

On the Throughput of Bluetooth Data Transmissions

Matthew C. Valenti
Lane Dept. of Comp. Sci. & Elect. Eng.
West Virginia University
Morgantown, WV 26506-6109

Max Robert and Jeffrey H. Reed
Mobile and Portable Radio Research Group
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061-0350, USA

Abstract— Analytical expressions for the throughput (in kbps) as a function of channel symbol signal-to-noise ratio (E_s/N_0) are derived for the six Bluetooth ACL packets that use automatic repeat request (ARQ). The analysis is exact under the assumptions that the outer CRC code provides perfect error detection and that the channel remains stationary for the duration of each packet. Using an expression for noncoherent correlated ($h < 0.5$) full response FSK signals, numerical results are provided for AWGN and quasi-static Rayleigh fading channels. These curves are an appropriate benchmark against which practical demodulators and custom error control techniques may be compared.

I. INTRODUCTION

Data in Bluetooth can be transmitted asynchronously using ACL packets, of which seven types are defined in the specification [1]. All but one of these packet types (AUX1) use a cyclic redundancy check (CRC) code and automatic repeat request (ARQ) for error control. Furthermore, three packet types (DM1, DM3, and DM5) use a (15,10) shortened Hamming code for forward error correction (FEC) prior to error detection. Frequency hopping (FH) is used in such a way that the radio is tuned to the same frequency for the entire duration of the packet, but then changes to a different frequency each time it transmits a new packet or retransmits an erroneous packet. Since the fading and interference in the new frequency channel will be significantly different than that of the previous one, the use of FH with ARQ provides an effective method of diversity.

Because the duration of a single packet is short relative to the coherence time of the channel and the frequency hops from packet to packet, it is appropriate to model the channel using a *quasi-static fading* assumption. With quasi-static fading, the envelope of the signal associated with the entire packet is multiplied by the same channel gain (which is typically Rayleigh or Rician distributed), while the channel gains vary independently from packet to packet. The quasi-static assumption is an effective ap-

proximation to many typical operating environments and admits tight analytical predictions of the performance of an isolated piconet. However, due to the asynchronous nature of Bluetooth piconets, the quasi-static assumption might not be appropriate when there is a second piconet located in close proximity to the piconet under consideration [2].

In this paper, we are concerned with measuring the throughput, in terms of maximum achievable one-way data rate, as a function of the channel signal-to-noise ratio (SNR) for the six ACL packets that use ARQ. Analytical expressions are provided for each packet type as a function of the error probability of the channel ϵ , which is a receiver dependent parameter that can be determined by measurement, analysis, or simulation. The analysis is exact under the assumptions that the symbol error rate is constant throughout the duration of a packet and that the CRC is always able to detect errors in the packet's payload. Using values of ϵ that correspond to noncoherent detection of nonorthogonal full response FSK signals, throughput curves for both additive white Gaussian noise (AWGN) and quasi-static Rayleigh fading channels are presented. While these curves do not account for the losses due to the intersymbol interference (ISI) induced by the partial response nature of the GFSK signal, they can serve as a benchmark against which practical demodulation methods may be compared.

The remainder of this paper is organized as follows: In Section II the features of Bluetooth that are relevant to this study are discussed. In Section III the probability of retransmission is derived, which is used in Section IV to compute the throughput. Throughput results are given for noncoherent demodulation, first for the AWGN channel and then for the quasi-static Rayleigh fading channel.

II. FEATURES OF BLUETOOTH

Bluetooth uses Gaussian frequency shift keying (GFSK) modulation with a time-bandwidth product of $BT = 0.5$, modulation index $0.28 \leq h \leq 0.35$, and symbol rate of 1 Megabaud. Time is divided into 625 μsec slots, and an

This work was supported by the Office of Naval Research under grant N00014-00-0655 and by the MPRG affiliates program.

