

# MYTV: A PRACTICAL IMPLEMENTATION OF TV-ANYTIME ON DVB AND THE INTERNET

A. McParland<sup>1</sup>, J. Morris<sup>2</sup>, M. Leban<sup>3</sup>, S. Parnall<sup>4</sup>, A. Hickman<sup>2</sup>, A. Ashley<sup>2</sup>,  
M. Haataja<sup>5</sup>, F. de Jong<sup>6</sup>

<sup>1</sup>BBC, UK, <sup>2</sup>Philips, UK, <sup>3</sup>University of Ljubljana, Slovenia, <sup>4</sup>NDS, UK and  
<sup>5</sup>Nokia, Finland, <sup>6</sup>NOB, Netherlands

## ABSTRACT

The technical goal of the myTV project is to marry TV-Anytime standards, local hard disk storage, DVB transmissions, DVB-MHP set-top boxes and the internet. This paper gives a view of the myTV implementation of the ideas and specifications of the TV-Anytime Forum. The paper includes descriptions of technical solutions for carriage of TV-Anytime data over DVB and the internet, the two separate myTV box implementations, the data services required to enable selection and acquisition of content and the system issues in implementing these standards. The applications written to utilise these services and demonstrate interoperability are then briefly described.

## INTRODUCTION

myTV [3] is an EU-funded collaborative project with the aim of developing personalised services for digital television. The aims of the project are:

- to develop new services exploiting local storage such as navigational aids, and try out new forms of content such as non-linear viewing of news or magazine programmes.
- to contribute to emerging technical standards in this area.
- to demonstrate the interoperability of open standards.

Put straightforwardly, the technical goal of the project is to marry TV-Anytime standards, local hard disk storage, DVB transmissions, DVB-MHP set-top boxes and the internet. Our wider goal is to understand how the introduction of home storage will affect a viewer's relationship with broadcast content, and how to provide services that the viewer wants.

The TV-Anytime Forum specifications define the environment and system architecture in which metadata, content referencing, and rights management are standardised. As a global forum, TV-Anytime has decided to be "agnostic" about transport and so is independent of broadcast protocols. However, to make an interoperable, end-to-end system the whole protocol stack must be standardised. The myTV project has built and demonstrated a complete TV-Anytime system based on DVB. It includes proposals for how TV-Anytime "objects" (Content Descriptions, Content References, Tables mapping Content References to locators etc) can be carried in a DVB system, how a receiver can find and identify these objects, and how linkages can be made between content and metadata.

This paper describes the background to the problem and the proposed method for carrying TV-Anytime data in a DVB broadcast system. The paper includes brief descriptions of an

IP-based solution, the two separate, interoperable myTV box implementations, the data services required to enable selection and acquisition of content and the system issues in implementing these standards. The paper describes downloadable applications that demonstrate new services, and the implementation of third-party metadata services.

## END-TO-END ARCHITECTURE

The project has implemented a complete TV-Anytime chain, with content and data about that content being delivered to the viewer. Figure 1 shows the myTV end-to-end architecture. We have implemented two DVB-MHP set-top boxes which use TV-Anytime data and provide hard-disk storage, several sources of TV-Anytime data, carriage of this data over DVB and the internet, and several applications to make use of this data.

