Research White Paper

WHP 150

July 2007

Video contribution over wired and wireless IP network – Challenges and Solutions

Y. Zheng, A. Lipscombe and C.Chambers

BRITISH BROADCASTING CORPORATION

Video contribution over wired and wireless IP network – Challenges and Solutions

Yuan-Xing Zheng, Andrew Lipscombe, Chris Chambers

BBC Research, Kingswood Warren, Tadworth, Surrey, KT20 7NP, U.K.

Keywords: Video Contribution over IP, WiMax, Pro-MPEG FEC

Abstract

Over the past few years, video transmission over IP network has become increasingly popular. However, the majority of research work in the area has concentrated on distribution, i.e. how to transmit TV signal to home users via internet. Video contribution over IP (VCIP), for use by broadcasters and program makers, is still a relatively new field and so far little research been carried out in this area.

This paper will introduce the VCIP work that has been done by BBC Research over the past few years. We will talk about the testing of various VCIP coders/decoders – problems unique to VCIP will be highlighted. We will also look at the video quality improvement that Pro-MPEG Forward Error Correction (FEC) has given; further equipment improvement is suggested.

We will then talk about the experience gained from our WiMax field trial, and also use it as a case study to further illustrate the importance of using FEC in real-time video transmission across lossy, high-jitter IP links.

1 Introduction

With the wide acceptance and fast deployment of IPTV, broadcasters have realised that video over IP is not only a great new tool for distribution, but potentially can also save a great deal of money on contribution links. From low bit-rate News-gathering applications to moving uncompressed HD video across studios, VCIP can be equally useful.

Although contribution and distribution shares many technical similarities, there are still a few big differences.

Firstly, for distribution, Internet Service Providers (ISPs) would like to squeeze in as many channels as possible for a fixed bandwidth on their network, therefore most of the video streams will operate with the emphasis on reducing the bit rate to the lowest possible. In comparison, for contribution, it is desirable to operate the streams at the highest bit rate available to enable the delivery of the highest quality pictures. Reducing the amount of compression used for video contributions also has the major benefit of enabling the use of greater compression for distribution while minimising the likely concatenation effects and so maintains the video quality.

Secondly, although low latency is important in both cases, for contribution, it is essential to minimise delay in order to allow live, two-way interviews with individuals who may not be familiar with working with long delays. For example, in a News-gathering situation, a two-way link for an on-air interview is the norm today - in this case it is accepted for most situations (with the exception of satellite links) that the link delay has to be less than 200ms to make an interview workable.

BBC Research began work on video and audio contribution over IP in 2001. Since then, a multicast-enabled network test bed has been built, with full Quality of Service (QoS) functionality. In this network, the core is a pair of 10GBit/s switches, with streaming servers and clients deliberately set in different domains to study various streaming related routing protocols. Recently we also expanded this network with a WiMax/WiFi combo edge network to explore the problems involved with contribution over a link with long delay and lower bandwidth characteristics. A simplified diagram for this test network is shown in Figure 1.

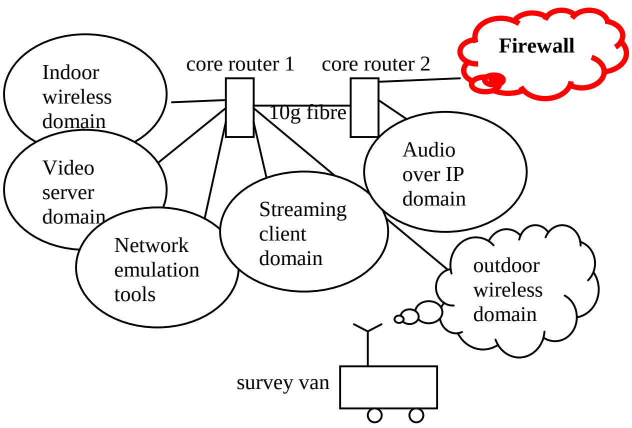

Figure 1. BBC Research video / audio over IP test network bed

With the help of this test network, a lot of work has been carried out to study multicast protocols such as PIM and IGMP v3, the impact of QoS on an end-to-end delivery chain and the capability of wireless networks to cope with real time applications.

