



R&D White Paper

WHP 088

June 2004

Audience monitoring of BBC Interactive services on DTV: a preliminary study into a possible technique

R.P. Marsden

Audience monitoring of BBC Interactive services on DTV: a preliminary study into a possible technique.

R.P.Marsden

Abstract

The BBC is keen to have a method of accurately measuring how much viewers use its Interactive services on digital television. This White Paper describes some preliminary investigations into a method for digital terrestrial television, based on the insertion of a “digital index” on each interactive page, created by the MHEG rendering engine in the set top box or integrated digital television receiver. This would be detected by a purpose-designed logging unit monitoring the composite signal on the SCART connection of the digital receiver. It is concluded that the method appears technically feasible, and although potentially visible on screen in some situations, would probably be acceptable both to our audiences and editorially. Further study would be needed to establish the total cost of such a system, which would need to cover not just the elements described here, but also mechanisms to assign and insert the digital index, and to collect and analyse the results.

Additional key words: Digital Television, Interactive Service, Audience Research, Monitoring

White Papers are distributed freely on request.
Authorisation of the Chief Scientist is required for
publication.

© BBC 2004. All rights reserved. Except as provided below, no part of this document may be reproduced in any material form (including photocopying or storing it in any medium by electronic means) without the prior written permission of BBC Research & Development except in accordance with the provisions of the (UK) Copyright, Designs and Patents Act 1988.

The BBC grants permission to individuals and organisations to make copies of the entire document (including this copyright notice) for their own internal use. No copies of this document may be published, distributed or made available to third parties whether by paper, electronic or other means without the BBC's prior written permission. Where necessary, third parties should be directed to the relevant page on BBC's website at <http://www.bbc.co.uk/rd/pubs/whp> for a copy of this document.

Audience monitoring of BBC Interactive services on DTV: a preliminary study into a possible technique.

R.P.Marsden

1 Introduction

The BBC is keen to have a method of accurately measuring how much viewers use its Interactive services on digital television. The immediate requirement is for digital terrestrial television, because technical solutions already exist for both digital satellite and digital cable, but a single technique that could be used on all three platforms has obvious advantages.

IT IS EMPHASISED THAT ANY FACILITIES THAT ENABLE PERSONAL VIEWING HABITS TO BE MONITORED ARE INSTALLED ONLY IN HOUSEHOLDS THAT HAVE GIVEN THEIR INFORMED CONSENT AND ARE SUPPLIED WITH THE SPECIAL LOGGING UNIT REQUIRED TO DETECT AND RECORD VIEWER USAGE.

Monitoring the viewing of Interactive content is not a trivial task and a number of other technical solutions were considered before arriving at the technique that is the focus of this White Paper.

Simply detecting the channel to which the television receiver is tuned is not sufficient (and in the digital age where each channel delivers a multiplicity of services within a multiplex, even this is not as straightforward as it may first appear). Also, unlike viewing pages downloaded over the Internet, where only the requested pages are transmitted, the entire set of Interactive data content is transmitted in a continuous “data carousel”, with the selection of what to display made inside the viewer’s receiver. The tight connection with the linear service being watched (and listened to) is no longer maintained, and indeed, the receiver can be instructed to retune to another multiplex or even show moving images from other services as part of the Interactive experience.

Pages may be opened very briefly, especially if the viewer is experienced in navigating through the service. A system that attempts to detect commands sent from the infrared remote controller can easily get out of step with the receiver itself, with very misleading consequences, and will not detect internally triggered page changes.

In principle, monitoring could be achieved with additional software in set top boxes to record which pages had been accessed by interrogation of the service information used by the receiver to access specific services and/or service components. This solution has been proposed for digital satellite services, where there is a vertically integrated market closely controlled by the network operator Sky. However, for the horizontal open market of digital terrestrial television, there is now such a wide range of receiver equipment available, particularly since the advent of Freeview, that this would not be a practical proposition. There is no guarantee that all types of set top box or IDTV

(integrated digital television receiver) would have enough spare memory or processor power, or that they would have a modem or other output to allow the viewing data to be collected. Swapping out viewers' equipment for a specially modified set top box would not be acceptable because it would inevitably distort their individual patterns of usage. The substituted box would almost certainly differ significantly from the equipment originally selected by the user, for example in its ease of use, its speed of respond to commands or its presentation of the Interactive content.