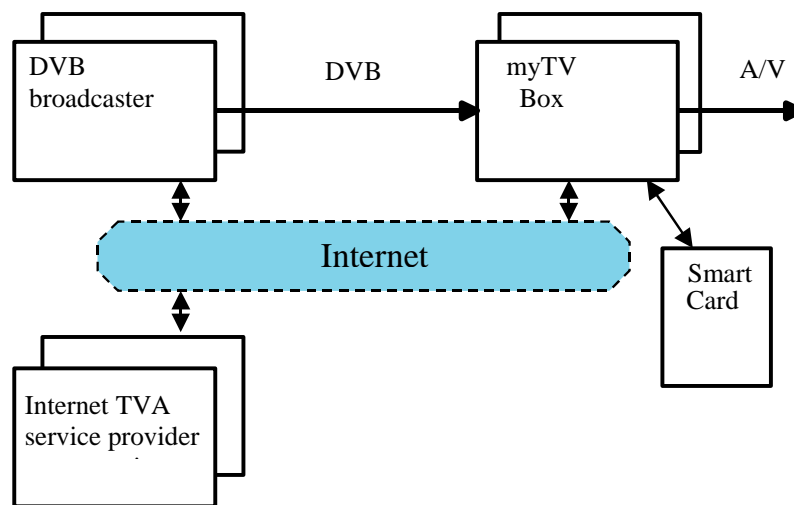


Figure 1. The myTV end-to-end architecture

## TV-ANYTIME OPEN STANDARDS

The TV-Anytime Forum recognised the need for a set of common specifications to support the processes needed to acquire and utilise content on a storage-based platform, dealing with metadata, content referencing and content protection. Common specifications are required, because there is a fundamental need to produce and transport content through a variety of transmission systems, to a broad range of boxes. Even vertically organised transmission systems with tightly controlled suppliers will benefit, as standards will help to ensure that generic and globally available content may be utilised.

The Content Referencing specification describes the separation of the reference to a content item (the Content Reference Identifier - CRID) from the information needed to actually retrieve the content item (the locator). The process of resolving a CRID to a locator (time & channel) is known as location resolution. The companion Metadata specification covers the data needed to enable the searching for, choosing of and navigating within, the content to be acquired – the 'attractors'. A systems document then ties these documents together to show how these specifications could be used to implement a TV-Anytime system. It is these three technical documents upon which the myTV project is based, implementing the transport-specific parts of the deliberately transport-agnostic TV-Anytime specifications.

