Experiences with multimedia streaming over 2.5G and 3G Networks

Julian Chesterfield, Rajiv Chakravorty, Jon Crowcroft
Cambridge University Computer Laboratory
Cambridge, UK
CB3 0FD
Email: {firstname.lastname}@cl.cam.ac.uk

Pablo Rodriguez
Mircosoft Research,
Cambridge, UK
CB3 0FD
Email: pablo@microsoft.com

Suman Banerjee
University of Wisconsin,
Madison, WI
53706, USA
Email: suman@cs.umd.edu

Abstract

As third generation technologies become more widely deployed, mobile data users increasingly experience ubiquitous global network access across a variety of heterogeneous 2.5 and 3G technologies. Such provision of higher bandwidth in mobile environments promises greater access to rich multimedia content delivered to handheld or mobile devices such as PDAs, laptops and vehicles. In this paper we provide a thorough evaluation of the performance of multimedia streaming across heterogeneous Wireless Wide Area environments through measurements taken from real networks (GSM, GPRS and UMTS) using our own unequal error protection architecture tool - vorbistreamer. *Our architecture is capable of aggregating lower capacity wide area data channels together in order to create a single, high bandwidth broadband multimedia channel. Furthermore we demonstrate that the key constraint for any streaming application is the level of interactivity required by the user, and we present performance benchmarks for different technologies under various levels of interactivity.*

1 Introduction

Wireless wide-area cellular networks such as GPRS and UMTS 3G networks in Europe and CDMA 2000 networks in America and Asia are being widely deployed, promising ubiquitous access to IP-based multimedia applications for end-users. The characteristics of the cellular wireless medium, however pose several new challenges to audio and video multimedia streaming applications. In this paper we consider the performance of such interactive multimedia streaming over cellular environments.

Furthermore, with the increase in potential data capacity, and the premium rate services for which users are prepared to pay, broadband audio/video streaming is becoming a reality. From 2.5 generation technologies such as GPRS through to the current third generation services such as UMTS and CDMA 2000, the mobile environment is becoming increasingly heterogeneous. UDP based multimedia streams, however are more restrictive than the traditional TCP type traffic such as web and email, and such applications are far more sensitive to underlying network performance variation.

This paper presents results from detailed studies of Wide Area Wireless Network (WWAN) streaming across various technologies, and describes experiences with the design and implementation of a WWAN optimised streaming application named *vorbistreamer*. The key contributions are:

- Real network measurement results collected from 2G, 2.5G and 3G data networks
- Actual modifications to the Vorbis codec to generate an Unequal Error Protection (UEP) architecture for streaming in cellular environments
- Real application measurements of performance in WWANs

Vorbistreamer is a streaming application that is capable of operating in any IP environment, is robust to link or stream failure and is optimised for efficient WWAN utilisation. It encompasses the following features:

- seamless roaming across cellular links, supporting both link layer and IP layer hand-off.
- Interactivity constraint gaurantees from fully interactive communication through to one-way video streaming
- Application level fine-grained error detection and recovery techniques
- WWAN channel aggregation to generate broadband streaming capability over the wide area
- Improved performance based on statistical channel diversity measurements

Our paper is structured into three sections. Section 2 presents a detailed evaluation of the link layer characteristics of WWAN links that impact multimedia streaming. We consider the benefits of link aggregation, focusing specifically on diversity measurements in the WWAN environment between parallel GPRS interfaces, and demonstrate the benefits of aggregating channel bandwidth. Furthermore we identify the streaming limitations of each technology, focusing specifically on one-way latency, jitter and consequently application buffering performance. Section 3 presents the design of *vorbistreamer*, and discusses the various optimisation techniques that have been implemented and tested over the WWAN environment. We present measurement results for each of the tests, and identify the relative benefits and disadvantages of each scheme. Finally in section 4 we present conclusions from the work and discuss future directions in our research regarding multimedia streaming.

2 Link and Media Characterisation and Measurement

As WWAN links become more widely available, users are growing increasingly dependent on the higher bandwidth they provide for multimedia communication such as videoconferencing, Voice over IP (VoIP) services and audio/video broadcasting. Hence it is important to design both adaptive applications capable of adjusting to their environment as well as networks capable of supporting real-time traffic. Unlike non-realtime activities such as web and email browsing over TCP, multimedia streaming, especially interactive communication, presents strict performance requirements from the end-to-end path. Such requirements are significantly more challenging to achieve in WWAN environments due to the vagaries of the cellular link caused by phenomena such as signal interference, distortion and fading.