2 VCIP encoder/decoder tests

One of the first problems we encountered for VCIP work is that there are not many real-time encoders available that are designed for outputting high data rate with high picture quality – most products are designed to produce the best picture at the lowest rate. Over the last twelve months, we tested five encoder products targeted at IP applications on the market today. In these tests, we used the free software encoder included within the open source media player VLC [1] as a comparison – any hardware coder should be significantly better than VLC before we think it is acceptable.

Encoding schemes tested included MPEG-2, MPEG-4 Simple profile, H.264 and some proprietary codecs.

Although we can't list manufacturer names in this paper, the most common problems we have found from tests are:

- Products are unable to encode with a data rate higher than 4Mbps
- Motion quality is generally poor
- Most of the hardware takes very long time to boot up, and is not broadcast resilient
- No product's delay was found to be lower than 200ms for H.264 images
- Not all of them supported multicast and none of them supported ProMPEG FEC
- Some of the end-to-end solutions are proprietary.

These are disappointing results, particularly as we still can't find a suitable standalone high data-rate real-time encoder. This makes it difficult to use an IP-based solution as a main contribution feed even though an IP pipe could be 100 times cheaper than the same bandwidth via a satellite link.

It should be pointed out that using IP as an output format will introduce extra overhead above the datarate of streams comprising simply of MPEG transport stream packets. This is because, for most situations, the maximum IP frame size is 1500 bytes, therefore one IP frame can take up to 7 transport stream packets without fragmentation. The overhead IP level introduced in is list in Table 1 (assuming a transport stream packet size of 204 bytes).

TS packets per IP frame	IP header overhead
1	26%
2	13%
3	8.7%
4	6.5%
5	5.2%
6	4.3%
7	3.7%

Table 1. Overhead vs. TS packets per IP frame

In summary, comparing the performance of the coders under test with VLC as a software encoder, most of them are better in one way or another, although the improvement is marginal. Considering the hardware encoders under test cost in the region of three to five thousand pounds per unit and that VLC is free software, the results are fairly disappointing. It is worth pointing out that the encoders under test were early versions and we would expect to see improvements in resilience and suitability for broadcast use as they mature. The market trend is increasingly for encoders to have IP (as well as ASI) outputs as standard, thereby increasing the choice of equipment for this specific application.

3 ProMPEG FEC

Due to the nature of IP network traffic, bursty errors are common and to be expected on any "real world" IP connection. To cope that, manufacturers have introduced FEC protection as recommended by Pro-MPEG Code of Practice 3 [2]. To test this feature, we purchased two network adaptation products that have this incorporated. Although they do not have encoding function, they take ASI signal inputs and re-packetise the transport streams into IP packets (up to 7 TS packets per UDP packet). One of the products can also split the IP stream across up to 3 independent routes and reassemble them at the receiving end into one ASI output stream. Both products provide FEC protection.

By using one or two independently transmitted FEC streams (IP port numbers for the FEC are N+2 and N+4 where N is the payload port number, usually 5004 for RTP traffic), the system can be more robust in dealing with errors at the price of increasing overhead and delay. The lower number port carries the "column" FEC data stream and the higher number port carries the "row" FEC data stream. Obviously these streams contribute a further overhead to the total stream bandwidth, but the intention is to send these 2 data streams at a different Quality of Service setting to the payload, so that a form of graceful degradation is possible when the bandwidth is restricted, with the payload having a greater priority.

Our tests show that the FEC function certainly improves received video quality, especially when the link error is high, such as in a wireless environment. However we were disappointed to find out that, in practice, some manufacturers only implement the one-dimensional (column) error correction, preventing the full range of error protection from being implemented.

Another disappointment is that although the equipment allows us to adjust the FEC protection ratio, the current generation of receivers give no information on the link statistics, therefore it is pure guess work to decide how much FEC is required for a specific link. If too little FEC is used, the protection is insufficient; if too much, it wastes precious bandwidth. We would therefore like to see some adaptive-feedback mechanism implemented in future equipment, or at the very least, a clear and easily interpreted display of link statistics.

