

Unequal Protection of Encoded Video Streams in Bluetooth EDR

R. Razavi, M. Fleury, and M. Ghanbari
University of Essex, Colchester, UK

Abstract—Bluetooth’s Enhanced Data Rate (EDR) modulation schemes can potentially support higher quality video streams. To avoid the impact of RF noise, this paper proposes adaptive modulation, previously unavailable to Bluetooth, based on content type importance at the picture and macroblock levels. Because, as a result, a lower bit rate for protected data may occur, buffer management is introduced to avoid transmit buffer overflow. The paper also proposes a tri-zone buffer with varying unequal protection (UP) policy in each zone. Application of these UP techniques results in up to 4 dB improvement in PSNR.

I. INTRODUCTION

Bluetooth [1], standardized as IEEE 802.15.1, is a short-range, radio frequency (RF) interconnect, which can be expanded to form a piconet, with one master node and up to seven active slaves. Version 2 of Bluetooth now includes two additional modulation types. In this paper, we introduce unequal protection (UP) of encoded video according to picture type through adaptive modulation. Applications include Bluetooth access networks, video sensor networks, and augmented reality over a personal area network [2]. Our primary method of protection is by adaptive modulation but transmit buffer management is applied as a necessary secondary mechanism. In combining these techniques, the paper reports an upper bound improvement in video quality (PSNR) of between about 2 and 4 dB from employing UP over the best fixed-modulation scheme without protection.

The UP policy is applied by monitoring transmitter buffer fullness, available through Bluetooth’s Host Controller Interface. The buffer zone boundaries are dynamically changed according to the relative size ratio of arriving I-, P- and B-pictures. Buffer fullness is responsive not only to buffer congestion from an arriving video stream but to an increase in buffer service time when piconet cross traffic is present. As buffer fullness reflects the congestion of the Bluetooth wireless channel, we use it to regulate the UP scheme, rather than other forms of noise estimation.

Because a lower bit rate is employed for priority packets, this runs the risk of buffer overflow at the transmit buffer compared to when transmitting all packets at the higher bit rate. Therefore, this paper proposes a tri-zone buffer with different UP policy for each zone. As the buffer fullness increases, packets begin to fill the second and then third zone, and accordingly the prioritization policy changes. Ideally, the output at the lower bit rate should decrease linearly with buffer fullness. However, to achieve this linear UP policy based simply on picture type will not work, because of a varying number of packets between the picture types. Consequently, the paper demonstrates how, in the

second zone of the tri-zone buffer, the number of P-picture packets offered protection can be governed by their importance according to macroblock type, given that there may be intra-coded macroblocks present as well as the normal predictive macroblocks and possibly SKIP macroblocks. Therefore, the principal contribution of the paper is the tri-zone buffer and its combination with adaptive modulation for UP. This scheme based for the most part on incremental priority is relatively simple to implement and maps onto Bluetooth’s available modulation types. A non-linear prioritization scheme taken across the picture types and involving a number of weighting factors (in the manner of a constrained optimization problem) would involve a computational burden difficult to support in real-time on Bluetooth devices or would involve encoder intervention or transcoding.

The remainder of this paper is organized as follows. Section II considers some related work. Section III is an overview of the Bluetooth UP system with adaptive modulation and tri-zone buffering. Section IV describes the features of the scheme in more detail. Section V contains illustrative results, Section VI is a discussion of wider issues resulting from the study, while finally Section VII draws some conclusions.

II. RELATED WORK

Unequal error protection for more important data through enhanced application or channel error coding is widely practiced, for example [3]. However, it is also possible to use modulation adaptation [4] for the same purpose without the need for coding complexity. For example, in [5], the video stream is partitioned through multi-description coding (with some redundancy) and each sub-stream is adaptively modulated and transmitted through an antenna array in a Multiple Input Multiple Output (MIMO) antenna system.

Adaptive modulation can also be applied [6] through multi-layering (at a cost in flexibility). In our paper, rather than partition a stream into separate streams, individual parts of the same stream are protected. Therefore, this paper’s techniques are applicable to single stream video.

In prior work involving one of this paper’s coauthors [7], motion vectors and other header data through H.264 data partitioning are prioritized through hierarchical Quadrature Amplitude Modulation (QAM) for Orthogonal Frequency Division Modulation (OFDM). However, that scheme was for Digital Video Broadcasting (DVB), and is not applicable to Bluetooth, as Ultra WideBand (UWB) is planned for version 3.

Four categories of MPEG-4 information occur in [8] for priority ARQ: header, I- and P-pictures with scene changes, shape and motion information in P-pictures, and fourthly

texture information in P-pictures. A finer level of data prioritization may also be applied by inspecting the number of intra-coded macroblock in an H.263 bitstream, though in [9] protected by ARQ and FEC.

