



R&D White Paper

WHP 096

September 2004

Archiving Interactive Digital Television

D. Chisholm¹, C. Gibson², J. Hunter¹, A. Pearce¹ and C. Smith³

¹ BBC R & D, ² BBC Technology Direction and ³ BBC Information & Archives

Archiving Interactive Digital Television

Duncan Chisholm¹, Chris Gibson², Jeff Hunter¹, Andrew Pearce¹, Cathy Smith³

¹ BBC Research & Development, UK, ² BBC Technology Direction, UK
and ³ BBC Information & Archives, UK

Abstract

Interactive Television is now considered an integral part of the viewer proposition for many digital television platforms. Increasingly interactive content is seen as a true peer of traditional linear content and subsequently there is a requirement for it to be treated as such in the broadcast chain. This includes archiving.

Archiving digital content can occur at different points in the broadcast chain. Generally these can be classified as either “pre-transmission” or “as-transmitted”. In both cases the archiving of traditional linear content, e.g. video and audio, is reasonably well catered for, though often relying on solutions developed for analogue broadcasting. However, at present there is little technology available to archive interactive content due to the diversity and nature of the deployed technologies.

This paper describes ongoing research into the archiving of as-transmitted interactive content.

This document was originally published in the Proceedings of the International Broadcasting Convention, September 2004.

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.

ARCHIVING INTERACTIVE DIGITAL TELEVISION

Duncan Chisholm¹, Chris Gibson², Jeff Hunter¹, Andrew Pearce¹, Cathy Smith³

¹ BBC Research & Development, UK, ² BBC Technology Direction, UK and ³ BBC Information & Archives, UK

ABSTRACT

Interactive Television is now considered an integral part of the viewer proposition for many digital television platforms. Increasingly interactive content is seen as a true peer of traditional linear content and subsequently there is a requirement for it to be treated as such in the broadcast chain. This includes archiving.

Archiving digital content can occur at different points in the broadcast chain. Generally these can be classified as either “pre-transmission” or “as-transmitted”. In both cases the archiving of traditional linear content, e.g. video and audio, is reasonably well catered for, though often relying on solutions developed for analogue broadcasting. However, at present there is little technology available to archive interactive content due to the diversity and nature of the deployed technologies.

This paper describes ongoing research into the archiving of as-transmitted interactive content.

INTRODUCTION

Interactive Television is now considered an integral part of the viewer proposition for many digital television platforms. Numerous operators have shown that when appropriately deployed it can attract viewers to the platform, reduce churn, increase audiences and generate revenue. However, the more Interactive Television content becomes a true peer of traditional linear content, e.g. video and audio, the more it needs to be treated as such in all aspects of the broadcast chain, including archiving.

Archiving systems are deployed for a number of reasons:

1. Re-transmission at a later date.
2. A historical record, i.e. a representative capture for posterity.
3. To fulfil legal obligations, i.e. a record of what was actually transmitted.

Over time it is expected that each of these functions will need to include interactive content, with the third of these becoming particularly relevant as Interactive Television comes under the increasing scrutiny of national regulatory bodies. However, to achieve this it will be necessary to modify, or even replace, existing archiving systems, many of which simply do not cater for Interactive Television content. Furthermore, any solution must be realistic in terms of both deployment and running costs due to the tight financial constraints under which most archiving units operate.

This paper describes ongoing research by the BBC, undertaken with the assistance of the BBC's operational archiving unit, to explore how Interactive Television content can be effectively incorporated into the archiving process.

REQUIREMENTS FOR ARCHIVING

Whilst there are a number of points within a typical broadcast chain where archiving of digital content can occur, they can generally be classified as either “pre-transmission” or “as-transmitted”.

Pre-Transmission Archiving

Archiving pre-transmission allows the broadcaster to store its content before it is suitably encoded for a particular digital television platform. This allows the broadcaster access to content for:

- Resale, where the content can be made available to a third party in the highest quality form possible;
- Distribution of content to programme makers in a form that allows them full artistic and editorial control, thus enabling them to generate material with the form and quality that suits their purpose.