ACL packet may occupy 1, 3, or 5 consecutive slots. Packets are composed of a 72 bit access code, a 18 bit header, and a variable length payload. The access code is used for synchronization, DC offset compensation, and addressing. The header is protected by a triple redundancy code (resulting in 54 symbols). The payload is composed of a 1 or 2 byte payload header, data, and a 16 bit CRC for error detection. For the DM1, DM3, and DM5 packet types, the payload is protected by a (15,10) shortened Hamming code.

Time division duplexing (TDD) is used and a transmitter may only begin a packet transmission every other slot. Master devices must begin their transmission on an even indexed time slot, while slave devices must begin their transmissions on odd indexed slots. Each packet is transmitted at a different frequency chosen pseudorandomly from the set of 23 (Japan, France, and Spain) or 79 (elsewhere) 1 MHz RF channels. In the slot immediately following the received packet, the destination radio sends a positive acknowledgment (ACK) or negative acknowledgment (NAK) back to the source radio in the header of a return packet. Unless the source receives a valid ACK, it will retransmit the packet during its next slot. Packets are unnumbered, but a flag in the packet header indicates if it is a retransmission.

The six ACL packet types for which we wish to compute the throughput are as follows:

DM1: Occupies 1 forward slot. Payload consists of a 1 byte payload header, up to 17 bytes of data¹, and 16 bit CRC, all of which are protected by the (15,10) Hamming code resulting in a maximum of 240 code bits.

DH1: Also occupies 1 slot and contains 240 total bits but is not protected by the Hamming code. The payload consists of a 1 byte header, 27 bytes of data, and 16 bit CRC.

DM3: Occupies 3 slots. Payload consists of a 2 byte payload header, 121 data bytes, and 16 bit CRC. The payload is protected by the Hamming code resulting in 1500 code bits.

DH3: Also occupies 3 slots but is not FEC encoded. Payload consists of a 2 byte header, 183 bytes of data, and 16 bit CRC, resulting in a total of 1496 bits.

DM5: Occupies 5 slots. Payload consists of a 2 byte header, 224 data bytes, and 16 bit CRC. It is FEC encoded producing 2745 code bits.

DH5: Occupies 5 slots but is not FEC encoded. Payload contains 2 byte header, 339 data bytes, and 16 bit CRC, resulting in a total of 2744 bits.

Note that all of the above packets are followed immediately by a return packet and thus the effective number of slots used is 2, 4, or 6.

¹Since we are concerned with computing *maximum* throughput, for the remainder of the discussion we assume that every packet contains the maximum amount of data.

III. PROBABILITY OF RETRANSMISSION

For any particular receiver implementation and instantaneous signal-to-noise ratio, it is possible to determine the probability of retransmission. Retransmission occurs when any of the following events occur:

A: The destination radio fails to synchronize with the access code of the forward packet.

B: The header of the forward packet is corrupted after the triple redundancy code is decoded.

C: The payload of the forward packet is corrupted after the optional Hamming code is decoded, thus causing the CRC check to fail.

D: The source is unable to synchronize with the access code of the return packet.

E: The header of the return packet is corrupted (after decoding).

A packet is retransmitted if any of the events A-E occur, and thus the probability of retransmission is:

$$P_r(\gamma) = 1 - P[\bar{A}]P[\bar{B}]P[\bar{C}]P[\bar{D}]P[\bar{E}], \quad (1)$$

where \bar{A} indicates the complement of event A, $P[A]$ is the probability of event A, and $\gamma = \{\gamma_f, \gamma_r\}$ is a two-tuple containing the instantaneous forward and return signal-to-noise ratios.

Synchronization is achieved by correlating the output of the demodulator with a stored copy of the access code. The packet becomes synchronized if the output of the correlator exceeds a threshold. The threshold is chosen such that the frame will be synchronized if T of the 72 bits in the access code are demodulated correctly, where the parameter T is chosen according to the desired probability of false alarm. Since the packet will be successfully synchronized provided that there are no more than $(72-T)$ errors in the received access code,

$$P[\bar{A}] = \sum_{k=0}^{72-T} \binom{72}{k} (\epsilon(\gamma_f))^k (1 - \epsilon(\gamma_f))^{72-k}, \quad (2)$$

where $\epsilon(\gamma_f)$ is the symbol error probability of the forward channel. Since the return packet also uses a 72 bit access code, error event D has the same form,

$$P[\bar{D}] = \sum_{k=0}^{72-T} \binom{72}{k} (\epsilon(\gamma_r))^k (1 - \epsilon(\gamma_r))^{72-k}, \quad (3)$$

where $\epsilon(\gamma_r)$ is the symbol error probability of the return channel.