Intra-coded macroblocks, as monitored in this paper, may appear in P-pictures as well as I-pictures, and may indicate scene changes, camera zooms or pans, and so on. The presence of intra-coded macroblocks, which is encoder implementation dependent, indicates important information in the encoded bitstream, though prior research in [9] did not associate them with the pictures themselves and did not employ adaptive modulation.

III. UP SYSTEM MODEL

A. Bluetooth video streaming

Like IEEE 802.11b,g (WiFi) Bluetooth operates in the unlicensed 2.4 GHz band. Bluetooth's short range (less than 10 m for class 2 devices), robust Frequency Hopping Spread Spectrum (FHSS) and centralized medium access control through Time Division Multiple Access (TDMA) and Time Division Duplex (TDD) means it is less prone to interference from other Bluetooth networks. Bluetooth employs variable-sized packets up to a maximum of five frequency-hopping time-slots of 625 μ s in duration. In TDD for transmit/receive separation, every Bluetooth frame consists of a packet transmitted from a transmitter node over 1, 3 or 5 timeslots, while a receiver replies with a packet occupying at least one slot, so that each frame has an even number of slots.

Bluetooth v.1.2 received comparatively limited investigation as a medium for streaming video. Research papers that do investigate Bluetooth streaming include [10]-[13]. However, B/T v. 2.0 [14] 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 MPEG-2 transport stream (TS), as used in DVB-T,H.

In fact, Bluetooth v. 2.0 supports both a gross air rate of 3.0 Mbps and 2.0 Mbps (mgup of 1.4485 Mbps), through respectively $\pi/4$ -DQPSK or 8DPSK modulation. This implies that, through adaptive modulation, a lower bit rate is available that can serve to give UP to some of the packets of the more important pictures, the Intra (I) and Predictive (P) anchor picture, as well as some bi-predictive (B-) picture packets, depending on circumstances.

This research was pleased to avail itself of the University of Cincinnati Bluetooth (UCBT) extension to the well-known ns-2 network simulator (v. 2.28 used). The UCBT extension supports Bluetooth EDR but is also built on the air models of previous Bluetooth extensions such as BlueHoc from IBM and Blueware.

B. Cross-layer interaction

In Fig. 1, showing the logical functionality of the unequal protection system, the encoded MPEG2-TS data enter a one-frame buffer, prior to Bluetooth packetization. The stream may be encapsulated as an arriving IP packet, though this is not a pre-requisite for the scheme to work and is simply a suggested application. Within the frame buffer, the UP system determines from the encoded data the type of frame, its size, and, if a P-picture, the ratio of intra-coded macroblocks to other macroblock types. The picture

information is passed to a decision unit that allocates the priority of the resulting Bluetooth packets as they are passed into the FIFO transmit buffer. The prioritizing decision is affected by the state of buffer fullness and the importance of the incoming Bluetooth packet. In an implementation, a packetization unit would break the frame up into Bluetooth packets. The packetization policy based on partition of the IP packets within the one frame buffer is described in Section III.D. The priority allocated by the decision unit would be marked within a temporary per-packet tag, serving as an instruction to the tri-zone buffer management unit, that no doubt would be present in an implementation.

Within the transmit buffer, priority marked Bluetooth packets are transmitted by either of the two modulation schemes, depending on priority. As already mentioned, low priority packets are sent at 3.0 Mbps, as this rate is subject to the largest risk of error from RF noise.

In the tests of Section V, an Additive White Gaussian Noise (AWGN) channel is modeled, with a Bit Error Rate (BER) of 10^{-5} corresponding at the higher data rate to an E_b/N_0 of 16 dB. This assumes an indoor environment and neglects fast fading and multipath. To avoid complicating the analysis, Bluetooth's Automatic Repeat Request (ARQ) mechanism was turned off. The default setting for Bluetooth ARQ is unlimited retransmissions, which is unsuitable for delay sensitive applications.