Pre-transmission archiving of video and audio is already well catered with both tape based, (e.g. DigiBeta) and disk based storage solutions in use. In addition, new digital container file formats exist that not only facilitate the exchange of high quality programme content but also its associated meta-data, for example the Advanced Authoring Format¹.

However there is currently little support for any kind of standardised mechanism of archiving interactive content pre-transmission. This is because there is no well-adopted standard for the pre-transmission description of interactive content. Instead, interactive content is usually provided either in some platform specific, ready-to-transmit form, i.e. a pre-built application, or is built on-the-fly in response to changing input data, i.e. the service description is inherent within the instructions of one or more build processes. A further complication is that there are no well-adopted standards for describing the temporal linking of interactive content with the flow of linear content such as video and audio. A number of standardisation activities, including the DVB's Portable Content Format² group are active in this space and may prove to be relevant to the archiving of interactive content at this stage.

As-Transmitted Archiving

Archiving the as-transmitted signal (typically received off-air) has several advantages over archiving pre-transmission. The key advantage is that it provides a record of exactly what was broadcast. This allows the following requirements to be met:

- A “legal” archive. This is where a broadcaster has an obligation to record exactly what has been broadcast in order to satisfy legal requests on exactly what content was received in the home. Now that interactive and ancillary components are a key ingredient in most programme propositions, it is becoming imperative that they are archived in addition to video and audio content.
- An “historical” archive. This allows key broadcasting events to be recorded for posterity. Recording events for posterity could include recording the first few hours of a particular interactive event (e.g. Interactive Wimbledon on UK Digital Terrestrial Television (DTT)).
- Operational management. The archive could be used to examine the viewer experience retrospectively, and so confirm if operational procedures have been followed correctly,

¹ See <http://www.aafassociation.org>.

² See the DVB's CM-MHP-PCF and TM-TAM-PCF for more information

e.g. a “press red” prompt did appear within 5 seconds of the start of a programme.

- Fault investigation. The archive could be used for retrospective analysis of any element of the digital broadcast in order to locate a reported fault. For example, the exact composition of a DSM-CC³ download message could be analysed to verify whether a certain carousel element was present or not at a particular time.

In addition to providing a definitive record of what was broadcast, as-transmitted archiving has the benefit that it maintains the inherent temporal link between linear elements and the subsidiary components (i.e. interactive, subtitles and audio description). Potentially, it also allows for an archiving architecture that does not require complex hardware or software elements with explicit knowledge of the interactive television middleware of a particular platform.

Currently, as-transmitted archiving is often achieved by recording the output of a typical receiver onto VHS; a relatively low quality solution for video and audio only. More subtle archiving solutions encode the video and audio output from a typical receiver (whether analogue or digital) into a compressed digital form. The key disadvantage of these two methods is that currently they do not capture any further components of the service, including interactive elements⁴.

A PROPOSED ARCHITECTURE FOR AS-TRANSMITTED ARCHIVING

The following architecture has been developed around a requirement to archive broadcast systems based on MPEG-2 transport streams. However, the underlying principles are applicable to digital broadcasting in general. Furthermore, they are not complex. In fact, some are already seen in the home in media devices such as Personal Video Recorders (PVRs) that directly record elements of an MPEG-2 transport stream onto a hard disk. A broadcaster’s archive might be seen as an “enterprise” version of a PVR!

The general principles involve:

1. Capture of the input digital stream;
2. Retrieval of meta-data about the piece of captured content (i.e. time of recording, current event name etc.) either in real-time or with post-analysis;
3. Storage of this meta-data in an index;
4. Provision of an interface to a client wishing to retrieve a portion of an archived store, normally based upon a given piece of meta-data (e.g. the client wishes to retrieve the content for “Eastenders”, broadcast at 1930 on BBC1 on Tuesday 11 May 2004).

The following diagram represents this architecture. Note that the architecture principally shows process and data-flow, i.e. it represents interfaces between system components. Critically, it does not specify a particular choice of hardware for these system components.

³ Digital Storage Media – Command and Control (ISO/IEC 13818-6). This is a standard used in the broadcast of interactive services in UK Digital Terrestrial Television.