The header of a packet will be corrupted if any of the eight repetition coded triplets are decoded incorrectly. Since the (3,1) code can correct one error,

$$\begin{aligned} P[\bar{B}] &= (3\epsilon(\gamma_f)(1 - \epsilon(\gamma_f))^2 + (1 - \epsilon(\gamma_f))^3)^{18} \\ P[\bar{E}] &= (3\epsilon(\gamma_r)(1 - \epsilon(\gamma_r))^2 + (1 - \epsilon(\gamma_r))^3)^{18}. \end{aligned} \quad (4)$$

The most likely error event is type C. It is assumed that the CRC decoder is able to detect all uncorrectable errors. For the DHx packet types, this occurs whenever any of the payload bits are received in error and thus,

$$P[\bar{C}] = (1 - \epsilon(\gamma_f))^m, \quad (5)$$

where $m = 240$ for DH1, $m = 1496$ for DH3, and $m = 2744$ for DH5.

The payload of DMx packets is protected by a (15,10) Hamming code, which is capable of correcting one bit error per 15 bit code block. The payload is correctly decoded provided that all code blocks contain one or fewer errors, i.e.

$$P[\bar{C}] = (15\epsilon(\gamma_f)(1 - \epsilon(\gamma_f))^{14} + (1 - \epsilon(\gamma_f))^{15})^M, \quad (6)$$

where $M = 16$ for DM1, $M = 100$ for DM3, and $M = 183$ for DM5.

IV. THROUGHPUT ANALYSIS

Define the (geometric) random variable N to be the total number of times a particular packet must be transmitted. The probability that $N = n$ is the product of the probability of failure during the first $n - 1$ trials and the probability of success during the n th transmission. Thus the pmf of N is:

$$p_N[n] = E \left\{ (1 - P_r(\gamma_n)) \prod_{i=1}^{n-1} P_r(\gamma_i) \right\}, \quad (7)$$

where γ_i is the instantaneous SNRs for the i th pair of forward and return packets. Assuming that the channel is uncorrelated from packet to packet (quasi-static fading),

$$p_N[n] = E_\gamma \{ 1 - P_r(\gamma_n) \} \prod_{i=1}^{n-1} E \{ P_r(\gamma_i) \}. \quad (8)$$

Assuming that the pdf of the SNR is the same from packet to packet,

$$p_N[n] = (1 - \bar{P}_r(\Gamma)) (\bar{P}_r(\Gamma))^{n-1}, \quad (9)$$

where $\Gamma = E\{\gamma\}$ is the average SNR, and $\bar{P}_r(\Gamma)$ is the corresponding average channel error probability.

For any particular type of packet, the data rate R is a function of N . In particular,

$$R = \frac{K}{(DN)(625 \times 10^{-6})}, \quad (10)$$

where D is the number of occupied slots per transmission including the return packet (2 for Dx1, 4 for Dx3, and 6 for Dx5) and K is the number of data bits in the packet

(136 for DM1, 968 for DM3, 1792 for DM5, 216 for DH1, 1464 for DH3, and 2712 for DH5).

The average throughput R_{avg} is found by taking the expected value of R with respect to N ,

$$R_{avg} = E_N \{ R \}. \quad (11)$$

A. Performance in AWGN

To compute the throughput as a function of E_s/N_o , which is the ratio of energy per symbol to one-sided noise spectral density, (9)-(11) must be evaluated by using the proper values of $\bar{P}_r(\Gamma)$. In AWGN, the instantaneous and average SNRs are the same, $\Gamma = \gamma = E_s/N_o$ and thus $\bar{P}_r(\Gamma) = P_r(\gamma)$. Furthermore, reciprocity of the channel implies that $\epsilon(\gamma_f) = \epsilon(\gamma_r)$.