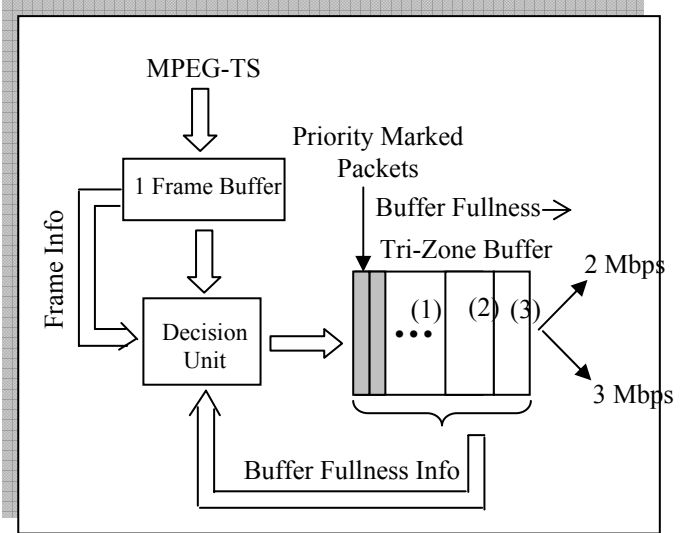


Fig. 1 Unequal protection system for video data

C. Buffer UP policy

In zone 1 of the buffer, all Bluetooth packets of type I- or P-picture are protected through dispatch at the lower bit rate. B-picture packets are protected in zone 1 according to the following procedure. A uniformly distributed random number in the interval [0,1] is generated and compared to the fraction $f = \text{zone packet occupation} / \text{zone capacity}$. If the random number is greater than f then that B-picture packet is also protected. The point of this common procedure is to establish a long term distribution according to ratio f , which would not otherwise occur if a deterministic decision was made.

As the buffer fullness increases and packets also occupy zone 2 of the buffer then a different prioritization policy for P-pictures is applied. I-picture packets remain protected within zone 2 of the buffer and B-picture packets are no longer protected. P-picture packets in zone 2 of the buffer

are protected according to the ratio of intra-coded macroblocks to other macroblocks within the picture, as detected when the picture was in the frame buffer. Again, the boundary between protected and unprotected P-picture packets is dynamically adjusted according to a past history of intra-coded macroblock ratios within P-pictures. Zone 2 buffer adjustment is further explained in Section IV.B.

Finally, in zone 3 of the buffer, when the buffer is at its fullest, no protection to any B- or P-pictures packets is applied. However, I-picture packets are protected according to the same policy applied for zone 1, i.e. by random number generation and comparison with a fraction f for zone 3.

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

D. Bluetooth packetization

As previously mentioned, a data frame across a Bluetooth link in asymmetric mode consists of a multi-time slot Asynchronous Connection-Less (ACL) packet. Because of packet quantization effects, the Bluetooth packet sizes become significant and their effect on user payload are summarized in Table I for a single master-slave ACL link for Bluetooth v. 2.0. As an example, a 3DH5 packet has a gross air rate of 3.0 Mbps and occupies five time slots. Data High (DH) contrasts with the Data Medium rate of Bluetooth v. 1.2, which was 1.0 Mbps and is retained for header transmission.

The normally assumed Bluetooth controller behavior is that, given a chosen 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 pre-set maximal scheme, then a reduced scheme is used. For example, the controller swaps from 3DH5 down to 3DH3 or even 3DH1. From Section III.B, IP packets accumulate within a one frame buffer before packetization after their priority has been determined by the decision unit of Fig. 1.

Unfortunately, if packetization takes place on a single MPEG-2 slice (one row of macroblocks) per Bluetooth packet, the possibility of many partially filled packets and many 1- or 3-slot packets is introduced. The result is a drop in throughput. Therefore, in [15] 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 [15] it is shown that for medium to good SNR, the overall video performance is superior. For poor SNR (below about 10 dB), smaller packets and a slice per packet may still be appropriate, but such conditions are outside the scope of this paper, as they require special measures [16]. Of course, though H.264 introduces variably sized slices, it is not possible to easily predict which slice size is optimal, given volatile channel conditions. In the experiments in Section V, the video Bluetooth packet size was either set to 3DH5 or 2DH5, depending respectively on whether a gross rate of 3.0 or 2.0 Mbps was chosen.

IV. METHODOLOGY

A. Buffer partition size allocation

In Fig. 2, for an MPEG-2 Common Intermediate Format (CIF)-sized video sequence (an episode of the situational comedy ‘Friends’) encoded at 1.2 Mbps at 25 fps, with Group of Pictures (GOP) structure $N=12$, $M=3$, the relative sizes of I-, P-, and bi-predictive B-pictures were monitored. In fact, as occurred in practice, averaging over 10 GOPs produces little change in the pattern. It will be seen that a static ratio of 6:3:2 for I-, P-, and B-pictures is a good fit [17], as the illustrative example from Friends demonstrates in confirmation of more general findings reported in [17]. However, the relative size of P-pictures and at the same time B-pictures may well increase in comparison to I-pictures, which is why adaptive partition adjustment is included to avoid reliance on statistical data from an inevitably limited set of example sequences.

