IEEE P802.11
Wireless LANs

GBT9 Throughput Efficiency

Date: January 16, 1998

Author:
Darrol Draper
Golden Bridge Technology
185 Route 36, West Long Branch NJ 07764
Phone: 732 870-8088
Fax: 732 870-9008
e-Mail: goldenbridge@mail.monmouth.com

Abstract

The need for a variable data rate for WLAN’s is discussed, and is considered crucial for maximum use of the 2.4 GHz band. A comparison is made to phone-line modems, which negotiate their data rate, for compatibility with older and less expensive modems while using the highest practical data rate.

Based on data presented in Doc: IEEE P802.11-98/28 , three tables show how maximum throughput is obtained by varying the modulation rate as the RF channel Trms changes. Because the overhead time for each packet is a real limitation on high throughput, a Short Header is proposed, and Table 3 shows the higher throughput.

Practical considerations of the channel estimate algorithm are discussed. This algorithm operates over a long period, and thus is not a large burden on a microprocessor, compared to channel equalizers.

The need for a variable and negotiated data rate

Every radio channel imposes constraints on maximum data rates. The maximum theoretical data rate and the possible data rate of actual equipment both depend on the channel characteristics. It is desirable to be able to use a high data rate when the channel is excellent, and fall back to a slower, more robust implementation when the channel is not so good. Seen in this light, the “best” system is one which gives the highest data throughput under “average” conditions. Excellent performance at one data rate may not allow for a high average data throughput, if this high data rate can only be used infrequently.

Throughput is net data transferred, as distinguished from the gross data rate. It does not include dropped packets, and sometimes does not include data which is only of use to a lower layer and is discarded by a higher layer protocol. This paper examines the effect of dropped packets and does not consider the physical layer header as data for determining the throughput rate.

Industry practice for phone-line modems is to negotiate the highest possible data rate, based on the characteristics of the channel and on the capabilities of each modem. So too, we believe, an evolved standard for wireless modems should allow for the negotiation of the highest practical data rate based on the actual characteristics of the RF channel and on the capabilities of each wireless modem.

The ability to negotiate a maximum data rate is, arguably, more important for wireless modems than it is for phone-line modems. First, there is a much greater variation in the maximum possible data rate for RF channels in typical WLAN environments, compared to most phone lines. So a negotiated data rate can increase the throughput for wireless more than it can for wired channels. Secondly, the RF channels are shared among users, and this places a limit on the number of simultaneous users and on the types of applications which are practical for RF links. Higher data rates mean less channel occupancy time for the same data messages, and permit more users on the same channel. This is a large issue in changing WLAN's from a niche into a widely used technology. 2.4 GHz spectrum is relatively limited. Higher data rates allow more users, and this generates a virtuous circle of volume production, lower costs, and even more end users. The third reason is that a maximum data rate, even when a channel is dedicated to a single pair of users, places a limit on the types of applications which are practical. The throughput of UTP cable has been increased beyond what engineers thought possible a decade ago, and today UTP has enough bandwidth to support any application, including video. This statement can not yet be made for wireless. However a negotiated data rate is one step in that direction.

Case One: Packet Length = 1000 bytes, 192 byte BPSK header & preamble

Shown in italics: optimum number of GBT9 codes, for maximum throughput, for inter-packet delays of 0, 0.1, 0.25, 0.5, 1 and 2 ms

