

Research Article

Unequal Protection of Video Streaming through Adaptive Modulation with a Trizone Buffer over Bluetooth Enhanced Data Rate

Rouzbeh Razavi, Martin Fleury, and Mohammed Ghanbari

Electronic Systems Engineering Department, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK

Correspondence should be addressed to Martin Fleury, fleum@essex.ac.uk

Received 1 March 2007; Revised 12 July 2007; Accepted 14 October 2007

Recommended by Peter Schelkens

Bluetooth enhanced data rate wireless channel can support higher-quality video streams compared to previous versions of Bluetooth. Packet loss when transmitting compressed data has an effect on the delivered video quality that endures over multiple frames. To reduce the impact of radio frequency noise and interference, this paper proposes adaptive modulation based on content type at the video frame level and content importance at the macroblock level. Because the bit rate of protected data is reduced, the paper proposes buffer management to reduce the risk of buffer overflow. A trizone buffer is introduced, with a varying unequal protection policy in each zone. Application of this policy together with adaptive modulation results in up to 4 dB improvement in objective video quality compared to fixed rate scheme for an additive white Gaussian noise channel and around 10 dB for a Gilbert-Elliott channel. The paper also reports a consistent improvement in video quality over a scheme that adapts to channel conditions by varying the data rate without accounting for the video frame packet type or buffer congestion.

Copyright © 2008 Rouzbeh Razavi et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Bluetooth [1], standardized as IEEE 802.15.1, is a short-range radio frequency (RF) interconnection, which can be expanded to form a piconet, with one master node and up to seven slaves. In this paper, we investigate unequal protection (UP) of encoded video data transmitted from master to slave, in the face of cross-traffic passing from slave to slave via the Bluetooth piconet master. In Bluetooth, there is no direct slave-slave communication, as all cross-traffic must pass through a Bluetooth master node. Such usage certainly occurs in Bluetooth personal area networks for wearable computers [2], whereas IEEE 802.11 wireless local area networks are less suitable for this purpose, for example, because of an order-of-magnitude higher-power requirement (100–350 mA as opposed to 1 mA). Providing differing levels of error coding to achieve UP is widely practiced. This is usually designated as unequal error protection (UEP) and not UP. However, it is also additionally possible to apply modulation adaptation to achieve UP, particularly in orthogonal frequency division multiplexing (OFDM) systems [3]. As an

example [4], adaptive modulation was traded against error coding. However, if data-link FEC is *not* available, it is still possible to apply adaptive modulation. In Bluetooth version 2.1, FEC is not implemented for enhanced data rate modes, possibly because low-cost devices could not cope with the computational requirements of coding at the higher data rates. On the other hand, Bluetooth EDR provides several forms of modulation, though not through OFDM.

Our main contribution is protection by adaptive modulation together with transmit buffer management to avoid packet loss from buffer congestion, with consideration of packet importance and wireless channel conditions. We propose trizone management of the transmit buffer for video stream packets, based on the relative content importance of the differing frame types. To the best of the authors' knowledge, no trizone buffer system of management based on video packet importance has been previously described. The combination of frame-packet-type and subsidiary-macroblock-type frequency counts provides a clear means of regulating the zones. The paper reports an upper bound improvement in video quality, reflected in peak

signal-to-noise ratios (PSNRs)¹ of about 2 to 4 dB employing UP over the best fixed-modulation scheme without protection additive white Gaussian noise (AWGN) channel and around 10 dB for a Gilbert-Elliott channel. The paper also improves a consistent improvement in video quality over a scheme that adapts to channel conditions by varying the data rate without accounting for the video frame packet type. The UP scheme involves no change to the Bluetooth version 2.1 specification [5], as we would wish to preserve the advantages of a Bluetooth single-chip, low-cost (<US \$5), and low-power implementation. We also do not assume FEC at the application layer, as this would reduce the generality of the solution as far as the video decoder is concerned. Single-layer video is assumed because most legacy content is in this form, though there are many good schemes such as [6] that rely on layering of some form (fine-grained, data-partitioning, wavelet coding, spatial/temporal scalability). Instead, UP by frame type and content importance is simpler to implement as a cross-layer system, avoiding the complexity that would militate against the positive features of Bluetooth.

Bluetooth v. 1.2 received comparatively limited investigation as a medium for streaming video. The potential performance of encoded video transmission was investigated in [7–9], but no error control measures were proposed. Hardware implementations are described in [10, 11], but error control is described by conventional MPEG-4 error resilience tools, though channel coding is not discounted. We also assume error resilience through slice resynchronization markers (see Section 3.4), except when the slice structure reduces packet throughput. Error concealment by previous frame replacement is a simple and standard means of error reduction [12] which we also assume to be present at the decoder. In [13], it was remarked that the default Bluetooth recommendation of automatic repeat request (ARQ) with unlimited repeats is unsuitable for video transmission and, therefore, a non-standard codec with built-in error resilience was assumed. While we agree with the former suggestion, using a nonstandard codec is only suitable for embedded applications and not for a Bluetooth access network for a possibly remote and anonymous server. The nearest similarity to our work is that reported in [14], which employs repeated transmission of intracode frames (rather than adaptive modulation). To avoid host intervention to control the number of retransmissions, the standard Bluetooth mechanism of setting the flush timeout is employed, which indirectly controls the number of retransmissions. However, the work in [14] does not consider frame types other than intracoded ones and does not report the impact on packet latency.

Bluetooth v. 2.0 increased the maximum gross user payload (MGUP) bit rate from a basic rate of 0.7232 Mbps to 2.1781 Mbps, which allows Bluetooth to carry an arriving

MPEG2 transport stream (TS). Bluetooth v. 2.1 [5] also includes near field communication, along with improvements to power consumption and security. It seems that the increase in bandwidth has decreased research in video transmission over Bluetooth, as very little consideration as a whole has been given to Bluetooth v. 2.0 or v. 2.1 in the research literature. In fact, Bluetooth v. 2.1 under EDR supports gross air rates of both 3 and 2 Mbps (MGUP of 1.4485 Mbps), through, respectively, $\pi/4$ -differential quadrature phase-shift keying (DQPSK) or eight-phase differential phase-shift keying (8DPSK) modulation.² This implies that, through adaptive modulation, a lower bit rate is available that can serve to give UP to some of the packets of more important frame types, the intra- (I-) and predictive- (P-) anchor frames, as well as some bipredictive- (B-) frame packets, depending on circumstances.

Because a lower bit rate is employed for priority packets, there is a risk of buffer overflow at the transmit buffer, compared to a situation in which all packets were sent at the higher bit rate. Therefore, a trizone buffer applies a different UP policy for each zone. However, it should be carefully noted that the fact that there are three zones does *not* mean that only I-frame packets occur in one zone, P-frame packets in another zone, and B-frame packets in the third zone. All packet types can occupy each zone, but the prioritization policy between each zone is different as a reflection of the greater fullness of the buffer as each successive zone is occupied. As the buffer fullness increases, packets of whatever type begin to fill the second and then the third zone, and the prioritization policy between frame-type packets changes accordingly. Ideally, the output at the lower bit rate should decrease linearly as buffer fullness increases. However, to achieve this, because of a varying number of packets between the frame types, a linear UP policy based simply on frame type will not work. Therefore, in the second zone of the trizone buffer, the number of P-frame packets offered protection is modulated by the content importance and its predominance within the arriving P-frames.

The buffer zone boundaries are based on the frequency within a video stream of I-, P-, and B-frames, and they dynamically change according to the relative size ratio of the arriving frame-type packets. In other words, the ratio of data allocated to each frame type within an arriving video stream dynamically determines the zone sizes, while the frame type determines the UP policy applied within the zone. Zone 1 is first occupied by arriving packets. In this zone, not all B-frame packets are protected, and B-frame packets are not protected in zones 2 and 3. As zone 1 is the only zone in which B-frames receive some protection, it makes sense to allocate the size of zone 1 according to the relative amount of data arising from B-frame packets. Doing otherwise would bias the zone size against B-frame packets. It should be noted that, because of the GOP structure, B-frame packets occur with greater frequency than other frame-type packets. In

¹ Specifically, $PSNR = 10 \log_{10} [p^2 / (1/n) \sum_{i,j} (Y_{ref}^{i,j} - Y_{prc}^{i,j})^2]$ dB, where p is the peak value for a given pixel resolution, for example, for 8-bit $p = 255$, n is the total number of pixels in a picture, i, j range over every pixel of the frame, and Y_{ref} is the luminance value in the original frame before transmission, while Y_{prc} is the pixel value in the frame after transmission, decoding, and display.

² In the paper, for ease of reference, these EDR modes are referred to by their gross rate.

zone 2, not all P-frame packets are protected and P-frame packets are not protected in zone 3. Therefore, in zone 2, when P-frame packets begin to lose the protection received in zone 1, the size of the zone determines, so to speak, how quickly they lose their protection. This rate is determined by the amount of P-frame type data within the stream to avoid unfairly biasing of the zone size against P-frame packets. Finally, a similar observation applies to zone 3. If there are packets occupying this zone, then the buffer would be at its fullest state and as a result not all I-frame packets are protected in zone 3.