To avoid this and other problems this White Paper proposes inserting a short data string on each Interactive page, detectable by a suitable electronic device intercepting either the connection between the set top box and the television, or the connection from an IDTV intended to link to a video recorder. Although we have concentrated on the digital terrestrial television platform, the functionality available on the digital cable and digital satellite platforms is adequate for this method to be used there too.

A very brief series of experiments has been conducted, to gauge the likely acceptability of the method, given that the digital index would be potentially visible on each Interactive page to all viewers, not just those with the monitoring facility installed. We also needed to see if the proposal might work in the wider context of the broadcast system before it is decided to put further effort into developing the concept.

We have therefore investigated how visibility might be mitigated, for example by placing the digital index near the top or bottom of the picture, or minimising its contrast. At this stage, we have written an Interactive application to create simple examples of the digital index. No equipment has been built to test detection. We believe that techniques such as those used to detect teletext could easily be adapted to do this.

2 A few words on terminology

When this idea was first proposed, it was described in terms of a "barcode" applied to each page. Most people are familiar with the small patch of black and white stripes on the label of almost everything they buy, and the ability of the barcode scanner at the checkout to identify the item down to the last detail, even from the huge range held in the typical modern hyperstore. This was a useful analogy with which to describe the idea to an audience containing non-technical people. In practice the identifying mark can be very much thinner than an ISO/IEC 16388 barcode, and is unlikely to share its encoding parameters or its method of detection. The term "digital index" will therefore be used.

3 The proposal in more detail

UK DTT set top boxes and IDTVs capable of displaying Interactive content do so by means of a "MHEG rendering engine" – effectively a software-based graphics generator already loaded in the set top box¹. The Interactive content is delivered in the multiplex as a "carousel" containing instructions and data from which the MHEG engine renders each page. Instructions in the broadcast

application determine the appearance of the display and the changes that occur as time goes by or when the viewer operates the remote control.

It is a relatively simple matter to include extra instructions and data for each page to form a “digital index” for that page. The MHEG engine can reproduce bitmap images, and also has a feature called “Dynamic Line Art” that can draw geometrical shapes. Dynamic Line Art was used to create the digital indices for these experiments.

It is important to realise that the digital index, or for that matter any other object forming part of interactive content is only ever generated or rendered in the receiver; it is only the description of the objects or content that is broadcast.

Set top boxes and IDTVs are consumer products, and must be produced at an acceptable price to the public. Their graphics capabilities are consequently restricted. This has implications for the characteristics of the digital index that can be created, either as a bitmap or with Dynamic Line Art.

The digital index must be created in the active picture area. This is because Interactive applications have access only to a 720 pixel by 576 line grid, nominally corresponding to the active picture area. (This is also true for the digital satellite and digital cable platforms). In any case, the set top box imposes its own regenerated blanking interval waveform, and will overwrite any signal put in the blanking intervals. Also placing the digital index in the active picture area removes the risk that the index data waveform will interfere with or confuse any other technical systems in the signal path, for example Teletext decoders (see Section 6).

Set top boxes typically provide only a limited palette of colours for Interactive applications to use. The MHEG palette contains 139 fixed “colours”, including 7 evenly spaced monochrome steps from black level to peak white. This restriction applies to graphics generated by Dynamic Line Art or a bitmap. This imposes a minimum amplitude on the digital index in the gamma-corrected domain of about 117mV between monochrome steps (although luminance differences less than 1 mV exist between some pairs of colours in the palette). Laying the index in as a watermark is therefore not an option.

An interactive page can be made up of graphics and images. It may be that the position chosen for the digital index falls within an image. MHEG provides no mechanism by which an application can determine the colour and brightness of a specific part of the screen, and working this out in advance would be a complex task. It is therefore very difficult to attempt to match the colour of the digital index to its surroundings. If the surroundings are made up of moving video it becomes impossible because in this case the MHEG instruction simply points to a MPEG-2 video stream within the multiplex.

Lastly, because set top boxes provide a pixel based display and have a limited grey scale, it is not straightforward to attempt to create a pulse train whose symbols are not an integral number of pixels. It is important to remember too that the creation of the digital index may increase the time taken to render the page, and increase the amount of interactive content data that has to be transmitted in the multiplex, so the method chosen must be as efficient as possible.