	0 ms		0.1 ms		0.25 ms		0.5 ms		1 ms		2 ms	
T rms	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s
14	*12*	14.01	*12*	11.87	*12*	9.66	*12*	7.37	*12*	5.00	*12*	3.05
16	*12*	13.63	*12*	11.55	*12*	9.40	*12*	7.17	*12*	4.87	*12*	2.96
18	*12*	13.07	*12*	11.08	*12*	9.01	*12*	6.88	*12*	4.67	*12*	2.84
20	*12*	12.42	*12*	10.53	*12*	8.57	*12*	6.54	*12*	4.44	*12*	2.70
22	*12*	11.57	*12*	9.80	*12*	7.98	*12*	6.09	*12*	4.13	*10*	2.53
24	*12*	10.63	*12*	9.01	*12*	7.33	*12*	5.59	*10*	3.81	*10*	2.36
26	*12*	9.78	*12*	8.29	*12*	6.75	*12*	5.15	*10*	3.50	*7*	2.21
28	*12*	8.79	*12*	7.45	*12*	6.06	*10*	4.64	*7*	3.26	*7*	2.10
30	*12*	7.84	*12*	6.64	*12*	5.41	*8*	4.21	*7*	3.04	*6*	1.98
32	*12*	7.20	*12*	6.10	*12*	4.97	*8*	3.96	*7*	2.89	*3*	1.91
34	*12*	6.50	*12*	5.50	*7*	4.65	*7*	3.77	*7*	2.73	*3*	1.86
36	*12*	5.89	*10*	5.02	*7*	4.31	*7*	3.49	*6*	2.56	*3*	1.81
38	*10*	5.29	*10*	4.56	*7*	3.95	*6*	3.21	*4*	2.42	*3*	1.76
40	*10*	4.91	*10*	4.23	*7*	3.67	*6*	3.00	*3*	2.35	*3*	1.71
42	*10*	4.52	*10*	3.90	*6*	3.40	*3*	2.83	*3*	2.30	*3*	1.67
44	*10*	4.23	*8*	3.68	*6*	3.21	*3*	2.78	*3*	2.26	*3*	1.64
46	*10*	3.92	*8*	3.43	*6*	3.01	*3*	2.73	*3*	2.21	*3*	1.61
48	*8*	3.58	*8*	3.15	*5*	2.82	*3*	2.67	*3*	2.17	*3*	1.57
50	*8*	3.25	*8*	2.87	*4*	2.79	*3*	2.61	*3*	2.12	*3*	1.54

T rms (in nanoseconds) represents the channel multipath (ref: doc IEEE P802.11-97/157r1), as a measure of impairment.

Mb/s is net throughput, so it does not include the header.

Table 1 Variable Data Rates for 1000 byte packets and 192 byte header

Case Two: Packet Length = 64 bytes, 192 byte BPSK header & preamble

Shown in italics: optimum number of GBT9 codes, for maximum throughput, for inter-packet delays of 0, 0.05, 0.1, 0.25, 0.5, and 1 ms

	0 ms		0.05 ms		0.1 ms		0.25 ms		0.5 ms		1 ms	
T rms	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s
14	*12*	2.35	*12*	1.91	*12*	1.60	*12*	1.09	*12*	0.71	*12*	0.42
16	*12*	2.32	*12*	1.88	*12*	1.59	*12*	1.07	*12*	0.70	*12*	0.41
18	*12*	2.27	*12*	1.84	*12*	1.55	*10*	1.05	*12*	0.68	*8*	0.41
20	*12*	2.21	*12*	1.79	*12*	1.51	*10*	1.03	*12*	0.67	*8*	0.40
22	*12*	2.14	*12*	1.74	*12*	1.46	*10*	1.00	*7*	0.66	*5*	0.40
24	*12*	2.06	*10*	1.68	*12*	1.41	*5*	0.98	*7*	0.64	*5*	0.39
26	*10*	1.98	*8*	1.62	*8*	1.37	*5*	0.96	*5*	0.64	*4*	0.39
28	*7*	1.91	*7*	1.57	*5*	1.35	*4*	0.95	*4*	0.64	*3*	0.39
30	*7*	1.85	*5*	1.54	*5*	1.32	*4*	0.93	*3*	0.64	*3*	0.39
32	*7*	1.80	*5*	1.51	*5*	1.30	*3*	0.93	*3*	0.63	*3*	0.39
34	*4*	1.77	*4*	1.49	*4*	1.28	*3*	0.92	*3*	0.63	*3*	0.38
36	*4*	1.75	*4*	1.47	*3*	1.27	*3*	0.91	*3*	0.62	*3*	0.38
38	*4*	1.72	*3*	1.45	*3*	1.26	*3*	0.91	*3*	0.62	*3*	0.38
40	*4*	1.69	*3*	1.44	*3*	1.25	*3*	0.90	*3*	0.61	*3*	0.37
42	*3*	1.67	*3*	1.42	*3*	1.24	*3*	0.89	*3*	0.61	*3*	0.37
44	*3*	1.67	*3*	1.42	*3*	1.23	*3*	0.89	*3*	0.61	*3*	0.37
46	*3*	1.64	*3*	1.40	*3*	1.22	*3*	0.88	*3*	0.60	*3*	0.36
48	*3*	1.63	*3*	1.39	*3*	1.21	*3*	0.87	*3*	0.59	*3*	0.36
50	*3*	1.61	*3*	1.37	*3*	1.19	*3*	0.86	*3*	0.58	*3*	0.36

