Taming the Cloud

— the use of IP networks in professional broadcast contribution and production applications

Rhys Lewis
Chief Enterprise Architect, BBC Future Media & Technology and
Chair of the EBU Network Technology Management Committee

In the past, professional broadcast networks that support media contribution, production and distribution have demonstrated a deterministic behaviour and have delivered a very high level of performance (availability, reliability etc.). IP networks[1] by contrast are inherently statistical in nature, usually adopting a best-endeavours approach to performance. The two types of network are chalk and cheese. Yet the economy, scale and flexibility of an 'IP Cloud' are extremely attractive to today's broadcasters, and the deployment of IP networks supporting professional broadcast operations is already well advanced.

This article draws on some of the presentations made at the recent EBU Network Media Exchange 2009 seminar[2] to illustrate some of the successes, highlight some of the problems and identify the areas where IP networks still need to 'do better'.

Introduction

The EBU Networks Technology Management Committee (NMC) has organized an annual seminar dealing with key topics in the network domain for almost ten years now. Over time the focus of interest has shown a marked change, with much greater interest being shown of late in the use of IP networks for the transport of media (high-quality video and sound) to support professional broadcast applications. More traditional solutions, such as those based on the SDH[3] or ATM[4] standards, have provided broadcasters with networks which provide:

1. Note that the term 'IP network' is used in this article as a generic term to describe networks based on the carriage of data packets through a 'cloud' of interconnections using the Internet Protocol (IP) and additional protocols standardized by the Internet Engineering Task Force (IETF).
2. "Network Media Exchange 2009", held in Geneva in June and jointly organised by the EBU (http://www.ebu.ch) and the VSF (http://www.videoservicesforum.org). A full report of the seminar is available on the EBU's website: http://www.ebu.ch/en/hr_training/training/technology/pastactivities/2009reports/networks_seminar09.php. In addition, EBU Members can access all the presentations by following the appropriate links.
3. Synchronous Digital Hierarchy – a common protocol for optical networks operated by telecom companies. For more information see http://en.wikipedia.org/wiki/Synchronous_Digital_Hierarchy.
4. Asynchronous Transfer Mode – a common protocol in the networks operated by telecom companies. For more information see http://en.wikipedia.org/wiki/Asynchronous_Transfer_Mode.

- Consistent, defined end-to-end performance, featuring:
  - time delays which are low and fixed;
  - guaranteed bandwidths;
  - very low signal impairment;
  - high availability – usually 99.999% of the time or better.
- Effective monitoring methods and tools to help identify and repair problems;
- Good management and control over network performance and configuration;
- Secure transport of data;
- Terminal equipment which is, to a large extent, interoperable;
- An easy means to connect one source to multiple destinations (multicast).

An IP network by contrast, at least judged by most people's experience of the best known IP network – the public Internet – struggles to achieve the above (see the text box "IP networks in brief…") and is generally characterized by:

- Performance which is very variable over time, having:
  - time delays which can be long and inconsistent;
  - no guaranteed bandwidth;
  - unreliable signal transport which can include duplication, loss or re-ordering of data;

IP networks in brief…

In a traditional or switched network, a direct connection with a given bandwidth is established between the source and one or more destinations. A full bandwidth path exists between the destination and source for the entire duration of the connection. There is no competition between the different users of the network for any proportion of the bandwidth allocated to a particular connection.

Switched networks employ a technique called admission control. When the connection is initially requested, bandwidth is allocated for the full duration of the connection. If the bandwidth is unavailable, the request is refused. The network operator needs to make sure that sufficient bandwidth is available to avoid too many refusals.

By contrast, the simplest IP networks are connectionless networks. There is no fixed connection between different places on the network – rather the network uses a series of routers to carry packets of data between source and destination. The packets are of variable length and each carries a destination address. The network 'cloud' is made up of many connections and nodes with a packet router situated at each node. The first router in the network accepts packets from the source and forwards them on to another router in the network. It makes a decision about which router this should be by inspecting the destination address, from its knowledge of the configuration of the network, and on how busy a particular connection may be (the first-choice path may be fully occupied in which case the router can choose to send the packet to an alternative device). Since a router typically accepts packets from many different sources, packets are queued, using data buffers, until the router has received the full packet and successfully found another router to which the packet may be sent. No packet is refused entry to the network, but there is no guarantee that the packet will traverse the network!