Detailed evaluations of non-realtime transport protocol and application performance over GPRS networks have been conducted. Studies on TCP under-performance [6, 2] have demonstrated that link under-utilisation does occur due to protocol design assumptions, however in general, application level issues are the primary cause of link under-utilisation [7]. Multimedia applications are typically unconstrained by such transport specific algorithmic decisions, however in this paper we demonstrate that significant application under-performance can also occur for multimedia streams through the lack of link/application cooperation.

2.1 Multimedia Traffic characterisation

Since certain frequency components in a signal have higher importance than others in terms of human perceptibility, a common approach towards digital encoding is to order data by importance such that easier application processing, whether to provide variable bitrate streaming or progressive decoding of a signal, is facilitated. Such ordering can be used to great effect in providing Unequal Error Protection (UEP) [15, 12, 13] for streams. The data produced by an encoder is ordered in terms of importance such that decoding can operate with only a minimal amount of information, but the residual quality is greatly enhanced by decoding all the data.

Current streaming media encoding formats typically produce a fairly constant bit rate over time utilising variable size samples. The reason for this is that signal analysis can evaluate the importance of a sample, such that certain samples benefit more than others from greater data representation, and a more perceptually pleasing result can be produced by maintaining flexibility with respect to the exact rate generated by the encoder. In all cases, however we find that the bandwidth produced by the encoder is not unrestricted, and will always range between a minimum and a maximum target. Any channel that cannot support the minimum data rate for constant periods up to the bounds of the playout buffer delay will cause decoder failure, and hence it is essential to maintain as predictable a channel performance as possible.

For our experiments we developed an UEP audio streaming tool, *vorbistreamer* utilising the Vorbis [14] codec. We engineered the codec to produce UEP data, providing information to the application regarding the content of the media to enable intelligent packetisation and streaming.

2.2 WWAN Network Performance

Unlike ‘staple’ Internet applications (e.g. web-browsing, e-mail etc.), where user tolerance is generally quite flexible, for multimedia applications, a user can not tolerate *any* flexibility beyond certain interactivity constraints. In video conferencing applications, for example, the delay constraints are tightly coupled with the human perception of interactive communication, whilst for one-way broadcast-style applications the delay tolerance is typically based on the willing-

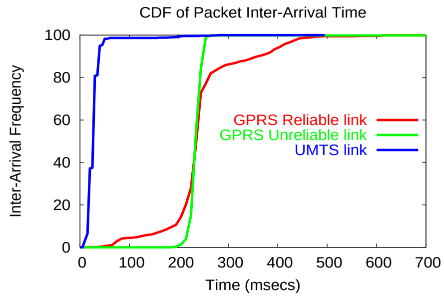

Figure 1. The inter-arrival time CDF of packets (1000 packet samples, 1371 bytes size) with ARQ enabled and ARQ disabled. The time values indicate one-way packet propagation latency, RTTs are significantly higher due to the lower bandwidth uplink.

ness of the listener to trade-off initial start-up delay versus long-term performance improvement through smart buffering and data retransmission. Additionally, there is an inherent interactivity constraint in any user-controllable media stream, e.g. using RTSP to provide pausing, rewinding, fast forwarding and segment skipping type functionality. Although we are unaware of any strict guidelines in this area, it is generally accepted that delays should be kept to a minimum, typically less than 500 milliseconds in order to facilitate user interaction.

In this section we evaluate the suitability of certain WWAN technologies for multimedia streaming under various interactivity constraints, and demonstrate that greater application intelligence can significantly improve the end-user experience of streaming in cellular environments. We measured three types of common network technologies for our tests; GSM data links, GPRS networks and UMTS, third generation networks. The characteristics of each are similar, demonstrating that although bandwidth capacity increases with the newer generation technologies, the characteristics of the wireless environment can cause a significant variation in performance.