T rms (in nanoseconds) represents the channel multipath (ref: doc IEEE P802.11-97/157r1), as a measure of impairment.

Mb/s is net throughput, so it does not include the header.

Table 2 Variable Data Rates for packets with 64 byte payload and 192 byte header

Case Three: Packet Length = 1000 bytes, Short Header & preamble

Shown in italics: optimum number of GBT9 codes, for maximum throughput, for inter-packet delays of 0, 0.05, 0.1, 0.25, 0.5, and 1 ms

	0 ms		0.05 ms		0.1 ms		0.25 ms		0.5 ms		1 ms	
T rms	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s	*codes*	Mb/s
14	*12*	18.08	*12*	16.20	*12*	14.67	*12*	11.44	*12*	8.36	*12*	5.44
16	*12*	17.59	*12*	15.76	*12*	14.27	*12*	11.13	*12*	8.14	*12*	5.29
18	*12*	16.87	*12*	15.11	*12*	13.69	*12*	10.67	*12*	7.80	*12*	5.08
20	*12*	16.03	*12*	14.36	*12*	13.01	*12*	10.14	*12*	7.42	*12*	4.82
22	*12*	14.93	*12*	13.38	*12*	12.12	*12*	9.45	*12*	6.91	*12*	4.49
24	*12*	13.72	*12*	12.29	*12*	11.13	*12*	8.68	*12*	6.35	*12*	4.13
26	*12*	12.62	*12*	11.31	*12*	10.24	*12*	7.98	*12*	5.84	*12*	3.80
28	*12*	11.35	*12*	10.16	*12*	9.21	*12*	7.18	*12*	5.25	*10*	3.48
30	*12*	10.12	*12*	9.06	*12*	8.21	*12*	6.40	*10*	4.72	*7*	3.26
32	*12*	9.29	*12*	8.32	*12*	7.54	*12*	5.88	*7*	4.41	*7*	3.11
34	*12*	8.38	*12*	7.51	*12*	6.80	*12*	5.30	*7*	4.16	*7*	2.93
36	*12*	7.61	*12*	6.81	*12*	6.17	*10*	4.85	*7*	3.86	*6*	2.74
38	*12*	6.80	*12*	6.09	*12*	5.52	*10*	4.41	*7*	3.54	*4*	2.56
40	*12*	6.16	*10*	5.57	*10*	5.11	*10*	4.09	*6*	3.29	*3*	2.47
42	*10*	5.64	*10*	5.13	*10*	4.71	*7*	3.83	*6*	3.07	*3*	2.41
44	*10*	5.28	*10*	4.80	*10*	4.40	*7*	3.60	*3*	2.96	*3*	2.37
46	*10*	4.90	*10*	4.45	*10*	4.08	*7*	3.37	*3*	2.90	*3*	2.32
48	*10*	4.46	*10*	4.05	*10*	3.72	*6*	3.14	*3*	2.84	*3*	2.27
50	*10*	4.05	*10*	3.69	*10*	3.38	*6*	2.90	*3*	2.77	*3*	2.22

T rms (in nanoseconds) is channel multipath (ref: doc IEEE P802.11-97/157r1).

Preamble is 32 symbols of 1 Mb/s BPSK. Header is 64 bits, encoded at 1.83 Mb/s (i.e. one code).

Table 3 Variable Data Rates for 1000 byte packets and 67 µs Short Header

Discussion of the Tables

Reference is made to data presented in Doc: IEEE P802.11-98/28 and which is not reproduced here. These represent the performance of GBT9 at a variety of Trms and data rates. From those, Tables 1, 2 and 3 were developed. These consider packets of 1000 bytes and 64 byte payloads. Note that these charts consider the case of multipath without thermal noise, so thermal noise may give slightly different results.

