b>Slot time</b>	9 $\mu$ s	20 $\mu$ s	20 $\mu$ s 9 $\mu$ s (optional)
<b>Preamble</b>	OFDM	Long Short (optional)	Long/Short/OFDM

*Table 1 Different WLAN system characteristics*

This whitepaper will explain, in as simple a way as possible, the impact of these various effects on the performance of an 802.11g network. The next section will address spectrum and channel issues that relate to 802.11g devices without consideration of legacy devices. The section following that brings in the complexity of co-existence between 802.11g and 802.11b devices. The final section presents Proxim's solutions to these issues.

## Simplest Situation: 802.11g Only Devices

Even without considering co-existence with 802.11b devices, it is true that 802.11g devices will not have equivalent performance to 802.11a devices simply because of the frequency band in which they operate. This section describes the effects of propagation and channel availability, and how they are different in the cases of 802.11a and 802.11g.

### Effect of Propagation

One of the most obvious differences between 802.11g devices and 802.11a devices is that they operate in different frequency bands. Because the size of the antennas used to transmit and receive signals depends on the frequency, for antennas with similar characteristics there is a frequency dependent effect on the reduction in signal strength as measured by two antennas. This effect is commonly referred to as frequency dependent path loss. In a free space environment, the path loss between isotropic antennas at both 2.4 GHz and 5 GHz is shown in Figure 1.



Figure 1: Free space path loss at 2.4 GHz and 5 GHz

As this figure shows, in the simple free space model there is approximately a 6 dB difference between propagation at 2.4 GHz and propagation at 5 GHz. ( $20 \times \log_{10}(2400) - 20 \times \log_{10}(5000) = -6.4$  dB.) In deployments in which signal range is the most important factor, this effect favors the 802.11g devices since, in principle the signals from those devices will propagate further with less loss.

Of course, the propagation conditions in the locations in which 802.11 wireless LANs are most often used are not likely to be well represented by the free space path loss model. Wireless LANs are typically used indoors, where multipath and shadowing effects will have a significant effect on the realized propagation. In order to more accurately represent the true environment, and in order to model more complicated effects, we have used a simulation with the following parameters.

- The propagation channel is modeled as a Rayleigh fading channel with a 50 ns rms delay spread.
- The fade margin is calculated based on 90% service availability. Shadowing is modeled as a log-normal distribution with an 8 dB standard deviation.

- For indoor radio environments, a two-slope path loss model with a breakpoint is used. Path loss slopes are modeled by two straight lines intersecting at the break point. Equation 1 describes the free space path loss before the break point.  $d$  is the distance between the transmitter and the receiver when  $d$  is less than the break point, and  $\lambda$  is the wavelength. For distances beyond the break point, the path loss is dependent on the environment and is shown in Equation 2 where  $\alpha$  is the path loss exponent. The break point is related to the first Fresnel zone clearance, antenna height, and the transmitted frequency. In this example we set the breakpoint to 10  $m$  and  $\alpha = 3.2$ . These values represent a medium obstructed office environment.

$$PathLoss = 20 \cdot \log_{10} \left( \frac{4\pi d}{\lambda} \right) \quad \text{Equation 1}$$

$$PathLoss = break\_point\_loss + 10 \cdot \alpha \cdot \log_{10} \left( \frac{d}{break\_point} \right) \quad \text{Equation 2}$$

$$break\_point\_loss = 20 \cdot \log_{10} \left( \frac{4\pi \cdot break\_point}{\lambda} \right) \quad \text{Equation 3}$$

- The rate fallback algorithm selects 10% PER (packet error rate) as a criterion. The current data rate scheme will be used until 10% PER is reached, at which point the system will reduce the data rate to next lowest available rate.

First, to see the effect mainly of the propagation environment, a Monte-Carlo simulation was used to model a single user environment. No packet errors are caused by collisions (since there is only one user on the network), that is, packet errors resulting in retransmissions are caused by insufficient SNR (signal to noise ratio) only. In this simulation the AP antenna height = 4  $m$  and the client antenna height = 1.5  $m$ . A total of 10,000 simulation runs were performed. In each simulation run, the user's location was generated randomly, uniformly distributed within the cell, and the SNR is calculated after applying shadowing and pathloss. The throughput is then derived based on received SNR. Figure 2 shows the average throughput for different operating modes against cell size. The average throughput is defined as the total throughput during 10,000 simulation runs divided by 10,000. After including the frequency-dependent propagation loss, 802.11g provides a higher average throughput than 802.11a in this single cell case. Even when the 802.11g throughput is degraded by the use of a 20  $\mu s$  slot time, the average throughput can still be comparable to 802.11a if the cell radius is over 70 meters. (The meaning of these slot time values will be discussed in detail later in this paper.)

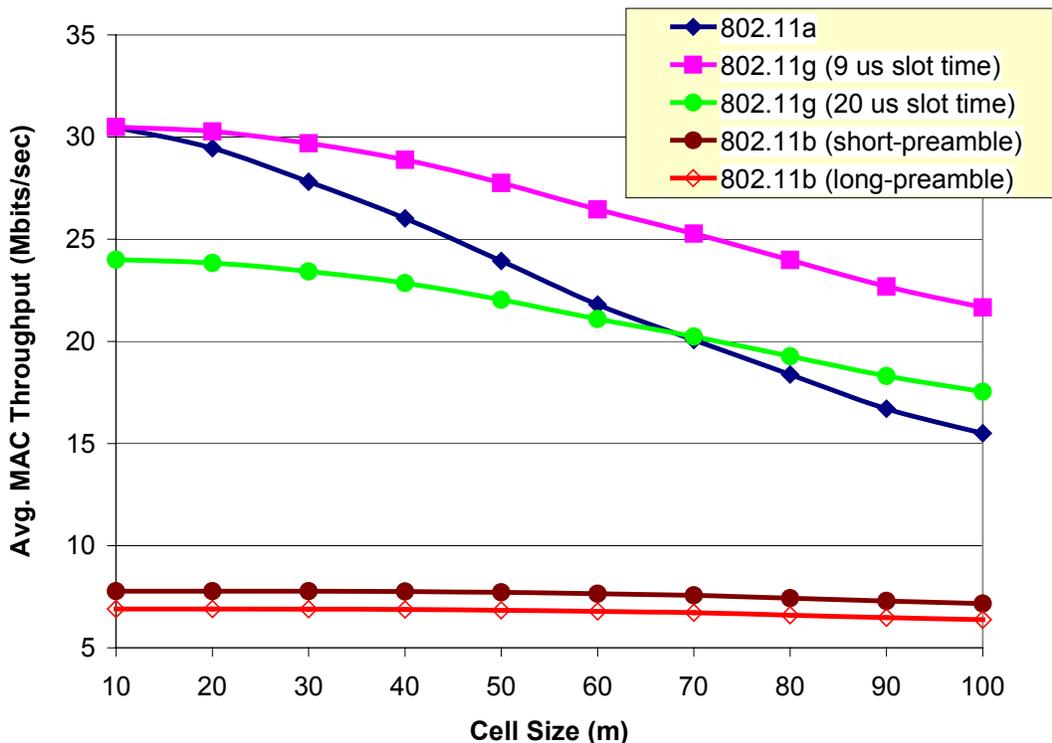


Figure 2: MAC layer throughputs for single cell networks

### Effect of the Frequency Band and the Number of Channels

While the propagation effect favors the 802.11g devices, the lower levels of interference in the 5 GHz band, and the number of channels available there, favors 802.11a devices.