4 Minimising visibility

Visibility of the digital index will be partly determined by its size, and it is useful to estimate how much information it will be expected to carry. It may be tempting to keep the information to the absolute minimum – currently the BBC's Interactive service contains typically about 1500 pages at any one time, which could theoretically be indexed with an eleven-bit code. However a realistic estimate needs to include other identification, error correction and possibly a run-in sequence to aid detection by the logging unit downstream of the digital receiver, as well as allowing for a much larger number of page codes to avoid too frequent re-use. At this stage it was considered better to over-estimate the requirement rather than under-estimate it, particularly if other broadcasters were to adopt the concept of a digital index and some form of standardisation were to be pursued. It is not necessary to populate all the fields that are defined.

A 160-bit (20 byte) code would allow information to be carried about the broadcaster, service, genre, platform, region or transmitter, time code (although a clock could be incorporated in the detector), as well as allowing for error correction, check bits, and a run-in code. This compares with 360 bits in one line of CEEFAX. Assigning two pixels per bit would lead to a digital index about 45% of the active line, and a bit rate on the video waveform of 6.75 Mb/s.

For the informal tests, a 90-bit index was generated where the number of pixels per bit could be varied from one to eight, and both the horizontal and vertical placing varied. Minor modifications to the MHEG instructions were needed to ensure the index was correctly reproduced on all the set top boxes; this is not unusual for Interactive applications on the UK DTT platform, where detailed conformance tests for receivers are only just becoming available. With Dynamic Line Art the mark-to-space ratio of the pulse train was not unity. All the set top boxes tested lengthened the marks at the expense of the spaces by various amounts, to the extent that two closed the spaces entirely with two pixels per bit. This would place a practical limit of four pixels per bit to generate a well-formed waveform, equivalent to 160 bits taking about 90% of the active line. This restriction would not apply to the bit map technique, where one pixel per bit is achievable (although alternating the value on every pixel would result in a frequency above the luminance bandwidth of the video signal).

Using the majority of the active line made the index more noticeable, as might be expected. In particular this seemed to be emphasised if the index ran towards the edge of the screen. This seemed to be because it drew the eye, as if the index was perceived to continue beyond it.

Rendering the digital index on alternate fields only, i.e. a constant line in the 625-line PAL System I sequence, was quickly abandoned. On a 50 Hz field-rate television, the type most commonly in use, the interlace flicker was clearly visible, even at the lowest amplitude of luminance steps. Using a stationary index value will also make it less likely to be noticed. If a time code is included, its resolution and therefore the rate at which it increments needs to be no higher than is essential for monitoring. Also, repeatedly changing the information in the digital indices will place undue demands on both the limited bit-rate available to broadcast the data carousel and the graphics rendering capabilities of digital receivers, and could degrade the overall performance of Interactive services.

The most effective method of reducing the visibility was to place the index near the top or the bottom of the picture. On most of the televisions tested, the index could be made to disappear off the screen. However, one potential problem was found. Tests on six types of terrestrial set top boxes chosen at random showed that not all of them placed the index, or indeed any of the Interactive content, on the correct line. Of the six boxes tested, only three were correct. One placed the index one line early, another two lines early, and the third three lines early. As a consequence, the digital index was not rendered correctly or at all on these boxes if an attempt was made to render it on the first (full) one, two or three active lines of the field respectively.

The digital index disappears off the screen despite being present on the SCART signal because most domestic television receivers overscan. This actually works to the advantage of the proposed technique, as it will be off-screen and hence not visible on the majority of television receivers and displays. In 1999 BBC R&D surveyed overscan on 64 domestic televisions belonging to staff and to BBC R&D itself. It might be expected that these would be better adjusted than the average domestic television, but in practice 95% were found to overscan by 8 lines or more at both the top and bottom of the picture (4 lines per field). These televisions were all fitted with cathode-ray tube displays. Our experience of both LCD and Plasma televisions is that these are also set up to overscan (although the scanning mechanisms in these devices permit them to be both line and pixel-accurate) presumably to avoid the possibility of the viewer seeing any anomalous timing or blanking effects from the production or distribution chain.

While in principle the digital index could be injected anywhere in the frame, using lines 27 and 339 on even and odd fields respectively at the top of the picture, and lines 307 and 619 at the bottom of the picture should be a good compromise between rendering and visibility. This allows for the possibility that there are also UK DTT set top boxes that insert the Interactive content late rather than early.