⁴ Further components could include subtitles, further video and/or audio components (e.g. audio description), or even EPG information (e.g. TV Anytime).

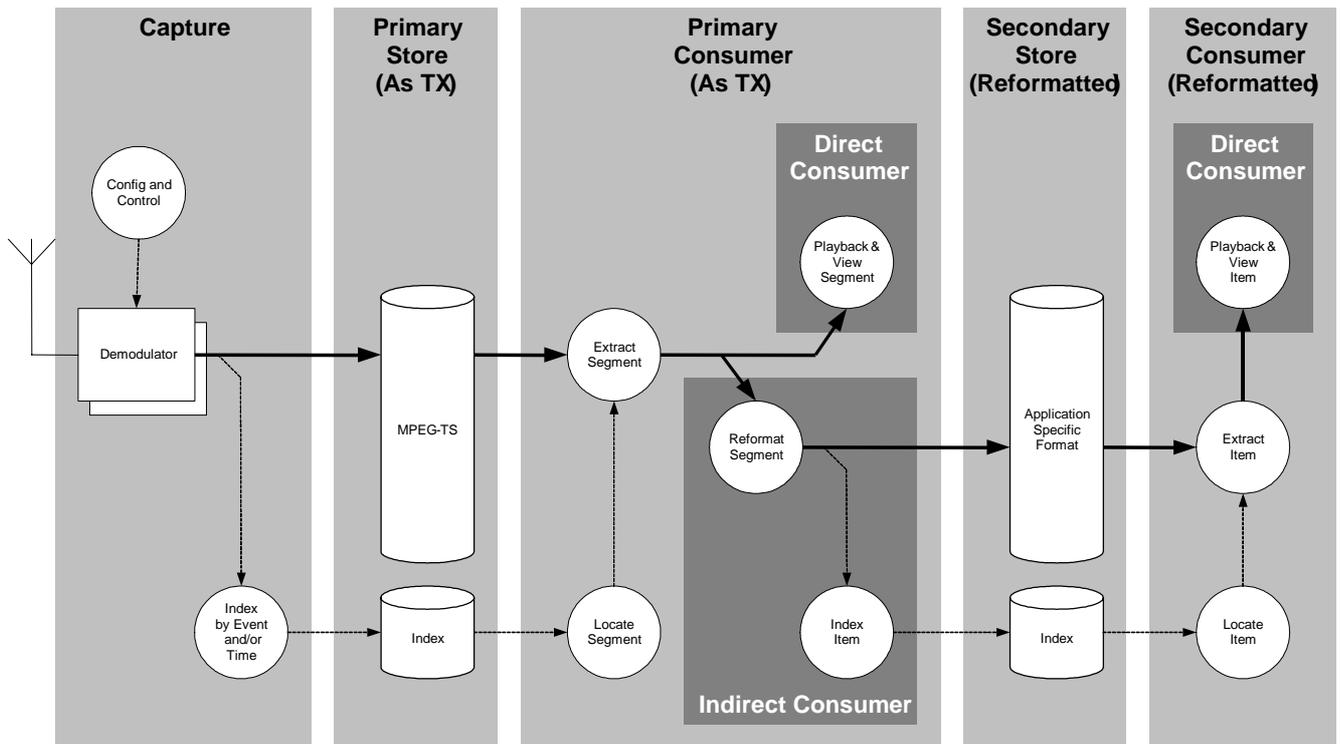


Figure 1: architecture for as-transmitted archiving

Capture

This stage of the architecture requires the direct capture of a single broadcast MPEG-2 transport stream, and the retrieval of meta-data from that stream that allows indexed access from the archive at a later date. Multiple instances of this stage can be used to archive more than one transport stream.

Capture of the broadcast transport stream inevitably requires some form of RF input and demodulation. In UK DTT, the input could be a standard rooftop aerial; for Digital Satellite it would be a satellite dish and Low Noise Block down-converter (LNB). Demodulation could be achieved with broadcast quality rack-mounted demodulators or even PCI cards in a standard desktop PC. Once the transport stream is available as a base-band signal it is possible to analyse it and so maintain an index into the archive. This analysis can be done on a standard desktop PC with an appropriate PCI based Asynchronous Serial Interface (ASI) card.

