SHARE IT! - A RIGHTS-MANAGED NETWORK OF PEER-TO-PEER SET-TOP-BOXES - SYSTEM ARCHITECTURE

J Walker[1], O J Morris[2], B Marusic[3]

[1]NDS UK, [2]Philips UK, [3]University of Ljubljana Slovenia

ABSTRACT

The *Share it!* project is developing and demonstrating a system of rights managed content sharing between homes connected by the broadband Internet in a peer-to-peer network of set-top boxes with storage. This paper looks at the system requirements, the design choices that have been made, and the overall system architecture that has been developed.

INTRODUCTION

Share it![1] is an EU-funded collaborative project [1]. It has the aim of developing an end-to-end system for easy access and transfer of content between people in their homes using the broadband network. It enables innovative services through a seamless combination of on-line, broadcast, and stored content, in a rights-managed world. It contributes actively to international standards enabling interoperability between content, service providers, and equipment manufacturers.

This paper outlines the context of the project, describes the architecture of the system that has been developed, and discusses the issues raised during the work of the project.

PROJECT GOALS

The goal of the Share it! project is to develop an end-to-end system enabling easy access and transfer of personal, stored and broadcast content, and in particular to:

- Develop usage scenarios and models for discovery and delivery of content and associated applications
- Enable seamless integration of broadcast, broadband and stored content
- Apply/extend/develop home-to-home, peer-to-peer technologies that enable these scenarios
- Experiment with peer-to-peer (P2P) scenarios and applications
- Develop an end-to-end suite of P2P protocols
- Pay particular attention to Digital Rights Management (DRM) issues

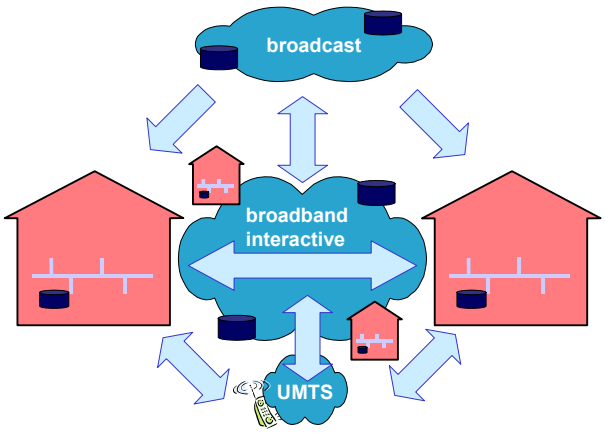

Figure 1 – End-to-end system

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

- Verify and demonstrate the system
- Contribute to DVB [2] and TV Anytime (TVA) standards [3].

SYSTEM REQUIREMENTS

In deriving the Share it! system requirements, we have adopted a user-centric top-down approach. Initially the project partners developed more than 50 usage scenarios. Examples of such usage scenarios include:

- Showing your home video to other named devices
- Sharing content amongst user groups
- Viewing and passing content rights

These usage scenarios were then analyzed for common properties, from these we derived user and system/content provider requirements, and then the technical requirements. These requirements were then prioritized according to whether we intended to support them in the demonstrator system or study the issues. The system requirements fell into the following categories:

- The Share it Set-Top Box
- Types of Content
- Content Distribution
- Security
- Content Creation
- Rights Management
- Discovery of content and User Groups
- Personalization

Rather than give a full detailed list of the system requirements [4] in this paper, a set of example applications is given in the following section which illustrates the kinds of applications and scenarios that the Share it! platform supports.

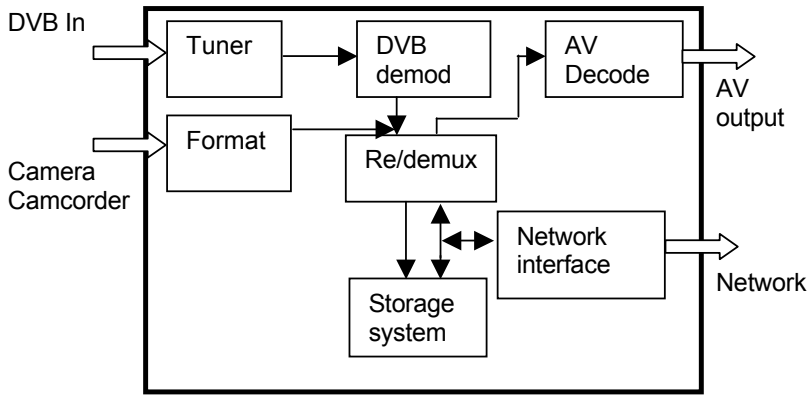

Figure 2 is a simplified view of the architecture of a Share it! terminal. In effect it is a set top box with a hard disc drive capable of storing content and metadata, making it a normal personal video recorder (PVR), with the addition of a network connection. The increase in cost over a PVR is small.

