

Defining and Improving Data Throughput in Wireless LAN

Yoram Solomon, Director of Business Development

Texas Instruments Wireless LAN Business Unit

Abstract

The amount of information that can be transmitted across an 802.11 interface is often interchangeably referred to as “signaling rate,” “data rate” or “throughput.” These terms all have different meanings and are not interchangeable. The 802.11b standard is considered to have a signaling rate of 11 Mbps, while the 802.11a and 802.11g standards are considered to have up to 54 Mbps.

In this article, a distinction will be made between these terms and show that the *signaling rates* of these standards are 11 Mbps and 54 Mbps, respectively, but that the actual *data throughput*, which is the real care about, is much lower than that. Suggestions for improving the data throughput will be presented, especially for the home network consumer space.

Signaling Rate

The signaling rate of 802.11 is defined as:

$$\text{Signaling Rate} = \frac{1}{\text{Bit Time}}$$

The bit time is the time it takes to transmit one bit of information. The signaling rate of 802.11b can range from as little as 1 Mbps up to 11 Mbps, depending on the specific transmission technique used. The 802.11a standard and 802.11g draft standard can offer signaling rates of up to 54 Mbps, five times faster than 802.11b.

Data Throughput

The 802.11 standard, with all its derivatives (the 802.11b, 802.11a, and the 802.11g draft) define physical layers (PHYs) as well as the protocol used for communications. In other words, the standard describes what format the information being transmitted needs to be in to be received and understood at the other end. These protocols include a significant amount of overhead. Overhead is a general name for all information and airtime used for anything that is not data. Overhead in an 802.11 system includes:

- Protocol headers that carry information required to control the device
- Frame Interval times that allow users of the frequency band to gain access to the communication medium (the air)
- Error and flow control to assure integrity of the transmissions
- Acknowledgment of received messages due to the relatively high probability of error messages and message collisions using the wireless medium.

Given that 802.11b includes both data and the overhead discussed above, the expected data throughput would be close to 5.5 Mbps rather than the signaling rate of 11 Mbps, and the expected throughput of an 802.11a or 802.11g product will be close to 30 Mbps rather than the signaling rate of 54 Mbps.

802.11 Standards

Parameter	802.11b	802.11a	802.11g
Highest signaling rate	11 Mbps	54 Mbps	54 Mbps
Mandatory modulation techniques	DSSS, CCK	OFDM	OFDM
Frequency bands	2.4 GHz ISM	5 GHz ISM and U-NII	2.4 GHz ISM
Bandwidth allocated	83.5 MHz	200 MHz for indoor use and 100 MHz for outdoor use. More being allocated	83.5 MHz
Number of non-overlapping channels available	3	12 (8 indoors plus 4 outdoors)	3
Backward compatibility	Baseline standard	None. Using different bands	With 802.11b products

There are currently three extensions to the original 802.11 standard, specifically 802.11b, 802.11a and 802.11g draft version. The original 802.11 standard was published in 1997. 802.11 provided data rates of only up to 2 Mbps. Two new, but incompatible versions of the standard, 802.11b and 802.11a, emerged in response to improve the data rate. 802.11b operates in the 2.4 GHz Industrial-Scientific-Medical (ISM) band and provides signaling rates of up to 11 Mbps. The “a” version works in the 5 GHz Unlicensed National Information Infrastructure (U-NII) band and provides signaling rates up to 54 Mbps. 802.11a equipment is not backward compatible with 802.11b equipment, meaning that all new equipment must be purchased and installed. The 2.4 GHz band is also used by many common household devices, including cordless phones, microwave ovens and baby monitors. These “other users” can cause interference and prevent an 802.11b user from being able to access the network. While

802.11b devices can transfer data at distances of up to 300 feet, 802.11a devices can only transmit up to 150 feet, due to the higher “path loss” characteristic of this higher frequency, requiring more access points in the same space. In addition, 802.11a devices require more power to transmit information in the 5 GHz band than an 802.11b device would require in the 2.4 GHz band, due to the lower efficiency of power amplifiers in higher frequencies.

