

Share it! – Content Transfer in Home-to-Home Networks

Boštjan Marušič, Andraž Jere, Marko Meža, Štefan Dobravec, Jurij Tasič
University of Ljubljana, Faculty of electrical engineering
Tržaška 25, SI-1000 Ljubljana
Slovenia
email: bostjan.marusic@fe.uni-lj.si

Abstract—This paper presents a possible solution for the introduction of a complementary distribution chain into the broadcast environment. In the described scenario personal video recorders with local storage are networked into an ad-hoc home-to-home network based on an emerging peer-to-peer protocol, enabling broadcast as well as home-produced content sharing. The paper presents the current system architecture as conceived within a consortium working on the Share it! project¹.

Index Terms—broadcast content sharing, P2P, TV-Anytime, Share it!, DRM

I. INTRODUCTION

POWERFUL signal processing, cheap mass storage and the increasing availability of always-on broadband networks are set to radically change models and habits of TV viewing. The Share-it! project examines how these changes may develop and how these technical factors interact with the world of digital rights management (DRM).

Powerful signal processing and cheap mass storage discs have led to the development of Personal Video Recorders (PVRs). PVRs are digital video recorders that allow much easier programming, instant access to recorded programmes, and methods of automating the capture of TV programmes. Early PVRs have been deployed in several markets (USA, UK, Japan) and have been met with both fear and enthusiasm by broadcasters and viewers. The experience of existing users is that broadcast times and schedules become much less relevant as the PVR can give easy access to the programmes you want to watch at the time you want to watch them. Often this is reported as leading to people watching more television, but at times of their choice. From a broadcaster's perspective, it is important to understand how PVRs create new ways to attract and keep viewers. For instance, PVRs automatically select and store TV programmes according to viewers' preferences by analysing content descriptions ('metadata') broadcast along with TV programmes and comparing these to the preference profile of the user. Well-authored metadata can facilitate this, and for instance help

viewers to record all the programmes in a series with a single action, or allow them to book a recording when the trailer is broadcast. Properly used, these tools can strengthen the branding of a channel.

Open standards, as being developed by the TV-Anytime forum [4], are essential for a strong development of this new business. TV-Anytime is a global standardisation forum consisting of organisations covering the entire broadcast chain (content creators, broadcasters, service providers, telecommunications companies and manufacturers of consumer devices). It is developing standards for the end-to-end signalling necessary to enable PVRs. This will enable consumers to explore and acquire a much wider range of TV content from a variety of sources – “what you want, when you want to watch it!” Further, DVB-MHP [5] will be extended to allow TV Anytime PVRs to download and execute “multimedia home platform” (MHP) applications that can access Internet servers, other devices and locally stored content. This combination has the potential to deliver a range of new options for the future of television.

The rapidly increasing availability to consumers of always-on broadband networks is another key factor that, in parallel with the evolution of the interactive television, will change television usage patterns. Access to a network is changing the habits and expectations of users. Besides being in control of when and what is watched, there is a trend towards people taking a more active role in content creation and distribution. By turning from a passive content consumer to a content producer, the user is challenging the traditional unidirectional broadcast model of distribution of TV.

As a consequence, a new content distribution model is evolving as a complementary distribution chain to the broadcast network. Enabled by broadband infrastructure availability and recent developments in peer-to-peer networking mechanisms, PVRs with massive local and network storage can be organized into a distributed content repository in which broadcast content as well as user provided content can be stored and accessed.

¹ The work described in this paper was funded by the European Commission (FP5 IST) under contract IST-2000-28703

The Share it! Project [1], funded by IST, is working actively to explore solutions and drive standards in this area. Share it! is building a system that enables content access/consumption scenarios similar but not limited to the following:

- An end-user can search for content regardless of its source (broadcast or network storage). If a programme has been pre-stored in the network the user can access it even if he missed the original broadcast.
- An end user connects to the system remotely and can view and manipulate stored content.
- An end user grants access to a limited set of his own content to another user of the Share it! system, possibly in another home.
- Personalized virtual channels can be created mixing broadcast local and network stored content
- Groups of users share content based on their interests.