The value of error probability is highly dependent on the implementation of the receiver. However, due to the low cost requirements of Bluetooth, noncoherent receivers are used almost exclusively. A lower bound on error probability can be found by considering the performance of the noncoherent detection of full response binary correlated ($h < 0.5$) FSK signals. This is a lower bound because it does not account for the additional losses that accrue due to the intersymbol interference (ISI) induced by the use of partial response GFSK signaling. However, an exact analysis requires knowledge of the receiver implementation and must take into account both predetection and postdetection filtering. The reader that is interested in an exact analysis of practical demodulators suitable for Bluetooth is referred to [3]. We have found that using the results of [3] in place of the performance of full response signaling degrades performance by 5-10 dB, depending on the receiver implementation. It should be noted, however, that this loss applies to all packet types and thus the relative performance of the packets remains the same.

When the modulation index of FSK is less than $h = 0.5$, then the signals are correlated with the correlation given by [4]

$$\rho = \frac{\sin(2\pi h)}{2\pi h}. \quad (12)$$

Now define two constants,

$$\begin{aligned} a &= \sqrt{\frac{\gamma}{2} (1 - \sqrt{1 - \rho^2})} \\ b &= \sqrt{\frac{\gamma}{2} (1 + \sqrt{1 - \rho^2})}. \end{aligned} \quad (13)$$

The error probability may be expressed as [4]

$$\begin{aligned} \epsilon(\gamma) &= Q_1(a, b) - \frac{1}{2} e^{(a^2 + b^2)/2} I_0(ab) \\ &= e^{-\gamma/2} \left\{ \frac{1}{2} I_0(ab) + \sum_{k=1}^{\infty} \left(\frac{a}{b} \right)^k I_k(ab) \right\}, \end{aligned} \quad (14)$$




Fig. 1. Average throughput of Bluetooth ACL packets in AWGN assuming noncoherent detection, modulation index $h=0.32$, no ISI, and synchronization threshold $T = 65$.

where $Q_1(a, b)$ is the Marcum Q-function whose definition is embedded in the expression above.

Fig. 1 shows the throughput performance in AWGN using the error probability for noncoherent detection given by (14). To compute this curve, the value of $\epsilon(\gamma)$ is determined for each particular value of $\gamma = E_s/N_o$ by using (12-14) with $h = 0.32$. Next, the probability of retransmission is computed using (1-6), where we have chosen a synchronization threshold corresponding to $T = 65$. Finally, throughput is computed using (9-11) and $\bar{P}_r(\Gamma) = P_r(\gamma)$. In order to bound the pmf of N , the maximum number of transmissions for any particular packet is limited to be contained within 96 consecutive slots (using a higher limit does not change the curves). If after 60 msec the transmitter has not received a valid ACK, then the system times out and the next data packet is transmitted.

At high SNR, the throughput converges to the maximum permissible value for each packet type (i.e. 108.8 kbps for DM1, 387.2 kbps for DM3, 477.9 kbps for DM5, 172.8 kbps for DH1, 585.6 kbps for DH3, and 723.2 kbps for DH5 [2]). It is interesting to note that at low SNR different packets achieve maximum throughput. Specifically, for $E_s/N_o < 9.67$ dB (or $R_{avg} < 87.6$ kbps) the DM1 packet provides the best throughput. For 9.67 dB $< E_s/N_o < 10.79$ dB (87.6 kbps $< R_{avg} < 304.5$ kbps) the DM3 packet provides maximum throughput. For 10.79 dB $< E_s/N_o < 13.52$ dB (304.5 kbps $< R_{avg} < 477.5$ kbps) the DM5 packet is best. Finally, for $E_s/N_o > 13.52$ dB the DH5 packet has highest throughput. This implies that in order to achieve maximal throughput, the packet type should be adaptively selected to match the channel SNR. Also, note that DH1 and DH3 never provide the top

throughput for any range of E_s/N_o .