Since 802.11g devices operate in the 2.4 GHz band, they are subject to many more interference sources than are 802.11a devices, which operate in the 5 GHz band. Since the 2.4 GHz band is an ISM (industrial, scientific, and medical band) from 2.400-2.500 GHz as well as an unlicensed band from 2.400-2.4835 GHz, wireless LANs in this band must deal with interference from devices like microwave ovens (more than 100 million in the US alone), medical uses of RF for heating of body tissues and magnetic resonance imaging, and of course other communications devices such as Bluetooth and cordless telephones. Since most of the 5 GHz allocations do not overlap an ISM band, and because microwave ovens do not operate in this band, there are fewer sources of external interference for 802.11a devices.

However, an equally problematic source of interference to wireless LANs is *self-interference*, which results from APs in a single network interfering with other APs on that same network operating at the same frequency. As the networks become larger and more APs are required for either coverage or capacity, or both, it becomes increasingly difficult to avoid interference between co-channel APs because there are such a limited number of channels from which to choose.

This effect is related to a classic problem in mathematics known as the Four Color Theorem. The Four Color Theorem dates back to 1852 when Francis Guthrie, while trying to color the map of counties of England noticed that four colors was sufficient for his task. He asked his brother Frederick if it was true that any map can be colored using four colors in such a way that adjacent regions (i.e. those sharing a

common boundary segment, not just a point) receive different colors. This theorem was finally confirmed by Appel and Haken in 1976, when they published their proof of the Four Color Theorem. What this theorem shows is that in any attempt to color a map in order to avoid having adjacent boundaries share a color (or, in the case of WLANs, avoid having adjacent cells share a channel) you need at least 4 colors (or 4 channels.)

The likely validity of this assertion can be seen by taking a simple example. Imagine a 4x3 grid, each section of which can represent an area that needs to be covered by a single AP. An example is shown in

Figure 3.

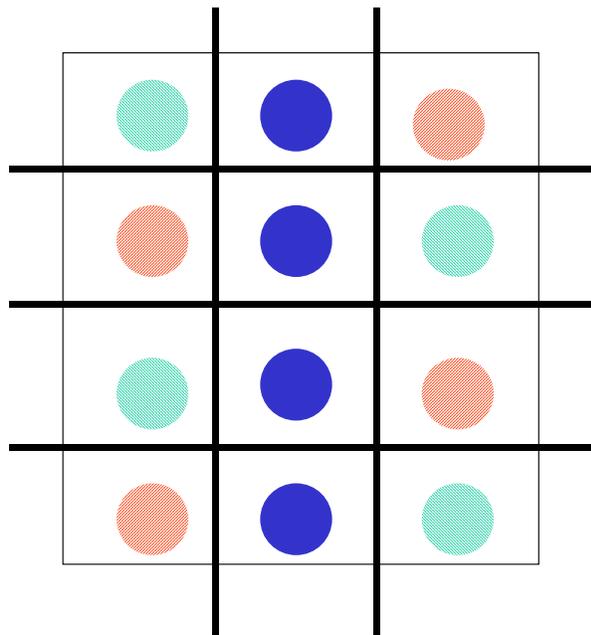


Figure 3: Illustration of co-channel neighbor cells in a 3-channel system

The three types of circles shown in this figure represent three different channels on which the AP can operate. The configuration shown in this figure will have significant interference problems since there are many cases in which the blue (solid) circles are adjacent to each other. There is no configuration of colors that avoids this problem with only 3 colors (channels) available. As shown in Figure 4, however, once the number of channels is increased beyond three, there are many ways to arrange the APs so that adjacent cells do not use the same channel. This leads to dramatic reductions in the self-interference effect.

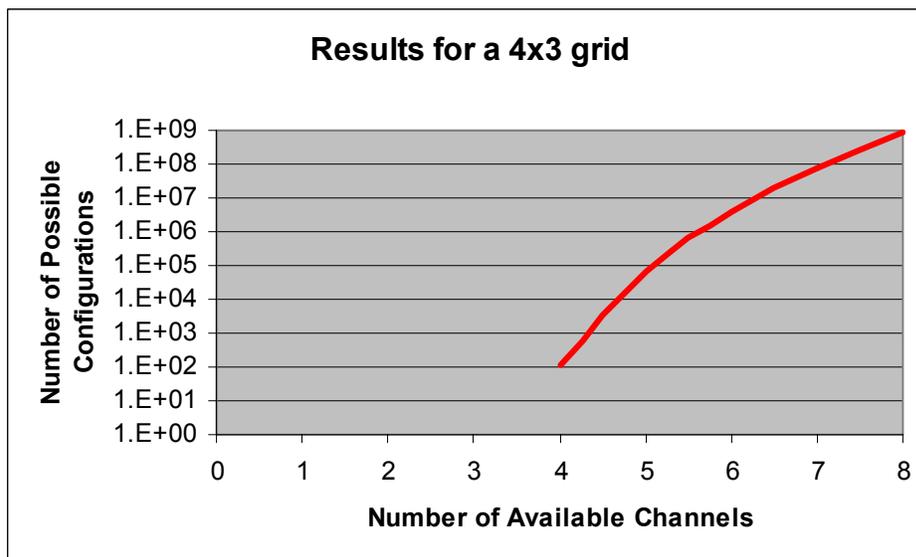


Figure 4: Demonstration of the number of possible configurations with no co-channel neighbor cells, as a function of the number of available channels

Thus with only three channels, 802.11g has insufficient channels to deploy access points on a grid layout without having to use the same channel on adjacent access points. 802.11a with its larger number of channels gives the wireless network designer much more flexibility in assigning channels. There is a possibility of additional spectrum adjacent to the current 2.4 GHz band being made available for unlicensed use in the US to increase the number of 802.11g channels from three to four. However, 802.11a still has considerably more flexibility in this area, and thus much larger network bandwidths can be achieved.

Using our simulation we can look specifically at the magnitude of this effect as it relates to the deployment of 802.11a and 802.11g networks.