In response to these issues between 802.11b and 802.11a, the IEEE started to work on the 802.11g extension in 2000. 802.11g operates in the same 2.4 GHz band as 802.11b, but provides the signaling rate of 802.11a using the same modulation technique used for 802.11a, Orthogonal Frequency Division Multiplexing (OFDM). 802.11g is also backward compatible with 802.11b, allowing users to continue to use older equipment without upgrading. Since 802.11g operates in the 2.4 GHz band like 802.11b, it too suffers from interference problems with other common household devices. During a period of transition from 802.11b to 802.11g, many WLANs will support the high speed signaling rate as a first option but fall back to 802.11b speeds when dictated by the equipment or conditions.

Mixed Mode 802.11g and 802.11b

There are currently some 20 million users of 802.11b networks and as the market moves to the higher data rates of 802.11g, these users need to be taken into account. The IEEE 802.11g draft standard intends to improve data throughput in the 2.4 GHz band. Task Group G adopted a draft that uses the same modulation technique as the 802.11a standard, OFDM, with the same, short preamble used for that standard. However, 802.11g is in the 2.4 GHz band. While the 802.11g draft requires backward compatibility with 802.11b standard, this compatibility is somewhat questionable. One would expect that “backward compatibility” would mean that both standards, using the same channel in the 2.4 GHz band, could co-exist. In fact, they cannot coexist without specific provisions that need to be used.

The reason for this inability to coexist is twofold: first, the “legacy” 802.11b receivers cannot decode OFDM data. Therefore, those receivers, whether access points or stations, will not be able to sense that an 802.11g transmitter is transmitting on the same channel. Second, the “short preamble” transmitted by an 802.11g transmitter, whether access point or station, is different than the one used with the 802.11b standard. A 802.11b device, wishing to transmit, will first “listen” to the channel, as required by the nature of the CSMA/CA protocol, and if it cannot detect a valid transmission, will assume that the airwaves are “clear” and transmit. However, since the 802.11b device will be unable to detect an 802.11g client transmitting due to the two reasons described above, it may incorrectly assume that the airwave is clear, and transmit “over” the transmission of the 802.11g client. Typically, both transmissions will be lost, and retries will be required by both devices, significantly lowering the system throughput. An 802.11g device cannot make the same “mistake,” since it is, by definition, backward compatible with the 802.11b standard, thus capable of sensing an

802.11b device and avoiding transmitting simultaneously with it. Due to the first problem, though, the aggregate system throughput will be hurt.

One possible solution is not to allow mixed mode 802.11g and 802.11b networks. This way, supporting both 802.11g and 802.11b will require the use of two channels. Given the fact that the 2.4 GHz band allows for only three non-overlapping channels, this will take 2/3 of the available spectrum. An even worse problem with this solution is that the cost of supporting both standards will significantly increase, since two radios will be required to support the simultaneous operation in two channels.

Improving Data Throughput

Another solution is to mandate a network coordination function, as proposed by the quality of service (QoS) IEEE 802.11 Task Group E. The 802.11e draft standard aims to improve the quality of service of wireless LAN such that it can be used to transfer real-time media, such as voice and video, without suffering from the “best effort” service level provided by the current 802.11 family of standards which were designed as data-centric. While this solution can solve the mixed-mode 802.11g-802.11b dilemma in the future, it cannot solve it for legacy 802.11b clients, not supporting the 802.11e draft standard.

Yet another solution exists within the basic 802.11 standard. The standard allows an access point to prevent stations from transmitting by requiring them to transmit a request-to-send (RTS) packet, and wait for a clear-to-send (CTS) packet from the access point before continuing with transmitting their payload data packets. This provision was included in the 802.11 standard to address a situation in which two stations are within coverage of the access points, but not within coverage of each other. In this case, while they cannot “hear” each other, they might transmit at the same time, and the access point will not be able to decode either message. Using this mechanism to solve the mixed-mode dilemma is viable. When used to protect 802.11g OFDM traffic from 802.11b transmissions, this technique is often referred to as “Protection Mechanism.” However, using it will create more system overhead (with the RTS-CTS packets), lowering the overall system throughput and once again defeating the benefits of using the higher-rate 802.11g.

So, what can be done now to improve the overall system throughput, while supporting legacy 802.11b clients, and maintaining a cost-effective solution? The answer really lies in segmentation of the wireless LAN market, and understanding the buying criteria and process in the different market segments. Interoperability and throughput are two driving factors in every market segment. Everyone wants their equipment to operate anywhere they want at the speed that they want. Interoperability refers to the ability of equipment from different manufacturers to work together in the same network.