In all of these scenarios a system of digital rights management is vital. The establishment of the connection is based on users' access and viewing rights to specific content, which may include (non-copyrighted) audio-visual data, different sets of metadata and other user-provided data. The networking of PVRs is capable of complementing and enhancing the traditional broadcast chain. The scenarios envisage peer-to-peer communication, where of course a central question is user authentication, access rights, control of the rights of content owners.

The aim of the Share it! project is to develop solutions that enable the parties involved to benefit from these significant changes in the world of TV:

- End users will benefit from a very rich audio-visual experience and may become content creators, either by putting audio-visual material up for sharing or by adding context to content.
- Service providers and broadcasters will benefit from the extended capabilities of a system using both on-line and broadcast content, enabling them to deliver new innovative revenue generating services and involve their consumers in the actual service creation process.
- Network operators, content creators, equipment manufacturers and other players in the chain will benefit from the increased demand for their products and services generated by the enhanced system capabilities.

The high-level architecture of the Share it! system can be seen in Figure 1.

II. PEER-TO-PEER NETWORKING ASPECTS OF SHARE IT!

Peer-to-peer networking is a new style of distributed computing with the following characteristics:

- no functional hierarchy between network nodes (peers) – for example a network node can be a client and server at the same time (although this is not entirely true since

specialised devices need to exist in the network in order to support the network functionality – for example JXTA rendezvous peers [7], centralised certification authorities).

- network nodes share resources with and provide services to other network nodes
- communication between network nodes is established on a hop-to-hop basis (network nodes forward messages to known neighbours until the message reaches the recipient).

Peer-to-peer networking has been recently popularised by file sharing systems such as Kaaza, Gnutella and Napster. These technologies lack two main properties in order to be able to provide an acceptable and interoperable solution for all the players involved in the content distribution chain: open standard protocols and DRM solutions that would enable content protection from unauthorized access.

The Share-It! network is being designed as a content sharing and distribution network. In other words content and associated metadata stored on the Share-It! device is made available to other users of the network subject to rights management conditions. The Share-It! user perceives the Share-It! network as a global content repository from which content can be accessed or where own content can be published.

The Share-It! network complements the broadcast network by providing persistent storage of content that was broadcast and thus providing an alternative acquisition mechanism. Furthermore each Share-It! user can become a content provider. The Share it! Network can be seen as a content sharing P2P network where content and associated metadata enters either through recording of DVB broadcasts on the PVR local storage or by adding content to the PVR directly (for home-grown content). Note that broadcasters can also inject content and metadata directly into the broadcast network beside broadcasting it through DVB systems.

A. Employed P2P protocols

In order to implement the required Share it! P2P functionality, Share it! employed an emerging protocol stack called JXTA [3]. The software layer developed on top of JXTA protocols enables content publishing and search, content recommendations and announcements etc. in the P2P network.

The JXTA project was initiated by Sun Microsystems, but in April 2001 the project was placed in the hands of the P2P community by adopting an open source license.

The vision of the JXTA project is to provide an open, innovative collaboration platform that supports a wide range of distributed computing applications and enables them to run

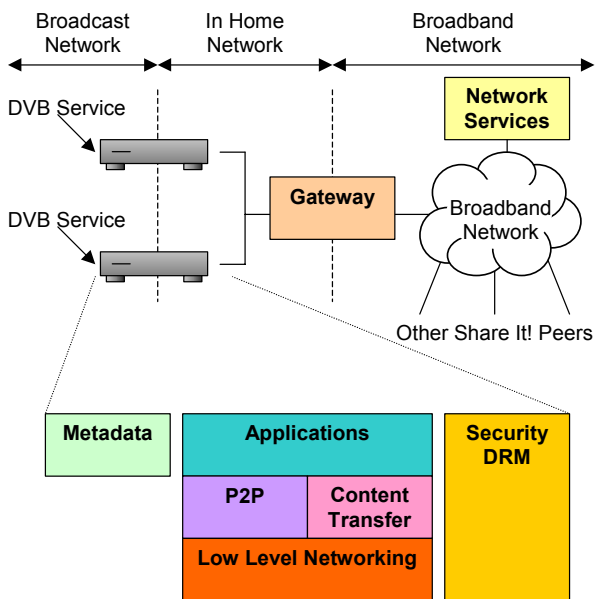


Figure 1: High-level architecture of the Share it! system.

on any device with a digital heartbeat.