By monitoring transmitter buffer fullness, available through Bluetooth host controller interface (HCI), an adaptive UP scheme is applied. It turns out that buffer fullness is an excellent indication of congestion within a Bluetooth piconet. Buffer fullness is responsive not only to buffer congestion from an arriving video stream but also to an increase in buffer service time when piconet cross-traffic is present. As buffer fullness reflects the congestion of the Bluetooth wireless channel, it can be used to regulate the UP scheme, and this is a feature of our proposal. The channel condition should also be ascertained. This can be achieved by received signal strength indicator (RSSI) [15] or we can rely on channel probing messages or channel condition feedback messages [16]. RSSI is an optional feature of Bluetooth implementations, though in [16] it was found that the RSSI reported that Bluetooth channel quality oscillated rapidly. This topic is otherwise outside the scope of this paper.

A range of packet types exists in Bluetooth according to the number of timeslots occupied by a packet (1, 3, or 5) and the modulation type. The classical Bluetooth channel quality-driven data rate (CQDDR) model assumes different packet types, and hence data rates are chosen depending on channel conditions. This model can be achieved by means of a lookup table (LUT) which effectively establishes the per-bit SNR boundaries between the differing packet types. Selecting the packet type by content type in addition to selection by channel quality overrides CQDDR. This is provided by offering *up* to some video packets when traffic on the shared Bluetooth channel permits it. When channel conditions deteriorate and/or traffic congestion across the Bluetooth piconet increases, then the trizone policy effectively converges upon the CQDDR model.

In the Bluetooth CQDDR model, retransmission after an automatic repeat request (ARQ) occurs until the packet arrives without errors. However, it is possible to set the “flush timeout” to a minimal value [5], which effectively turns off ARQ. The details of what this value should be and possible side effects from setting it are discussed in Section 3.1. As unbounded retransmissions may well lead to missed display deadlines when transmitting video frames, some such action is advisable. Otherwise, packets may not be lost over the wireless channel, but they are dropped by the decoder. The sender informs the receiver of a change in the default flush timeout by a logical link control and adaptation protocol (L2CAP) command message [5], with no alteration to the Bluetooth packet header being required. A consequence of abandoning CQDDR in some circumstances for video is that the choice between the two EDR modes is no longer bi-

nary. It is on this observation that the UP adaptive modulation scheme is founded.

The proposed scheme has no implications for the Bluetooth EDR standard such as changing the form of modulation. Priority packet marking can take place above the HCI boundary within the host’s software, which is available in open source form, such as the Bluez stack for the Linux operating system. However, firmware modification would be required at the data-link layer in order to recognize marked packets and apply adaptive modulation.

The remainder of this paper is organized as follows. Section 2 considers related work on UP of video streaming over wireless channels. Section 3 describes how the UP system is modeled in the paper. Section 4 details the application of the UP system, while Section 5 presents the evaluation of the system. Finally, Section 6 draws some conclusions.

2. RELATED WORK

This section employs a simple division into research on UP for multistream and single-stream videos (with UEP being considered by us as a subset of UP). A more complete taxonomy might also account for wireless technology capability and performance according to channel conditions. For example, with respect to the wireless technology, Bluetooth v. 1.2 has only one form of modulation, Gaussian frequency-shift keying, Bluetooth v. 2.1 has two additional forms, whereas IEEE 802.11a has eight modulation modes. Any protection scheme should take account of these differing capabilities.

2.1. Multistream video UP

In [17], the video stream is partitioned through multiresolution coding (with some redundancy), and each sub-stream is adaptively modulated and transmitted through an antenna array in a multiple-in multiple-out (MIMO) system. The solution in [17] is, of course, unsuitable for Bluetooth because of the assumption of MIMO. Adaptive modulation can also be applied [18] through multilayering, but, as remarked in Section 1, this is at the expense of flexibility. OFDM systems such as IEEE 802.11a lend themselves to a combination of FEC and adaptive modulation [15, 19]. In [15], layering occurs through fine-grained scalability in which a progressive intracoded enhancement layer is employed. Vertical integration of protection means, including adaptive ARQ and FEC, is applied. However, the (N, K) Reed-Solomon (RS) coding of [15] is not particularly unsuitable for Bluetooth, as RS codes have a $K(N - K)\log_2 N$ complexity. Adaptive ARQ for Bluetooth [20] is a promising alternative to adaptive modulation. Similarly, in [21] in work by one of the coauthors, motion vectors and other header data through H.264 data partitioning are prioritized through hierarchical quadrature amplitude modulation (QAM) for OFDM, intended for a digital video broadcasting (DVB) system. In [22], horizontal FEC coding across packets was applied, so that the initial data within each packet was afforded greater protection than later data, though this scheme was actually applied to the fixed Internet.

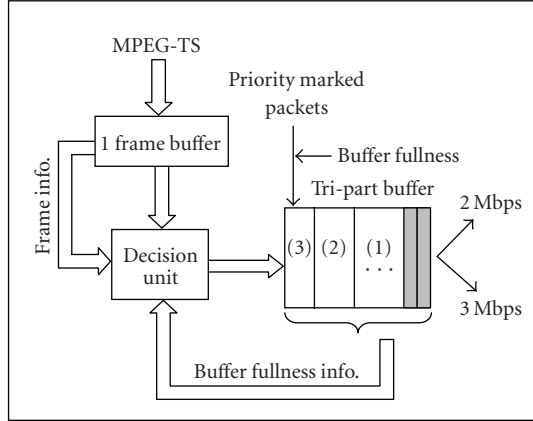


FIGURE 1: Unequal protection system for video data.

2.2. Single-stream video content-importance UP

In our paper for single-layer video, individual parts of the stream are protected according to the content importance. In comparison, [23] takes four categories of MPEG-4 information: header, I- and P-frames with scene changes, shape and motion information in P-frames, and fourthly texture information in P-frames. The scheme in [23] employed priority-based ARQ combined with data-link FEC protection of retransmitted packets, that is, a form of type-1 hybrid ARQ. A finer level of data prioritization may be applied [24] by inspecting the number of intracoded macroblocks in an H.263 bitstream, though in [24] they are protected by ARQ and FEC, rather than adaptive modulation. Intracoded macroblocks, as monitored by us, may appear in P-frames as well as I-frames and may indicate scene changes, camera zooms or pans, and so on. The presence of intracoded macroblocks, which is encoder implementation-dependent, indicates important information in the encoded bitstream, though prior research in [24] did not associate them with the frames themselves and did not employ adaptive modulation. For an MPEG-4 bitstream, in [25] packets are reorganized into fixed-size segments containing data of differing importance. The intention was to reduce side-information overhead by avoiding the need to indicate data-type boundaries. The side information is needed for adaptive ARQ at the wireless link. However, again this was a UEP scheme not a UP one, with RS coding forming the protection. On the other hand, [26] does rely on side information, namely, an error propagation rating found at the encoder.

3. UP SYSTEM MODEL

3.1. Cross-layer interaction

In Figure 1, prior to Bluetooth packetization, the encoded MPEG2-TS enters a one-frame buffer. The stream may be encapsulated as an Internet protocol (IP) packet arriving, say, by DVB-T (digital video broadcasting for terrestrial transmission) or Internet protocol TV (IPTV), or directly from, say, a DVD. Within the frame buffer, the UP system deter-

mines the type of frame, its size, and, if a P-frame, the ratio of intracoded macroblocks within the encoded data. The frame information is passed to a decision unit that allocates the priority of the resulting Bluetooth packets when they are passed into the first-in first-out transmit buffer. The prioritizing decision is affected by the state of buffer fullness and the importance of the incoming Bluetooth packet. The tri-zone buffer configuration is further explained in Sections 3.2 and 4.1 Within the transmit buffer, priority-marked Bluetooth packets are transmitted by one of the two modulation schemes, depending on the packet's priority. As already mentioned, low-priority packets are sent at 3 Mbps, as this rate is subject to the largest risk of error.

As mentioned in Section 1, Bluetooth default ARQ mechanism (unlimited retries) is effectively turned off by altering the flush timeout to avoid excessive packet delay, which would result in missed display or decoded deadlines at the receiver. The flush timeout value is set in multiples of 625 microseconds. As this is the Bluetooth timeslot period, no packet transmission can be shorter than 625 microseconds. In fact, as part of Bluetooth time division duplex (refer to Section 3.4), a mandatory reply is always sent from the receiver to the sender. Therefore, setting the flush timeout to two timeslots (1250 microseconds) serves the same purpose. In our Bluetooth simulation model, we assume that, once a flush timeout has occurred, the link controller sends no further handshake packets to the receiver. Resetting the flush timeout value will affect all other communication streams as well as the video stream. However, in practical terms, this is avoided by setting the packets in the other communication streams as nonflushable and in our Bluetooth simulation model by intervening at the buffer level to distinguish between flushable and nonflushable packets.