In a typical deployment the metadata that would be extracted from the transport stream includes:

- Broadcast time information.

In DVB systems this can be obtained by analysing Time and Date Tables (TDTs) or Time Offset Tables (TOTs) in the transport stream. The following table shows an example of what information might typically be stored in a "time record" for a number of fragments of a transport stream:

CPU Time	Broadcast Time	Bitrate (bps)	URI
22:01:21 11 May 2004	22:01:24 11 May 2004	24128342	file://archive/1.ts

22:31:21 11 May 2004	22:31:24 11 May 2004	24128342	file://archive/2.ts
----------------------	----------------------	----------	---------------------

Table 1: example time records for two archived transport stream fragments

When recording time information, it is not necessary to analyse every TDT or TOT provided the bit-rate of the captured transport stream is known. Once the time has been verified for a particular point in a data fragment it is possible to interpolate/extrapolate to any other “time” within the captured data. It may be necessary to analyse more than one TDT or TOT both in order to ensure that TDT or TOT is correct, and to remove any jitter seen on reception of these tables.

- Event information.

In DVB systems this can be obtained by analysing the “present” event from the Event Information Tables (EIT) “present/following”. The following table shows an example of what information would have to be indexed in order to retrieve an event from the archive:

Event Name	Containing Service	Start Time and Date	Duration	Description
Eastenders	BBC Three	22:29:15 11 May 2004	00:30:00	Soap Opera
Newsnight	BBC Two	22:31:27 11 May 2005	00:45:00	News programme

Table 2: example event records for two events

Information extracted from EIT may not always be completely reliable as the basis for an index into the archive. Firstly, EIT information is intended for presentation to the viewer and much of the information is generated accordingly. For example, the encoded value for the (billed) duration of a news programme might be 45 minutes and 00 seconds, which is rounded up from the less viewer friendly scheduled (intended) duration of (say) 44 minutes and 33 seconds. And if the news programme is live it might overrun so that its actual duration is (say) 46 minutes and 14 seconds.

An alternative is to track the transitions of the EIT present event. However, even this is not always reliable, particularly for reactive channels. It is worth noting that many of these issues have been explored in detail by DVB during recent work on broadcast signalling to support PVRs and that if introduced such signalling could be used to provide more accurate indexing for archiving systems.

Primary Store

The Primary Store is where the actual as-transmitted archive and associated index is written. The Store’s size and type will generally be a function of the time period to be archived and the number of transport streams being recorded. For example, archiving a single transport stream in UK DTT, which is broadcast at 18Mbps, for a period of 90 days requires a store that is approximately 17 Tbytes in size. For this kind of solution it is expected that a dedicated “network-attached store” of some kind would be used. However, to record one 18Mbps transport stream for just one day requires only 194 Gbytes - achievable using a standard IDE PC disk.

The transport stream archive itself is unlikely to be stored in the Primary Store as a continuous single file. In practice it may be necessary to split the archive into “fragments” whose size is dependent on factors such as:

- Capabilities of the archiving process/system writing the file;

- Number of clients that are allowed to simultaneously access the archive;
- Acceptable error of the time records for each transport stream fragment;
- How the archive size is to be managed. If a ring-buffer technique is to be used (i.e. when the archive size is near storage capacity, the oldest fragment is removed each time a new fragment is added) then the step change in the availability in the archive is equal to the size of the fragment. Thus the fragment size will impact the availability of content at the extremes of the archive.

This kind of detail should be a feature of the implementation of the Primary Store and hidden from any accessing process. However, it does create some implementation challenges in terms of retrieval.

Note that extracting the actual data from the archive may be more complex than just reading from the Primary Store; there may be an additional proxy required that enables access authorisation and manages the transfer of the content in a known and system-dependent manner.

Indexing relevant to the archive could be stored in a variety of ways. In its simplest form this would be a flat-file(s), i.e. the time records are written into a file in text-form, with another equivalent file for the event records. However it is expected that a fully implemented solution would provide access to database records through a fully featured Relational DataBase Management System (RDBMS).