The Home / Consumer / Retail Market Segment

In the home market, wireless networking equipment is typically purchased in “kits.” The buyer will decide to install a wireless network, and in order to avoid potential incompatibilities, will buy both the access point and network cards at the same place, at the same time, and from the same vendor. Retail chains will typically offer kits that include both access points and network cards in either one package, or under one promotion plan. While this may solve the interoperability problem in the actual home environment, the consumer will definitely want to be able to use the wireless card in other locations (typically in the public space), and will definitely want to use the wireless card installed in his or her “work computer” by the IT department, in their home environment.

Throughput is becoming an increasingly important requirement in the home market as consumers look to connect rich content devices at the highest possible speed. For example, photographs and high-quality video streaming will be downloaded over the WLAN from a camera to a PC. Having a high data throughput is critical in applications like this.

The Enterprise Market Segment

One key selection criterion for the enterprise market is interoperability. The reason lies in the fact that the access points are bought separately from the network cards. Once the wireless access point infrastructure is installed, the selection of the network cards for the different devices does not have to adhere to the same process. The network cards will not necessarily be bought from the same vendor as the access points, and will in-fact typically be acquired from a different vendor. Furthermore, the wireless access point infrastructure may be upgraded from time to time, so even if both access points and network cards were originally bought from the same vendor, after such upgrade they will not necessarily remain from the same vendor. Interoperability is therefore a key selection criterion for the enterprise, and the Wi-Fi Alliance is working hard at assuring it across vendors.

Throughput is once again an important factor in the enterprise market segment because of the sheer volume of traffic typically seen on the network. Most enterprise IT managers will opt for providing the highest guaranteed speed possible, per user, for any communication equipment.

The Public Access Market Segment

Much like the enterprise market, but for different reasons, the public access market is mostly concerned with interoperability. There is no way in the public access market to ensure that both access points and network cards are purchased from the same vendor. Users will use the same network cards they have installed, from multiple different vendors, regardless of the vendor that was selected by the public-access service

provider as a source for the access points. Any attempt of the service providers to demand that the users use network cards from specific vendors will only limit market growth for public access wireless LAN. Furthermore, different service providers will most likely use different access point vendors, thus adding another dimension of inconsistency between access point vendors and network card vendors. Interoperability, therefore, is even more important in the public access market segment than in the enterprise market segment.

Data throughput will be important for the public access market as well. The majority of the users will be using these networks almost solely for Internet access requiring a high level of data throughput.

An Interoperable Solution for the Home Environment

In the following, methods are proposed to improve the “real throughput” in the *home* environment in a cost-effective manner, but in a way that will not hurt interoperability. The characteristics of such solutions must be:

- (a) The ability to interoperate in any 802.11b environment without special requirements and provisions such as quality of service;
- (b) The ability to provide higher throughput in the home environment, while allowing other “legacy” 802.11b network cards to interoperate in such network without significant throughput degradation that can be expected with a mixed-mode 802.11g-802.11b network;
- (c) The ability of the wireless card to roam among other service locations, such as “hot spots” of public access, with full interoperability with such services, and without degrading the service level or throughput at those locations;
- (d) An assumption that can be made is that both the access point and network card were purchased from the same vendor, and thus may include some “proprietary technology” or optional features not mandatory for the standard implementation in both ends, subject to the above two characteristics / conditions.

Packet Binary Convolutional Coding (PBCC) was developed as a method to increase the signaling rate for wireless networks from the 11 Mbps supported by 802.11b to 22 Mbps, while maintaining full co-existence and interoperability with legacy 802.11b systems. This modulation/coding technique was offered to the 802.11g task group, and while it was rejected as the mandatory high-speed modulation technique for that draft standard, it was accepted as an optional one. It should be noted that even the 802.11b standard has PBCC, although limited to 11 Mbps, as an optional modulation method. Sometimes PBCC is referred to as 802.11b+ in the marketplace.

A key characteristic of PBCC is that it uses exactly the same header as the legacy 802.11b device, unlike 802.11g which uses an OFDM header, undecodable by legacy 802.11b devices. This means that the legacy client can decode the PBCC packet header. This header has a field, as defined in the original 802.11 standard, called the “Duration” field. This field has the sole purpose of letting a receiver know how long (in microseconds) the currently transmitted packet will be. In other words, even if the receiver cannot decode the rest of the packet, whether due to interference, or simply because the data is encoded using PBCC or another modulation technique unfamiliar to this receiver, it will know how long to wait, and when it can transmit again, such that it will not collide with this current transmission.