In the tests of Section 5, an AWGN channel is modeled, with a bit error rate (BER) of 10^{-5} at the higher rate of 3 Mbps, corresponding to an E_b/N_0 of 16 dB. This value of SNR is convenient as it lies within the range for which five slot packets are optimal (refer forward to Section 3.5), thus simplifying the interpretation. However, to judge the response of the UP scheme to different channel conditions, a Gilbert-Elliott [27, 28] two-state discrete-time ergodic Markov chain is also employed to model the wireless channel error characteristics. By adopting this model, it was possible to simulate burst errors, which are typical of practical channels. Though Bluetooth v.1.2 adopts an adaptive frequency hopping (AFH) scheme, the Gilbert-Elliott model is still used herein to model the channel, because AFH is of limited benefit to audio/video applications [29], especially when interference occurs across the unlicensed 2.4 GHz industrial scientific medical (ISM) band. The mean duration of a good state, T_g , was set at 2 seconds and that of a bad state, T_b , to 0.25 seconds. In units of 625 microseconds (the Bluetooth timeslot duration), $T_g = 3200$ and $T_b = 400$, which implies from

$$T_g = \frac{1}{1 - P_{gg}}, \quad T_b = \frac{1}{1 - P_{bb}} \quad (1)$$

that, given that the current state is good (g), P_{gg} , the probability that the next state is also good (g), is 0.9996875 and P_{bb} , the probability that the next state is also bad (b), given

that the current state is bad (b), is 0.9975. At 3 Mbps, the bit error rate (BER) during a good state was set to 10^{-5} and during a bad state to 10^{-4} in 3 Mbps mode. The transition probabilities, P_{gg} and P_{bb} , as well as the BER, are approximately similar to those in [30], but the mean state durations are adapted to Bluetooth. The two states result in SNRs of, respectively, 16.00 and 14.70 dB. The first value is chosen to provide a point of comparison with the single-state model, while the second SNR value lies within the range in which a rate of 2 Mbps is optimal (refer forward to Section 3.5). In subsequent experiments, the already high BER is made worse by linearly modifying the bad-state BER. For SNRs below 10 dB (see Table 2), only protected basic rate packets are suitable, while the UP adaptive scheme is appropriate for EDR modes.

This research applied the University of Cincinnati Bluetooth (UCBT) extension (download is available at <http://www.ececs.uc.edu/~cdmc/ucbt>) to the well-known NS-2 network simulator (with v. 2.28 being used). The UCBT extension supports Bluetooth EDR, but it is also built on the air models of previous Bluetooth extensions such as BlueHoc from IBM and Blueware. Specification details at both the baseband and the above such as L2CAP are simulated in UCBT, including connection setup and multislot packet-type negotiation. UCBT also takes clock drift into account, to allow for accurate simulation of synchronization and scheduling. However, clearly any implementation of Bluetooth may differ from the simulation and, in particular, the speed of switching between EDR modulation modes may differ if a longer guard interval is applied to separate the modes.

3.2. Buffer UP policy

An overview of the buffer zone UP policy has been given in Section 1. In zone 1 of the buffer, all Bluetooth packets of I- or P-frame type are automatically protected through dispatch at the lower bit rate. B-frame packets are only protected in zone 1 if they pass the following test. A uniformly distributed random number in the interval $[0,1]$ is generated and compared to the fraction f , zone packet occupation/zone capacity. If the random number is greater than f , then that B-frame packet is also protected. This test is adopted so that the number of B-frame packets that are protected linearly changes with zone-1 buffer fullness.

As the buffer fullness increases and packets also occupy zone 2 of the buffer, a different prioritization policy for P-frames is applied. I-frame packets remain protected within zone 2 of the buffer, and B-frame packets are no longer protected. P-frame packets in zone 2 of the buffer are protected according to the ratio of intracoded macroblocks within the frame, as detected, while the frame is in the frame buffer. Again, the boundary between protected and unprotected P-frame packets is dynamically adjusted according to a past history of intracoded macroblock ratios within P-frames. Section 4.2 further explains zone-2 adjustment of the buffer.

Finally, in zone 3 of the buffer, when the buffer is at its fullest state, no protection to any B- or P-frame packets is applied. However, I-frame packets are protected according to




FIGURE 2: Spatial information change over time.

the same policy applied for zone 1, that is, by random number generation and comparison with a fraction f for zone 3.

Notice that in zones 1 and 3, the UP policy approximates a linear regime. This is because the allocation function f grows linearly with buffer fullness for B-frame packets in zone 1 and I-frame packets in zone 3. However, the P-frame UP policy is nonlinear, as it is based on a tradeoff between content importance and buffer fullness. By compensating for buffer fullness, the actual P-frame packet output is actually adjusted to approach once more a linear regime.

3.3. Dynamic variation of frame content

In Section 3.1, it was found that it is necessary to dynamically adjust the ratios between the zones. In general, this is due to the following. Firstly, the spatial content varies over time, which will impact upon I-frame size. Secondly, the temporal content will also vary over time, which will affect B- and P-frames in approximately equal measure. In [31], for the purpose of selection of suitable video sequences for subjective testing, two measures were provided for judging the spatial and temporal information, respectively. In the spatial measure, the luminance is Sobel-filtered for each frame under test, and subsequently the standard deviation (SD) is taken over all pixels in a frame. The measure takes the maximum, but in our illustration the SDs are simply plotted (see Figure 2). Figure 2 represents the spatial content in successive frames of part of the *Italian Job* (European-formatted standard interchange format (SIF), 352×288 pixel resolution, 25 frames/s (fps), encoded at 2 Mbps), a film with many scene changes owing to the action in the film. For the temporal measure, the difference in luminance value is computed between the current frame and the previous one for all pixels in the current frame. The per-frame SD is taken from the temporal information of all pixels in each frame. Figure 3 plots the temporal SDs over time for the same video sequence. In both Figures 2 and 3, the variability in spatial and temporal information is evident, justifying the need to vary the buffer zone sizes over time.




FIGURE 3: Temporal information change over time for the same sequence as in Figure 2.

TABLE 1: Bluetooth packet types: user payload and bit rates.

Packet type	User payload in bytes	Asymmetric maximum rate (kbps)
DM1	0–17	108.8
DM2	0–121	387.2
DM5	0–227	477.8
DH1	0–27	172.8
DH2	0–183	585.6
DH3	0–339	723.2
2DH1	0–54	345.6
2DH3	0–367	1174.4
2DH5	0–679	1448.5
3DH1	0–83	531.2
3DH3	0–552	1776.4
3DH5	0–1021	2178.1

Length and master-to-slave bit rates for a single ACL master-slave logical link, with DM = data medium rate (FEC enabled) and DH = data high rate (no FEC). 2-DH3 is a 2Mbps modulation three-timeslot packet.

3.4. Packetization policy

A data frame across a Bluetooth link in asymmetric mode consists of an asynchronous connectionless (ACL) packet occupying one, three, or five timeslots and at least a single slot reply, with either master or slave as receiver. Because of packet quantization effects, the Bluetooth packet sizes become significant and their effects on user payload are summarized in Table 1 for a single master-slave ACL link for Bluetooth v. 2.1. Packet types at the basic rate (DH1-5, DM1-5) are not part of EDR, but they are included because the data medium (DM) packets are effective at low SNRs. The DM packets employ data-link FEC through an expurgated (15,10) Hamming code.

The normally assumed Bluetooth controller behavior is that, given a maximal Bluetooth packetization scheme, for example, 3DH5 or 3DH3, packets up to the maximum user payload will be formed. However, if the arriving data or IP packets do not justify the preset maximal scheme, then a

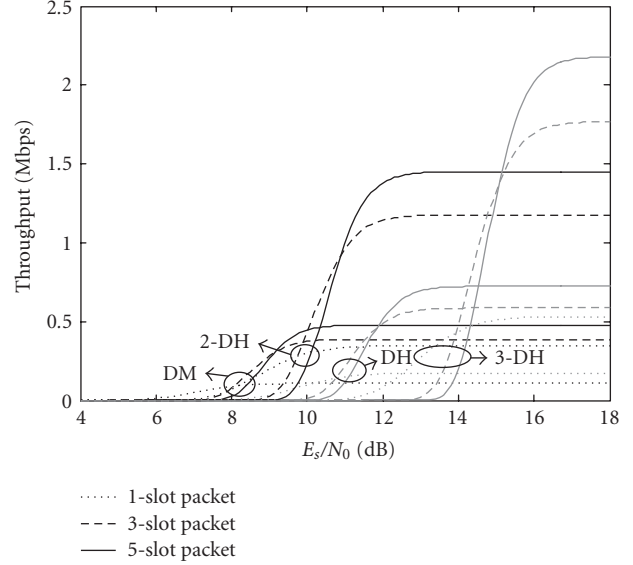


FIGURE 4: Throughput versus SNR for different Bluetooth packet types.

reduced scheme is used. For example, the controller swaps from 3DH5 down to 3DH3 or even 3DH1.

Unfortunately, if packetization takes place on a single MPEG2 slice (one row of macroblocks) per Bluetooth packet, this behavior introduces the possibility of many partially filled packets and many 1- or 3-slot packets. The result is a drop in throughput. Therefore, in [32], fully filled Bluetooth packets were formed, regardless of slice boundaries. While this results in some loss in error resilience, as each MPEG-2 slice contains a decoder synchronization marker, in [32] it is shown that the overall video performance is superior. In the experiments in Section 5, the video Bluetooth packet size was set either to 3DH5 or 2DH5, depending, respectively, on whether a gross rate of 3 or 2 Mbps was chosen.

3.5. CQDDR model

As introduced in Section 1, the CQDDR model adapts the Bluetooth packet type to channel conditions. The pure CQDDR model does not account either for packet content or the congestion level of the network, whereas this paper's scheme accounts for both through the trizone buffer. Figure 4 plots the throughput of the Bluetooth packet types of Table 1 for an AWGN channel. It will be seen that certain Bluetooth packet types never provide optimal throughput. Table 2 shows the SNR boundaries between the optimal packet types. The expurgated (15,10) Hamming code is capable of double adjacent error correction (DAEC) [33], as well as single error correction (SEC). An SEC-DAEC decoder involves no additional complexity in its implementation. However, as much research on Bluetooth such as [34] has assumed an SEC decoder, Table 2 includes SNR boundaries for both types of decoder, while Figure 4 assumes an SEC-DAEC decoder.