Primary Consumer

A Primary Consumer is one that can use the as-transmitted data in the Primary Store, e.g. the transmitted transport stream. Such a consumer will typically request an item of content from the archive via a proxy. The proxy's responsibility is to locate and extract the relevant segment from the archive, and present it to the Primary Consumer.

Primary Consumers can be further categorised into Direct and Indirect Consumers.

A Direct Consumer will take the supplied transport stream and directly generate the corresponding viewer experience. An example of this is a commercially available set-top box, fed via a DVB-ASI output interface and a modulator.

An Indirect Consumer will take the supplied transport stream and convert it into some other form suitable for a Secondary Consumer. This converted version of the content is stored in a Secondary Store.

Secondary Store and Secondary Consumer

A Secondary Consumer is one that does not use the as-transmitted data in the Primary Store, e.g. the transmitted transport stream. The reasons for this will typically be:

- The encoding used for components of the as-transmitted broadcast require decoding resources that are not easily made available to the end users;
- The bitrate used for encoded components of the as-transmitted broadcast are too large to be able to deliver the content in an effective manner.

As a good example, consider the scenario where various programmes that have been archived need to be made available to all staff in an organisation via a company-wide Intranet. Providing all staff with suitable resources to make them a Primary Consumer, e.g. a set-top box and modulator feed, is likely to be both expensive and complex to manage and may significantly change the requirements used for the specification of the Primary Store. Further, if all staff can access the content on demand then capacity requirements could be

untenable given the relatively high bitrate used for the encoding of broadcast material, typically 10 times greater than that used for streaming over the Internet.

However, solutions that provide a large number of users on a shared network with access to streamed content on-demand are relatively commonplace, e.g. MPEG-4, Windows Media, Real. A better approach is to use these technologies as the basis for any Secondary Consumers and ensure that the required content is available in a suitable form in the Secondary Store, i.e. by a Primary, Indirect Consumer as described previously.

In this case, a Primary, Indirect Consumer would make a request from the Primary Store for each of the programmes to be made available, the result of which would be a set of transport stream segments. The Indirect Consumer would then extract the relevant components⁵ and trans-code them to a form that is presentable by the Secondary Consumer. And here lies one of the more complex issues. Whilst the trans-coding of video and audio can be achieved using a number of readily available tools, the trans-coding of interactive content is far more complex and not well supported. One approach is to only trans-code the video and audio and require the Secondary Consumer to support the as-transmitted format used for interactive content. This remains an open issue.

It is quite likely that content in the Secondary Store is not managed as a continuous archive of what was transmitted but rather as a library of items of content, e.g. programmes, that need to be accessed.

One of the key advantages of the separation of functions between the Primary and Secondary Stores is that they can be built to reflect their different specific requirements. The Primary Store will, as a priority, need to store continuous input from the broadcast stream reliably, whilst simultaneously serving a number of output/retrieval requests. On the other hand, the Secondary Store will, as a priority, need to service a potentially large number of output/retrieval requests, whilst supporting the upload of new content to the store. Aggregating these requirements into a single storage block may prove difficult to implement.

CONCLUSION

At the moment there is little or no available technology for enterprise-level archiving of all components of the digital output from a broadcaster. In particular, there is little support for archiving interactive content and other ancillary components, which are playing an increasingly important role in the viewer proposition.

Archiving the as-transmitted digital broadcast, such as an MPEG-2 transport stream, appears to be a viable approach to solving this problem, addressing many of the requirements for archiving. It is independent of the specific interactive television technology being used and is therefore applicable to all digital television platforms. It also preserves the temporal relationship that the interactive content had with video and audio.

Capturing MPEG-2 transport streams is technologically simple and can generally be achieved by using off-the-shelf hardware and even standard PCs. The software required to accurately store, index and maintain the archive is not complex. Critically the proposed architecture does not require any specific hardware; it defines "interfaces". Subsequently the archiving system is not constrained to any particular platform or device, and the cost of implementation is a factor determined by the archive's storage requirements.

⁵ It would do this by analysing the transport stream Service Description Table (SDT), Program Association Table (PAT), and finally the containing service's Program Map table (PMT).