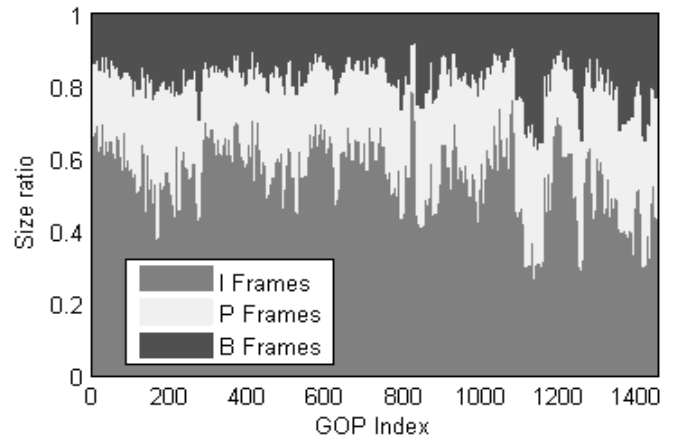


Fig. 2. Example distribution of frame sizes by frame type per GOP for an MPEG-2 video sequence.

To consider how the buffer zone boundaries are allocated, firstly take the static size ratio of 6:3:2 between the different picture types. Within a GOP structure of $N=12$ and $M=3$, the frequency of picture types is in the ratio 1:3:8. Therefore, by simple multiplication of the three ratios the buffer zone sizes would be in the ratio 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. As the actual monitored ratio of sizes varies dynamically, Fig. 2, in the implemented algorithm, for every N GOPs, where N is a pre-set constant, the average size ratio of the different pictures types is found and the zone allocation is accordingly adjusted. By observation, a reasonable value of N was found to be 10 GOPs. In other words, the size ratios was wide-sense stationary in time over this GOP interval, while larger values result in a worse fit and smaller monitoring intervals are inefficient.

TABLE I.
BLUETOOTH EDR PACKET TYPES: USER PAYLOAD AND BIT RATES

Packet type	User payload in bytes	Asymmetric max. rate (kbps)
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

B. P-picture macroblock type prioritization

While I-pictures are formed entirely by intra-coded macroblocks in MPEG-2, P-pictures, apart from macroblocks of type predictive and SKIP (no update of matching macroblocks from the prior picture), may also include intra-coded macroblocks. Fig. 3 plots the ratio of intra-coded macroblocks within P-pictures for a ‘Football’ sequence. The ‘Football’ sequence has the same GOP structure as the ‘Friends’ sequence and was CIF-sized at 25 fps. It is chosen as an illustration, as there is rapid motion and, between P-pictures indexed 65, Fig. 3 (b), and 66, Fig. 3 (c), a scene change occurs from a wide view of the pitch to a close up of players. The plot in Fig. 3 (a) shows a sharp peak in the ratio of intra-coded macroblocks for these P-picture indices, and for others. As matching macroblocks in subsequent pictures (after P-picture index 66) depend for coding on these macroblocks, until the arrival of the next I-picture, it is important that they are delivered intact to the decoder. Notice that in general the distribution of P-pictures with a high intra-coded ratio is dependent on film genre and motion content, and Fig. 3 should not be taken as typical.

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

Fig. 4 plots the ratio of intra-coded macroblocks in P-pictures for the ‘Friends’ sequence of Section IV.A. Fig. 5 shows the resulting distribution over the P-pictures, grouped into the ten categories used by the current algorithm (though for 1000 P-pictures in this example rather than the 100 used in practice). The derived mapping function is plotted in Fig. 6 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 Fig. 6.

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-pictures that have more than 62.4% of their macro-blocks intra-coded are protected. At any time, if the current number of packets in zone 2 and the ratio of intra-coded macroblocks of a given picture is known, the decision can be made easily.

The mapping function is formed by taking the set of ten probabilities, such as those in Fig. 5, and projecting them onto the zone-2 capacity. For example, in Fig. 5 the 0.1 ratio

of intra-coded macroblocks has a probability of approximately 0.25.





Fig. 3. Example distribution of macroblock types within P-frames, with (a) Frequency of intra-coded macroblocks (b) Frame 65 macro-block types (c) Frame 66 macroblock types, with grey circles = predictive-, black = SKIP, and white = intra-coded macroblocks.

Therefore, $0.25 \times 50 = 13$ packets for a zone-2 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 Fig. 5. Data points are connected in piece-wise linear fashion.




Fig. 4 Intra-coded macroblock ratio for successive P-frames.

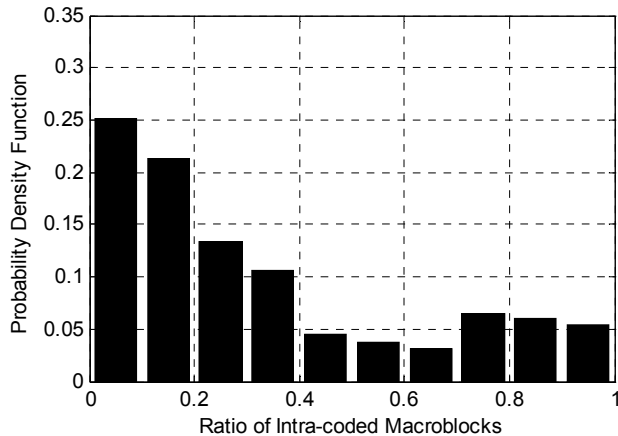


Fig. 5. Distribution of the ratios of intra-coded P-frame macroblocks from Fig. 3.