Horizontal placement of the index is less critical, although subject to the visibility criteria already mentioned. It was noted that most of the set top boxes blanked several nominally active pixels at each end of the line, although it is unlikely that the amount of information to be carried would justify using the full length of the active line. There was also a horizontal timing variation between set top boxes tested of about 400 ns (measured between the leading edge of the line synchronising pulse and an arbitrary fixed point in the digital index). This confirms that a run-in code would help detection. This timing uncertainty led us to dismiss another idea, that of putting the index vertically down the side of the picture as a line of single or a small number of pulses.

5 Effect of widescreen switching and aspect ratio changes

This has yet to be investigated fully, but we believe that set top boxes do not clip Interactive content in any aspect ratio mode, nor do they move it when letter-boxing 16x9 pictures for displaying on 4x3 receivers (the interactive content being displayed full height over the letter-boxed picture). Any aspect ratio conversion performed in the television receiver itself will not affect the signal being intercepted on the SCART connection from the set top box.

6 Interference with other technical systems

The instructions and data for the digital index will conform to the MHEG specification and will be transmitted through the broadcast chain in exactly the same way as other Interactive content. The one situation where the digital index might conceivably be misinterpreted is when interactive pages are rendered in the studio, for transmission as conventional video, for example to promote the Interactive service to viewers watching the linear services. Even then, it should not be possible for any problems to arise. The digital index is entirely within the active picture area, and its amplitude and frequency range will be within those occupied by the normal video signal.

7 Detection

It is possible to make graphical content semi-transparent, which would reduce its visibility, but this would make the digital index difficult or impossible to detect if overlaid on still pictures. It is expected that the digital index would always be substituted entirely for any Interactive content normally occupying the chosen part of the screen.

Detection on the composite video signal is proposed because this is the only form of video output provided on some set top boxes. Otherwise there is no reason why the digital index should not be detected from the red, green and blue signals, either individually or in combination, or from an S-video output, where available. A coloured rather than monochrome index could also be considered, and this might prove less noticeable.

It should be noted that the detector will for much of the time be intercepting normal pictures (unless the viewer is a particularly avid Interactive user), and must discriminate between conventional video and the digital index. This will be assisted by the different characteristics of the index and typical video. The video waveform will usually excure most of the amplitude range from black to peak white, while the digital index will probably be injected at low level, for example between black level and the darkest grey in the MHEG palette (about 117 mV). Also, the digital index should appear on a known line in each field, certainly within a small window of lines. If the detector does extract what it believes to be an index from video, it is statistically unlikely that this will be declared valid because the error correction and check bits can be designed to heavily restrict the set of legal values. The digital index will also be stationary over a number of fields.

8 Considerations for a complete system

This White Paper has concentrated on the technical aspects of creating a suitable means of labelling each page of an Interactive service with an index, and some aspects of detecting that index. A complete system would require a numbering scheme to be devised, an automatic means, preferably, of passing these numbers to the Interactive content-authoring systems and for adding the extra instructions and data to each page. It would also require a means of extracting the viewing data collected in homes, possibly an auto-dialling modem in the detector unit using the domestic telephone line, and a means of analysing the results, and correlating the digital indices with the

pages they represent. All these would incur costs and raise design and management issues that are outside the scope of the work reported here.

9 Conclusions

The BBC is keen to have a method of accurately measuring how much viewers use its Interactive services on digital television, and a technical system based on the insertion of “digital indices” on Interactive pages has been proposed. Preliminary investigations including some informal tests of the visual impact of the index suggest that the method is technically feasible, and could be made acceptable both to our audiences and editorially. The technical aspects described in this document are only part of the solution, and it still needs to be established that the total costs are acceptable, in particular the costs associated with assigning and inserting the digital indices, and collecting and analysing the results.

10 References

1. Digital Terrestrial Television MHEG-5 Specification Version 1.06, BBC on behalf of TDN and Digital TV Group, 2003

11 Acknowledgements

I would like to acknowledge the contributions of colleagues Chris Poole, who wrote the code to create the digital indices for these tests, Neale Davison who carried out the set top box tests, Richard Salmon, who carried out the overscan survey, and Richard Stevens, for advice on modern television displays.