From the above description it should be clear that both the route followed and the transit time of each packet through a simple IP network is a statistical process and that all users/packets compete for the bandwidth available in different parts of the network. Packets usually, though not always, manage to find a route through the network cloud. But there are likely to be significant variations in packet delay and packet ordering (a late packet may take a longer route through the network and arrive after one sent earlier). In many cases, routers are designed to deliberately 'drop' or delete packets which arrive if their queuing buffers become too full, requiring the destination to signal back to the source for the missing packet to be resent if complete data reception is required.

A more detailed tutorial on the mechanics of IP was given by Andy Leigh of the BBC at a session during the 2005 networks seminar, see

http://www.ebu.ch/CMSimages/en/NMC2005report_FINAL_tcm6-40551.pdf

- circuit availability which can be very variable over time;
- multicast transmissions which rely on immature protocols and which can be difficult to set up and manage.

- Immature monitoring methods and tools for assessing real-time signals;
- No central management and little ability to control network configuration;
- Difficult to control who can access the data being transported;
- A level of complexity which makes it difficult to establish interoperability across equipment from different manufacturers, both within the network and for terminal equipment (e.g. video and audio codecs) connected to the network.

On the face of it, this leads to a conclusion that IP networks should be avoided at all costs for professional applications. But this would be to overlook the inherent benefits of IP networks which include their ubiquitous nature, flexibility to carry multiple signal types, the ability to readily extend a broadcaster's internal 'office' networks and systems out 'into the field' and the availability of relatively low-cost networks and terminal equipment. Much of the current work on IP networks looks to take advantage of these benefits whilst finding ways to deal with the performance issues listed above.

This article draws on some of the material presented at the 2009 NMC seminar to illustrate some of the successes and challenges associated with the deployment of IP networks to support professional broadcast contribution and production applications.

Let's start with a couple of the success stories…

A false prediction

Andy Rayner of BT Design kicked off the seminar by reminding the audience that, at IBC in 1998, he had predicted that *"IP [would] never be suitable for use in professional broadcast networks"*. Fittingly, his presentation then went on to describe a worldwide IP network built by BT to support professional broadcast applications. The irony was not lost on Andy – or the audience!

But what's changed in the 10 years since Andy made his original prediction? How has the inherent unpredictability of IP networks been tamed by BT?

In his presentation 'Building an MPLS network to support broadcasters', Andy described the key elements of BT's approach to delivering the stringent requirements of professional broadcasters. These include:

- The use of protocols which 'tie down' many of the uncertainties inherent in the network – in effect establishing a connection in a connectionless network. Examples of such protocols include Multiprotocol Label Switching (MPLS) to fix routes through the 'cloud'; RSVP-TE (a resource reservation protocol) to reserve bandwidth, and the use of Point-to-Multipoint (P2MP) signalling to provide reliable multicast capability.
- The use of these protocols is supplemented by a central network control and scheduling system, implemented by the network operator, which can flexibly configure network routes and paths. Network resources can be booked by clients through a web portal from days in advance down to five minutes or less before the planned use.
- Service monitoring which measures service quality from input to output. The monitoring system interfaces with the scheduling system to switch between alternate routes through the network (e.g. a main and reserve path) when necessary. The monitoring and path switching is designed to provide largely glitch-free automatic switching to a reserve circuit in case of failure of the main route.
- Working closely with the router and end-equipment manufacturers to ensure that buffering and throughput delays are as deterministic and as low as they can practically be made. By careful design, the throughput of routers has been reduced to around 60µs, and the end-to-end jitter (a

measure of the delay variation through the network) has been kept below 1ms. The end result is that the total end-to-end delay for video transmission is set largely by the choice of video codecs rather than network delays.

Fig. 1 illustrates the extent of BT's Media and Broadcast IP network. The network is fully operational and proving its effectiveness in examples such as the transport of UK Digital Terrestrial Television feeds from broadcasters to transmitters and supporting studio-to-studio and outside broadcast contribution feeds for commercial and public-service broadcasters in the UK and elsewhere.