We were pleased to see that both systems tested adhered to the code of practice in as far as interoperability is concerned. That is to use the FEC streams if available but to carry on with best effort if neither FEC stream is delivered to the receiver.

We would like to see the Pro MPEG FEC offered as an option on the IP outputs of more video encoders to enable a choice of one-box solutions for outside broadcast, news and sport applications. We would also like to see it implemented in software solutions such as VLC (both for transmission and reception).

4 Video contribution over WiMax network – A Case Study

Compared with WiFi, WiMax can provide links over longer distances and higher, more managed bandwidth; it has therefore increasingly attracted the attention of broadcasters as well as Internet service providers. There are however significant challenges to its use in a broadcast environment and, as ever, the marketing hype surrounding the headline throughput rates creates an unrealistic expectation for instant solutions from early-to-market equipment.

BBC Research first started tests with WiMax in 2005, using the first generation of WiMax equipment operating at the 'licence free' 5.8GHz band. The maximum range we have achieved with this set of equipment is approximately 20 miles with very clear line-of-sight (Epsom Downs to BBC Television Centre). We have also discovered that for the 5.8GHz frequency band, trees have a very strong negative impact. For example, one small tree with juicy leaves directly in front of a subscriber station introduced a 20dB SNR loss, compared with an unobstructed subscriber station at the same distance from the base station.

However due to the lack of QoS support, we found the first generation WiMax equipment was not suitable for any real time streaming applications. So we had to wait until the approval of the IEEE 802.16e standard, which targets mobile applications and therefore includes QoS protection.

In 2007, one year after the release of IEEE 802.16e standard, we purchased a set of point-to-multipoint WiMax equipment for a field trial. It consists of one base station (BS) and four subscriber stations (SS); again, all of them working in the 5.8GHz 'license free' band.

The test network is set up centred on the BBC Research antenna range with the base station in the middle of the field and four subscriber units dotted around the site on various buildings, as shown in Figure 2.

Figure 2. BBC Research WiMax network

In our arrangement, one subscriber station takes video feeds from the existing on-site coding and multiplexing test-lab, where any required bandwidth or format of video sources can be provided (effectively acting as a broadcast centre). The other subscriber stations act either as network throughput test nodes or as video decoder points, enabling us to better understand the end-to-end network performance and subjective video quality over a WiMax link. This can be done as either a single hop or double hop as appropriate.

Our initial experiences are that, firstly, due to the higher jitter nature of the wireless link, even when the link is protected by its own high QoS setting, there is little chance of getting a clean error-free picture. Therefore ProMPEG FEC generated at source is a must-have for this type of wireless transmission. However, we have to bear in mind that WiMax is already a relatively long delay link, so any extra delay added by the reception and decoding of the FEC streams is not good news for live two-way interviews.

Secondly, at the 5.8GHz band, the maximum available channel size for a TDD WiMax base station is 10 MHz, therefore in reality the maximum data rate for all subscriber stations associated with one base station is about 20Mbps, shared between the uplink (SS to BS) and downlink (BS to SS). Although a WiMax system allows us to adaptively allocate channel bandwidth between the up and downlink, a total of 20Mbps capability still means it is only suitable for medium to low video quality transmission. This bandwidth, however, is still very useable for news-gathering and other outside broadcast links.

At the end of May 2007, OFCOM (the UK government regulator governing, amongst other things, frequency spectrum issues) made a change to the maximum permitted output power for Wimax use in the 5.8Ghz band, from 2 watts to 4 Watts EIRP across all geographic areas [3]. This will be very useful for increasing range, throughput and suitability for more challenging outdoor RF environments, but because of the licence free nature of the band, the potential for interference increases as well.

5 Conclusions

Compared with IPTV, video contribution over IP is still a relatively new field. Although it shares many similarities with video distribution over IP, the need for higher bandwidth and very low delay/jitter still imposes an important challenge to many broadcasters, especially with growing interest in the use of wireless systems such as WiMax.

References

[1] www.videolan.org
[2] http://www.pro-mpeg.org/publications/pdf/Vid-on-IP-CoP3-r2.pdf
[3] http://www.ofcom.org.uk/consult/condocs/powerlimits/power/statement.pdf