2.2.1 Latency

The packet propagation latency exhibited across cellular links tends to be much greater than that experienced in the wired environment. This is typically due to the hardware delays imposed by the encoding and interleaving processes, the aggressive ARQ retransmission mechanism provided at the link layer as well as the channel bandwidth limitations which restrict the modulation speed of the data across the air interface. Figure 1 demonstrates the distribution of transmission speeds of data (1400 byte packets) across the cellular link, demonstrating that latency and consequently packet interarrival jitter is reduced as link speeds increase. For multimedia streaming over GPRS, however the large propagation time presents a significant challenge since many interactive applications require minimal round trip time propagation delays in order to maintain perceptually acceptable communication (the ITU recommendation for voice communication over telecommunication links [1] is around 400 milliseconds, RTSP functionality also limits RTTs to around half a second).

For GSM data, we have measured significantly longer packet propagation times, proportional to the modulation rate of the link. The average bitrate in our tests was around 8kbit/s (the total GSM voice bandwidth channel provides a maximum of 9.6 kbit/s including link level headers and RLP traffic), hence 1400 byte packets take an average of 1460 milliseconds to propagate.

UMTS links provide a higher throughput, and consequently lower propagation delays, making interactive communication easier. In these tests we measured full link utilisation at an average rate of 386.3 kbit/s (the official capacity of UMTS is 384kbit/s, we consider the minor additional data rate measured in our experiments to be due to smoothing effects caused by hardware buffering), indicating that there was no bandwidth contention with other users in the cell. Typical inter-arrival times for 1400 byte packets as measured in our tests are around 30 milliseconds but range as high as 300 milliseconds, illustrating that strict link layer reliability is always preferable.

2.2.2 Jitter

The packet propagation delay for end-to-end communication is further aggravated by the variation in delays, or the inter-arrival jitter. The CDF graph in figure 1 further demonstrates the variation in jitter experienced across WWAN cellular links. In the case of GPRS, once again we find that the inter-arrival variation is more pronounced than the higher speed UMTS links, ranging anywhere from 80 milliseconds to 500 milliseconds. By disabling reliability, we notice that the effects of the ARQ retransmission delays are removed, causing a smooth and predictable jitter bound on all traffic, in exchange for a higher level of packet loss (in these tests, the reliable mode produced no packet loss, the unreliable mode produced 3%). UMTS jitter on average is much lower due to the shorter propagation time of the link, and consequently faster retransmission periods, however, as indicated previously, our tests did uncover some extreme inter-arrival gaps ranging as high as 300 milliseconds.

2.2.3 Capacity Variation

We also consider the effect of the packet propagation latency and jitter variation on the instantaneous throughput

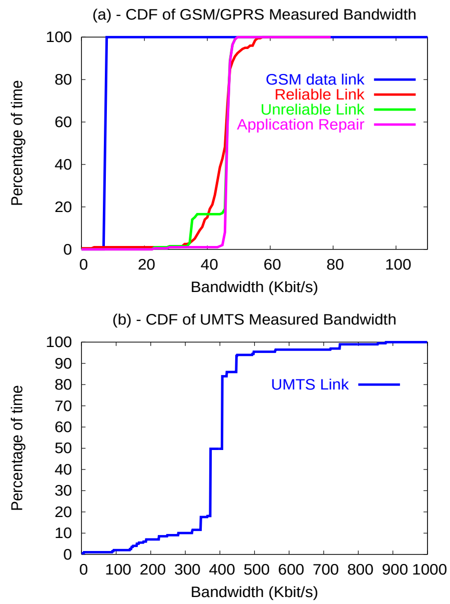

Figure 2. The CDF of instantaneous measured link capacity showing a) GSM and GPRS links and b) UMTS

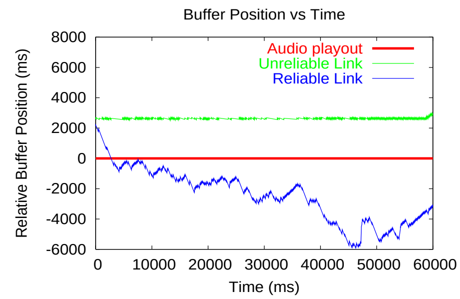

Figure 3. The challenge of trying to maintain a sufficient buffer in advance of the media playout over a reliable GPRS link compared to transmitting data over an unreliable link.