The goals of the JXTA project are to provide a "general-purpose" network programming and computing infrastructure. Its goals are:

- *Interoperability*: by enabling inter-connected peers to easily locate each other, participate in community-based activities and offer services to each other seamlessly across different peer-to-peer systems and different communities.
- *Platform independence*: JXTA is designed to be independent from programming languages (such as Java or C), system platforms (such as Linux or MS windows) and network platforms (such as TCP/IP or Bluetooth)
- *Ubiquity*: JXTA is designed to be implementable on every device with a digital heartbeat.

Peer groups are the core of JXTA's infrastructure. A peer group is essentially a partition of the world of peers for communication, security, performance and other reasons. A single participant can be in multiple groups at one time. JXTA provides core protocols for peer discovery, peer group membership and peer monitoring. JXTA uses asynchronous uni-directional communication channels, called pipes, for sending and receiving messages. All data interchange in JXTA is in the form of XML formatted documents.

1) *JXTA architecture*

JXTA is a framework with a set of specifications [7] for protocols that support any imaginable peer-to-peer application. The JXTA protocols are defined in terms of XML-based messages and are not rigidly defined, so their functionality can be extended to meet specific needs.

JXTA is made out of three distinct layers. The first is the platform, the JXTA Core. The platform contains the core functionality, which is used by services, which are the second layer. Services provide access to the JXTA protocols. Finally, there are applications that use services to access the JXTA network and utilities.

The layered architecture enables JXTA to easily incorporate new protocols and services to support a growing number of users.

The JXTA protocols are specifically designed to traverse gateways and firewalls. This firewall transparency allows the same protocols to be used within the home as those used outside the home.

B. *Content search in Share it!*

The inclusion of a decentralized content distribution network into the traditional content distribution chain requires a new resolution approach that cannot be implemented by the centralized TV-Anytime "resolution authority" concept [4]. The term resolution authority refers to a centralised server mechanism that resolves a content reference identifier (CRID) associated with a particular content to a broadcast schedule (time/channel of broadcast) or to video on demand server locations (URLs). This by no means excludes the use of the resolution authority concept in the broadcast environment. The Share-It! network is a complementary content distribution system (with its own specific properties) and as such it requires a specific resolution approach.

Due to the non-deterministic (devices come and go, new content appears and disappears randomly, availability of content changes, etc.) behaviour of distributed networks performing location resolution based on decentralized peer-to-peer search seems the most appropriate solution. This means that a particular content is searched by sending out a discovery message (query) with searched CRID or other metadata information. The query is propagated throughout the network on a hop by hop basis reaching all relevant peers on the network. The peers that have the requested content signal back the content location as well as metadata. The method of query propagation heavily affects system parameters such as response time and scalability. This is an important area in which the Share-It! system architecture has to provide elements for optimising query propagation in the network.

The distributed storage of content in the Share-It! network can be seen as a complementary content distribution chain.

Share it! uses a JXTA mechanism to propagate queries and responses through the network called the Peer Resolver Protocol (PRP) [7][8]. Share it! queries and query responses are formatted as specified in TV-Anytime specification on Bi-directional metadata [9]. These XML formatted messages are wrapped by the JXTA PRP message envelopes in order to be propagated through the network. Example of Share it! query/response messages and their transport within PRP messages can be found in figures 2, 3 and 4 .

```
<?xml version="1.0"?>
<!DOCTYPE get_Data>
<get_Data maxPrograms="10">
  <QueryConstraints type="AND">
```

```

<Predicate fieldID="tvaf:Title"
fieldValue="Santa Claus" test="contains">
</Predicate>
<Predicate fieldID="tvaf:Synopsis"
fieldValue="Christmas" test="contains">
</Predicate>
</QueryConstraints>
< RequestedTablesType>
  <Table type =
"ProgramInformationTable" />
  <Table type = "ProgramLocationTable" />
  <Table type = "ProgramReviewTable" />
  <Table type = "ShareItSpecificTable" />
< /RequestedTablesType>
</get_Data>