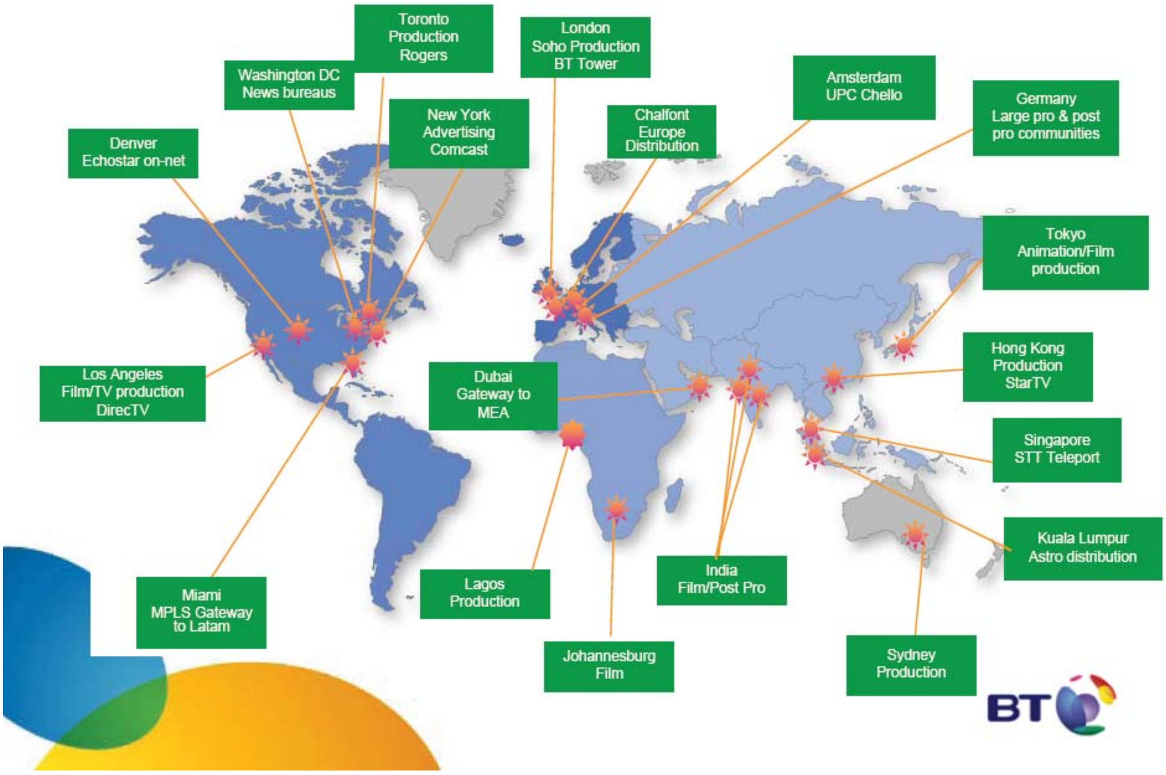

Figure 1
BT's Worldwide IP network
(Credit: Andy Rayner, BT)

HD over IP

In a paper entitled "Moving towards HD on an IP network", Anders Nyberg of SVT in Sweden described their use of an already well-established IP network. The original network had been built in 1999 as an affordable solution to carry locally-recorded news items from various locations around Sweden back to Stockholm. The application was for time-shifted playout in a 24-hour news channel. It was based on a 34 Mbit/s star topology using point-to-point leased lines from the central point to each regional office.

The network was expanded in 2003 when live broadcasts were also transported over the network from tape-less 'light newsrooms' in each region. This phase of work replaced the earlier videotape-based operation with a file-based workflow and real-time contribution circuits.

The latest phase of work moved SVT to a fully network-based operation and included building an off-premises backup and recovery facility, the ability to play out from a number of sites as well as making SVT's entire legacy video-tape archive available via a browsable, on-line content library. With over 180,000 hours of material it is the largest digital library in Sweden.