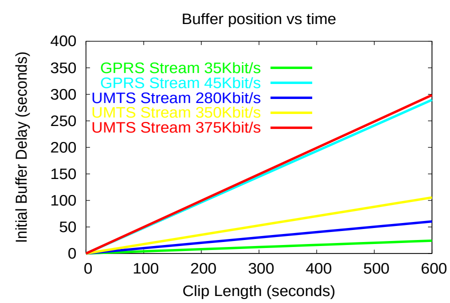

Figure 4. The initial buffer delay that must be inserted prior to media playout in order to support various streaming rates.

capacity of the channel. Figures 2 (a) and (b) demonstrate the fluctuation in capacity over time and as a cumulative distribution. Results from GSM data, GPRS operating in reliable mode, GPRS in unreliable mode and GPRS sub-packet repair are shown. The fourth result is particularly significant, since this highlights the benefits of sub-packet error detection as first presented in [4]. In earlier work it was demonstrated that sub-packet error detection is an efficient means of utilising useful portions of packets experiencing errors, when full reliability at the link level cannot be used due to delay and variation constraints. This work built upon the concept of UDPLite [5, 9, 3, 11] partial packet protection, optimising the technique through *bucket alignment* for GPRS links, enabling errors to be detected and isolated to their exact length through an application-level inference approach. Errors occur in the GPRS environment in multiples of the RLC Coding Scheme (CS) size (typically either 20 or 30 bytes for CS-1 or CS-2).

2.2.4 Buffer Performance

The impact of variability on a streaming media application depends upon the constraints of the application itself. It is clear that a complete and reliable stream is always preferable for a receiver if there is sufficient flexibility to support the variation in inter-arrival times and fluctuating channel capacity. Hence we evaluate the threshold at which reliability operating at the link layer can be considered unsuitable to the application requirements by modelling the relationship between the length of the media that is to be transmitted and played, the target bitrate of the application and the tolerance of the user to an initial startup delay. We do not consider the impact of RTSP functionality here, however it is worth noting that any non-cached data being requested, e.g. due to a rewind or fast forward operation, would require the application to re-buffer the stream every time. Under such circumstances, however it is likely that the data source is a non-realtime generated stream such that any excess throughput can also be utilised to replenish buffers while streaming.

For these data results it was assumed that in addition to streaming and playing in real-time at the receiver, the original media source was also being generated in real-time, e.g. a television or radio station feed, and consequently the target data rate is both a minimum and a maximum streaming rate, since no advance buffering can be achieved.

Figure 3 demonstrates the challenges associated with buffer management for real-time media streaming over WWAN links. This particular test demonstrates the effect of link fading due to poor signal quality triggering multiple retransmissions and consequently causing the application buffer to underrun. A 2 second initial playout buffer was generated in these tests for a 60 second music clip, transmitting data at an average rate of 35 - 40kbit/s. This was targetted as a feasible data rate for the average capacity of the link, and highlights the sensitivity of an application to fluctuations in capacity.

Based on our earlier results presented in figure 2 we evaluated the percentage of capacity available over time and consequently the ability of the link, under these circumstances, to support a minimum target streaming rate. Figure 4 indicates the theoretical buffer requirements of our application to target various streaming rates over different periods of time based on the media clip length in a GPRS environment. We see that in order to support a higher streaming rate closer to the link bandwidth limit, our theoretical buffer requirement grows exponentially, requiring significantly more relative buffer delay as the bitrate increases.

Figure 4 also provides measurement results for the UMTS environment. We see similar effects but on a larger scale for the high end of the UMTS capacity. As the target rate gets close to the capacity of the link (here 375 kbit/s out of a possible 384 kbit/s) the initial buffer startup delay also reaches 50% of the stream length (the equivalent measurement in GPRS is 46 kbit/s out of a possible 53.6 kbit/s).