```

Figure 2: An example Share it! query where user searches for The movie titled “Santa Claus” where the synopsis contains the word “Christmas”. Note that the query also specifies which TV-Anytime information it wishes to receive in responses to this query (see the RequestedTables tag).

```

<?xml version="1.0" encoding="UTF-8"?>
<jxta:ResolverQuery
xmlns:jxta="http://jxta.org">
<HandlerName> . . . </HandlerName>
<Credential> . . . </Credential>
<QueryID> . . . </QueryID>
<SrcPeerID> . . . </SrcPeerID>
<Query> Payload </Query>
</jxta:ResolverQuery>

```

Figure 3: The JXTA PRP query message in which Share it! queries are propagated through the network as payload of the <Query> field.

```

<?xml version="1.0"?>
<!DOCTYPE get_Data_Result>
<get_Data_Result version="1" truncated="true">
  <TVAMain>
    //Content advertisement for matching
    content #1
  </TVAMain>
  //Content advertisement for matching
  content #2
  <TVAMain>
    //Content advertisement for matching
    content #3
  </TVAMain>
  ...
</get_Data_Result>

```

Figure 4: The Share it! query response message. Note that the message is mainly a list of <TVAMain> fields that represent content metadata as specified by TV-Anytime.

C. Content transfer in Share it!

Share it! query responses contain also locators from where the actual resources can be downloaded. Share it! has limited itself to support only MPEG 2 content. The content itself can be either downloaded using the HTTP protocol or real-time streamed using RTSP. In other words, conventional and standard mechanisms are employed for the transfer, while only search and location is achieved by means of the P2P

network. Note that Share it! locators are HTTP or RTSP URLs.

III. DIGITAL RIGHTS MANAGEMENT

A key aspect of Share-it!, and one that has a pervasive impact on most aspects of the project, is rights management. This section is an explanation of our basic approach to this issue by giving a high-level review of the rights processes that are required, and then a rather more detailed look at some of the implications.

Note that Share it! enables different content distribution models, that based on DRM aspects can go into one of the four categories:

1. Traditional Broadcast Content. Traditionally, the broadcaster only has the rights to broadcast a given piece of content in given territorial domain (typically national, or multi-national. The current distribution of content by broadcasters is to implicitly managed populations of consumers (managed by virtue of the broadcaster platform, conditional access system and broadcast footprint).
2. Global distribution and payment. In this model the copyright holder distributes content globally and derives revenue. An example is the direct distribution of niche interest material on video. Share-It! offers a new distribution path for such content. In this model a transaction must be made and content rights adhered to even though there may not be any ongoing business relationship between the content provider and user.
3. Freely distributed content. In this model, the rights holder does not demand any payment (i.e. they are effectively not asserting their rights when the content is published), and the content may be freely distributed anywhere.
4. Private content. This is content where the rights owner wishes to restrict distribution and even visibility. An example is home produced content for distribution to a users family only.

A. A seven-step approach to rights management

This section identifies and describes seven basic steps necessary to search for and successfully acquire content in a rights compliant Share it system. The seven steps can be summarised:

1. Query - response
2. Rights resolution
3. User interaction and content selection
4. Rights exchange transaction
5. Secrets acquisition
6. Content acquisition
7. Adherence to usage rights

The process steps are described in more detail below with the assumptions that:

- All peers knows some basic rights related information about themselves, e.g: UK based, Sky platform
- Content shared from a broadcast will have some basic

rights related information available, eg UK, BBC1.

- Online rights resolution authorities maybe available to resolve or broker rights requests.

1) *Step 1: Query-response process*

Queries can include a simple rights profile to allow peers to respond with content that fits the profile of the requesting peer. Responses to queries should indicate some simple content related rights information to allow a simple content rights – user rights match. This match can be used to help the requesting peer filter or select content already within its known usage profile. A response may include a rights resolution authority reference.