The network now has more than 5000km of fibre, providing 2.5 Gbit/s of network capacity in the core, and carries multiple video feeds using MPEG-2 over IP contribution circuits running at between 9 and 24 Mbit/s. Towards the end of 2006, the network was upgraded to carry HD signals using JPEG-2000 codecs, preferred for their short coding delay, at 130 Mbit/s between a number of fixed sites as well as well as from two transportable HD/SD 'flight-packs'. The flight-packs allow full network connectivity to be established between the OB site and the main production centre including file-transfer of media, a return video path, 4-wire communications and office traffic.

What had started life as a news-only network has now grown to support HD-production of Sports, Drama, Documentary and Factual material and has turned SVT into an almost entirely IT-based organization.

So much for the successes, now let's take a look at one of the challenges…

What's happened to my files?

"Where are my files" by Luc Andries (who has the enviable job title of a 'Top Expert' with VRT in Belgium) proved to be one of the most well-received papers of the seminar. Luc described the adverse impact that the statistical nature of an IP connection can have on network-based production workflows. His recurring theme was that the design of an IP-based system was more akin to Einstein's quantum physics than the classical Newtonian thinking which could be applied to switched networks!

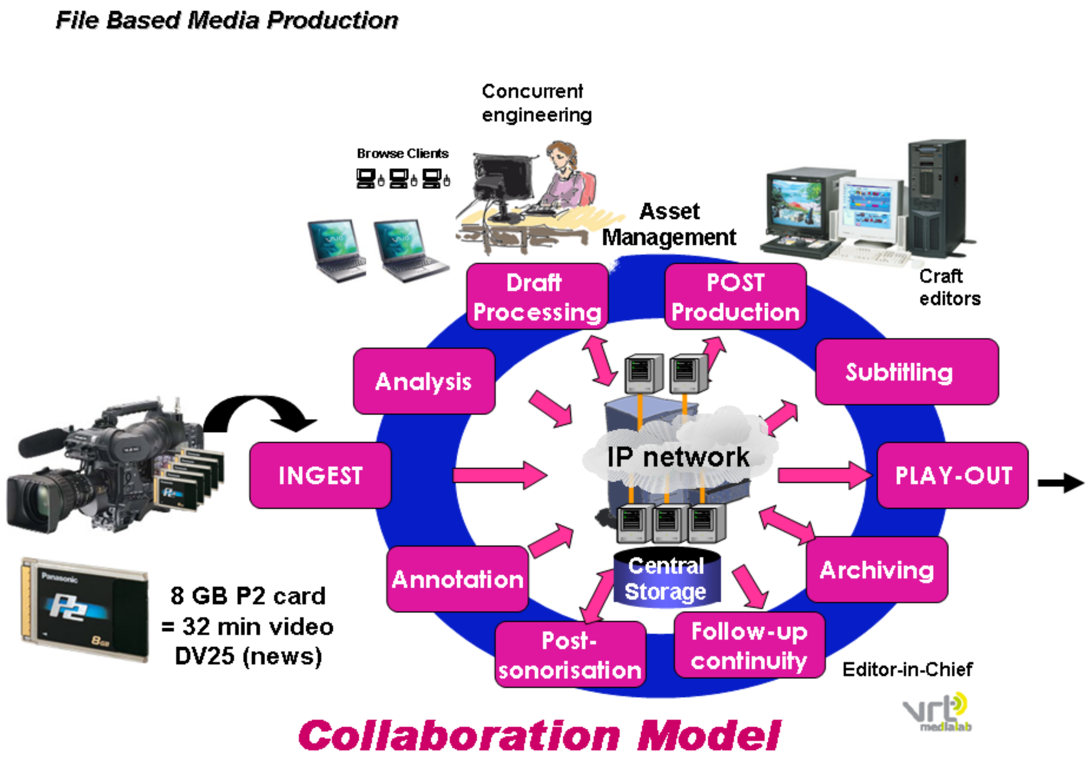

Figure 2
Illustration of the collaborative approach enabled by networked production
(based on a slide presented by Luc Andries and used with his permission)

VRT has installed a large file-based news production system – the Digital Media Factory (DMF)[5] – at its headquarters in Brussels. The previous and largely sequential tape-based workflows have been replaced by a collaborative tape-less approach as illustrated in *Fig. 2*. This is a collaborative approach with archival and work-in-progress file storage accessed from multiple process areas.