2.2.5 Channel Aggregation and Diversity

WWAN links, particularly between different providers, exhibit strongly uncorrelated behaviour [10]. In the design and evaluation of *vorbistreamer* we have leveraged the benefits of link diversity to support higher streaming rates through channel aggregation. In this manner, wide area broadband multimedia streaming channels can be created. We have focused specifically on the ability to emulate UMTS type streaming rates across multiple GPRS links to create a truly heterogeneous broadband streaming environment. For our tests we assumed a target data rate of 160kbit/s, emulating a high quality surround sound audio system experiment supporting 5 channels at a rate of 32kbit/s per channel (e.g. MPEG-2 audio). We considered an environment in which a mobile user roams between UMTS coverage, e.g. in an urban area, and GPRS only coverage, and consequently measured the ability of a GPRS channel aggregation system with 4 GPRS interfaces and 1 UMTS link to support such a streaming rate. The results have indicated that when operating in reliable mode, the GPRS aggregated channel requires a significant initial startup buffer compared to UMTS (e.g. for a 6 minute clip

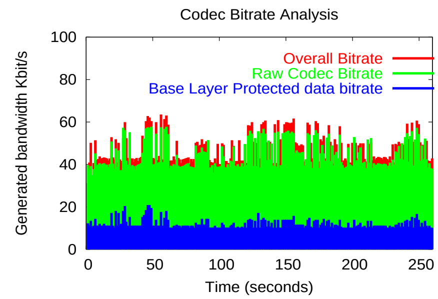

Figure 5. The division of protected to unprotected data as generated by *Vorb istreamer*. Figure 5 also indicates the amount of data overhead generated by packet and payload headers.

UMTS would buffer for only 5% of the clip length versus 20% for the GPRS case). When operating in unreliable mode, the streaming rate is much more stable, and in fact the aggregated GPRS channel outperforms UMTS operating in reliable mode by a factor of 5, requiring only 1% of the clip length to be buffered.

3 Vorbistreamer and UEP encoding

We designed an application, *vorbistreamer*, that is specifically optimised for multimedia streaming across broadband cellular links. *Vorbistreamer* is a bitrate adaptive application that optimises the encoding of data based on information inferred about the underlying link. It is capable of operating seamlessly across any IP based environment, the Forward Error Correction scheme applied to the data stream being dynamically adjusted based upon *end-to-end* error detection and measurement techniques. It supports mobility both at the link level (e.g. intra-provider hand-off between UMTS and GPRS networks) and at the IP level (e.g. inter-provider hand-off between UMTS and WLAN), utilising RTP for the data transport protocol, and implementing IP-based data striping in order to aggregate channel bandwidth and create WWAN broadband links.

3.1 Application Design

The Vorbis codec [14] provides high quality encoding of raw audio data. Upon initialisation, various parameters can be specified such as Average, Constant or Variable rate coding with an additional quality or bitrate settting. The encoder initially generates a base layer ‘spectral floor’ which defines the basic ‘shape’ of the signal, followed by residual higher frequency components. The codec has the useful property that any combination of residual layers can be decoded to improve the quality of the signal. The encoded data therefore lends itself well to UEP formatting with a single base layer data stream, and any combination of residual layers received over a 'best effort' channel. For the purposes of these tests, a hybrid version of the codec was produced which allows data to be directly categorised as it comes out of the encoder, in essence producing a UEP-ready version of the data. This allowed intelligent packetisation of the data prior to streaming. The codec also benefits from the optimised design of pre-streaming codebooks which must be transmitted reliably to the decoder, allowing a higher degree of signal compression for all components of the bitstream.

We selected the vorbis encoding format as a suitable high quality audio codec, and instrumented the encoder to our requirements in order to generate UEP content. Figure 5 illustrates the typical ratio of essential/protected data to higher level residue unprotected data. The ratio of data rates varies depending upon the maximum target bitrate of the encoder, with the rates closer to equal as the overall bitrate drops to its lowest threshold. The lowest overall bitrate achieved from the codec in separate tests has been around 30 - 35kbit/s including packetisation overhead, at which point the residue data rate and the base layer data rate are approximately equal. The codec ratio identified in figure 5, is representative of vorbis streaming rates in general as the overall data rate increases above approximately 40Kbit/s. The maximum data rate generated by the encoder is around 500kbit/s.

Each vorbis sample, as generated by the encoder, includes a header which specifies, among other things, the length of the sample and the position of the sample in the continuous audio stream. This allows any sample to be decoded independently of another, without being impacted by loss of earlier or subsequent data. This is an important property for our streaming environment since we anticipate that samples will be lost on occasions. Decoded vorbis sample sizes are always a factor of 2, and range between 64 to 8192 bytes. Typically, however for lower bitrate streaming, as required over GPRS, the sample sizes tend to be equivalent to 32 millisecond segments of audio, ensuring that any single gap occuring in the decoder stream as a result of loss is not too significant.