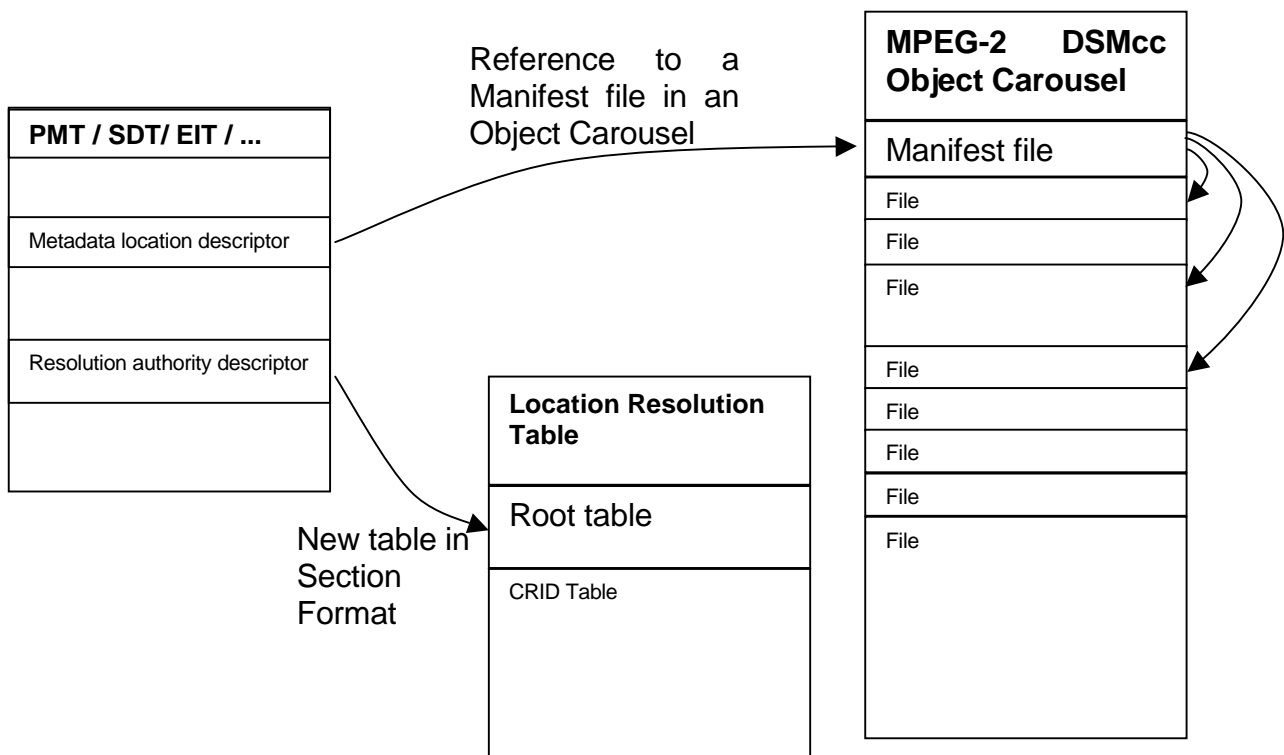
## DELIVERY OF TV-ANYTIME DATA

### DVB Implementation

We have implemented a method of delivering TV-Anytime Content Referencing and Metadata information over a DVB system to enable experiments with a complete TV-Anytime system. A design aim has been to develop a proposal that reuses existing standards as much as possible. At the time of writing, DVB is in the requirements phase defining a standardised method of transporting TV-Anytime data.

### Overview

Within MPEG-2, the Systems group is defining an amendment to the MPEG-2 standard that will create the basic framework of signalling needed to transport TV-Anytime data [0]. MPEG-2 has defined a set of low-level descriptors that enable the system to identify where metadata is carried within the broadcast stream. MPEG-2 has not, and will not define *how* the information is to be carried. That will be left to regional standards organisations. Figure 2 shows the basic signalling solution for transporting essential TV-Anytime data.



**Figure 2. Basic signalling structure for Content Referencing and Metadata information**

The following sections outline the basics of the system, how the different types of information are transported and how a system can find the information it needs within a Transport Stream.

## **Carriage of Content Referencing Information**

There are several ways in which location resolution information might be carried, such as in a private stream, an MPEG-2 transport stream section, in an object carousel or as part of the existing schedule information (i.e. EIT for DVB systems). After careful consideration, we have chosen to carry the information as MPEG-2 sections. They are lightweight, well suited to carrying the amount of information required and are flexible enough to support different update-rate, carousel rate, and bandwidth trade-offs. The system used in myTV is described in more detail in [1].

### **Resolution authority descriptor**

A new descriptor needs to be standardised to carry the TV-Anytime resolving authority record (RAR) (see [1]). The resolving authority record carries the information that identifies the locations where content reference resolution information can be found. In the case where the content resolution tables are carried within the DVB stream, the RAR points to the component that carries the Location Resolution Tables.

The descriptor may be included in the PMT in the descriptor loop for either the program or an elementary stream, but could be contained in SI tables such as the SDT, EIT etc. The advantages and disadvantages of different choices of where to place this information need more consideration.

### **Location resolution table**

The location resolution table allows a CRID to be mapped to a “DVB locator” that identifies the channel, time and event\_id of the broadcast. A location resolution stream is made up of a number of sections. The first type of section is called the root table, which lists all of the authority names that this location resolution stream can resolve. The second type of section is called the CRID table and provides the mapping from the CRID to be resolved to its locators. There shall be one section for the root table and one or more sections for the CRID table.

In order to save bandwidth, the CRID is split into two parts, namely the authority name and the data portion. In most situations there will be many CRIDs from the same authority in a location resolution stream. We assign a stream-unique numeric identifier to each authority, and use the numeric version of the identifier in the location resolution tables, thus avoiding repetition of the authority portion. More details can be found in [1].

Of course the exact time at which a programme is broadcast often differs from the intended time because of schedule slip caused, for example, by over-running live events. Within the DVB-SMI group it is proposed that DVB event\_id transitions are used for fine-control of the start and end of recordings, and that is what is used within the project.

## **Carriage of content description and segmentation information**

The methods used to carry metadata are similar to those used for Content Referencing information – a metadata locator descriptor is also required. A service can carry an MPEG-2 defined descriptor (the metadata locator descriptor) in the PMT, SDT, EIT or other table that identifies the location or locations of the relevant metadata service(s).

The metadata itself will be carried as files within an Object Carousel. The metadata locator descriptor will indicate the carousel and file where a *manifest file* is carried. The manifest file has a defined format (described in [2]) that describes the files where content description and segmentation metadata are carried. To make filtering of metadata easier, the manifest file is

carried in the root directory of the Object Carousel.

The reason to choose the Object Carousel is that it is already an MPEG-2/DVB standard implemented in deployed DVB systems. It has the functionality needed to define a versioned, hierarchical file-system “on the wire”. It includes an efficient compression system and is well-suited to the task.