5. http://www.ibc-show.com/IBC2007/Candice/IBC2007/PAPERS/1205%20J%20Hoffman/1205%20J%20Hoffman.pdf

However, the user's experience of the collaborative process has been less than ideal with a significant proportion of the file transfers either being significantly delayed or not happening at all. This ultimately led to a very public failure of a popular VRT News website – which is when the Top Expert was called in to investigate!

After his investigation, Luc concluded that there were two fundamental causes of the problems, both related to the statistical nature of packets in the IP-network 'cloud':

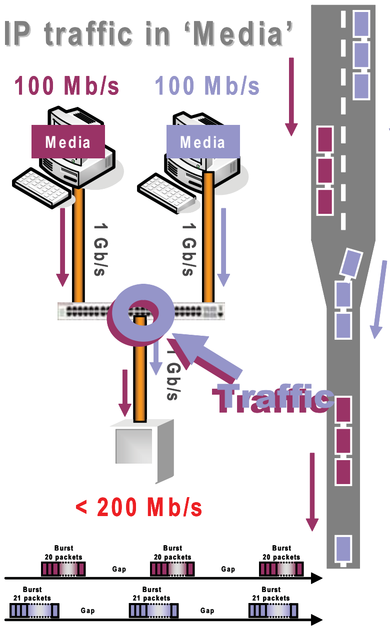

1. Two media sources generate bursts of packets at an average bitrate of 100 Mbit/s.

2. When the packet bursts arrive at a router at the same time, the buffers in the router start to fill and packets are queued or may even be rejected.

3. The result is that the average bitrate leaving the router may be substantially less than 200 Mbit/s and some packets are 'lost'.

Figure 3
The statistical nature of the media traffic causes congestion and packet-loss
(extract from a slide presented by Luc Andries and used with his permission)

- Competition between data packets for network bandwidth rather than shared use. Just because the network has a capacity of 200 Mbit/s there is no guarantee that it can reliably carry two 100 Mbit/s packet streams. This problem is exacerbated by the common practice of employing large packet sizes for video signals which makes sharing more difficult *(see Fig. 3)*.
- Complex data-flows between different elements in the collaborative network resulting in network congestion. Luc nicely illustrated this with an example of how a simple ingest of a short video file from a memory card resulted in 36 file transfers across the DMF network. Some additional simulation work had shown that at busy news periods up to 500 simultaneous large file transfers could be taking place on the network.

The combination of these two factors led to a high proportion of packet losses (slowing down the transfer of files) and in extreme cases led to the file transfer being aborted or even to an application or server failing completely and having to be re-started – to the considerable frustration of the users.

Luc's solution was two-fold:

- The design of the network needs to be fine-tuned to avoid the packet loss caused by the contention for resources such as bandwidth of buffer capacity. Because this contention takes place at an 'atomic' timescale (of the order of the duration of the individual packets), existing network-analysis tools were not of much use – the traditional concepts of average bandwidth and available capacity are insufficient to guarantee the service quality required.
- Systems must be architected so that the workflows, and their associated data flows are carefully managed to avoid placing excessive demands on the underlying network. Techniques such as traffic management operating at the network level, or reducing the reliance on long packets for video, are also helpful in avoiding excessive short-term demands being placed on the routers in the network.

Interestingly, there is a common lesson emerging from Luc's presentation and the success stories presented earlier; treating IP networks as commodity technology that can be rolled out in professional applications is a recipe for failure – rather, the design of the system must be carefully thought out (architected in other words) in order to ensure that problems are avoided. Design rather than hope should always be the guiding principle.

As well as demonstrating the successes and challenges associated with the use of IP networks, the conference also looked at the issue of interoperability across terminal equipment from different manufacturers…

An interoperability challenge…

In introducing the session on Interoperability, Chris Chambers of BBC Research laid down a challenge…"I believe we must have product interoperability for ALL key elements of our systems to reduce risk as well as deliver future flexibility and choice to our businesses!". During the session, two of the papers illustrated the work that the EBU has undertaken to try and ensure that broadcasters and manufacturers can deliver to Chris' aspiration.