TABLE 2: Optimal Bluetooth packet types by SNR boundaries.

Optimum packet type	SNR range in dB for receiver with double adjacent error correction functionality	SNR range in dB for receiver without double adjacent error correction functionality
DM1	SNR < 8.06	SNR < 8.15
DM3	8.06 < SNR < 9.13	8.15 < SNR < 9.20
DM5	9.13 < SNR < 10.03	9.20 < SNR < 10.02
2HD3	10.03 < SNR < 10.88	10.02 < SNR < 10.88
2DH5	10.88 < SNR < 15.14	10.88 < SNR < 15.14
3DH5	SNR > 15.14	SNR > 15.14

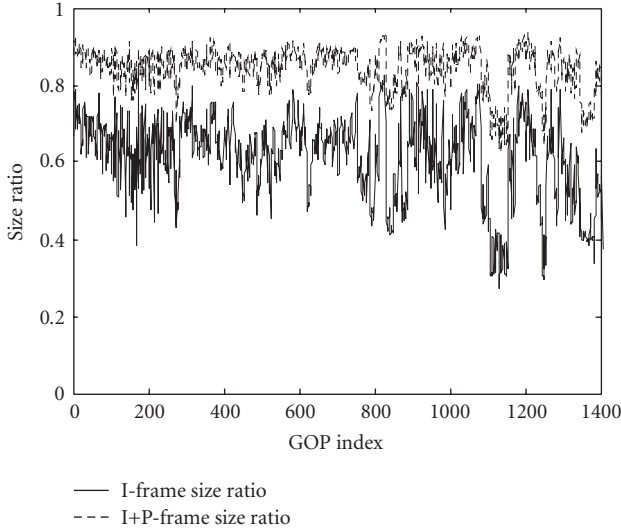


FIGURE 5: Example measured distribution of frame ratios by frame type per GOP for an MPEG-2 video sequence.

4. METHODOLOGY

4.1. Buffer zone size allocation

In Figure 5, for an MPEG-2 SIF-resolution video sequence (an episode of the situational comedy *Friends*) at 25 fps, with group of pictures (GOP) structure³ of $N = 12$ and $M = 3$, the relative sizes of I-, P-, and bipredictive B-frames were monitored. In fact, as occurred in practice, averaging over 10 GOPs produces little change in the pattern. It will be seen that though a static ratio of 6 : 3 : 2 for I-, P-, and B-frame sizes is a good fit [35], the relative size of P-frames and at the same time B-frames may well change in comparison to I-frames.

To consider how the buffer zone boundaries are allocated, firstly take the static size ratio of 6 : 3 : 2 between the different frame types. Within a GOP structure of $N = 12$ and $M = 3$,

³ N determines the number of frames from one I-frame before another one occurs. M determines the number of frames before a further anchor frame (I- or I-frame) occurs. $M = 3$ implies that there are 2 B-frames before each anchor frame.

the frequency of frame types is in the ratio of 1 : 3 : 8. Therefore, by simple multiplication of the three ratios, the buffer zone sizes would be in the ratio of 6 : 9 : 16. For a total buffer capacity of 50 packets divided in this last ratio, the zone allocation is (10, 15, 25), with zone 1 being 25 packets, zone 2 being 15 packets, and zone 3 being 10 packets. The zone allocation was adjusted accordingly by a P -order linear prediction filter (LPF) [36], with an eight-order filter resulting in very little difference between the predicted and the actual ratios of Figure 5. Ratio values were predicted by the P -order LPF previously mentioned. Specifically, the I- to P-frame and P- to B-frame ratios were predicted. The P -order linear prediction filter is represented by

$$X(m+1) = \sum_{k=1}^P w_k \cdot X(m-k+1), \quad (2)$$

where $X(m+1)$ is a predicted ratio value estimated from P previous values over sample instances m , while the w_k are the P adaptive filter weights indexed by k . The weights are estimated [36] through

$$\mathbf{w}(m+1) = \mathbf{w}(m) + \frac{e(m) \cdot \mathbf{X}(m)}{\|\mathbf{X}(m)\|^2}, \quad (3)$$

where \mathbf{w} is the length- P column vector of weights and \mathbf{X} is a length- P column vector of ratio measurements over time as in:

$$\mathbf{X}(m) = [X(m), X(m-1), \dots, X(m-P+1)]^T \quad (4)$$

when T represents the vector transpose. The variable $e(m)$ is the error between the measured and the predicted ratio value. The system was initialized with a ratio of 6 : 3 : 2, which, as previously mentioned, is a good fit for the relative sizes of I-, P-, and B-frames. Figure 6 then represents the predicted values over time, bearing out the claim that the predicted values differ little from those in Figure 5.

4.2. P-frame macroblock-type prioritization

In MPEG-2, while I-frames are formed entirely by intracoded macroblocks, P-frames, apart from macroblocks of predictive type and SKIP (no update of matching macroblocks from the prior frame), may also include intracoded macroblocks. Figure 7 plots the ratio of intracoded macroblocks

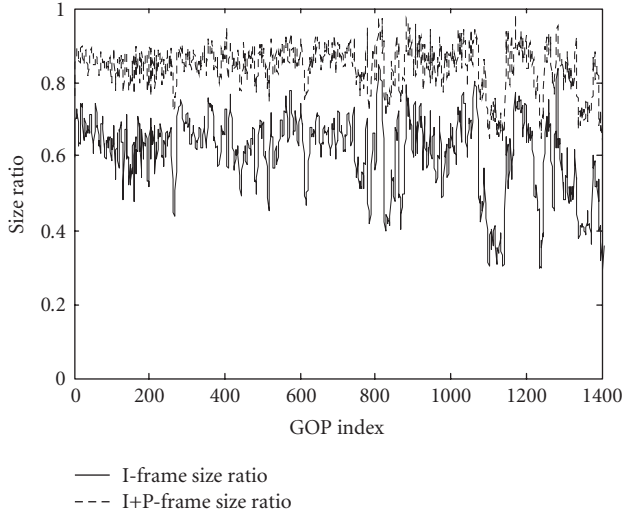


FIGURE 6: Predicted distribution of frame ratios by frame type per GOP for an MPEG-2 video sequence.

within P-frames for a *Football* sequence. The *Football* sequence has the same GOP structure as the *Friends* sequence, and it is again an SIF-resolution sequence at 25 fps. It is chosen as an illustration, as there is rapid motion, and between P-frames indexed as 65 (see Figure 7(b)) and 66 (see Figure 7(c)), a scene change occurs from a wide view of the pitch to a close-up of players. The plot in Figure 7(a) shows a sharp peak in the ratio of intracoded macroblocks for these P-frame indices, and for others. As matching macroblocks in subsequent frames (after P-frame index 66) depends for coding on these macroblocks, until the arrival of the next I-frame, it is important that they are delivered intactly to the decoder. Notice that in general the distribution of P-frames with a high intracoded ratio is dependent on film genre and motion content, and Figure 7 should not be taken as typical.

In the buffer zone-2 algorithm, every M P-frame, for some constant M , is sampled to determine the distribution of intracoded macroblocks. Depending on that distribution, the policy of protecting P-frame packets within zone 2 of the buffer is adjusted and applied to the next M P-frames. During the application of this protection policy, the next M frames are similarly inspected. A size of $M = 100$ frames was chosen assuming that the video characteristics are wide-sense and time-stationary over this interval.

Figure 8 plots the ratio of intracoded macroblocks in P-frames for the *Friends* sequence of Section 4.1. Figure 9 shows the resulting distribution over the P-frames, grouped into the ten categories used by the current algorithm (but for 1000 P-frames in this example rather than 100 used in practice). The derived mapping function is plotted in Figure 10 for two different illustrative buffer zone-2 capacities. The mapping function is quantized according to the integer-valued number of packets on the horizontal axis of Figure 10. Using this mapping function enables a linear change in the number of protected P-frame packets versus buffer occupation of zone 2.


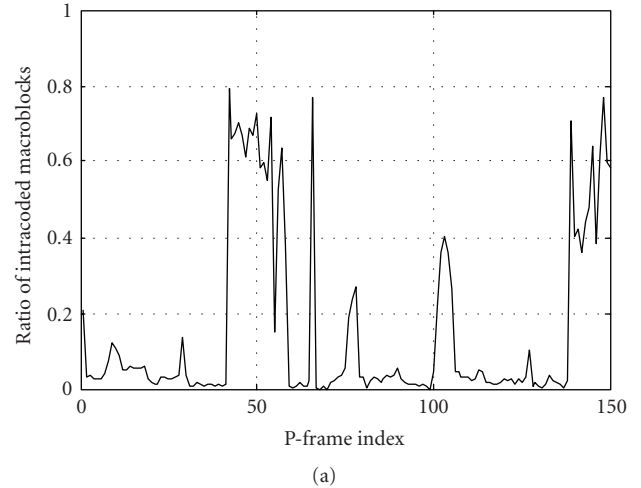


FIGURE 7: Example distribution of macroblock types within P-frames, with (a) frequency of intracoded macroblocks, (b) frame-65 macroblock types, and (c) frame-66 macroblock types, with grey circles = predictive macroblocks, black = SKIP, and white = intracoded macroblocks.