3.2 UEP encoding Techniques

Here we identify some different techniques that have been used in applying redundancy to the data in order to generate receiver-based repair of the stream. All these techniques are optimised for WWAN links that operate in transparent mode, causing errors to occur within the bitstream in multiples of the RLC coding size. The motivation for disabling reliability is to maintain tight jitter bounds for interactive communication.

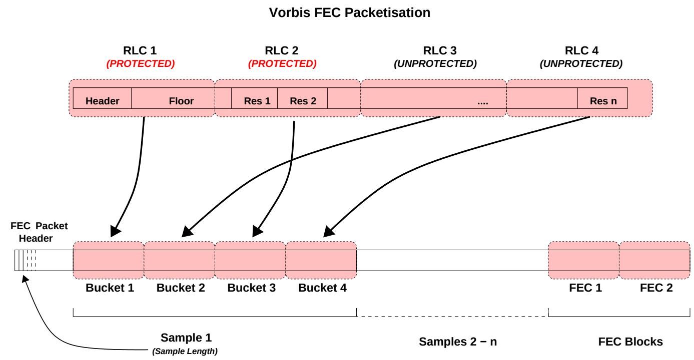

Figure 6. The packetisation approach for utlising intra-packet error checking, and generating Unequal Error Protection for media such as Ogg Vorbis encoding.

3.2.1 Intra-packet redundancy

Intra-packet redundancy is a technique that leverages the nature of error bursts across the wireless interface, utilising fine-grained error detection to isolate errors down to the exact length of the burst and adding small amounts of additional redundancy to every packet. Since error events occur randomly but individual errors within each event occur in bursts, we packetise the data within a frame to distribute the impact of an error burst over multiple data segments. By interleaving important with unimportant data, the effect of the error is spread across the sample stream, thereby decreasing the perceptual impact of any single burst. To further improve the goodput of the link, however we generate an intra-packet UEP architecture to bias the probability of arrival of important data segments over unimportant data segments.

The output from the vorbis encoder is in UEP layers of base and residual data to which we assign different levels of priority. We divide our base layer into multiples of the RLC coding size minus the CRC hash, padding where necessary with residual data, and then apply a reed-solomon FEC code [8] to generate parity data which is included in the packet. For further information on the RLC coding optimisation strategy see [4]. Reed Solomon(RS) codes have the property that for k source data symbols, an (n, k) code can be constructed where any k symbols out of the n will reconstruct the original data. Figure 6 illustrates the UEP scheme adopted by *vorbistreamer* depicting a sample with 4 RLC bucket lengths worth of data. The first 2 blocks are protected by the RS codes, and the second 2 blocks are un-

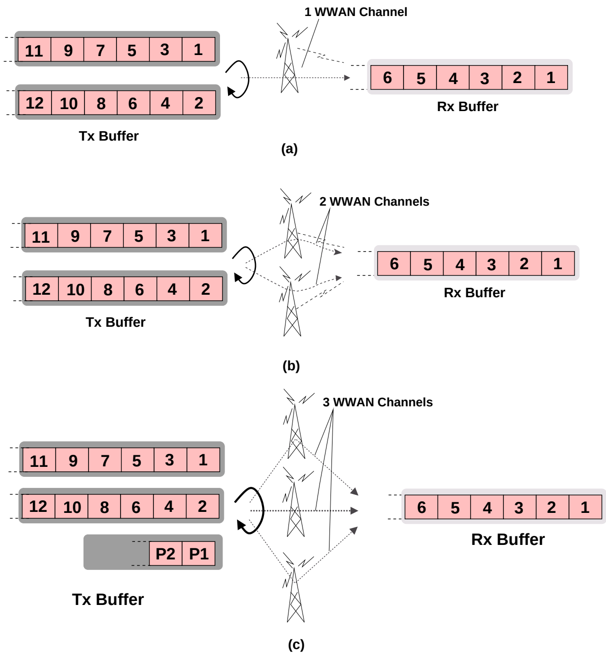

Figure 7. The effects of interleaving over (a) 1 channel, (b) 2 channels and (c) 2 channels + 1 parity channel

protected. The figure also illustrates that multiple samples are stacked into one packet in order to minimise the amount of header overhead, the RS encoding process is then applied over the k protected data segments throughout the packet, generating $(n - k)$ parity symbols of data appended to the packet. Normally in wireless environments, it is beneficial to maintain a lower packet length in order to keep retransmission times, and consequently network jitter low, however where retransmission and error checking is disabled this has no impact, and therefore we minimise the header and FEC overhead by generating larger packets.