Figure 2 – Simplified view of the Share it! box

EXAMPLE APPLICATIONS

Three broad categories of user scenarios were identified in the project: basic content sharing and exchange; network services and opportunities for adding value to broadcast services; and publishing home produced content.

Content sharing and exchange

In this scenario users search the Share it! network to find content of interest (for example, maybe they are interested in content about 'Ferrari cars'). The Share it! box user interface needs to facilitate easy entry of search criteria in a TV-centric 'lean-back' manner (most current PC-based peer-to-peer clients are somewhat 'lean-forward' in nature). The box will issue the query to the network, peer boxes with matching content will return relevant results. These results include rights data so that the querying peer can establish whether it has the rights to acquire and access the content, and if not how it might acquire those rights. These results are presented to the user who can select and acquire the content they desire, and if necessary, the rights. The content published and searched for could include both originally broadcast, and consumer-created content.

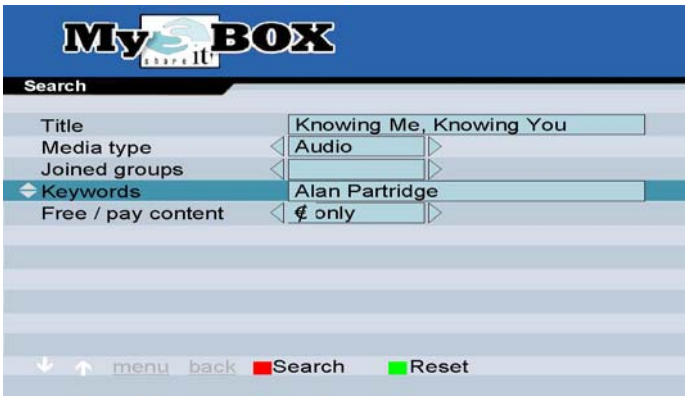

Figure 3 – Searching for content

Users are also able to publish content, which may be of interest to other users, to specific Share it! User groups if the content rights allow them to do so.

Network services

By placing a server supporting Share it! protocols in the Share it! network, service providers could offer a number of services:

- Premium content
- Managed user groups tied in to a particular content concept or theme (e.g. the official 'Friends' user group)
- Better quality of service (for example enabling real-time streaming of content, faster downloads, or enhanced / richer metadata).
- Rights brokering, enabling users to access a wider range of content, and content providers to find a wider audience.
- Sophisticated audience measurement functions, including tracking super-distributed content.

An important aim in providing these services is to encourage viewer loyalty and to extend their interest in the programming and related content.

Creating, editing and publishing by consumers

The Share it! platform enables users to easily distribute content that they have created. For example, users could:

- Send images or videos to a friend – images from a digital camera, or video from a camcorder can be easily loaded onto the Share it! box and sent to friends.
- Creating a complex content bundle – users can bundle related items of content and publish this bundle of content.

- Annotating content – users can add text or image annotations to an existing piece of A/V content (rights permitting) and make this available for friends or family.

Figure 4 – Sharing photographs

END-TO-END ARCHITECTURE

Having developed the system requirements we needed to design the end-to-end system architecture. In order to do this we started off by defining a layered model. We conducted a survey of existing candidate technologies, and mapped these to the previously identified layers. We then progressively detailed the choices, driven by need for early choices on issues that impact hardware, and identifying gaps in existing solutions where we needed to invent new technologies or customize existing technologies.

The resulting system architecture at the highest level is as follows:

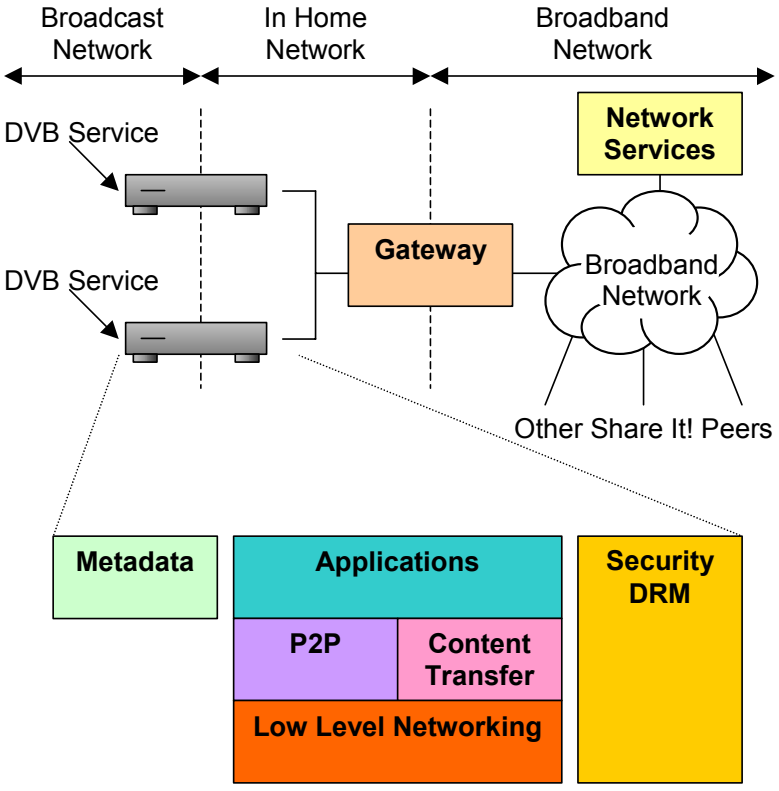

Figure 5 – System architecture

System concepts and identifiers

A central concept of many peer-to-peer networks is that of bringing together groups or communities of peers that are interested in similar things. Within Share it! we felt this concept was still valid, however we developed this into the idea of user groups. User groups have three attributes which define how users can interact with that user group. These

attributes are:

- Membership – restricted or unrestricted
- Publishing – moderated or unmoderated
- Visibility – hidden or visible

Various combinations of these attributes give different types of group, for example:

- The unrestricted unmoderated visible group, e.g. Ferrari interest group. This is a simple self-managing group of people interested in a given topic. All members are able to indicate content of interest to the group
- The unrestricted moderated visible group, e.g. BBC EastEnders group. This is a simple group whereby a single entity (e.g. the BBC) is able to indicate public content that is of interest to the group
- The restricted unmoderated hidden group, e.g. family group. The group itself is invisible to non-members. Any member can publish content to the group.

Within Share it! we need to support a range of content sharing models, for example broadcast originated content and home produced content. These have different visibility requirements, for example, home produced content may have some desired privacy, which is not required for broadcast content. To this end we have defined public content and private content – public content can be freely discovered by users, whilst private content cannot (note that even if you can discover public or private content, you may not have the rights to acquire, consume or share it).

Public content is available from any type of user group, whereas private content is only available from restricted membership user groups. Typically broadcast originated content would be public whilst home originated content could be private or public. In principle private content might be used as a way of illegally distributing content (much as today's unmanaged peer-to-peer networks are used), so an audit trail mechanism has been included for private content that indicates where it originated in the network.

Many identifiers are needed within the system to identify various entities e.g. users, devices, pieces of content etc. In particular the identification of content has proven to be a problem area. Within the demonstration system we have adopted the TV-Anytime Content Reference Identifier (CRID) [5], however this is to some degree stretching the CRID beyond its original intended application. The CRID has shortcomings when applied to peer-to-peer content; primarily as it has been designed as a managed entity for use in broadcast networks rather than as an unmanaged entity to be released into a distributed network.

Peer-to-peer protocols, discovery and search for content

Candidate peer-to-peer protocols identified in the survey phase of the project included JXTA [6, 7, 8, 9], UPnP, and Gnutella. Whilst UPnP is a popular protocol for in-home networks, its difficulties with scalability made it unsuitable for our home-to-home network architecture. The Gnutella protocol is basic compared to JXTA and so we essentially had the choice of extending Gnutella to include the functionality we needed, or using JXTA (a more complete but 'heavier' solution). We chose JXTA, which is an open-source peer-to-peer protocol initiated by Sun Microsystems.

JXTA provides a platform for distributed applications, enabling inter-connected peers to locate each other, participate in community-based activities and offer services to each other across peer-to-peer systems. JXTA is designed to be independent from programming languages, system and network platforms.

Peer groups are the core of JXTA's infrastructure. A peer group is essentially a sub-set of all the peers in the world. 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.

JXTA provides much of the required Share it! peer-to-peer functionality, with a fairly direct mapping between JXTA peer groups and Share it! user groups. We have developed a software layer on top of JXTA protocols and defined new JXTA peer and peer-group services for Share it!:

- Publishing, Searching, Recommendation and Announcement Services (built on top of the JXTA Resolver service using TVA compliant messaging formats where applicable)
- User Group Membership Services (both centralised and distributed)
- Rights Negotiation Service
- Chat Service with address book functionality
- Hidden User Group Recommendation Service

Share it! has defined the formats of the queries and responses using the TV Anytime specifications and wrapped them in JXTA formats so that we can use JXTA mechanisms to propagate them through the network.