Table 3 considers the possibility of a Short Header, using a 32 microsecond BPSK preamble, and a 64 bit header with a 1.83 Mb/s modulation rate. This rate is provided by using only one code, and is the lowest data rate for the GBT9 system. Thus, the header would represent a lowest common denominator for all 802.11 High Speed systems. Of course, backward compatibility with existing low speed 802.11 equipment would use the existing header. After the units have use the low speed header to establish a session, the subsequent packets can use the Short Header.

This preamble and header require a total of nearly 67 microseconds This does not represent the shortest possible header. A comparison between Tables 1 and 3 shows that considerable improvement in throughput is possible, especially in less impaired channels.

These tables illustrate that there is a “best” data rate, for each combination of channel Trms, packet length, and overhead in the form of headers and dead time between packets. Higher rates are more susceptible to channel impairments, and more packets are dropped. To some degree, this is not a problem. A system which has, say, a 50% higher data rate, can lose 33% of its packets and still have performance superior to the slower system which might only have a 5% PER.

802.11 has a special consideration, in that there are 192 microseconds of preamble and header on every packet. This creates a large incentive to not drop packets, as it lowers the throughput. This incentive looms larger for the higher rate systems. This effect can be seen in the tables by comparing throughputs for any constant Trms, but with overhead times which increase from columns on the left to the right.

A Forward Error Correction code (FEC) might improve results, but this is not included in the results. Higher data rate systems can afford the increased overhead of the FEC, as the net benefit is a lower PER.

What is an average channel Trms ?

Given that the goal is the highest average data rate over a variety of RF channels, and the need to choose a single high speed modulation method which will give the best performance most of the time, the question then becomes one of how to model the distribution of channel Trms's. Considering the range of environments over which 802.11 is expected to operate, the general form of a single best system is

$$K = \sum P_i (R_i) / n$$

K is the average throughput rate, in Mb/s
i represents discrete samples of the range of Trms.
n is the number of i samples of Trms.
P_i is the probability that each Trms will occur, and varies from 0 to 1.
R_i is maximum net data throughput possible, in Mb/s, for each Trms.

K is a figure of merit for the modulation method, as it defines throughput. K is more relevant if we place greater emphasis on, or consider only those RF which present higher likelihood of congestion by a large number of end users with high volumes of traffic.

There are two ways to consider R_i. One way considers the non-ideal nature of the channel estimate algorithm. This adds another layer of complexity. The other way simplifies things by ignoring the contribution of this algorithm, on the assumption that the algorithms are a separate issue and should not be a significant part of the evaluation of the best modulation method. To the extent that the possible modulation rates vary smoothly, the contribution of non-ideal algorithms does not, at first sight, appear to greatly favor one modulation system over another.

In a future paper, Golden Bridge Technology expects to review the literature on surveys of distribution of Trms, to draw some conclusions on the distribution of P_i.

Implementation of Channel Estimate Algorithms

GBT anticipates that channel estimate algorithms is an area for gradual development, and that the standard would allow for improved algorithms to be implemented without a change to the physical layer specifications. Actual algorithms will only approach making the best prediction of actual channel Trms. For example, the estimates will be approximations. Also, algorithms may have their own delay in tracking a changing RF environment. No doubt this will be an area for improving algorithms in the future. The important issue is that the High Speed 802.11 standard should accommodate such future growth, by providing for these possibilities.

It is important to note that channel estimating algorithms, as envisioned here, typically must examine at least a few packets, so these algorithms operate over a relatively long span of time. Channel equalizers, by comparison, must operate in real time and at a much higher speed. Therefore, channel estimate algorithms can be orders-of-magnitude less computationally intensive.

Most of the processing to determine the data rate for each mobile user would most likely be done in the Service Access Point. One reason is that the SAP has a greater volume of wireless traffic and thus a greater knowledge of the RF channels. Another is that this is more economical than having this software reside with every mobile user.

- end -