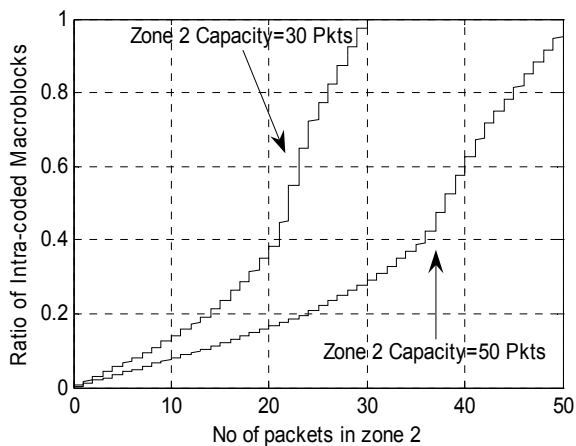


Fig. 6. Protection mapping function based on two different buffer zone 2 capacities.

C. Piconet congestion and buffer fullness

Fig. 7 shows the simulation configuration for the results of Section V. The encoded 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, Fig. 8. The Bluetooth standard does not specify the queue service discipline, and, along with Bluetooth implementations, this paper assumes pure round-robin (1-limited) scheduling. The work in [18] 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 [10], it is suggested that for congestion control the input packet rate to the shared RF channel should be increased (decreased) rate 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 [19], packet delay recorded at a Bluetooth receiver was found to be a better indicator of congestion than packet loss but resulted in oscillations in both video quality and delay in packet delivery when used as input for congestion control.

On the other hand, Figure 9 shows the ability of buffer fullness to track both variations in direct traffic (M to S1 in Fig. 7) and in cross traffic (S2 via M to S3 in Fig. 7). In [19], it was also shown that buffer fullness when applied to congestion control significantly reduces delay and improves PSNR. The video traffic rate plot in Fig. 9 reflects a fixed constant bit rate (CBR) cross traffic at 200 Kbps and packet size 800 B. Notice that this implies an effective bit rate of 400 Kbps across the shared channel, as the CBR traffic transported by UDP makes two hops to reach its destination. Equally, the packet size implies less than optimal use of the bandwidth capacity. The video traffic source was a 40 s MPEG-2 CIF-sized 25 fps ‘Newsclip’ (moderate motion) with GOP structure N=12 M=3, with fully filled packets. As its rate passes a threshold of around 1.6 Mbps, buffer fullness sharply climbs as the saturation rate of the Bluetooth link at 2.1 Mbps is approached. Similarly, with the MPEG-2 source rate fixed at 1.25 Mbps, when the CBR rate approaches channel saturation there is a sudden increase in buffer occupancy. It will be seen that the distraction of storing and forwarding higher rate cross traffic has a more severe effect at the master than increasing the locally source traffic.




Fig. 7 Bluetooth piconet with cross-traffic

V. RESULTS

A. UP behavior without cross-traffic

In Fig. 10, 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 tri-zone buffer characteristics. The traffic source was 4000 frames of the ‘Newsclip’ of Section IV.C but in this instance, to achieve maximum or saturation throughput, fully filled packets were sent back-to-back. Buffer adjustment refers to changing the number of protected P-picture packets in zone 2 according to the policy of Section IV.B.




Fig. 8. The buffering model for Bluetooth

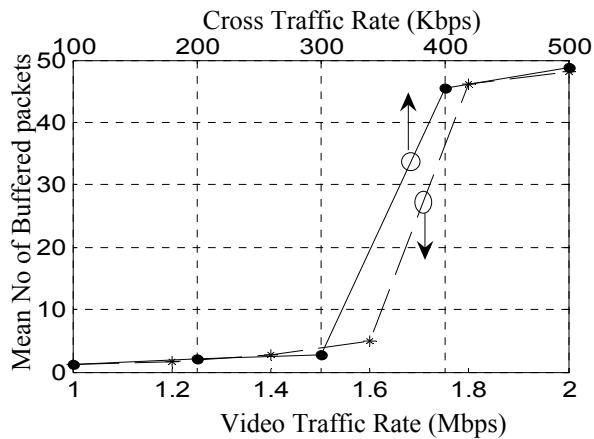


Fig. 9. Buffer fullness against varying cross traffic, and varying video rate.

For the plot without buffer adjustment, the boundaries between zones were set statically according to the size ratio 6:3:2 and a linear UP mapping function was applied instead of the non-linear mapping function of Fig. 6.