3.2.2 Inter-packet redundancy

Building upon our earlier observation that channel diversity can be very beneficial in improving the statistical performance of the WWAN link [10], we also leverage the benefits of diversity for multimedia streams. By applying FEC codes across WWAN channels, rather than within a single packet of data, we increase the statistical recovery performance of the channel. For environments which do not have channel aggregation flexibility, we also consider the statistical improvement of applying cross-packet redundancy where separate packets carry redundant repair data. Our tests have demonstrated that there is a higher statistical probability of consecutive packets within a channel receiving errors compared to parallel packets transmitted across 2 separate channels.

- *Cross-channel coding* Figure 7(c) illustrates the architectural approach for applying cross-channel redundancy. For an aggregated link of n channels, every $n - 1$ packets sent across separate channels are also accompanied by a parity packet on the nth channel. Each parity code is the length of an RLC block, and is applied across parallel protected data segments. In this manner, UEP coding is achieved, where the level of redundancy is a factor of the number of channels. In assigning data packets to channels, the physical channel selection is not important, the significant aspect is that the parity data is transmitted over a separate WWAN path.

- *Cross-packet coding* Where multi-channel coding is not available, e.g. due to lack of channel resources or hardware, Cross-packet coding is an effective approach to counter the effects of burst errors on the channel. By sending parity data in separate packets from the original data the statistical probability of samples being unrecoverable is increased. The greater the distance between original data and the parity data the more effective the recovery performance becomes, but at the expense of increased recovery delay for errored packets.

3.2.3 Sample Interleaving

For protected samples that are unrecoverable through the application level RS decoder, a gap in the audio stream is created since the sample must be discarded. In order to mask the effects of the loss, receiver based repair techniques can be applied. The easiest, and most effective methods are packet repetition (the previous sample is repeated), packet interpolation (a sample is derived through inference based on the two samples on either side) or noise insertion. The effectiveness of any repair technique is significantly increased by minimising the length of the gap created as a result of data loss. Therefore, we also measured the effects of packet level interleaving. Since error events do not statistically occur close together, even though bursts of errors occur within an event, we find that distributing the sample stream in time makes an important difference to the audio quality perception. At the expense of increased buffer delay to recover from sample re-ordering, the proximity of packet losses is a parameter that is minimised as much as possible within the buffer delay bounds.

We considered three techniques for interleaving, similar to the UEP approach; intra-packet sample interleaving where samples within a packet are re-arranged to ensure they do not occur consecutively, inter-packet sample interleaving across packets but within the same channel, and inter-packet sample interleaving across channels. Figure 7 demonstrates the interpacket interleaving approach over a single channel (a) and across multiple channels (b). The results from the experiments are presented in the following section.

3.3 Comparison of UEP Techniques

FEC Scheme	Recovery (%)	Overhead (%)	Goodput (%)
1 Channel FEC 0	0	0	96.2
1 Channel FEC 1	39.5	2.84	97.7
1 Channel FEC 3	92.17	6.59	99.7
1 Channel FEC 5	99.7	12.93	99.9
cross-packet FEC	99.7	12.39	99.9
cross-channel FEC	99.9	12.39	99.9

Table 1. Redundancy Scheme Performance

Interleaving Scheme	Average Interleave Gap
Single Channel, 1 buffer	1.2
Single Channel, 2 buffers	1.08
Round-robin Scheduler	1.04

Table 2. Interleaving Scheme Performance

The FEC schemes perform relatively similarly to each other, as indicated in the results presented in table 1. The best recovery performance has been measured using cross-channel redundancy due to the statistical diversity properties between channels as outlined previously [10]. The overhead generated by each scheme is low, the highest value measured being 12% of the original data, however these values are variable based on the recovery percentage required, and the amount of protected data that is generated. As the raw codec output increases, the ratio of protected to unprotected data widens in favour of the protected stream, generating a lower redundancy amount relative to the overall stream.