Figure 5 plots the throughput performance as affected by co-channel interference when the frequencies are reused in a cellular configuration. Only 6 1st-tier co-channel interferers are considered. (Only seven channels for 802.11a were considered.) However, the worst-case probability of co-channel interferer activity, i.e. 100%, is assumed. Since the interference impact is also related to the type of WLAN link traffic (the reference WLAN device can be an access point or a client), 20% of the traffic is assumed to be uplink traffic (client to access point) and 80% of the traffic is assumed to be downlink traffic (access point to client).

Figure 5 shows the effect of less dense frequency reuse possible due to the larger number of channels available in the 5 GHz band. The 802.11a networks have much higher average throughput in these multi-cell environments.

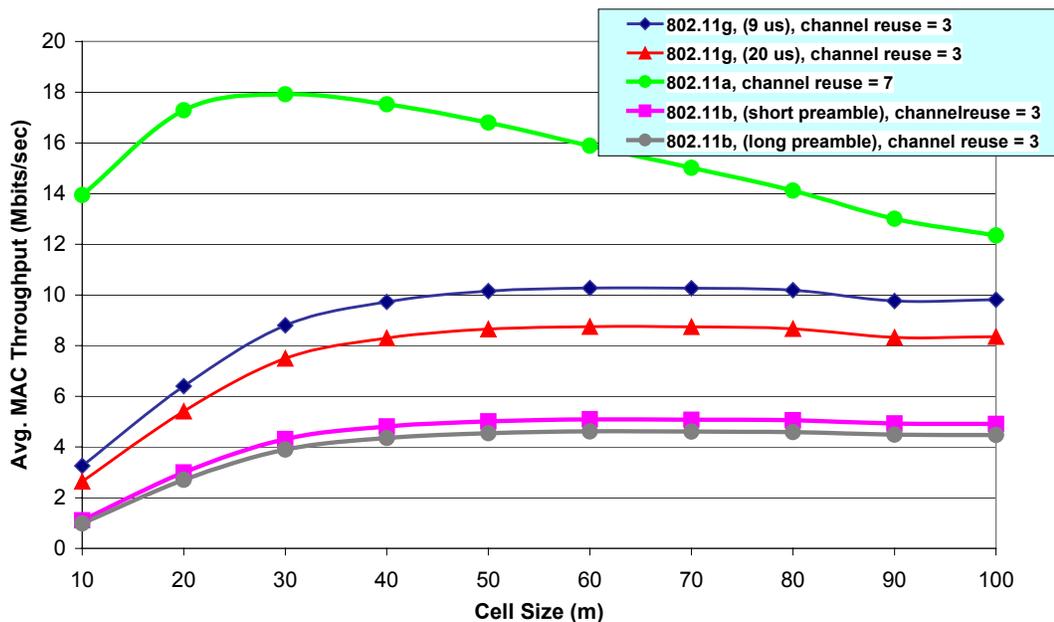


Figure 5: MAC layer throughputs for multi-cell networks

As mentioned above, some of the other aspects of this figure (preambles and slot sizes) will be discussed in detail below.

## Adding Complexity: Co-existence with 802.11b Devices

One of the benefits of 802.11g is that the standard is being prepared so that 802.11g devices have built-in backwards compatibility with 802.11b devices. Simultaneous operation of these two types of devices, however, affects the overall network performance, and we discuss that issue in this section. First we describe some necessary technical terms, and then we describe what happens in the case of overlapping and mixed mode networks. Finally, we discuss the protection mechanisms, and their effect on network performance.

### Technical Terms

As we have discussed above, there are differences that a user can expect between the performance of an 802.11a and an 802.11g network based solely on the frequency bands in which they operate and the number of non-overlapping operating channels available in those bands. However, there is a much more complex issue that will affect the performance of 802.11g devices – co-existence with 802.11b devices. In order to discuss these complexities, it is unfortunately necessary to understand some of the details of the 802.11 physical layer and media access layer protocols.

### Framing and the Preamble

The first detail that requires explanation is how data is transmitted within an 802.11 network. While it is relatively well known that the 802.11 standards define a data modulation that determines the rates at which data can be transmitted (for example, in 802.11b, the highest data rate of 11 Mbps is achieved using CCK modulation), less well known are the “overheads” that accompany each of these data packet transmissions. The term “overhead” refers to information that must be sent in order for the

network to operate properly, but which is not part of the data that the user is transmitting. Since the network must still spend time to send this information, but since it is not recognized by the end user as part of his traffic stream, the total throughput of the network will appear to be lower than the peak data rate that is quoted. It is partly for this reason that even the highest performing 802.11 networks do not achieve user throughput results much in excess of 50% of the peak data rate (about 5 Mbps for an 11 Mbps 802.11b network using long preambles, and 30 Mbps for a 54 Mbps 802.11a network.) The calculations that follow, both in this section and in the rest of this paper, are intended to give the reader an understanding of the magnitude of the effect being discussed. There are many overheads that, for clarity, are not being specifically discussed in this paper. Also, while actual throughput measurements are taken using the result of some higher layer application (like FTP for example), we will confine our discussion to the 802.11 protocol layer.

A very simplified figure of an 802.11 data packet is shown in Figure 6.



*Figure 6: Simple illustration of 802.11 preamble overhead*

The areas shown in gray make up what is known as the PLCP (Physical Layer Convergence Protocol) field in 802.11. There are three different preambles that will be relevant in this discussion. For 802.11b there is the long preamble (192  $\mu$ sec) and the short preamble (96  $\mu$ sec), and for 802.11a and 802.11g there is the OFDM preamble (20  $\mu$ sec). The effect of this preamble can be estimated in the following way. If 1500 bytes of data are sent at 11 Mbps, this takes 1091  $\mu$ sec ( $1500 \times 8 / 11$ ). If this data were sent with the long preamble attached, it would take an extra 192  $\mu$ sec for a total of 1283  $\mu$ sec. Sending 1500 bytes of data in 1283  $\mu$ sec means that the user will see a throughput of 9.5 Mbps, a reduction of 13.6% from the quoted 11 Mbps rate.

On the other hand, if the short preamble is used, then 1500 bytes is sent in 1187  $\mu$ sec, for an effective rate of 10.2 Mbps. The use of the long preamble lowers the effective data rate by about 7%. The importance of this for mixed mode 802.11g/802.11b networks will be discussed later.

### **Channel Access and the Slot Time**

The section above touches on how the data in 802.11 is packaged for transmission. This section describes how, once it is created, a data packet like that shown in

Figure 6 gets transmitted in 802.11. The 802.11 media access protocol attempts to avoid having different devices send their data packets at the same time. When this does happen it is known as a collision, and it usually results in neither packet being correctly received. This requires the data to be retransmitted, and the effect is a lowered average system throughput.