For the plot with buffer adjustment, the zones were set according to the actual ratio of sizes between the picture types, averaged over the sequence. In that plot, within zones 1 and 3 the plot is linear. A small non-linearity 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-picture 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 as the zone 1 to 2 boundary is crossed. This is because more P-picture 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-picture 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 takes place. This implies too many packets are being protected as 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 overflows, with certain types of cross-traffic. Conversely, had a policy of no buffer adjustment been applied to a

monitored zone boundary setting, the result would be an influx of P-picture packets at the higher bit rate. This in turn leads to a greater number of packets with error and, consequently, lower received video quality.




Fig. 10. The effect of size and content aware UP policy on throughput.

B. UP behavior with cross-traffic

In this Section, cross traffic is applied according to the scenario of Fig.7, while the 'Newscip' sequence of Section IV.C forms the main MPEG-2 video stream. The noise model is the one described in Section III.B.

In the first set of simulations, the cross traffic was CBR at a rate of 200 Kbps and payload packet size of 800 B. In Fig. 11 (a) the UP scheme was applied with both dynamic zone boundary changing and zone-2 buffer adjustment.

Compared to Fig. 11 (b), when 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 on buffer fullness). In Fig. 11 (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 until after this period (and similarly in Fig. 11 (a) and (c)). The CBR source was not turned on initially simply in order to better judge the effect when it is turned on. The video quality gain over the sending at a constant 3.0 Mbps, Fig. 11 (c) is less easy to discriminate by visual inspection. However, Table II 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.

Corresponding buffer fullness during the CBR cross traffic simulations of Fig. 11 is recorded in Fig. 12. Once the CBR cross traffic starts after 6 s, in Fig. 12 the buffer fullness with UP applied buffer fullness settles to a constant level, more than 10 packets below the 50 packet buffer capacity.

At a gross air rate of 2.0 Mbps, with SNR at 16 dB packet loss from RF interference is minimal. However, when all packets are protected, the buffer remains close to capacity and consequently packets are lost. Finally, transmitting all packets at the highest rate brings no risk of packet loss, but as in Table II, packet loss still occurs through RF interference. To illustrate the generality of the adaptive modulation scheme, Table II also includes PSNR and packet loss figures from the two other clips discussed in


Sections IV A and B. The same initial buffer boundaries were set for all of the sequences under test.

In the second set of simulations, under the same conditions as the previous set, the cross traffic was Web traffic (HTTP over TCP) from a Web server, with mean inter-page request time being 2 s 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. The Web traffic source was not turned on for about the first 150 video frames, again simply to better judge the effect. For this typical web traffic source, Fig. 13 reports the impact upon video quality. The pattern of PSNR results broadly follow those for CBR cross traffic. Table III, summarizes the results, from which it is apparent that less loss occurs due to buffer overflow at the 2 Mbps when Web traffic is present. Again, to illustrate the generality of the scheme, Table III also includes illustrative figures for two more complex video sequences than ‘Newsclip’. In Fig. 14 (a) buffer fullness is reported for the UP scheme 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 Fig. 14 (b) it is clear that transmitting exclusively at 2.0 Mbps exposes the video packets to an increased risk of being dropped from the transmit buffer. At the higher transmission rate all packet loss is due to the impact of the AWGN channel, as the buffer is under utilized.


Comparing Tables II and III, it is apparent that the ranking of the fixed modulation schemes change in ranking in respect to delivered video quality. As cross traffic characteristics are not generally known in advance this further disadvantages the fixed schemes without UP.

TABLE II.
MEAN VIDEO QUALITY WITH CBR CROSS TRAFFIC


Video sequence	Transmission scheme	PSNR (dB)	Packet Loss (%)
Newsclip	Adaptive modulation	38.06	5.08
	2 Mbps	33.15	12.10
	3 Mbps	34.05	9.53
Football	Adaptive modulation	37.46	6.31
	2 Mbps	32.24	14.67
	3 Mbps	33.98	10.94
Friends	Adaptive modulation	38.30	4.57
	2 Mbps	33.19	11.66
	3 Mbps	35.09	7.07



(a)



(b)



(c)

Fig. 11 Video quality with CBR cross traffic (a) with the full UP scheme (b) without UP at 2.0 Mbps (c) without UP at 3.0 Mbps

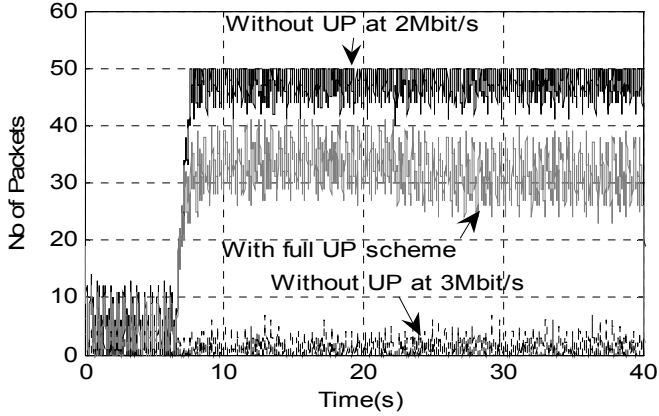


Fig. 12. Buffer fullness with CBR cross traffic.