FIGURE 8: Intracoded macroblock ratio for successive P-frames.

As an example, assume the total capacity of zone 2 to be 50 packets, then when there are 40 packets in the buffer, only those P-frames that have more than 62.4% of their intracoded macroblocks are protected. At any time, if the current number of packets in zone 2 and the ratio of intracoded macroblocks of a given frame are known, the decision can be made easily.

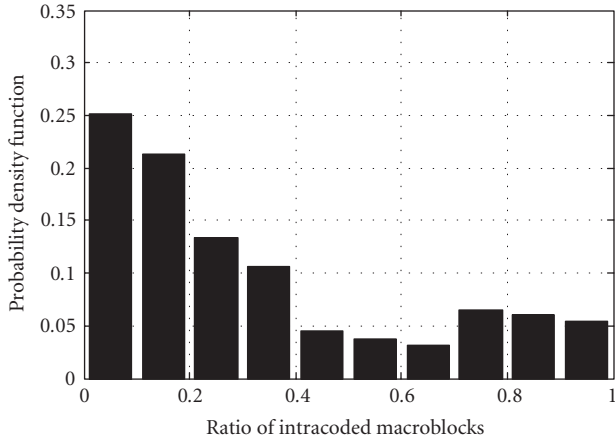


FIGURE 9: Distribution of the ratios of intracoded P-frame macroblocks from Figure 5.

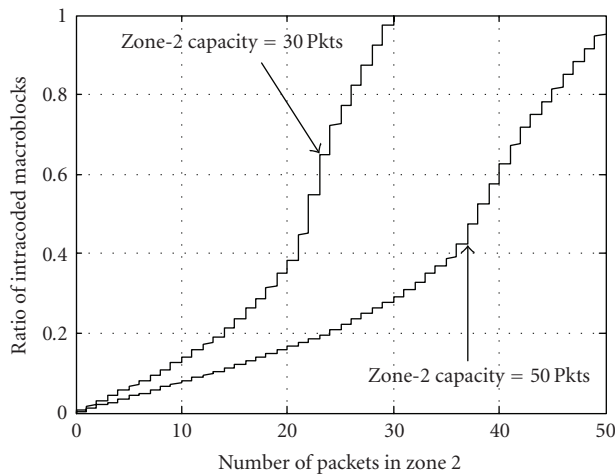


FIGURE 10: Protection mapping function based on two different buffer zone-2 capacities.

The mapping function is formed by taking the set of ten probabilities, such as that in Figure 9, and projecting them onto the zone-2 capacity. For example, in Figure 9, the 0.1 ratio of intracoded macroblocks has a probability of approximately 0.25. Therefore, there are 13 (0.25×50) packets allocated for a zone-2 with capacity of 50 packets. The same calculation is repeated for the next data point at a ratio of 0.2, but with aggregated probability of $(0.25+0.21)$ from Figure 9. Data points are connected in piecewise linear fashion.

4.3. Piconet congestion and buffer fullness

Figure 11 shows the simulation configuration for the results of Section 5. The MPEG-2 video stream is sent from the Bluetooth master node to slave S1, while slave S2 acts as a traffic source to slave node S3. As already mentioned, there is no direct slave-slave communication, and therefore a master maintains separate queues for each master-to-slave link (see Figure 12). The Bluetooth standard does not specify the queue service discipline, and along with Bluetooth




FIGURE 11: Bluetooth piconet with cross-traffic.




FIGURE 12: The buffering model for Bluetooth.

implementations, this paper assumes pure round-robin (1-limited) scheduling. The work in [37] showed that 1-limited servicing performed better under high load than an exhaustive queue discipline.

Various metrics have been considered to monitor congestion, which can be caused by cross-traffic or traffic from a local source (which we call self-congestion). In [6], it is suggested that for congestion control, the input packet rate to the shared RF channel should be increased (decreased) when the loss rate is below 5% (higher than 15%), based on periodic feedback from the receiver. Unfortunately, packet loss rates of 10% or more are likely to lead to a drastic reduction in video quality. In [38], packet delay recorded at a Bluetooth receiver was found to be a better indicator of congestion than packet loss, but it resulted in oscillations in both video quality and delay in packet delivery when used as input for congestion control.

On the other hand, Figure 13 shows the ability of buffer fullness to track both variations in direct traffic (M to S1 in Figure 11) and in cross-traffic (S2 via M to S3 in Figure 11). In [38], it is also shown that buffer fullness when applied to congestion control significantly reduces delay and improves PSNR. The video traffic rate plot in Figure 13 reflects a fixed constant bit rate (CBR) cross-traffic at 200 Kbps and packet size of 800 B. Notice that this implies an effective bit rate of 400 Kbps across the shared channel, as the CBR traffic makes two hops reach its destination. Equally, the packet size implies less-than-optimal use of the bandwidth capacity. The video traffic source was a 40-second MPEG2 CIF-sized 25 fps *Newsclip* (moderate motion) with GOP structure of $N = 12$ and $M = 3$, with fully filled packets. As its rate passes a threshold of around 1.6 Mbps, buffer fullness sharply climbing as the saturation rate of the Bluetooth link at 2.1 Mbps is approached. Similarly, with the MPEG2 source rate fixed

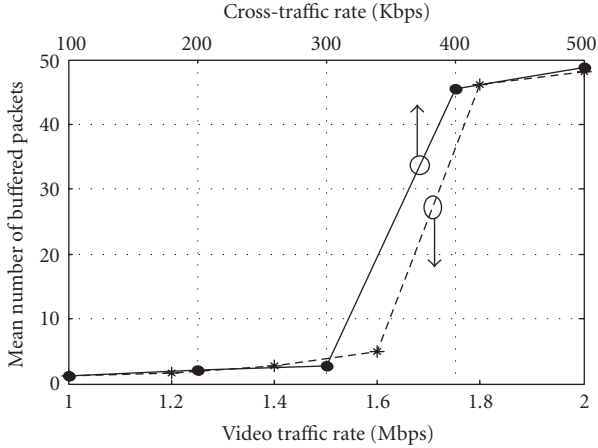


FIGURE 13: Buffer fullness against varying cross-traffic and varying video rate.

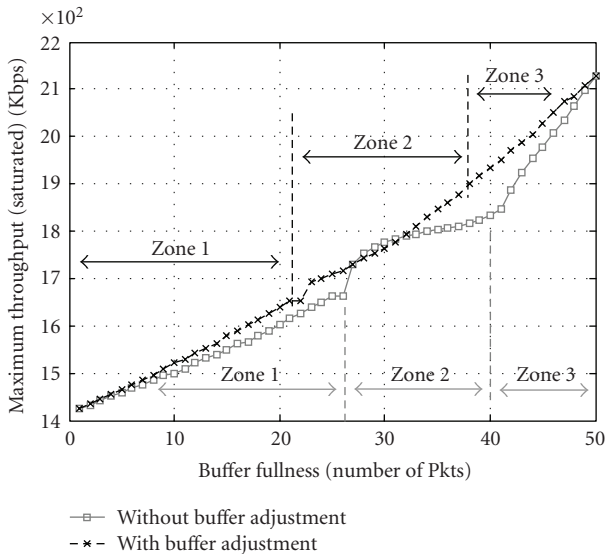


FIGURE 14: The effect of size- and content-aware UP policy on throughput.

at 1.25 Mbps, when the CBR rate approaches channel saturation, there is a sudden increase in buffer occupancy.

5. RESULTS

5.1. UP behavior without cross-traffic

In Figure 14, total buffer fullness is plotted across the horizontal axis for a 50-packet Bluetooth transmit buffer. Maximum achievable bit rate is plotted with and without dynamically changing trizone buffer characteristics. The traffic source was 4000 frames of the *Newsclip* from Section 4.3, and to achieve maximum or saturation throughput, fully filled packets were sent. Buffer adjustment refers to changing the number of protected P-frame packets in zone 2 according to the policy of Section 3.2.

For the plot without buffer adjustment, the boundaries between zones were set statically according to the size ratio of 6 : 3 : 2, and a linear UP mapping function is applied instead of the nonlinear mapping function of Figure 10.

For the plot with buffer adjustment, the zones were set according to the actual ratio of sizes between the frame types, averaged over the sequence. In that plot, within zones 1 and 3, the plot is linear. A small nonlinearity is present as buffer fullness crosses the boundary between zone 1 and zone 2 because of the quantization effect of taking ten categories of P-frame macroblock ratio. However, in general, zone-2 maximum throughput, when buffer adjustment is applied, is linear.

This is not the case if no buffer adjustment is applied, as a sudden increase in throughput occurs when the boundary between zones 1 and 2 is crossed. This is because more P-frame packets are sent at the higher bit rate, thus increasing the overall throughput. No account is taken of a relative increase in the number of arriving P-frame packets that are eligible for protection when no buffer adjustment takes place.

It should be noted that the overall throughput under the static zone boundary plot is down on that when buffer adjustment and monitored boundary setting take place. This implies that too many packets are being protected, because the lower bit rate is used more often. However, a consequence of this is that the buffer occupancy is increased, which is likely to lead to greater packet loss through buffer overflow for certain types of cross-traffic. Conversely, had a policy of no buffer adjustment been applied to a monitored zone boundary setting, the result would have been an influx of P-frame packets at the higher bit rate. This in turn leads to a greater number of packets with errors and consequently lower received video quality.