In order to attempt to avoid these collisions, 802.11 devices access the channel in the following way. While any one device is transmitting, no other devices will transmit. The other devices will wait until the channel is clear before transmitting. But after they determine that the channel is clear, and before they transmit their data packets, they will wait a random period of time in order to decrease the statistical likelihood that any two devices that have been waiting for the channel to become clear will transmit at the same time. The period of time that a device waits before accessing the channel (The

“backoff time”) is counted in integer units of what is known as the “slot time”. The slot times are shown schematically in Figure 7.

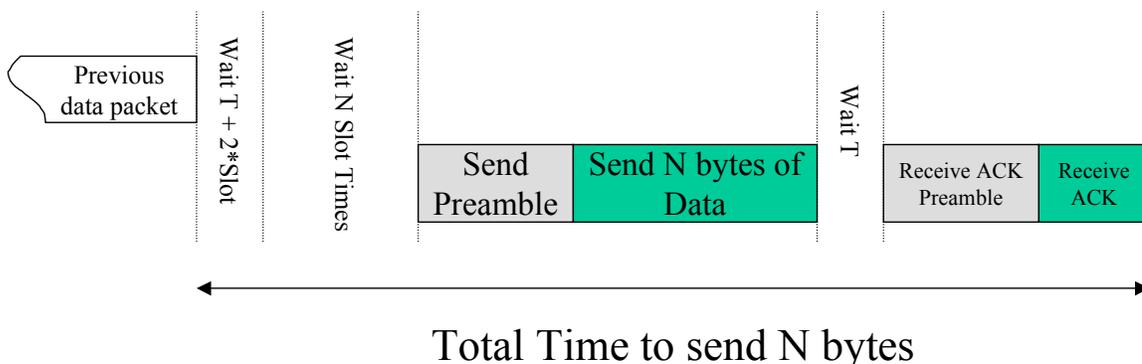


Figure 7: Simple illustration of 802.11 preamble, slot, and acknowledgement (ACK) frame overheads

In the 802.11b standard, the slot time is defined to be 20  $\mu$ sec. In 802.11a, the slot time is defined to be 9  $\mu$ sec. The effect of these times can be estimated in the same way that we estimated the effects of the preamble, above. If a device is transmitting a 1500 byte packet at 11 Mbps, this will take 1091  $\mu$ sec. In 802.11b, the slot time is 20  $\mu$ sec, and the number of slots used is a random number between 0 and 31. Using 15 slot times as an average, this adds an additional 300  $\mu$ sec to the transmission time, for an effective throughput of 1500 bytes in 1391  $\mu$ sec, or 8.6 Mbps. So the backoff times reduce the throughput by about 22% in this case. As can be seen, the magnitude of this effect will change depending on the data rate, because while the time required to send a packet will change, the time spent backing off will not.

This will be a very important issue in understanding the throughput of 802.11g devices in the presence of 802.11b devices. Because 802.11g is based on the 802.11a standard, it uses a 9  $\mu$ sec slot time with a random number of slots between 0 and 15. Using the same calculation as above, this can be seen to reduce a 54 Mbps data rate by 22% to 42 Mbps.

The channel access mechanism cannot operate correctly unless all devices in a network use the same slot timing to determine when they should send their packets. If different slot times were used, the devices with the smaller slot time would have preferential access to the wireless medium. (If an 802.11g device chose a backoff value of 10 slots, and an 802.11b device in the same network chose a backoff value of 5 slots, the 802.11g device would *still* access the channel first if it was using 9  $\mu$ sec slots [90  $\mu$ sec total] while the 802.11b device was using 20  $\mu$ sec slots [100  $\mu$ sec total].) Therefore 802.11g devices operating in the same network as 802.11b devices must use a 20  $\mu$ sec slot time. The same calculation as above shows that this effect alone will reduce a 54 Mbps data rate by 39%, to 33 Mbps. This slot time effect, therefore, has the effect of reducing the data rate perceived by the user by about 21%.

The rough calculations discussed above are used to help understand the magnitude of these effects. There are other overheads and effects that are not being taken into account in an effort to make the description understandable. For a more reliable calculation, a simulation of the 802.11 protocol was used to gauge the magnitude of these effects taking into account the full complexity of 802.11. In figure Figure 8 we show the effect on an 802.11g user of using the 20  $\mu$ sec slot time versus using the 9  $\mu$ sec slot time.

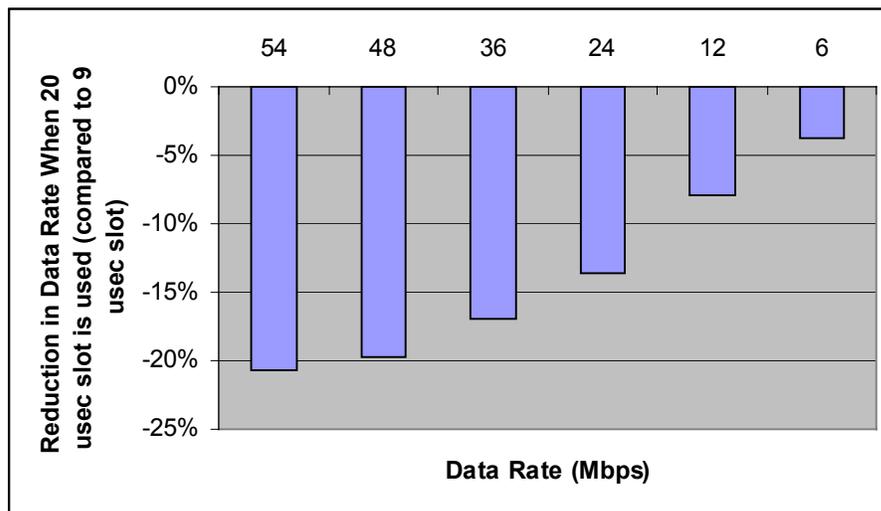


Figure 8: Reduction in data rate caused by the use of a 20  $\mu$ sec slot time

The agreement with the simple calculation shown above is excellent, with the perceived throughput being reduced by about 20% at the 54 Mbps data rate. The results for other data rates are shown as well.

### Overlapping BSS Case

We now have the tools to begin to analyze some of the issues surrounding 802.11g system deployment in the presence of 802.11b devices. First we take the case of Overlapping BSSes on the same channel. This refers to the situation, depicted in

Figure 9, in which two systems, unrelated to each other, are in close enough proximity to hear each other.