TABLE III.
MEAN VIDEO QUALITY WITH WEB CROSS TRAFFIC

Video sequence	Transmission scheme	PSNR (dB)	Packet Loss (%)
Newsclip	Adaptive modulation	39.11	2.19
	2 Mbps	37.61	6.42
	3 Mbps	33.98	11.13
Football	Adaptive modulation	38.87	3.18
	2 Mbps	37.21	8.59
	3 Mbps	33.09	14.27
Friends	Adaptive modulation	38.89	2.65
	2 Mbps	37.66	6.11
	3 Mbps	34.08	9.88

VII. CONCLUSION

Due to the fragility of encoded video data, it is often necessary to protect the most important content types. Unequal protection through adaptive modulation in Bluetooth transmission has been shown by us to achieve a significant improvement in delivered video quality over fixed modulation schemes, depending on cross traffic conditions. The paper shows that an unequal protection scheme ought to be dynamic, as the content type importance characteristics change within a video sequence. The scheme ought to account for a varying ratio of picture type data sizes and of intra-coded macroblocks within P-pictures, owing to the occurrence of scene changes, rapid motion, and camera pans and zooms. Varying the modulation type has an important effect on buffer fullness, with the risk of packet loss through buffer overflow. To counter this it seemed logical to impose a zoned system of buffer occupation, which, for video transmission, reflects the picture and macroblock types. High-quality video, at around 40 dB for a TV series clip of CIF pixel size at 25 fps, encoded at an average rate of 1.2 Mbps, is delivered through a combination of adaptive modulation and zoned buffer management. The default Bluetooth error protection scheme at the higher data rates, infinite ARQ without Forward Error Correction, is inadequate for this purpose despite the widespread deployment of Bluetooth transceivers in mobile phones and other mobile devices.

This paper has aimed at a system which is relatively simple to implement, capable of a real-time response. Given the low energy and computational power of Bluetooth wireless devices this may be preferable. The usual approach in video coding that is familiar from rate-distortion analysis would be to treat the adaptive modulation/tri-partite buffer system as a constrained optimization problem, for which the method of Lagrangian multipliers is a topical approach. This type of solution has the advantage of greater generality. For example, it would aid adaptation of the scheme to situations where there were more than three priority levels (given by the frame types). Again, in this paper, the practical perspective is that most content exists in single-layer form but this observation does not preclude future development of the scheme within a more formal structure.