B. Performance in Quasi-Static Fading

In a quasi-static fading channel, the SNR of all symbols in a frame is constant throughout the duration of the frame, but varies from frame to frame. Since the return packet is at a different frequency than the forward packet, the return SNR γ_r is different than the forward SNR γ_f . The average probability of retransmission $\bar{P}_r(\Gamma)$ is found by taking the expected value with respect to the joint pdf of $\gamma = \{\gamma_f, \gamma_r\}$,

$$\begin{aligned} \bar{P}_r(\Gamma) &= E_\gamma \{P_r(\gamma)\} \\ &= 1 - E_\gamma \{P[\bar{A}]P[\bar{B}]P[\bar{C}]P[\bar{D}]P[\bar{E}]\}, \end{aligned} \quad (15)$$

where $P[\bar{A}]$, $P[\bar{B}]$, and $P[\bar{C}]$ are random variables that are functions of γ_f while $P[\bar{D}]$ and $P[\bar{E}]$ are functions of γ_r . Since the fading on the forward and return channels are uncorrelated,

$$\begin{aligned} \bar{P}_r(\Gamma) &= 1 - E_{\gamma_f} \{P[\bar{A}]P[\bar{B}]P[\bar{C}]\} E_{\gamma_r} \{P[\bar{D}]P[\bar{E}]\} \\ &= 1 - \int_{-\infty}^{\infty} f(\gamma_f) P[\bar{A}]P[\bar{B}]P[\bar{C}] d\gamma_f \\ &\quad \int_{-\infty}^{\infty} f(\gamma_r) P[\bar{D}]P[\bar{E}] d\gamma_r, \end{aligned} \quad (16)$$

where $f(\gamma_f)$ is the pdf of the SNR of the forward channel and $f(\gamma_r)$ is the pdf of the SNR of the return channel.

Evaluation of (16) requires the bit error probabilities $\epsilon_f(\gamma_f)$ and $\epsilon_r(\gamma_r)$ and the fading pdfs $f(\gamma_f)$ and $f(\gamma_r)$. If both the master and slave have the same type of demodulator, then $\epsilon_f(\gamma_f) = \epsilon_r(\gamma_r) = \epsilon(\gamma)$. Furthermore, reciprocity of the channel implies that $f(\gamma_f) = f(\gamma_r) = f(\gamma)$. In a Rayleigh fading channel, the SNR is exponentially distributed [4],

$$f(\gamma) = \begin{cases} \frac{1}{\Gamma} \exp\left(-\frac{\gamma}{\Gamma}\right) & \text{if } \gamma \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

Using the values for noncoherent demodulation given by (12-14), the throughput in a quasi-static Rayleigh fading channel was computed and plotted in Fig. 2. The integrals in (16) were computed numerically, and the x-axis of Fig. 2 represents the *average* SNR Γ . Since the channel error probability is for full response signaling, this plot represents an upper bound on throughput when considering practical means for demodulating the partial response GFSK signals used by Bluetooth. The loss in SNR due to using practical GFSK demodulators such as those analyzed in [3] is the same as it was for the AWGN case (i.e. 5-10 dB).

As with the AWGN case, different packets have maximum throughput at different values of E_s/N_o . Specifically, DM1 has highest throughput for $E_s/N_o < 4.02$ dB (or $R_{avg} < 3.35$ kbps), DM3 for 4.02 dB $< E_s/N_o < 7.17$




Fig. 2. Average throughput of Bluetooth ACL packets in quasi-static Rayleigh flat fading assuming noncoherent detection, modulation index $h=0.32$, no ISI, and synchronization threshold $T = 65$.