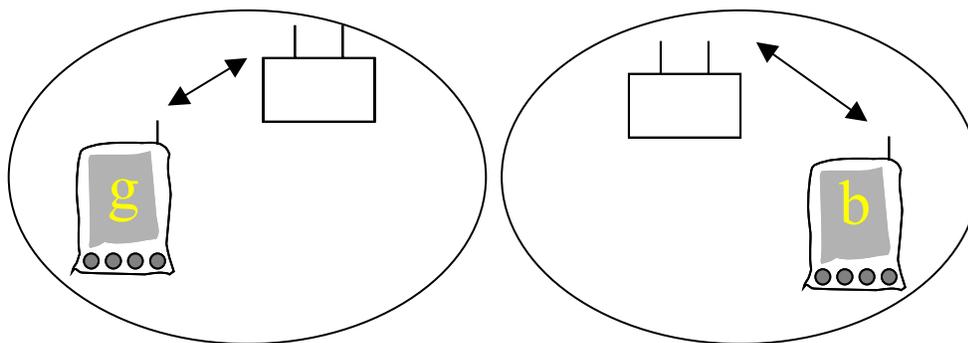


Figure 9: Schematic diagram of an overlapping BSS situation

As shown, we make this situation as simple as possible and designate that one of the networks has only 802.11g devices in it, and the other network has only 802.11b devices. Given the ideal circumstances, what would be the effect of this case?

In the ideal case, the 802.11b devices and 802.11g devices will not transmit at the same time since that would cause packet collisions. Rather, the devices would sense each other and would only transmit when the other devices were not transmitting. This is shown schematically in Figure 10.

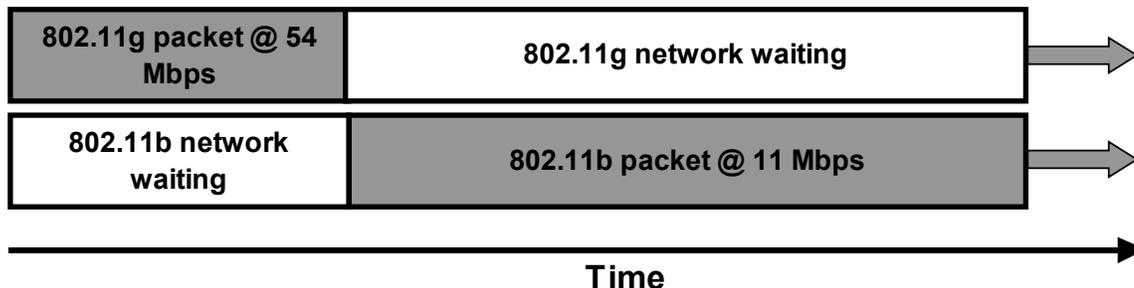


Figure 10: Demonstration of device throughput when devices back off to each other

Take a simple example of the effective throughput in this case, as above. Ignoring overheads (as mentioned several times, the calculations shown here indicate the magnitude of the effect, but do not include the full complexity of the 802.11 protocol), the 802.11g network at 54 Mbps can send a 1500 byte packet in 222  $\mu$ sec. The 802.11b network at 11 Mbps can send a 1500 byte packet in 1091  $\mu$ sec. Therefore, from the point of view of each network (or from the point of view of a device on each network) it will require 1313  $\mu$ sec to transmit a packet of 1500 bytes. From the point of view of the 802.11b device this translates into an effective throughput of 9.1 Mbps, or about a 17% reduction in throughput due to the presence of nearby 802.11g devices.

From the point of view of the 802.11g device, however, the picture is far worse. That device also sees that it takes 1313  $\mu$ sec to transmit 1500 bytes, so its throughput is the same as that of the 802.11b device. (This should actually be obvious from a glance at Figure 10. It takes the same amount of time for a device on each network to transmit the same amount of data, so the effective data rate must be the same.) The 802.11g devices, therefore, see a reduction in their effective data rate by about 83% due to the presence of nearby 802.11b devices.

### Mixed-Mode Networks

In mixed mode networks the situation is even more complicated due to the fact that the 802.11b devices and the 802.11g devices will operate in a compatible mode. An example of a mixed mode network is shown in

Figure 11.

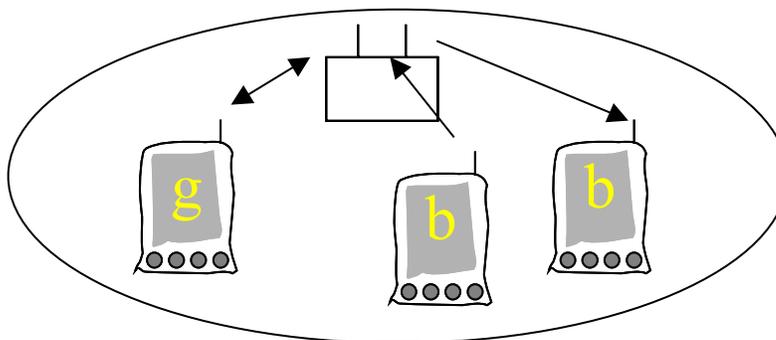


Figure 11: Schematic diagram of a mixed-mode configuration

In this network, not only will the devices, ideally, back off to each other creating the same situation as described above, but there will be other effects as well. First, in order for the devices to operate in conjunction with each other, the slot times must be equivalent. Since the 802.11b devices can only operate with a 20  $\mu$ sec slot time, the 802.11g devices must use this as well. Figure 8 shows the magnitude of this effect, which can be as large as a 20% reduction in the throughput of the 802.11g devices, depending on the data rate at which the devices are operating. Requiring that 802.11g and 802.11b devices to back off in the presence of one another means, as shown above, that the effective throughput for any devices can be limited by the slowest device on the network.

### **Protection Mechanisms; The Final Complication**

The situation we have described at the moment is that, in an environment with both 802.11g and 802.11b devices,

- The devices will ideally defer to each other, causing the effective throughput of any given device to be determined by the slowest active device in the area at a given time.
- In order to operate in a compatible mode, the 802.11g devices must use the longer slot time of 802.11b, causing a further degradation to their throughput performance.

In fact, there is one further complication that we have not yet described. As we have mentioned numerous times, the 802.11b and 802.11g devices will ideally defer to each other in order to avoid having their packets collide over the air. These devices, however, may not defer to each other, because the 802.11g OFDM modulation may not be received properly by the 802.11b receivers. The result of completely uncoordinated devices operating within range of each other will be a reduction in effective throughput for a different reason – over the air packet collisions. One example of this is shown in Figure 12.