2) *Step 2: Rights resolution*

For cases where there is no direct rights profile match from the simple content rights metadata, the listed rights resolution authorities may be used. This rights resolution might comprise

1. Submit a rights query to the rights resolution authority-including:

- Content Instance descriptor incorporating simple content rights info
- Requesting peer ID for rights resolution including simple user rights currently held.
- ID of server peer offering content.

2. Redirection between rights brokers to find appropriate server based on user locale, content ownership etc.

3. Response, if positive to include:

- Offer – selection of usage models
- Prices
- Terms and conditions

3) *Step 3: User interaction and content selection*

For any content returned from a search the user must be presented with the options in a simple manner so selection can be made.

4) *Step 4: Rights exchange transaction*

After the decision has been taken to commit to the offered terms and conditions, the signalling of this acceptance, acknowledgement and associated transaction takes place. This may be done locally on the PDR, if the content is designated accordingly. Otherwise this transaction will be server based or peer to peer based. The outcome of this process is that the user is granted rights to receive and consume the content.

5) *Step 5: Secrets acquisition*

Having obtained the rights for a given user to access the content in a specified way, the PDR then has to acquire the necessary secrets that enable the content to be transferred, used or both (encryption keys etc).

6) *Step 6: Content acquisition*

Once the rights and necessary secrets have been acquired, then the content can be transferred. N.B this transfer could in principle take place at any point after step 3.

7) *Step 7: Usage adhering to rights.*

Within the constraints of the precise rights granted to the user with respect to the specific content during the rights resolution process, the PDR should allow and control the use

of the content.

Depending on the type of distribution the rights acquisition process can be termed as light touch or heavy touch. Light touch refers to the case where the entire content acquisition process occurs exclusively between peers, the content and rights are transferred between the devices without the involvement of any external rights brokers. Usually the peer providing the rights information is the peer providing the content itself. To enable some models of rights managed content transfer requires the involvement of trusted third party rights resolution authorities, that provide the necessary rights related information. This is the so called ‘heavy touch’ approach to content acquisition. In this it is assumed that the content acquisition involves the payment of a fee.

IV. CONCLUSION

In this paper an approach to the introduction of a complementary distribution chain into the broadcast environment based on P2P networking of PVRs was presented. Both the mechanism for discovering content in such a decentralised repository as well as the mechanism for ensuring authorized usage of content were introduced and discussed from a high level perspective. We feel that the described developments can have deep impact on the future of interactive TV and therefore represent a challenge as well as an opportunity for all the players in the broadcast chain.

ACKNOWLEDGMENT

The authors would like to thank all members of the Share it! Consortium for their cooperation. This includes the teams at Philips Research Labs Redhill, Philips Research Eindhoven, British Broadcasting Corporation, NDS, Nederlands Omroepproductie Bedrijf N.V., KPN, University of Linköping, FhG-Forschungszentrum Informationstechnik GmbH and Elisa Communications Corporation.

REFERENCES

- [1] The Share it! project web page, http://www.extra.research.philips.com/euprojects/share_it/
- [2] Report on the general infrastructure and results from feasibility tests, Share it! deliverable 5, http://www.extra.research.philips.com/euprojects/share_it/
- [3] The JXTA project web page, <http://www.jxta.org>
- [4] The TV-Anytime forum web page, <http://www.tv-anytime.org>
- [5] The web page of the Digital Video Broadcasting project, <http://www.dvb.org>
- [6] The web page of the MHP project, <http://www.mhp.org>
- [7] JXTA v2.0 Protocols Specification, <http://spec.jxta.org/v1.0/docbook/JXTAProtocols.html> Protocol
- [8] Brendon Wilson: JXTA, <http://www.brendonwilson.com/projects/jxta/pdf/JXTA.pdf>
- [9] The TV-Anytime Forum: Metadata Services over a Bi-directional Network, Specification Series: S-6, <http://www.tv-anytime.org>
- [10] The HTTP protocol specifications, HTTP 1.0 (RFC 1945), HTTP 1.1 (RFC 2616), <http://www.ietf.org>
- [11] Real Time Streaming Protocol (RTSP), <http://www.ietf.org/rfc/rfc2326.txt>