5.2. UP behavior with cross-traffic

In this section, cross-traffic is applied according to the scenario of Figure 10, while the *Newsclip* sequence from Section 4.3 forms the MPEG2 video stream. The single-state and two-state noise models are those described in Section 3.1.

In the first set of simulations, the cross-traffic was CBR at a rate of 200 Kbps and payload packet size of 800 B. The transport protocol for CBR was set as UDP. As introduced in Section 1, PSNR is the normal objective metric for comparison of video quality. As PSNR is a relative metric, it is reliable when making comparisons between the PSNRs for the same video clip. The higher the PSNR is, the better will be the quality, with a level around 40 dB presenting excellent quality for mobile communication, while levels below 25 dB are probably unwatchable. Though some fluctuations in quality are unavoidable, fluctuations in quality are subjectively disconcerting, especially when the level drops below 25 dB. The reader is referred to [39] for further comparisons of video quality under wireless communication.

The channel noise model was initially set to the single-state model of Section 3.1. In Figure 15(a), the UP scheme was applied with both dynamic zone boundary changing and zone-2 buffer adjustment. Compared to Figure 15(b), when




FIGURE 15: Video quality with CBR cross-traffic, and a single-state channel model (a) with the full UP scheme, (b) without UP at 2 Mbps, and (c) without UP at 3 Mbps.

all packets are protected on the RF channel, video quality is clearly improved both in the overall PSNR level and in the

fluctuation in quality. The drop in quality is due to packet loss through buffer overflow (see later comments in this section on buffer fullness). In Figure 15(b), it is apparent that there is an initial burst of high-quality video reception at 40 dB, and this is because the CBR source was not turned on till after this period. Figure 15(c) is less easy to discriminate by visual inspection, but summary results presented shortly show the advantage of UP with adaptive modulation.

The channel noise model was now set to the two-state model of Section 3.1. A comparison is additionally made with the CQDDR scheme of Section 3.5. In Figure 16(a) for the UP scheme, the video quality over time does not differ greatly from that of Figure 15(a). In Figure 16(b), the drops in quality owing to packet losses are more severe compared to those when there is a single-state AWGN channel (see Figure 15(b)). Because of the more severe channel conditions during bad states, a pure 3 Mbps rate results in a severe drop in video quality, as illustrated by Figure 16(c). Lastly, though a CQDDR model is certainly an improvement to a single sending rate policy, it is apparent from Figure 16(d) that the average video quality is below that of the UP scheme.

Table 3 shows that adaptive modulation with buffer management achieves superior video quality, as more packets are lost due to RF interference when transmitting exclusively at the higher bit rate. Table 3 also includes the results of simulations with the *Friends* and *Football* sequences, confirming that adaptive modulation maintains its advantages for different types of video stream. This is the case whether one- or two-state channel model is assumed. In the two-state model, packet losses at the 3 Mbps rate increase owing to the increased likelihood of packet error on the channel. Though CQDDR has its advantages, for video the UP scheme is superior as it also takes into account the packet content as well as traffic conditions.

Corresponding buffer fullness during the CBR cross-traffic simulations of Figure 15 is recorded in Figure 17. Once the CBR cross-traffic starts, after 6 seconds (see Figure 17), the buffer fullness with UP applied settles to a constant level, more than 10 packets below the 50-packet buffer capacity. At a gross rate of 2 Mbps, with SNR at 16 dB, packet loss due to RF interference is minimal. However, when all packets are protected, the buffer remains close to the capacity, and consequently packets are lost. Finally, transmitting all packets at the highest rate without UP brings no risk of packet loss due to buffer overflow, but as in Table 3 packet loss still occurs through RF interference. Thus, in this example, the 2 Mbps rate without UP cannot cope with the arrival rate of the video, causing the buffer to become saturated. The 3 Mbps rate without UP can cope with the arriving video stream, causing the buffer to scarcely be filled, but this rate is prone to RF interference. Employing UP with adaptive modulation allows for a choice between these two extremes.

Corresponding buffer fullness during the CBR cross-traffic simulations of Figure 16 is recorded in Figure 18. The 3 Mbps rate still causes the buffer to be emptied without risk of overflow, thus confirming that the drop in video quality at this rate is due to packet loss through RF interference. Both the UP adaptive modulation scheme and the CQDDR scheme suffer from potential packet loss through

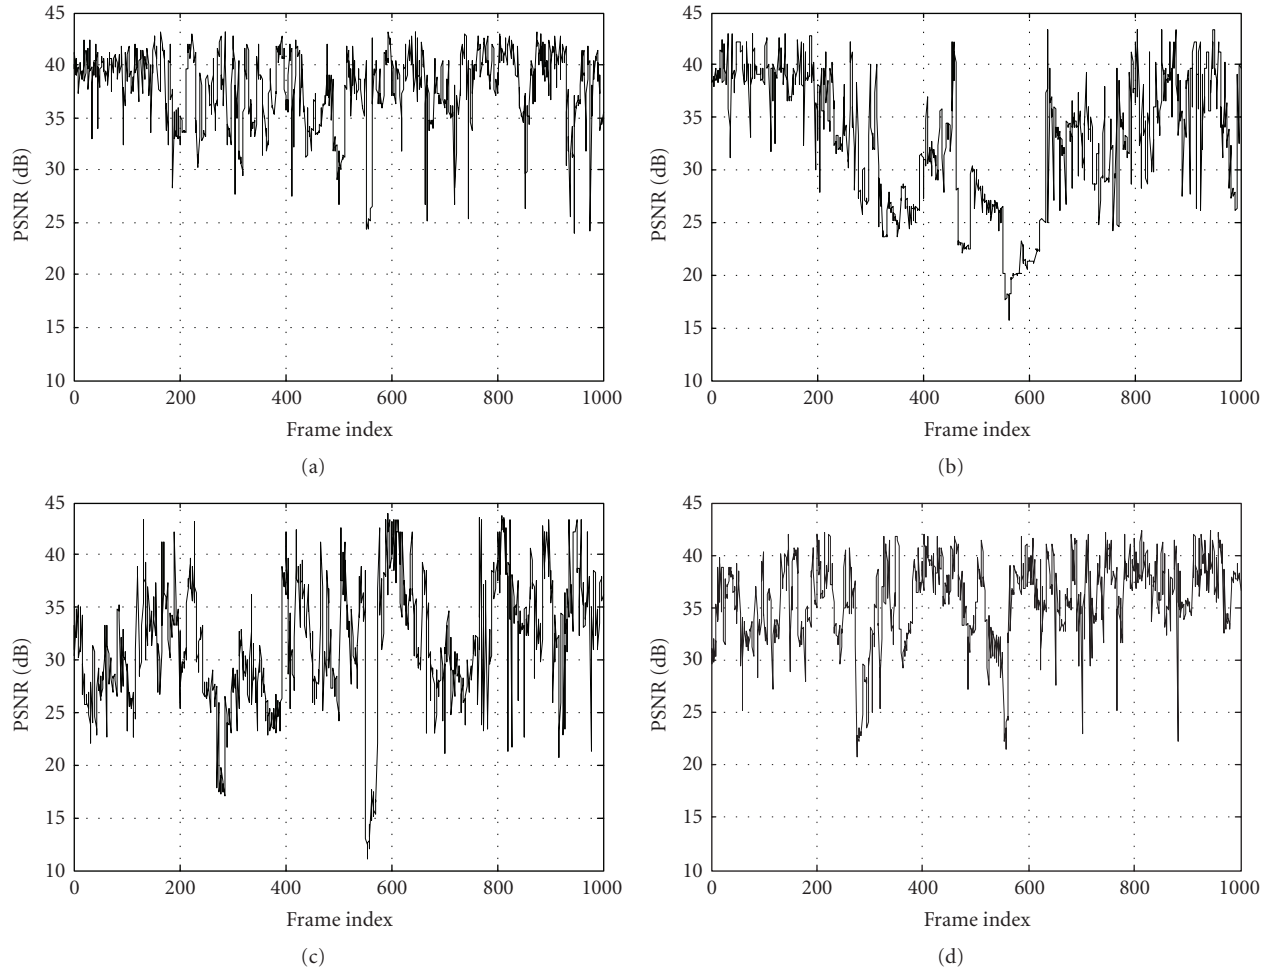


FIGURE 16: Video quality with CBR cross-traffic, and a two-state channel model (a) with the full UP scheme, (b) without UP at 2 Mbps, (c) without UP at 3 Mbps, and (d) with CQDDR.

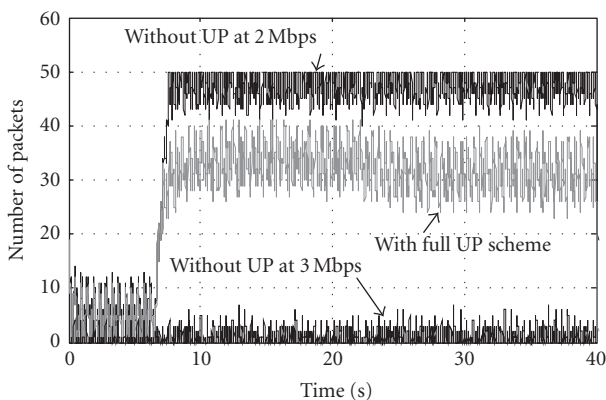


FIGURE 17: Buffer fullness with CBR cross-traffic for a one-state channel model.