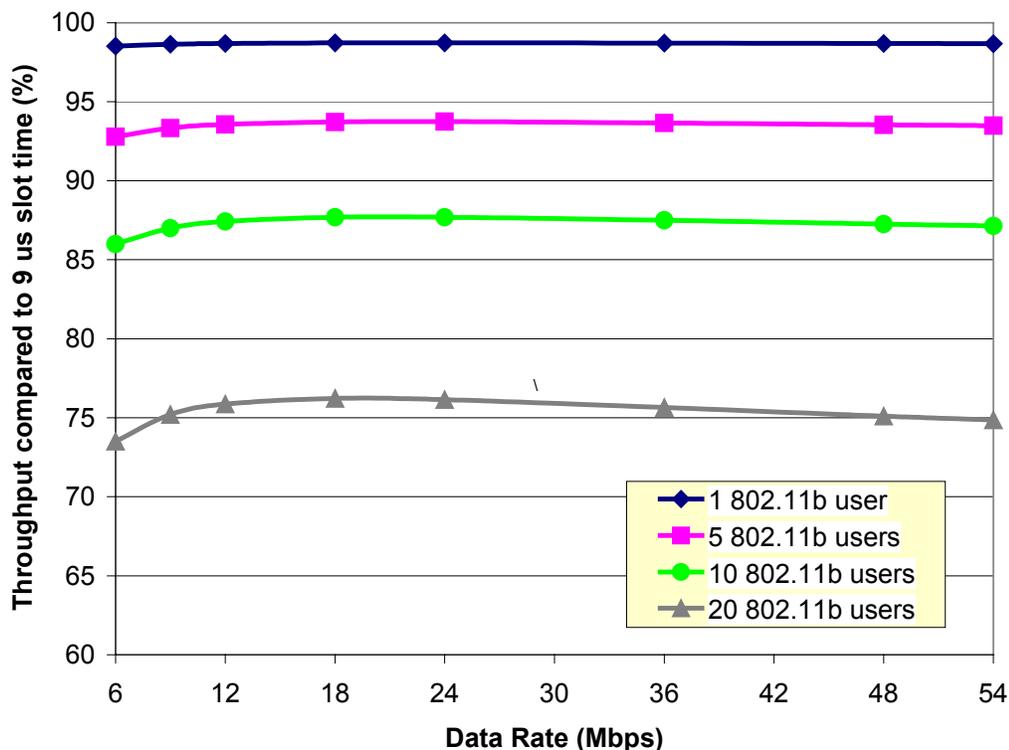


Figure 12: Effect on 802.11g device when uncoordinated 802.11b devices send probes to the AP

To generate this figure, we have simulated the following scenario. An 802.11g network is operating in its optimum mode, using only 802.11a-like parameters (slot time, preamble, etc.). When an 802.11b device is brought into the area, that device will not be able to recognize, nor will it be recognized by, the 802.11g network. According to the 802.11 protocol, this 802.11b device can begin to send messages called probe requests, trying to locate an 802.11b network. Those requests (using 802.11b signaling) will collide with the 802.11g devices, and will cause the throughput degradation shown here. As more and more devices enter the area, the increased number of probes will degrade the 802.11g performance further.

It is important to recognize that, in this simulation, the 802.11b devices are not even sending any data. They are only sending 802.11 probe requests, and already they have degraded the 802.11g network throughput. If they were sending data packets in an uncoordinated fashion, the collisions would be far worse causing even more degradation.

There is a mechanism for dealing with this problem as well, and this is the protection mechanism known as RTS/CTS (Request-to-Send and Clear-to-Send.) RTS can be thought of as a reservation request sent by a device on the network. CTS is a response to this message, informing the device making the request that its request has been received, and that it is OK to send its data packet. In its traditional mode, RTS/CTS was used in the following way. When there are devices on the network that cannot hear other devices on the network (the so-called “hidden node problem”), those devices do not transmit their data to the AP immediately when they sense that the channel is quiet. Rather, it sends an RTS message to the AP. If the AP receives that RTS message, it sends the CTS, which will be heard by all of the devices on the network, even those that the RTS-sending device may not be able to hear. (By definition, there are no devices in an infrastructure-based network that cannot be heard by the AP.) All the devices hearing the CTS will know to cease transmissions for a period of time (defined by the CTS) and this will result in fewer collisions. The mechanism is shown in

Figure 13.

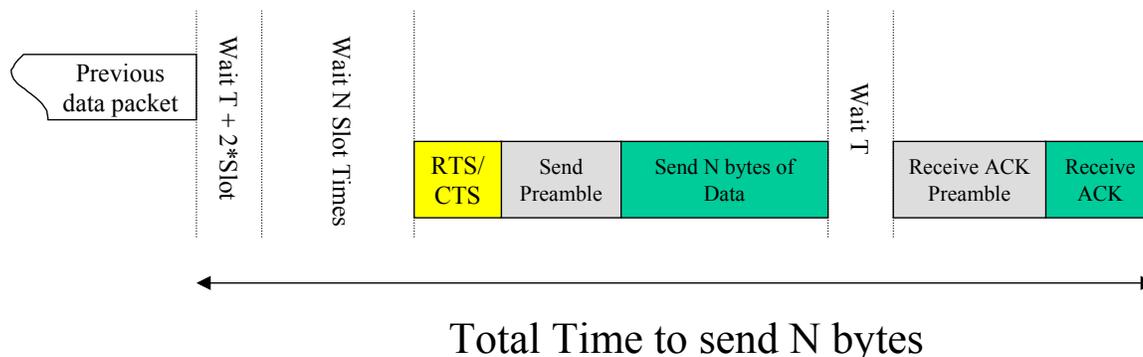


Figure 13: Simple illustration of 802.11 preamble, slot, acknowledgement (ACK), and RTS/CTS frame overheads

The benefit of reduced packet collisions does not come at zero cost, however. As is clear in

Figure 13, it now takes more time to transmit the same amount of data, this reducing the effective throughput seen by the user.

In the case of 802.11g/802.11b co-existence, there is a provision which allows for the sending of only a CTS packet by the AP; that is, the CTS packet can be sent by the AP when it needs to send a downlink packet, and that CTS does not need to be in response to an RTS. We refer to this protection mechanism as “CTS-Only”.

The effect on user throughput of these protection mechanisms is also easily estimated (with the standard disclaimer that there are many other overheads not being considered here):

Time to send 1500 bytes at 54 Mbps	= 222 µsec
OFDM preamble time	= 20 µsec
<u>Effective throughput</u>	<u>= 49.5 Mbps</u>
CTS (14 bytes) at 11 Mbps	= 10 µsec
CTS (CCK) header (short)	= 96 µsec
CTS + Data total time	= 348 µsec
<u>Effective throughput</u>	<u>= 34.4 Mbps</u>

So, the inclusion of a CTS, sent at 11 Mbps using the short preamble, before the sending of 54 Mbps data packets will reduce the user perceived throughput by 30% (from 49.5 Mbps to 34.4 Mbps).

Simulations bear out this back-of-the-envelope calculation. We have simulated the effect of using both RTS/CTS and CTS-only protection mechanisms on 802.11g traffic.