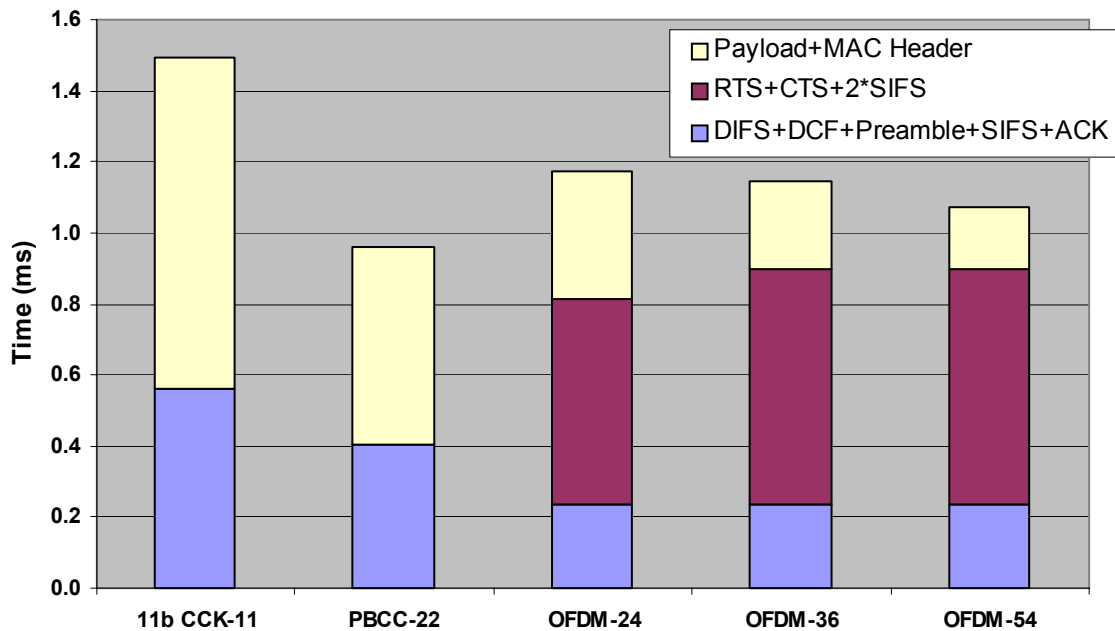
PBCC meets requirement (a), since it allows legacy 802.11b stations, “PBCC enhanced” 802.11b stations, and 802.11g stations using the optional PBCC mode to interoperate without requiring special quality of service or coordination function provisions. This is achieved by using the same header, decodable by all three types of stations (legacy 802.11b stations, “PBCC enhanced” 802.11b stations and 802.11g stations using the optional PBCC mode), and the “Duration” field within the header. It meets requirement (b), by offering higher throughput through an increased signaling rate, without hurting such “mixed-mode” operation. It also meets requirement (c), since a PBCC-enabled network card will use the standard, interoperable 802.11b mode when communicating outside of the PBCC-enabled home network. This can all be achieved based on assumption (d), since in the consumer market both the access point and wireless cards will be most likely purchased from the same vendor, typically using chipsets from the same vendor, and thus offering PBCC in both ends of the application.

The last question to be answered is whether using PBCC offers a comparable throughput to what 802.11g-based products offer, when supporting a mixed-mode operation with legacy 802.11b products. In the following charts, created by Texas Instruments Wireless LAN Lab, two simulations of the different alternatives for mixed mode operation are offered. The charts compare the time it takes to transmit a 1,000-Byte packet using different techniques. The times include not only the “payload” data, but also all the overhead required to transmit them, and co-exist with legacy 802.11b equipment in the network. The shorter the time bar is, the faster it can transmit the same amount of data, and the highest real data throughput is achieved by the network. The first chart presents the results in an environment using a long preamble (used to synchronize receivers to transmitters), while the second chart presents the results when using a short preamble. It should be noted that some legacy equipment might support only long preamble, and thus supporting it may require the use of long preamble throughout the network.