An interoperability success story

A project team of the EBU, N/ACIP – Audio Contribution over IP, has been in place for 3 years. Led by Lars Jonsson of Swedish Radio, the group has worked closely with a range of manufacturers and has established, as well as demonstrated, that a high degree on interoperability can be achieved between equipment from different manufacturers. In his presentation "Hints on the Selection of Network Parameters for Audio over IP", Lars summarized the progress that had been made and went on to give a few tips, based on his experience, on how to get the best from audio contributions made over IP networks.

Figure 4
Audio codec testing in Munich
(Credit: Lars Jonsson, Swedish Radio)

An interoperability specification has been published[6]. The document defines a set of parameters including audio-coding algorithms, signalling methods and IETF protocols which manufacturers must implement in order to establish interoperability between equipments. These requirements had been developed collaboratively by EBU Members and manufacturers. Earlier in 2009, 11 of those manufacturers and a number of EBU representatives had attended a 'plugfest' in Munich *(see Fig. 4)*. This was the second such event and the benefits of the collaboration were very visible. Whilst the first event had highlighted many interoperability issues, the second event established that most of those issues had now been tackled and resolved.

But Lars also reminded the audience that the solutions were not yet fully mature. Issues which remained included:

- The robustness and reliability of the equipment and network connections were not yet good enough – a live stream of a prestige concert to many millions of listeners was not the same thing as a single telephone call.
- Monitoring and control software for IP networks were not yet able to meet the needs of broadcasters.
- Whilst acceptable Quality of Service was achievable on well-specified corporate or private networks, getting a good enough service on public networks remained problematic.

6. EBU Tech Doc 3326-2007 "Audio Contribution over IP: Requirements for Interoperability". Available at: **http://tech.ebu.ch/docs/tech/tech3326.pdf**

Finally, Lars gave his tips for successful contributions using audio over IP.

- File transfer generally gives better results than streaming, which should be reserved for direct, live transmissions.

- The use of quality of service tagging techniques, such as Diffserv[7], to give some types of services priority passage through the network could be beneficial. But there were problems with preserving the tagging through concatenated networks.

- There were some benefits in optimizing the packet size used for audio over IP. Packet lengths of greater than 250 bytes were beneficial but the packet length should not be increased too much or excessive delays would be introduced.

- Error correction should be used sparingly, to avoid excessive delays, and must also be matched to the characteristics of the network. Simple error-correction methods such as packet duplication with error concealment in the receiver give good results.

- Lars concluded with his observations on the best choice of audio-coding format. PCM or EAPT-X audio coding had given the best results for fixed network transmission in the range 3-6 Mbit/s. AAC LD or EAPT-X for radio-frequency terminals at bitrates between 384 kbit/s and 2 Mbit/s. The use of AAC HE had only been beneficial on wireless links at bitrates below 100 kbit/s.

…and now for the video

Whilst the work of the ACIP group is nearing its close, the group which is seeking to establish interoperability between video codecs is just setting out: EBU N/VCIP – Video Contribution over IP. The chair of the group, Markus Berg from IRT, along with Alexander Hetfleisch from ORF, outlined the progress being made by the group in his presentation.

In common with the experience of the ACIP group, interest from manufacturers in establishing interoperability had been low at the outset but was now growing thanks to the efforts of the group. In addition, the liaison with VSF, who have a common interest in this area, was also proving fruitful.

The plethora of different video and IP network standards is a significant difficulty in this area and, in order to reduce the number of possible permutations, the group had defined four interoperability profiles. The four are expected to meet the vast majority of broadcaster requirements. A short-form description of each is given below.

Profile 1 – Uncompressed video over IP

Designed to allow the most common SD and HD formats to be carried over IP using the IETF-defined Realtime Transport Protocol (RTP).

Profile 2 – MPEG transport stream over IP

This profile defines the minimum set of transport-stream protocol variants for which interoperability between different manufacturers must be supported. This is the most well-developed profile definition and is likely to be published shortly.