Manifest files have a standard format defined in [2] starting with a header followed by a manifest file body. The header identifies the version, author, a brief description of the metadata, and a list of the CRID authorities that are supported by this set of metadata. The manifest file body is a list of CRIDs and file references where metadata directly associated with the CRID can be found within the object carousel.

The files containing metadata are TV-Anytime XML instance documents. Each file contains a single XML instance document, uncompressed, containing metadata about a single CRID. It is permitted to have more than one instance document for a single CRID. For example separate files could be used for Segmentation Data and for Program Information.

The metadata files are broadcast within the carousel as plain XML. Since Object Carousel uses zlib compression for compression of data within modules, and because compression of XML can be treated as a separate, non-interfering function, it has not been a priority to define a specific compression system for metadata.

### **Dynamic metadata for trailers**

Some of the TV-Anytime requirements and business models require that a CRID can be carried along with an item of programming in such a way that the CRID can be used to form a link to another programme. A typical example would be to insert a CRID into the stream during the showing of a trailer. The receiver can then allow the user to decide to record the content that is being promoted by the trailer.

This functionality is supported within the myTV system by carrying CRIDs within a new section format table. The component (PID) carrying this table is referenced from the PMT of the service.

### **Time synchronisation**

Segmentation is another important set of functionality defined within TV-Anytime. Segmentation allows a service provider to refer to different sub-parts of programmes and define a sequence order in which they should be played back.

Segments in TV-Anytime metadata reference sub-parts of the programme by time on an unambiguous, continuous time-line defined for the programme. The myTV project proposal is that DVB should use MPEG-2 DSMCC NPT (Normal Playtime) for these time lines. This of course requires that both head ends and receiving equipment can handle NPT accurately.

### **Internet Protocol Services**

The Internet Protocol (IP) Suite provides a well understood mechanism for providing set top boxes with programme information from a number of different metadata owners. By implementing such an IP-based end-to-end system, we have demonstrated the following:

- A TV-Anytime compliant content referencing service and metadata service, based upon IP.

- Interoperability enabled by TV-Anytime data. Three metadata providers (servers) were accessed by two different set top boxes (client devices).
- An embedded client device successfully handling large amounts of delivered XML.

### **The design of the system**

The HTTP protocol is used in the application layer, with the servers returning XML content in the body of the HTTP response. To specify the request the client forms a GET request with the URL query string encoded in the usual fashion (i.e. that commonly generated by an HTML CGI form). For example, to resolve two CRIDs the URL might be,

`http://www.movies.com/resolve?CRID=crid://movies.com/12&CRID=crid://movies.com/74`

whilst to search for content based on metadata requirements the URL could be,

`http://www.movies.com/tva.cgi?Actor=Tom+Cruise&Actor=Nicole+Kidman&Genre=thriller`

The second example shows the underlying mechanism by which a user can search for content. The field names in the query string are standardised and can be repeated, whilst undesired values (`Genre=horror`, say) can be explicitly excluded. The same mechanism is used to provide metadata for a CRID or CRIDs.

The format of the XML response is similar to that given in the TV-Anytime specifications - the differences are minor and only present because the TV-Anytime specifications were not finalised at implementation time. Since all responses must adopt the same XML grammar, which is well defined by the specified schema, the set top box can guarantee being able to interpret the data, whatever its origin.

### **Services and Applications**

Both the BBC and NOB have implemented a location resolution service and metadata service for their broadcast material. In order to simulate the different possible roles of a TV-Anytime third party, the University of Ljubljana implemented two services: the provision of broadcaster's metadata, which has been enhanced (contains extra information about movies); and the provision of detailed metadata for a specialist interest area (Champion's League football matches). Each service used a completely different scripting technology and database backend. For example, the University of Ljubljana used Java Servlets combined with a mySQL database, the NOB used a combination of PHP and mySQL whilst the BBC used Perl CGI and Java DOM.

The client devices were a Philips and a Nokia set top box with completely different hardware and software stacks. Both boxes contained a Java Virtual Machine which handled the networking and XML parsing. The boxes cached the data in a different manner. In other words, we have clearly demonstrated that there is considerable flexibility in the underlying hidden implementation on both the client and the server. It is clear that retrieving TV-Anytime data over IP using existing protocols can be realistically implemented on today's set top boxes.