FIGURE 18: Buffer fullness with CBR cross-traffic for a two-state channel model.

buffer overflow during channel bad states. However, CQDDR evidently chooses the higher 3 Mbps gross rate more frequently, leading to an emptier buffer but an increased risk of loss of more important anchor frame packets. This ex-

plains the resulting lower video quality of CQDDR recorded in Table 3.

TABLE 3: Mean video quality with CBR cross-traffic.

Video clip	Scheme	Single-state channel model		Two-state Markovian channel model	
		PSNR (dB)	Packet loss	PSNR (dB)	Packet loss
<i>Newsclip</i>	UP	38.06	5.08%	37.85	6.24%
	CQDDR	—	—	35.41	9.03%
	2 Mbps	33.15	12.10%	32.71	12.81%
	3 Mbps	34.05	9.53%	31.35	13.33%
<i>Football</i>	UP	37.46	6.31%	37.19	7.51%
	CQDDR	—	—	35.83	8.42%
	2 Mbps	32.24	14.67%	32.01	15.03%
	3 Mbps	33.98	10.94%	32.19	14.67%
<i>Friends</i>	UP	38.30	4.57%	37.92	5.83%
	CQDDR	—	—	36.11	8.94%
	2 Mbps	33.19	11.66%	32.87	12.06%
	3 Mbps	35.09	7.07%	32.39	12.54%

TABLE 4: Mean video quality with Web cross-traffic.

Video clip	Scheme	Single-state channel model		Two-state Markovian channel model	
		PSNR (dB)	Packet loss	PSNR (dB)	Packet loss
<i>Newsclip</i>	UP	39.11	2.19%	38.52	4.45%
	CQDDR	—	—	37.86	6.12%
	2 Mbps	37.61	6.42%	37.21	7.21%
	3 Mbps	33.98	11.13%	31.30	13.39%
<i>Football</i>	UP	38.87	3.18%	37.65	5.47%
	CQDDR	—	—	37.11	7.53%
	2 Mbps	37.21	8.59%	36.24	9.24%
	3 Mbps	33.09	14.27%	32.44	15.01%
<i>Friends</i>	UP	38.89	2.65%	38.10	5.12%
	CQDDR	—	—	37.87	6.20%
	2 Mbps	37.66	6.11%	37.22	6.89%
	3 Mbps	34.08	9.88%	32.41	12.49%

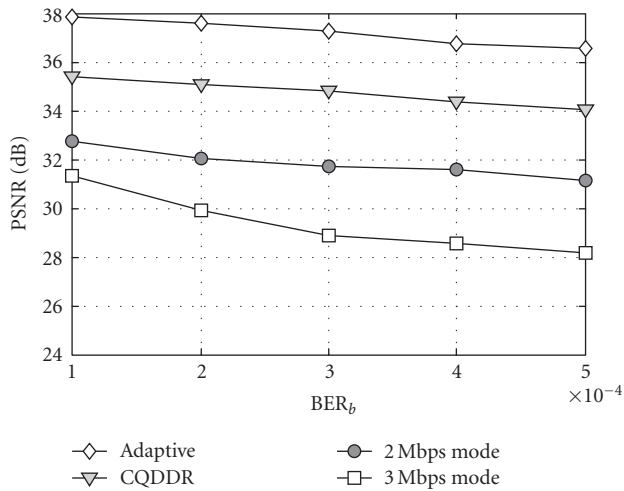


FIGURE 19: Mean video quality with CBR cross-traffic for a two-state channel model with varying bad-state BER (for a 3 Mbps gross rate).

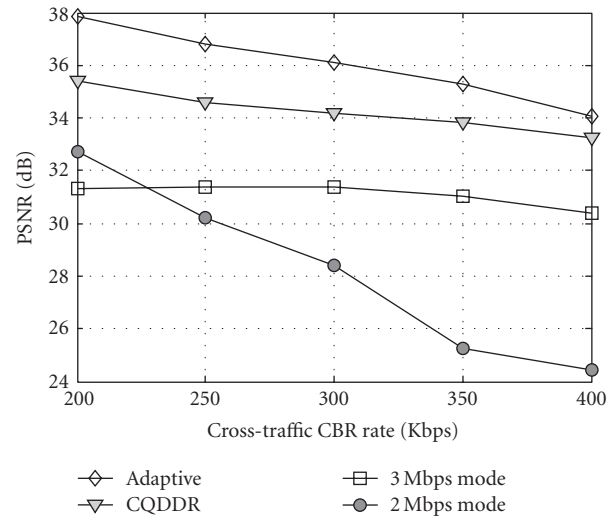





FIGURE 20: Mean video quality for different CBR cross-traffic intensities for a two-state channel model.



(a)




(b)




(c)

FIGURE 21: Video quality with Web cross-traffic (a) with the full UP scheme, (b) without UP at 2 Mbps, and (c) without UP at 3 Mbps.


To further judge the impact of channel conditions, the BER for a 3 Mbps gross rate in the bad state of the two-state Gilbert-Elliott model was varied as $i \times 10^{-4}$, with $i = 1, 2, 3, 4, 5$, while the remaining model parameter settings of Section 3.1 were retained. In Figure 19, for the *Newsclip* video, the mean PSNR deteriorates with increasing BER, as one might expect. The 3 Mbps rate mode suffers relatively severely from packet loss due to RF interference compared to



(a)



(b)



(c)

FIGURE 22: Buffer fullness with Web cross-traffic (a) with the full UP scheme, (b) without UP at 2 Mbps, and (c) without UP at 3 Mbps.

that of the 2 Mbps rate. The superior performance of the UP adaptive modulation scheme compared to CQDDR is confirmed across the range of BERs.

The impact of increasing the intensity of the CBR background traffic was also simulated. From Figure 20, it is apparent that, as the CBR rate increases, the delivered video quality of the UP adaptive modulation scheme and CQDDR starts to converge. This is because the UP scheme is increasingly

more likely to lose packets through buffer overflow. This risk is highlighted by the impact on the 2 Mbps rate. Because the service rate of the send buffer is reduced by the presence of more CBR packets, an increasing number of packets are discarded from the buffer, leading to a rapidly deteriorating delivered video quality. More importantly, with this increase in buffer occupancy, a smaller number of packets are eligible to be protected by the UP scheme, and so the performance of UP scheme starts to converge to the CQDDR scheme.

In the second set of simulations, under the same conditions as those of the previous set for the single-state channel model, the cross-traffic was from a Web server. HTTP over TCP transport was set in the NS-2 simulations. The Web traffic had a mean interpage request time of 2 seconds with an exponential distribution. A mean of 5 embedded objects within each page was set, with the number again being exponentially distributed. The mean object size was 20 KB, with a Pareto distribution with shape factor set to 1.2. Again, the Web traffic source was not turned on for about the first 150 video frames.

For this typical Web traffic source, Figure 21 reports the impact upon video quality. The pattern of PSNR results broadly follows that for CBR cross-traffic. Table 4 summarizes the results, from which it is apparent that less loss occurs due to buffer overflow at the 2 Mbps rate when Web traffic is present. Again, the results for the other video sequences under test are included in Table 4, to demonstrate that the result for the *Newsclip* is not an isolated result. Table 4 also includes a set of results for the two-state channel model. These follow the trends of the one-state model, though in all cases there is deterioration in mean PSNR. The UP scheme remains superior to CQDDR in terms of delivered video quality.

In Figure 22(a), buffer fullness is reported for the UP scheme for the single-state channel, when it is apparent that the buffer rarely reaches a level (50 packets when completely full) such that packet loss can occur. However, due to the slower transmission rate, from Figure 22(b) it is clear that transmitting exclusively at 2 Mbps exposes the video packets to an increased risk of being dropped from the transmit buffer. At the higher transmission rate (see Figure 22(c)), all packet loss is due to the impact of the AWGN channel, as the buffer is under utilization.

Comparing Tables 3 and 4, it is apparent that the fixed modulation schemes change in ranking with respect to delivered video quality. As cross-traffic characteristics are not generally known in advance, this further disadvantages a fixed scheme without UP. Two adaptive schemes were compared, but CQDDR without content-type awareness underperforms compared to the UP adaptive modulation scheme. Unfortunately, for video, this difference would be noticeable to the viewer, especially when the quality drops significantly owing to error bursts, which may give rise to “freeze frames.”

6. CONCLUSION

For delay-sensitive applications such as video streaming, reliable data delivery cannot simply be achieved by retransmission of packets. Due to the fragility of encoded data, it is also necessary to protect the most important information. Un-

equal protection in Bluetooth streaming has been shown by us to achieve a significant improvement in delivered video quality over the best fixed bit rate scheme according to cross-traffic conditions. In terms of delivered video quality, the UP scheme also consistently outperforms a classic Bluetooth CQDDR scheme in which the data rate is adjusted according to channel conditions, though without consideration of packet content. The paper shows that an unequal protection scheme ought to be dynamic, as the content-importance characteristics change within a video sequence. The scheme introduced accounts for varying ratios of frame-type sizes and intracoded macroblocks arising from the occurrence of scene changes, rapid motion, camera pans, zooms, and so on. While high-quality video, at around 40 dB for a TV clip of CIF pixel size at 25 fps, is delivered through unequal protection, a single bit rate option will result in an overall drop in quality, and furthermore it will behave differently depending on the cross-traffic present. A CQDDR scheme is preferable, but for video over Bluetooth it is suboptimal.