7. See http://en.wikipedia.org/wiki/Differentiated_services

Profile 3 – MXF encapsulation over IP

The MXF standard[8] is emerging as a popular means of transporting media within a production environment. The ability to carry MXF over an IP infrastructure has a number of key advantages, including reduced cost and the ability to minimize the number of video re-encoding processes which preserves picture quality.

Whilst existing standards from the SMPTE and IETF provide much of what is needed to support the carriage of MXF over IP, the VCIP group has identified two missing links which are necessary to join up the MXF and IP standards – a definition of the minimum set of MXF features required for interoperability (including options for error correction) and a definition of how the MXF data is to be encapsulated in the IP stream. Two proposals have been drafted in conjunction with the VSF and are being submitted to the relevant organizations.

It almost goes without saying that, without interoperability between different manufacturers, the flexibility offered to broadcasters by using MXF as an interchange standard will only partly be realized.

Profile 4 – Low bitrate video over IP

In many ways this is the most difficult area to standardize. There is a multiplicity of proprietary solutions used to carry video over IP in the 3 to 15 Mbit/s range, and the requirement for interoperability is not yet apparent to broadcasters. Typical applications in a broadcast environment include news contributions from journalist in the field and connections between remote studios.

Some progress has however been made in establishing a minimum set of requirements, though it is as yet unclear whether there is sufficient interest from end users in establishing interoperability in this area.

The group's emphasis at present is to improve the engagement of manufacturers and seek agreement on the various interoperability profiles. Future work would include signalling formats (to signal which parameters are being used by the source to the destination), additional profiles for contribution use and scrambling or encryption of the signal.

Never mind the quality, feel the experience

"Quality of Experience (QoE) Issues in Video over IP" proved to be one of the most thought-provoking presentations. Stefan Winkler, a Principal Technologist with Symmetricom in the USA, gave an outline of some work under way in the VSF, aimed at providing a more meaningful measure of the impact of impaired pictures on viewers. Current methods rely heavily on network Quality of Service (QoS) metrics such as delay, packet loss or jitter. The new measures seek instead to assess the quality of experience to try and provide much better correlation with the viewer's perceptions and taking account of all sources of impairment between source and viewer. To illustrate this, Stefan showed the two examples of *Fig. 5*. The network impairments give identical measurements in both cases – but the impact on the viewer is substantially different.

In order to measure QoE, an understanding of the human vision system is required and algorithms which model the behaviour need to be produced. In order to demonstrate some of the challenges of this, and how the eye and the brain can sometimes conspire to distort reality, Stefan presented a number of optical illusions. You can see the one which had most of the audience rubbing its eyes in disbelief at: **http://www.creativebits.org/amazing_optical_illusions**

A number of standards bodies, including the VSF, the ITU and some IPTV forums, are actively engaged in developing QoE measurement standards. The general approach is to define a set of

8. See **http://en.wikipedia.org/wiki/Material_Exchange_Format**

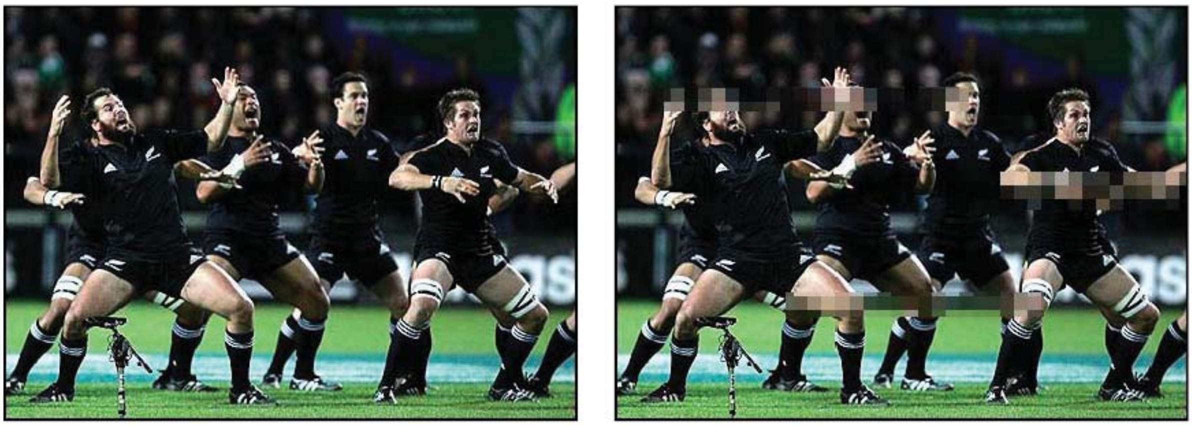

Figure 5
Same network impairment – but a very different impact on the viewer
(Credit: Stefan Winkler, Symmetricom)