It is immediately apparent from the charts that only in the legacy 802.11b devices and PBCC devices does the transmission of data take more than 50% of the overall transmission time. In the OFDM devices (rates simulated are 24, 36, and 54 Mbps) the data takes a much smaller portion of the overall transmission time. In fact, when transmitting OFDM at 54 Mbps, it appears that less than 20% of the time is used to

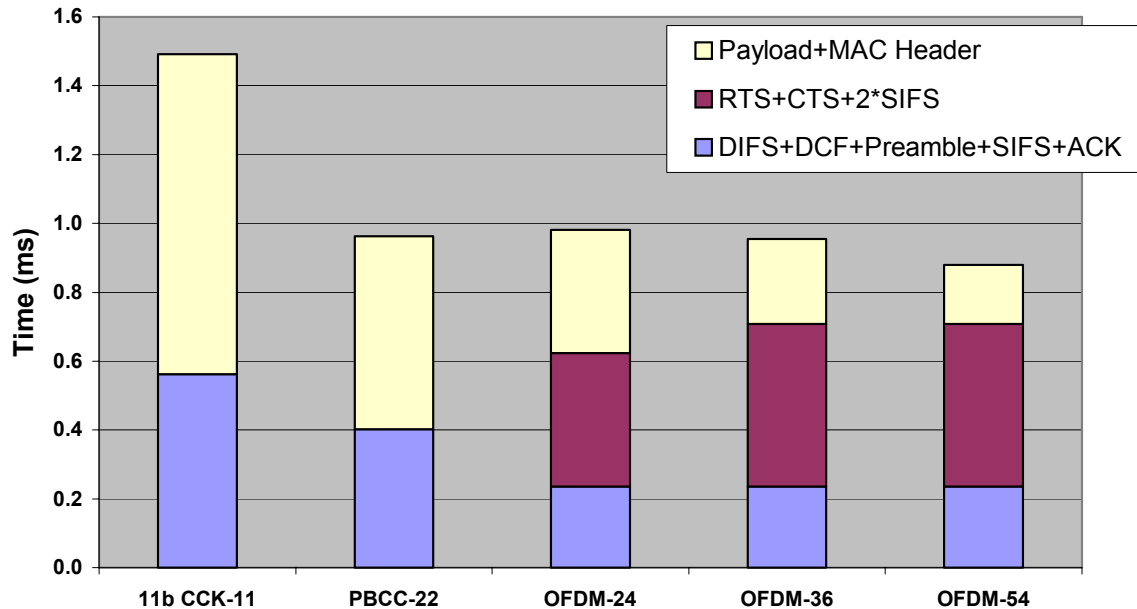
transmit “payload” data, whereas the rest of the time is used for the overhead and protection mechanisms (not required for PBCC devices).

Chart 1: Time to transmit 1,000-Byte blocks using a long preamble



It can be seen that only when using PBCC 22 does it take less than 1 millisecond to transmit one 1,000-Byte data block, representing a data throughput of over 8 Mbps, compared to an 802.11g device using 54 Mbps and RTS-CTS protection mechanism, taking approximately 1.1 millisecond, representing just over 7 Mbps.

Chart 2: Time to transmit 1,000-Byte blocks using a short preamble



When using short preamble, assuming that the legacy 802.11b devices associated with this network are capable of supporting them, there is a slight advantage for using 802.11g OFDM rate of 54 Mbps over PBCC, but this advantage is less than 5%, and gets nullified when the signaling rate falls back to 36 Mbps (typically due to increased distance between the access point and the station).

Summary

If you try to send a 54-Mbit file over a 54-Mbps 802.11a or 802.11g link, it will take more than a second. The real data throughput of the wireless media is significantly lower than the signaling rate and should not be confused. However, something can be done about this, to get the data throughput as close to the signaling rate as possible. While the enterprise and public access market segments prevent any improvements of the protocols due to the interoperability requirement, the home market segment can, as long as it meets the requirements described in this article. The most important requirement being that it maintains interoperability in different environments. PBCC increases the signaling rate for wireless networks from the 11 Mbps supported by 802.11b to 22 Mbps, while maintaining full co-existence and interoperability with legacy 802.11b systems and without the throughput degradation due to the use of “protection mechanisms” required for 802.11g OFDM devices. Finally, as the WLAN standards continue to evolve and new technological innovations are made, providing that the main focus is on real throughput and not only the signaling rate, data throughput will continue to improve.

© 2003 Texas Instruments Incorporated

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.