The interleaving scheme results outlined in table 2 also demonstrate similar properties, with the cross-channel diversity producing the best results, cross-packet interleaving with a single channel being marginally more correlated (i.e. generating a greater average sample gap) and intra-packet interleaving generating the largest sample gap.

4 Conclusions and Future Work

In this paper we presented detailed analysis of the environmental impact of Wide Area Wireless networks on multimedia streaming. We identified the constraints imposed on the network by different interactivity levels as required by applications. We further measured the ability of WWAN links with different reliability parameters to support various targeted streaming rates.

In order to generate Wide Area Wireless broadband streaming capability we designed and built an application named *vorbistreamer* that is capable of operating seamlessly across heterogeneous IP environments. We designed a system to aggregate 2.5G WWAN channels and produce high throughput data links suitable for transmitting broadband multimedia content.

For highly interactive applications that require the lowest possible jitter and propagation delay, we also considered the benefits of utilising sub-packet error detection in order to extract useful data that would otherwise be discarded. *Vorbistreamer* has been optimised for the wide area wireless environment, utilising a variety of application-based error recovery techniques in order to minimise the impact of errors on the data stream. We presented some analysis of the benefits of each recovery technique, demonstrating that there is a great deal to gain in leveraging the performance diversity across channels to improve the goodput of the link.

Future work will entail further investigation of cross-channel coding schemes and fully interactive videoconferencing communication using adaptive UEP codecs. Furthermore, we intend to conduct detailed analysis of WWAN and WLAN mobility and handover performance, evaluating the benefits of various optimisation proposals to create a true seamless mobile multimedia environment across heterogeneous networks.

References

[1] ITU-T recommendation g.114, (05/00): One-way transmission time. Technical report, ITU-T, May 2000.

[2] R. Chakravorty and I. Pratt. Performance issues with general packet radio service. In *JCN, 2002*, 2002.

[3] F. Hammer, P. Reichl, T. Nordstrom, and G. Kubin. Corrupted speech data considered useful. In *Proceedings of the First ISCA Tutorial and Research Workshop on Auditory Quality of Systems, Mont Cenis, Germany*, Apr. 2003.

[4] J.Chesterfield, R.Chakravorty, S. Banerjee, P.Rodriguez and I.Pratt. Transport Level Optimisations for Interactive Media Streaming Over Wide-Area Wireless Networks. In *WiOpt'04: Modelling and Optimization in Mobile, Ad Hoc and Wireless Networks*, 2004.

[5] L. Larzon. The udp-lite protocol. Internet Draft, Internet Engineering Task Force, Dec. 2002. Work in progress.

[6] R. Ludwig, A. Konrad, and A. Joseph. Optimizing the end-to-end performance of reliable flows over wireless links. In *MobiCOM*, 1999.

[7] R.Chakravorty, J.Chesterfield, P.Rodriguez and S. Banerjee. Measurement Approaches to Evaluate Performance Optimizations for Wide-Area Wireless Networks. In *Passive and Active Network Measurement (PAM 2004) 5th International Workshop*, 2004.

[8] I. Reed and G. Solomon. Polynomial codes over certain finite fields. In *J. Soc. Industrial Appl. Math, volume 8*, pages 300–304, 1960.

[9] R. Reine and G. Fairhurst. Mpeg-4 and udp-lite for multimedia transmission. In *PGNet 2003, Liverpool, UK*, June 2003.

[10] P. Rodriguez, R.Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee. Mar: A commuter router infrastructure for the mobile internet. In *MOBISYS - 2nd International Conference on Mobile Systems, Applications and Services on Network Computing and Applications*, June 2004.

[11] A. Singh, A. Konrad, and A. Joseph. Performance evaluation of udp lite for cellular video. In *Proceedings of NOSSDAV*, June 2001.

[12] T. Stockhammer. Uep - useful or useless error protection?

[13] T. Stockhammer and C. Buchner. Progressive texture video streaming for lossy packet networks. In *Packet Video Workshop, Korea*, page 57, 2001.

[14] The Ogg Vorbis multimedia coding format, forming part of the Xiph.org's Ogg Multmedia Project, http://www.xiph.org/ogg/vorbis.

[15] Y. Wang, A. Ahmaniemi, D. Isherwood, and W. Huang. Content-based uep: A new scheme for packet loss recovery in music. In *To Appear in Proceedings of ACM Multimedia 2003*, 2003.