dB ($3.35 \text{ kbps} < R_{avg} < 56.7 \text{ kbps}$), DM5 for $7.17 \text{ dB} < E_s/N_o < 12 \text{ dB}$ ($56.7 \text{ kbps} < R_{avg} < 265 \text{ kbps}$), and DH5 for $E_s/N_o > 12 \text{ dB}$. Again, the DH1 and DH3 packets never offer maximum throughput. It is interesting to note that at very low E_s/N_o (i.e. for $E_s/N_o < 10.1 \text{ dB}$ and $R_{avg} < 185 \text{ kbps}$) the throughput of the quasi-static channel is actually higher than that of the AWGN channel. This is because all the packets in an AWGN channel have the same SNR, and thus at low SNR all transmitted versions of the packet are equally likely to be bad. However, in a quasi-static fading channel, some packets will have SNR that is larger than the average SNR. Thus when the packet is repeated a number of times, there is a significant probability that one of the transmissions will be at a SNR substantially higher than the average.

V. LATENCY

The analysis presented in the previous section can be easily extended to provide expressions for both *average latency* and *latency jitter*. In particular, if a packet must be transmitted N times, the associated latency is:

$$\tau = (DN)(625 \times 10^{-6}), \quad (18)$$

where D is again the number of occupied slots per transmission including the return packet. The average latency $\bar{\tau}$ is found by taking the expected value of τ with respect to N , i.e.

$$\bar{\tau} = E_N \{ \tau \}. \quad (19)$$

Likewise, the latency jitter σ_τ is found by computing the standard deviation of τ , i.e.

$$\sigma_\tau = \sqrt{E_N \{ \tau^2 \} - \bar{\tau}^2}. \quad (20)$$

VI. CONCLUSION

The throughput of the six Bluetooth ACL packets that use ARQ was derived in this paper. Exact expressions were found for each packet type by first finding the probability of retransmission and then determining the average throughput. The resulting expressions are a function of the error probability of the channel and the frame-by-frame SNR of the channel. Plots are given for the case of noncoherent detection of full response FSK signals in both AWGN and quasi-static Rayleigh fading. While not attainable in practice due to the ISI caused by Gaussian pulse shaping, these curves can serve as a benchmark for actual demodulation techniques used by Bluetooth devices.

One interesting result of this analysis is that for both AWGN and quasi-static fading channels, the DH1 and DH3 packets never achieve maximum throughput. Thus, we suggest that these two packet types should be avoided and only used if latency or data-length requirements dictate that they must be used. Furthermore, it is noted that while the DM1 packet offers the best energy efficiency in AWGN for data rates less than 87.6 kbps, in the quasi-static Rayleigh fading channel it only has best energy efficiency for rates less than only 3.35 kbps. Thus in Rayleigh fading channels, DM1 frames are of limited utility and should be reserved only for the harshest of channel conditions.

For future work, we wish to study the implications of partial response signaling further. In addition, other types of fading distributions, such as Rician and Nakagami fading, should be considered. Finally, our main motivation in undertaking this work is that we wish to find methods to improve throughput at low SNR by using customized error control techniques. The key to implementing custom error control in Bluetooth is the use of the AUX1 packet type which does not use ARQ. By using the AUX1 packet, it is possible to use a custom code that is more powerful than those in the Bluetooth standard. Because ARQ is turned off on the Bluetooth device, it can be implemented off chip in a DSP coprocessor or in the software of the host computer itself. The results contained in this paper serve as a benchmark against which our custom error control technique and those of others can be measured.

REFERENCES

- [1] Bluetooth SIG, "Specification of the Bluetooth system," *Core Version 1.1*, Feb. 22, 2001.
- [2] S. Zúrbes, "Considerations on link and system throughput of Bluetooth networks," in *Proc. IEEE Personal Indoor and Mobile Radio Commun. Conf.*, (London), pp. 1315–1319, Sept. 2000.
- [3] M. Shimizu, N. Aoki, K. Shirakawa, Y. Tozawa, N. Okubo, and Y. Daido, "New method of analyzing BER performance of GFSK with postdetection filtering," *IEEE Trans. Commun.*, vol. 45, pp. 429–436, Apr. 1997.
- [4] J. Proakis, *Digital Communications*. New York, NY: McGraw-Hill, Inc., fourth ed., 2001.