### **Remote Programming**

The myTV Internet connection can also be used for remote programming of the set-box. Remote access to the myTV box has more security implications than downloading data and additional security measures need to be considered. For demonstration purposes, we have implemented basic authentication with usernames and passwords. Secure protocols like SSL for the Internet and WTLS for WAP would be required for commercial devices on the

market.

The myTV set-top box can be remotely accessed from any computer connected to the Internet with a Web browser, or from any WAP phone. The most attractive functionality is that of programming a recording remotely. By connecting to a 3<sup>rd</sup> party server supporting the TV-Anytime specifications, a user can send a record command to his set-top box simply by specifying a hostname of his set-top box and clicking on a title of a programme in a programme guide. A Web browser on a computer or a WAP browser on a WAP phone sends an HTTP request with a CRID of the programme to the myTV set-top box, which accepts the request as a record command after authentication of the user. To ensure the interoperability between different servers and different set-top boxes, the format of the command sent as a URL should be specified in an open standard (for example, TV-Anytime). For the myTV project we used the URL:

```
http://<Settop-box-IP-address>/3p?Method=Mark_for_recording&
      CRID=crid://bbc.co.uk/foo[&CRID="..."]
```

The same functionality would be available if the user had a direct connection to his set-top box, and then viewed the navigator running on the box remotely. Because no 3<sup>rd</sup> party could send a damaging command (e.g. delete instead of record) is involved in this case, all other commands are available for remote control of the myTV set-top box. The commands are sent to the set-top box as HTTP requests, which do not need to be standardised because they are sent back to the set-top box from where the Web or WAP pages have been downloaded.

## **SET-TOP BOX IMPLEMENTATION**

### **DVB-MHP**

The DVB Multimedia Home Platform (MHP) is an open standard platform for interactive digital television and multimedia services, which has been developed against a comprehensive set of commercial requirements that address the needs of both commercial and public sector broadcasters. The primary goal of the MHP is to enable the birth of horizontal markets for digital television and multimedia services, where there is open competition between content providers, network operators or platform manufacturers at each level in the delivery chain. It addresses this goal by specifying an interoperable application format, independent of any specific operating system or hardware technology, which allows MHP applications developed by different broadcasters to run on any MHP compliant receiver.

DVB has adopted Java as the interoperable application format for the MHP and has developed a version called DVB-J that includes the core of the standard Java language and provides extensions appropriate to the broadcast TV environment. The current specifications support enhanced broadcasting, interactive broadcasting and Internet Access. DVB is currently considering the possibility of extensions to support the provision of mass storage within the receiver platform.

The myTV implementation uses a subset of the MHP [7] which is supported by both Nokia and Philips to ensure interoperability. The public myTV-API is built over MHP. The functionality of the myTV-APIs are divided into three common packages: storage, location resolution and database access. These packages are briefly described below.

## **Storage: Recording and Playback**

The boxes must have means for managing their disk space, to ensure that requests to record are met or that failures are handled gracefully. The boxes must allow recording of arbitrary streams and playback of this same content.

It was agreed that there is no need for direct record commands, as we can use a request to record a CRID to provide the functionality that application writer's desire. For example, if an application needs to record the current broadcast, it can find out the CRID of the current broadcast and make a request to record it.

To play back an item from disk, the Content Reference Identifier (CRID) is returned by the API, which is passed then to the location resolution API. The locator returned by this can be passed to the JMF player APIs for playback. Playback control is based on the JMF 1.0 APIs, as defined in MHP.

## **TV-Anytime data access**

The boxes are able to receive and decode TV-Anytime data delivered via DVB and IP. The boxes store (cache) some of this data to provide immediate local access

## **Location resolution**

To fit with other ideas from MHP, the resolution of CRIDs, to CRIDs or locators, is based on using an extended version of the javax.tv.LocatorFactory API. When a CRID resolves into locators, the locator class has to implement the javax.tv.locator API. An advantage of this approach has been that, within the myTV project, we have not had to specify a format for locators pointing to content stored on the hard disk. An application can resolve a CRID into javax.tv.locator locators without needing to know any details about the underlying implementation. Furthermore, the LocatorFactory can be given a time range and a DVB channel and it will return a list of (CRID, locator) pairs of the programs that will be broadcast within the given time range and channel.