video metrics which can be extracted from the video and analyzed to give a score which is hopefully a good correlation with subjective viewer ratings of the same pictures. Examples of the metrics being considered include: frame-losses, average quantizer scale values, video freeze rates and so on.

Full reference methods, where the metrics from the source and received pictures are available for analysis, give the best correlation between the measured rating and viewer judgement but are essentially only suitable for laboratory work. No-reference methods, where analysis is only carried out on the received pictures, are the most challenging to make work but also have some potentially interesting applications. For instance, payment for the network capacity could be based on the QoE provided to viewers. Some work is also going on to look at a compromise approach, the reduced-reference method, which does not rely on the measuring device having full access to source pictures. Rather the essential features of the measurement at the source, and also possibly the receiver, are extracted and transmitted to the measurement equipment through an additional low-rate data channel.

Current standardization activities are at an early stage and aimed at refining the measures of subjective quality and choosing the best metrics. The focus of the work is on providing a static, rather than a dynamic, measurement. Although there is considerable interest in using these techniques for video transmitted over IP, the tools are not yet well enough refined to deal with these applications.

Conclusions

This article was written for a general audience and has drawn on some specific examples presented at the 2009 NMC seminar to illustrate some of the successes, but also some of the challenges, of deploying IP networks to support broadcast contribution and production applications. Peter Brightwell, of BBC Research, has written a companion article, based on a presentation he made at the seminar, which provides more technical details on some of the tools and techniques available to help draw the best out of IP networks supporting file-transfer applications[9].

If your appetite has been whetted by this report then why not plan to attend the next seminar which takes place in Geneva in mid-May 2010 – or better still join one of the project teams of the Network Technology Management Committee (NMC)[10]. In the words of Brad Gilmer, Executive Director of the VSF and co-chair of the event, “we’re stuck with IP – for better, for worse, for richer, for poorer, in sickness and in health”. We might as well try to make the best of it!

9. “High performance file transfer over IP networks”, Peter Brightwell, EBU Technical Review 2009-Q4 **http://tech.ebu.ch/docs/techreview/trev_2009-Q4_IP-Networks_Brightwell.pdf**

10. More information on the activities of the NMC at: **http://www.ebu.ch/en/technical/organisation/nmc/index.php**

Rhys Lewis leads the BBC's Enterprise Architecture team – a technology governance team working to align the BBC's business and technology strategies.

He has spent most of his working life researching and developing new broadcast technologies and was closely involved with the launch of many of the BBC's digital services; from early work on the Radio Data System for FM Radio to the launch of Digital Radio (DAB) and multi-platform Digital Television services. After a spell as the BBC's Head of Technology Strategy, he became their Chief Enterprise Architect in 2005.

Mr Lewis is also closely involved with the European Broadcasting Union and currently chairs their Network Technology Management Committee, looking at the technical future of the Eurovision network and the impact of network technology on public-service broadcasting in Europe.

Acknowledgements

I would like to thanks the various presenters who supported the 'Networked Media Exchange 2009' seminar – the quality and range of presentations was exceptional this year but, inevitably, I have had to be very selective in choosing the presentations to summarize in this article. I also have to give a particular thank-you to Jean-Noel Gouyet (I have drawn heavily on his report of the event which is available on the EBU website) and to Julian Cable for suggesting a number of improvements to the original text.

This version: 17 November 2009

Published by the European Broadcasting Union, Geneva, Switzerland

ISSN: 1609-1469

Editeur Responsable: Lieven Vermaele

Editor: Mike Meyer

E-mail: tech@ebu.ch

EBU·UER

The responsibility for views expressed in this article rests solely with the author