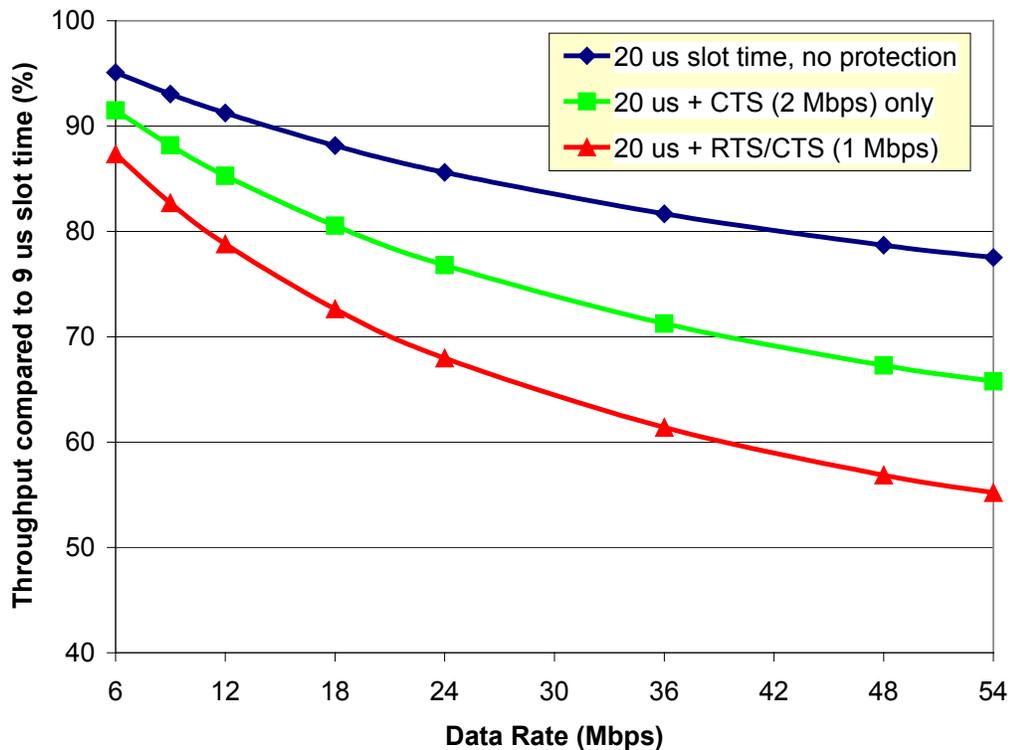


Figure 14: Illustration of the throughput reduction due to the use of the CTS-only and RTS/CTS protection mechanisms

## Proxim Solutions – Taming the Complexity

As we have described above, there are many issues that will affect the performance of an 802.11g network, and which make the deployment of these networks different from the deployment of an 802.11a network. These all are related to the fact that 802.11g operates in the 2.4 GHz unlicensed band, while 802.11a operates in the 5 GHz unlicensed band.

- Propagation issues tend to favor 802.11g in that signals can propagate further.
- The greater number of channels tends to favor 802.11a in that self-interference can be reduced by proper channel assignments.
- Nearby 802.11b networks, or mixed mode 802.11b/802.11g networks need to deal with:
  - Lower rate 802.11b devices which consume bandwidth
  - Long and short preambles which consume bandwidth
  - Different slot times which consume bandwidth
  - RTS and CTS transmissions which consume bandwidth

Several of the operational parameters in these networks are configurable, however, and it will be difficult for many system managers to know how to configure these parameters to achieve the best performance for their system. For that reason, Proxim provides certain static configuration

parameters, and certain dynamic system capabilities, that will help system managers to optimize the performance of their networks. The summary of Proxim's various modes is shown in

Table 2.

Mode	Parameters	Value
802.11b Only	Slot Time	Static (20 $\mu$ sec)
	Preamble	Use Long or Short
	Protection Mechanism (CTS-to-Self)	Disabled
802.11g Only	Slot Time	Static (9 $\mu$ sec)
	Preamble	Long and Short Preamble Not Applicable
	Protection Mechanism (CTS-to-Self)	Disabled
802.11b/g	Slot Time	Dynamic (9 or 20 $\mu$ sec)
	Preamble	Use Long or Short
	Protection Mechanism (CTS-to-Self)	Enabled

Table 2: Available modes in the Proxim 802.11g products

### 802.11b Only Mode

When in 802.11b mode, the AP will always operate with a 20  $\mu$ sec slot time. This is required since 802.11b clients can only recognize the 20  $\mu$ sec slot time. In this mode the Proxim APs will be able to respond to clients using either the long or the short preamble. The long preamble is only used when a client joins the network that can only use the long preamble.

### 802.11g Only Mode

In the 802.11g mode, only 802.11g modulation (OFDM) will be used, restricting associations to 802.11g clients. Therefore, the short slot time of 9  $\mu$ sec will be used and the only preamble that will be relevant is the OFDM preamble. And, finally, there is no need for a protection mechanism because the traditional 802.11 medium access mechanisms will manage the sharing of the wireless media. This is the configuration that will provide the maximum throughput for 802.11g devices because, as described above, there are no throughput degradations caused by long slot times, 802.11b devices, or protection mechanisms.

### 802.11b/g Mixed Mode

The 802.11b/g mixed mode is the most complex of the three modes, because it must accommodate all of the complexities described above. Proxim's 802.11b/g mixed mode provides the optimal functionality for the case in which both types of wireless devices are present at any given time.

If the AP detects any 802.11b devices, it will automatically fall back to the use of 20  $\mu$ s slot times. Since it provides the maximum throughput, slot times of 9  $\mu$ s will be used when no 802.11b devices are detected. In the same way, the AP will only use the CTS-only protection mechanism when

802.11b clients are present in the network. Again, when there is no mixture of clients, the CTS-protection mechanism provides an unnecessary overhead. Finally, when there is a mixture of 802.11b and 802.11g clients, the network will still use both the long or short preamble, using the long preamble only when a client unable of supporting the short preamble joins the network.

### Mode Performance

In Figure 15 we illustrate the performance of these various modes.