REFERENCES

- [1] J. Haartsen, “The bluetooth radio system,” *IEEE Personal Communications*, vol. 7, no. 1, pp. 28–36, 2000.
- [2] R. Razavi, M. Fleury, and M. Ghanbari, “Low-delay video control in a personal area network for augmented reality,” in *Proceedings of the 4th Visual Information Engineering*, pp. 1245–1300, London, UK, July 2007.
- [3] S. Ye, R. S. Blum, and L. J. Cimini Jr., “Adaptive modulation for variable-rate OFDM systems with imperfect channel information,” in *Proceedings of the 55th IEEE Vehicular Technology Conference (VTC '02)*, vol. 2, pp. 767–771, Birmingham, Ala, USA, May 2002.
- [4] M. Sajadieh, F. Kschischung, and A. Leon-Garcia, “Modulation-assisted unequal error protection over the fading channel,” *IEEE Transactions on Vehicular Technology*, vol. 47, no. 3, pp. 900–908, 1998.
- [5] “Core specification of the bluetooth system, version 2.1 + EDR,” July 2007, <http://www.bluetooth.com>.
- [6] Q. Li and M. van der Schaar, “Providing adaptive QoS to layered video over wireless local area networks through real-time retry limit adaptation,” *IEEE Transactions on Multimedia*, vol. 6, no. 2, pp. 278–290, 2004.
- [7] A. Iyer and U. B. Desai, “A comparative study of video transfer over bluetooth and 802.11 wireless MAC,” in *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC '03)*, vol. 3, pp. 2053–2057, New Orleans, La, USA, March 2003.
- [8] C. H. Chia and M. S. Beg, “Realizing MPEG-4 video transmission over wireless bluetooth link via HCI,” *IEEE Transactions on Consumer Electronics*, vol. 49, no. 4, pp. 1028–1034, 2003.
- [9] R. Kapoor, M. Kazantzidis, M. Gerla, and P. Johansson, “Multimedia support over bluetooth piconets,” in *Proceedings of the 1st Workshop on Wireless Mobile Internet*, pp. 50–55, Rome, Italy, July 2001.
- [10] J. Y. Khan, J. Wall, and M. A. Rashid, “Bluetooth-based wireless personal area network for multimedia communication,” in *Proceedings of the 1st IEEE International Workshop on Electronic Design, Test, and Application*, pp. 47–51, Christchurch, New Zealand, January 2002.

- [11] C. Scheiter, R. Steffen, M. Zeller, R. Knorr, B. Stabernack, and K.-W. Wels, "A system for QOS-enabled MPEG-4 video transmission over bluetooth for mobile applications," in *Proceedings of International Conference on Multimedia and Expo (ICME '03)*, vol. 1, pp. 789–792, Baltimore, Md, USA, July 2003.
- [12] S.-G. Miaou, C.-Y. Huang, K.-J. Ho, and N.-C. Tu, "Quality degradation and improvement of H.263 video transmitted in bluetooth packets under the interference of wireless LAN," in *Proceedings of IEEE Global Telecommunications Conference (GLOCOM '02)*, vol. 2, pp. 1738–1742, Taipei, Taiwan, November 2002.
- [13] M. F. Tariq, P. Czerepiński, A. Nix, D. Bull, and N. Canagarajah, "Robust and scalable matching pursuits video transmission using the bluetooth air interface standard," in *Proceedings of IEEE International Conference Transactions on Consumer Electronics*, vol. 46, pp. 673–681, Los Angeles, Calif, USA, August 2000.
- [14] R. Kapoor, M. Cesana, and M. Gerla, "Link layer support for MPEG video over wireless links," in *Proceedings of International Conference on Computer Communications and Networks*, pp. 477–482, Dallas, Tex, USA, October 2003.
- [15] S. Krishnamachari, M. van der Schaar, S. Choi, and X. Xu, "Video streaming over wireless LANs: a cross-layer approach," in *Proceedings of the 13th International Packet video Workshop*, Nantes, France, April 2003.
- [16] L.-J. Chen, R. Kapoor, M. Y. Sanadidi, R. Lee, and M. Gerla, "Audio streaming over bluetooth: an adaptive ARQ timeout approach," in *Proceedings of the 24th International Conference on Distributed Computing Systems*, vol. 24, pp. 196–201, Tokyo, Japan, March 2004.
- [17] C. Ru, L. Yin, J. Lu, and W. Chen, "A new UEP scheme for robust video transmission in MIMO System," *China Communications*, vol. 4, no. 5, pp. 102–108, 2006.
- [18] Y. Pei and J. W. Modestino, "Multi-layered video transmission over wireless channels using an adaptive modulation and coding scheme," in *Proceedings of IEEE International Conference on Image Processing*, vol. 2, pp. 1009–1012, Thessaloniki, Greece, October 2001.
- [19] X. Xu, M. van der Schaar, S. Krishnamachari, S. Choi, and Y. Wang, "Fine-granular-scalability video streaming over wireless LANs using cross layer error control," in *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP '04)*, vol. 5, pp. 989–992, Montreal, Canada, May 2004.
- [20] R. Razavi, M. Fleury, and M. Ghanbari, "Deadline-aware video delivery in a disrupted bluetooth network," in *Proceedings of IEEE Sarnoff Symposium*, Princeton, NJ, USA, April-May 2007.
- [21] B. Barmada, M. M. Ghandi, M. Ghanbari, and E. V. Jones, "Prioritized transmission of data partitioned H.264 video with hierarchical QAM," *IEEE Signal Processing Letters*, vol. 12, no. 8, pp. 577–580, 2005.
- [22] J. Goshi, R. Ladner, A. E. Mohr, E. A. Riskin, and A. Lippman, "Unequal loss protection for H.263 compressed video," in *Proceedings of the Conference on Data Compression (DCC '03)*, pp. 73–82, Snowbird, Utah, USA, March 2003.
- [23] Y. Shan, "Cross layer techniques for adaptive video streaming over wireless networks," *EURASIP Journal on Applied Signal Processing*, vol. 2005, no. 2, pp. 220–228, 2005.
- [24] P. Batra and S.-F. Chang, "Effective algorithms for video transmission over wireless channels," *Signal Processing: Image Communication*, vol. 12, no. 2, pp. 147–166, 1998.
- [25] J. Cai, Q. Zhang, W. Zhu, and C. W. Chen, "An FEC-based error control scheme for wireless MPEG-4 video transmission," in *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC '00)*, vol. 3, pp. 1243–1247, Chicago, Ill, USA, September 2000.
- [26] C.-M. Chen, C.-W. Lin, and Y.-C. Chen, "Packet scheduling for video streaming over wireless with content-aware packet retry limit," in *Proceedings of the 8th IEEE Workshop on Multimedia Signal Processing (MMSP'06)*, pp. 409–414, Victoria, Canada, October 2006.
- [27] E. N. Gilbert, "Capacity of burst-noise channel," *Bell System Technical Journal*, vol. 39, no. 8, pp. 1253–1265, 1960.
- [28] E. O. Elliott, "Estimates of error rates for codes on burst-error channels," *Bell System Technical Journal*, vol. 42, pp. 1977–1997, 1963.
- [29] N. Golmie, N. Chevrollier, and O. Rebala, "Bluetooth and WLAN coexistence: challenges and solutions," *IEEE Wireless Communications*, vol. 10, no. 6, pp. 22–29, 2003.
- [30] R. Fantacci and M. Scardi, "Performance evaluation of preemptive polling schemes and ARQ techniques for indoor wireless networks," *IEEE Transactions on Vehicular Technology*, vol. 45, no. 2, pp. 248–257, 1996.
- [31] ITU-T Recommendation P.210, "Subjective video quality assessment methods for multimedia applications," September 1999.
- [32] R. Razavi, M. Fleury, E. Jammeh, and M. Ghanbari, "An efficient packetization scheme for bluetooth video transmission," *Electronic Letters*, vol. 42, no. 20, pp. 1143–1145, 2006.
- [33] N. B. Abramson, "A class of systematic codes for non-independent errors," *IEEE Transactions on Information Theory*, vol. 5, no. 4, pp. 150–157, 1959.
- [34] J. H. Yoon, S.-B. Lee, and S.-C. Park, "Packet and modulation type selection scheme based on channel quality estimation for bluetooth evolution systems," in *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC '04)*, vol. 2, pp. 1014–1017, Atlanta, Ga, USA, March 2004.
- [35] M. Ghanbari, *Standard Codecs: Image Compression to Advanced Video Coding*, IEE Press, Stevenage, UK, 2003.
- [36] M. K. Honig and D. G. Messerschmitt, *Adaptive Filters Structures, Algorithms, and Applications*, Kluwer Academic, Boston, Mass, USA, 1990.
- [37] Y.-Z. Lee, R. Kapoor, and M. Gerla, "An efficient and fair polling scheme for bluetooth," in *Proceedings of IEEE Military Communications Conference (MILCOM '02)*, vol. 2, pp. 1062–1068, Anaheim, Calif, USA, October 2002.
- [38] R. Razavi, M. Fleury, and M. Ghanbari, "Detecting congestion within a bluetooth piconet: video streaming response," in *London Communications Symposium*, pp. 181–184, London, UK, September 2006.
- [39] A. H. Sadka, *Compressed Video Communications*, John Wiley & Sons, Chichester, UK, 2002.