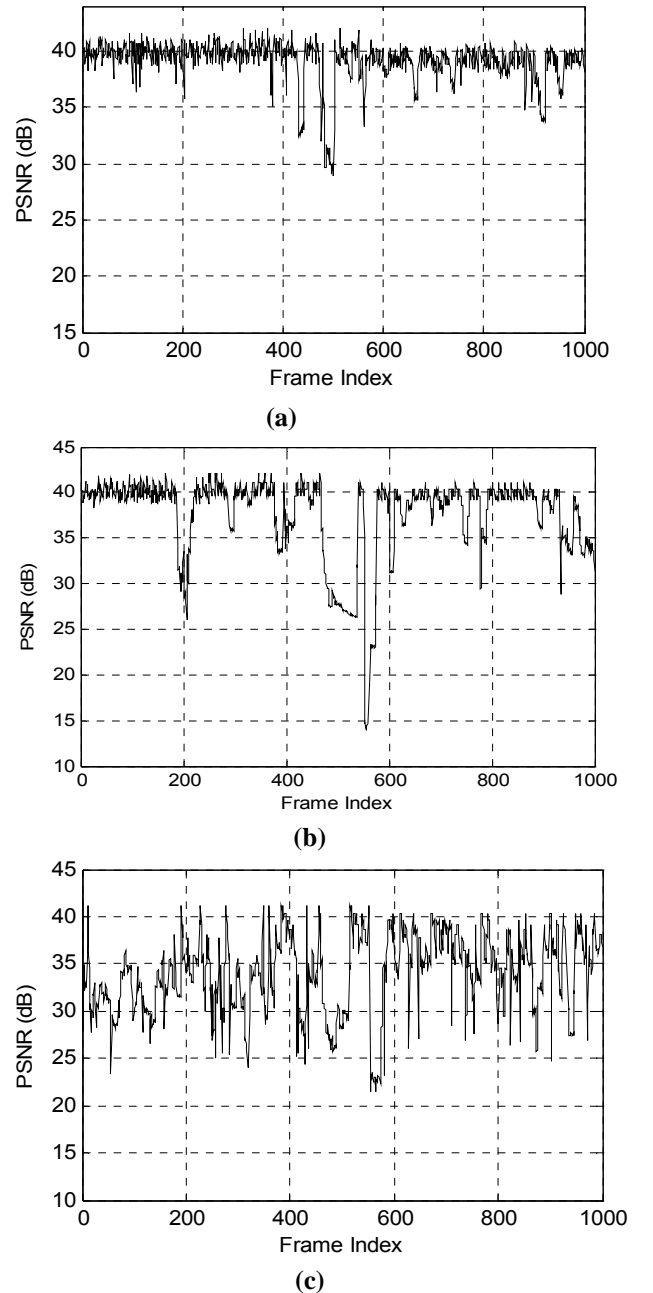


Fig. 13. Video quality with Web cross traffic (a) with the full UP scheme (b) without UP at 2.0 Mbps (c) without UP at 3.0 Mbps

REFERENCES

- [1] J. Haartsen, "The Bluetooth Radio System", *IEEE Personal Comms.*, Vol. 7, No. 1, pp. 28-26, 2000.
- [2] R. Razavi, M. Fleury, and M. Ghanbari, "Low-Delay Video Control in a Personal Area Network for Augmented Reality", *IET Visual Information Engineering*, July 2007.
- [3] Q. Qu, Y. Pei, J. W. Modestino, X. Tian, "Source-Adaptation-Based Wireless Video Transport: A Cross-Layer Approach", *EURASIP J. on Applied Signal Processing*, no. 5, pp. 1-14, 2006.
- [4] M. Saladih, F. Kschischung, and A. Leon-Garcia, "Modulation-Assisted Unequal Error Protection over the Fading Channel", *IEEE Trans. on Vehicular Tech.*, vol. 47, no. 3, pp. 900-908, Aug. 1998.
- [5] C. Ru, L. Yin, J. Lu, and W. Chen, "A New UEP Scheme Based on Adaptive Modulation for Robust Video Transmission in MIMO System", *China Communications*, pp. 102-108, Oct. 2006.
- [6] Y. Pei and J. W. Modestino, "Multi-layered Video Transmission over Wireless Channels using an Adaptive Modulation and Coding Scheme", *IEEE Int. Conf. on Image Processing*, pp. 1009-1012, 2001.
- [7] 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, Aug. 2005.
- [8] Y. Shan, "Cross Layer Techniques for Adaptive Video Streaming over Wireless Networks", *EURASIP J. on Applied Signal Processing*, vol. 2, pp. 220-228, 2005.
- [9] P. Batra and S.-F. Chang, S.-F. Chang, "Effective Algorithms for Video Transmission over Wireless Channels", *Signal Processing: Image Communication*, vol. 12, no. 2, pp. 147-166, Apr. 1998.
- [10] R. Kapoor, M. Kazantzidis, M. Gerla, and P. Johansson, "Multimedia support over Bluetooth piconets", *1st Workshop on Wireless Mobile Internet*, pp. 50-55, 2001.
- [11] A. Iyer and U.B. Desai, "A Comparative Study of Video Transfer over Bluetooth and 802.11 wireless MAC", *IEEE Wireless Comms. and Networking Conf.*, pp. 2053-2057, Mar., 2003.
- [12] C. H. Chia and M. S. Beg, "Realizing MPEG-4 Video Transmission over Wireless Bluetooth link via HCI", *IEEE Trans. on Consumer Elect.*, vol. 49, no. 4, pp. 1028-1034, 2003.
- [13] M.F. Tariq, P. Czerepinski, A. Nix, D. Bull, and N. Canagarajah, "Robust and Scalable Matching Pursuits Video Transmission using the Bluetooth Air Interface Standard", *IEEE Int. Conf. on Consumer Electronics*, pp.673-681, 2000.
- [14] Specification of the Bluetooth system --- 2.0 + EDR, Available online at <http://www.bluetooth.com>, 2004.
- [15] 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.
- [16] M. C. Valentii and M. Robert, "Improving the QoS of Bluetooth through Turbo Coding", *IEEE Military Conf.*, pp. 1057-1061, 2002.
- [17] M. Ghanbari, *Standard Codecs: Image Compression to Advanced Video Coding*, IEE Press, Stevenage, UK, 2003.
- [18] Y.-Z. Lee, R. Kapoor and M. Gerla "An efficient and fair polling scheme for Bluetooth", *MILCOM*, vol. 2, pp. 1062-1068, Oct. 2002.
- [19] R. Razavi, M. Fleury, M., and M. Ghanbari, "Detecting Congestion within a Bluetooth Piconet: Video Streaming Response", *London Communications Symposium*, pp. 181-184, Sept. 2006.




Fig. 14. Buffer fullness with Web cross traffic (a) with the full UP scheme (b) without UP at 2.0 Mbps.