## **Access to Metadata**

The underlying database API has been abstracted in order to hide the underlying implementation. Two APIs are defined, one for making queries on the database inside the set top box, and another API to query a metadata server on the Internet. A downloaded application can only read from the database, as being able to write to the database was felt to be a security risk. Any application that wants to store information can store its data in its own files. If applications were allowed to write to the database, it is not clear how different applications would be able to find their own metadata whilst skipping other applications' private metadata extensions.

The metadata can be accessed using either local cache or a remote server depending on how up-to-date the metadata needs to be. The local cache is accessed using the mytv.database.TVADatabase class and the remote server is accessed using the mytv.database.TVARemoteDatabase class. Both classes have two different functionalities. Firstly, all the metadata associated with one CRID can be retrieved and secondly all the CRIDs matching some metadata criteria can be retrieved.



## **APPLICATIONS**

Several applications have been written to exercise the system, to test interoperability of TV-Anytime data and to demonstrate TV-Anytime-focussed MHP applications. They are briefly explained here, but are considered in more detail in the companion paper [3].

The other applications show some of the possibilities of TV-Anytime ideas and test interoperability at the application level.

### **Resident Navigators**

Resident navigators can use all of the capabilities of the manufacturer's platform, and the TV-Anytime data resources that are available. Different approaches to the design of Resident Navigators have helped to identify the issues in using metadata to search for content, and allow us to demonstrate the interoperability of TV-Anytime data by using the use the same data source with a different look and feel, and other abilities.

### **Downloadable Navigators**

Downloadable navigators are MHP applications that provide the user with a view of programme data and the available content. Downloadable navigators cannot hope to have as full access to the capabilities of the platform as a resident navigator, however they can provide the best interface to particular features of a TV-Anytime data set generated by the authors of the navigator. Downloadable navigators developed within the myTV project test the TV-Anytime APIs and explore issues of access to the disk and other services.

### **Segmented Latest News Application**

Within the myTV project NOB (Dutch Broadcasting Services Corporation) has developed a 'Segmented Latest News' application. It demonstrates how local storage can be used to enhance an important type of programme, and it is used as a vehicle for testing TV Anytime Segmentation metadata and extensions to MHP APIs.

### **Trailer Selection**

Trailer selection allows a user can request the advertised programme to be recorded. The project has implemented this feature by allowing CRIDs to be inserted into the stream during the broadcast of a trailer. The implementation of this feature has raised issues such as how to cope with multiple short trailers in a row, and generates ideas such as a real or virtual trailer channel for pointing to new or interesting programmes.

## **CONCLUSIONS**

The project has taken a leading role in the development of the TV-Anytime standards, and fed back lessons learned from implementing a TV-Anytime system. We have implemented IP-based services, with multiple metadata service providers and multiple users of the services. We have developed protocols for carrying the data over DVB and implemented them. We have developed two boxes that implement TV-Anytime to demonstrate the attractive features of the system, and in doing so have demonstrated interoperability of an open standard. We have developed extended APIs for the DVB-MHP to implement TV-Anytime services, for both storage- and data-based facilities. Several applications were written to take advantage of these APIs and, interoperably, run on both boxes. We have

also demonstrated that open standards allow a third party to provide useful and interesting user services without requiring any access to the broadcast chain. In short: we have implemented TV-Anytime, and it works.

## **REFERENCES**

0. MPEG, 2001. PDAM 1 to ISO/IEC 13818-1:2000. ISO/IEC N3845.
1. myTV, 2001. Proposal for MPEG-2 Carriage of Content Referencing Information. TV-Anytime contribution AN250.
2. myTV, 2001. Proposal for a metadata location record. TV-Anytime contribution AN235.
3. Tol. et al., 2001. TV-Anytime metadata: Attractors and segmentation – a myTV view. IBC 2001.
4. TV-Anytime, 2001. Content Referencing Specification. SP004v1.1.
5. TV-Anytime, 2001. Metadata Specification. SP003v1.0.
6. TV-Anytime, 2001. System Specification. SP002v1.0.
7. DVB, 2000. MHP Specification. Tam232r15.

## **ACKNOWLEDGEMENTS**

The authors would like to thank their colleagues for their contributions to this work. We would also like to thank our respective companies for permission to publish this paper.