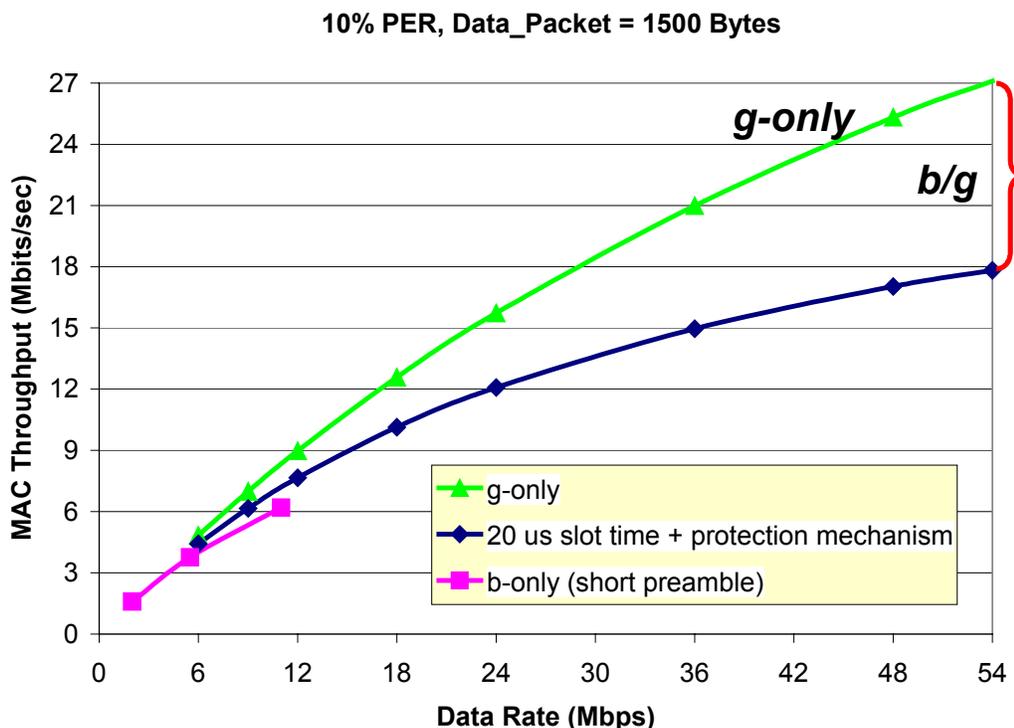


Figure 15: Performance of clients when operating in each of the three Proxim modes

This figure shows that Proxim’s g-only mode offers dramatically improved performance relative to b/g mode, so that in the case in which there are only 802.11g clients on the network this mode offers enhanced performance.

### Two Slot Architecture Migration

The Proxim two-slot access point allows a single access point to offer two channels of service, in a variety of configurations. One of the important advantages that this type of access point architecture has offered in the past is the ability to help WLAN users migrate from one WLAN standard to another. For example, when 802.11a access points became available, customers with 802.11b networks could get 802.11a coverage by using one of each type of technology in the two access point slots. The two-slot access point can play an important role in helping customers incorporate 802.11g technology into their 802.11b networks.

The two-slot architecture deployment model is simple to understand, and is illustrated in Figure 16.

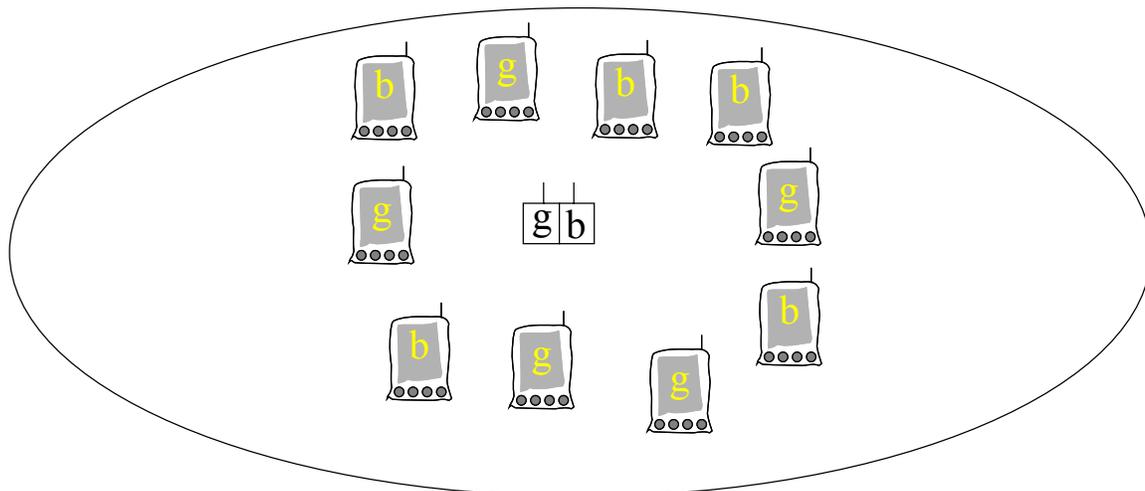


Figure 16: Schematic diagram of 802.11b and 802.11g devices being serviced by a single, dual-slot access point

In deployments in which 802.11g is being used to increase the total system throughput, a configuration like that shown in

Figure 16 can be used to get the most benefit out of the 802.11g technology. A simple calculation explains why.

As shown in Figure 15 above, an 802.11g device operating in mixed mode operates at a maximum of 18 Mbps, and for an 802.11b device using the short preamble, the equivalent rate is about 6 Mbps. The weighted average of the 11b and 11g mixed-mode rates (since the 802.11b device takes 3× longer on the channel) is  $(3 \times 6 + 18) / 4 = 9$  Mbps. Therefore, if both slots of a dual-slot access point were to be operated in mixed-mode, the total network throughput would be about 18 Mbps.

On the other hand, the maximum data rate for an 802.11g device in g-only mode is 27 Mbps, so if one slot were operated in b-only mode, and the other slot were operated in g-only mode as shown in the figure, the total throughput would be closer to 33 Mbps, an increase by a factor of 1.8×. 802.11b clients would only be permitted to associate with the 802.11b side of the AP, and similarly for the 802.11g clients.

The power of this dual-slot access point architecture can truly be seen when examining the enhanced performance of the 802.11g devices. In the mixed mode case, the 802.11g devices share a fraction of a 9 Mbps channel. In the dual-slot architecture case, they share a fraction of a 27 Mbps channel, which means that they see their data rate increased by 3×. The dual slot architecture allows users to take full advantage of the new technology, while still supporting legacy users as well.

## Summary

Customer reaction to pre-standard 802.11g products has already been very favorable, which augurs well for the continued adoption of 802.11g devices as the standard becomes ratified. The deployment of this new technology is not as straightforward as it may first appear, however. With an understanding of some of the complexities it is possible to offer certain suggestions as to how users can get the most benefit out of 802.11g.

- In controlled networks, in which it is not necessary to support legacy devices, access points should be configured to run in Proxim's 802.11g-only mode. This mode will offer the highest throughput scenario.
- In uncontrolled networks, or networks in which legacy devices need to be supported, two operational modes are possible.
  - If throughput is not a major concern, and the best performing, least expensive, solution is desired, Proxim's single-slot 802.11g access point should be used in its mixed, b/g mode. This mode will support both 802.11g and 802.11b devices, and will automatically adjust its use of preambles, slot times, and protection mechanisms to get the best performance depending on the client population it is supporting.
  - When throughput is a primary concern, Proxim's dual-slot access point should be used to segregate the 802.11b traffic from the 802.11g traffic. While this configuration still supports both new and legacy users, it offers as much as a 3× improvement in system throughput over a mixed-mode AP configuration.