As for every peer-to-peer network, performance scalability is a key issue, and when considering a potentially large deployed network of consumer peer-to-peer devices this needs to be considered. The JXTA architecture includes two solutions to address scalability; these include an overlay network of super-peers called rendezvous peers that provide an effective means of caching and routing system messages, as well as the previously mentioned grouping concept. The application of the JXTA peer group mechanism provides added value to the user results, but more importantly provides an effective mechanism for scoping queries to peers which are likely to have relevant content, although users need to know which peer group to search. Although potentially more effective in terms of guaranteeing scalability, search hubs offered by JXTA were not considered due to their centralised nature.

The full reference JXTA implementation does have a large footprint when considering resources in typical current generation set top boxes, and we will go on to look at how it could be profiled to better suit a current STB device within the project.

Content transfer

Of the several well-known IP content transfer protocols, we decided to use HTTP for content download and RTP/RTSP for real-time streaming. Given typical current in-home and broadband between-home connection bandwidths it is likely that streaming would only be feasible for in-home and network-server to home connections (as most broadband connections have lower rate uplink than downlink connection bandwidths).

Whilst in principle multiple video codecs can be supported within the Share it! system e.g. MPEG2, MPEG4 part 2, MPEG4 part 10, within the demonstrator we have focussed on MPEG2 video as this is currently the only codec supported by DVB. Clearly box-based transcoding to more efficient codecs and/or lower bit-rates potentially enables faster content transfers over rate-constrained links. This is enabled in our architecture but we have not made this a focus within our implementation.

We have defined a ‘content-bundle’ which allows multiple media elements to be distributed as a single file. Such elements might include an MHP Xlet, still images etc.

Metadata and metadata management

As far as possible we have adopted the TV-Anytime standards for metadata [10] within the Share it! project. TV Anytime metadata forms the basis of the JXTA content advertisements used to publish and search for content on the peer-to-peer network. We have also implemented TV Anytime SP006 compliant [11] metadata servers to enable Share it! boxes to query and retrieve metadata for broadcast schedule.

Rights management and security

Digital Rights Management (DRM) is a very complex area, as perhaps indicated by the number of standards groups tackling this topic! Our approach with Share it! has been to focus on appropriate models to enable home-to-home sharing, extending current broadcast rights models.

We have therefore defined two models suitable for different circumstances; known as the light touch and heavy touch approaches.

In the light touch model, rights negotiations and verification are performed directly between trusted peers. Typically the requesting peer already has the rights to consume the desired content, it has to satisfy the serving peer that it has appropriate rights before it can acquire the content.

In the heavy touch model, rights negotiations include a trusted third party in a rights-brokering role. Typically the requesting peer does not have rights to consume the desired content, and has to obtain these from the third party broker before then satisfying the serving peer that it has appropriate rights and then acquiring the content.

Within the system, we have defined both usage and sharing rules. Usage rules govern each peers ability to consume content, whilst sharing rule govern each peers ability to share content with other peers. The goal of the system is to be able to support suitable distribution models for the full range of possible content, from premium movies through to home produced content. The inclusion of the heavy touch approach enables complex cross territory business models to be addressed and provides one route for payments to be collected. The specific models to be applied to a given piece of content will be defined by the content rights owner. It might be that in the future that the rights managed peer to peer sharing of content is a new distribution path that content owners will wish to embrace.

Within the project demonstrator we are implementing the DRM protocols using simple XML and JXTA services however, in principle the DRM data could be encapsulated in various rights expression languages.

Although our focus has been more on DRM than on security, we are mindful of security and privacy issues (particularly content protection and robustness of the DRM solution to attack), and are conducting a threat analysis of the platform.

Within our system architecture we have presumed that each in-home network’s Internet connection is protected by a gateway, and as such our peer-to-peer and content transfer protocols must be able to cope with gateways and Network Address translation (NAT).

One of the key issues of building a rights managed secure platform is that of trust. Managing and maintain a network of trustworthy devices is a prerequisite to enabling rights managed content sharing. Within a typical vertical TV platform the platform operator and their security technology provider could clearly perform a trust management role. Within a horizontal

market, it is less clear who might perform this role, however there could be new opportunities for providing such services.

APIs

Within the Share it! box implementation, a number of Java API's have been defined to allow interactive application writers to take advantage of Share it! system functionality. These extend the MHP stack and include support for:

- TVAnytime location resolution
- Content access, control and acquisition
- Metadata queries (local database, peer-to-peer and TVA SP006 servers)
- Share it! user group access (note that rather than expose the full JXTA implementation, a simplified abstraction is available to applications)
- Address book and user to user messaging

SEARCHING AND ACQUIRING CONTENT AND RIGHTS – AN ILLUSTATIVE EXAMPLE

Having reviewed the architectural topics and issues, an illustrative example of searching and acquiring content and rights is given below. This shows the high level steps that are taken in order to search for and acquire content and rights from another peer. In this example, the user of Box A is searching for content, which is available on Box B.

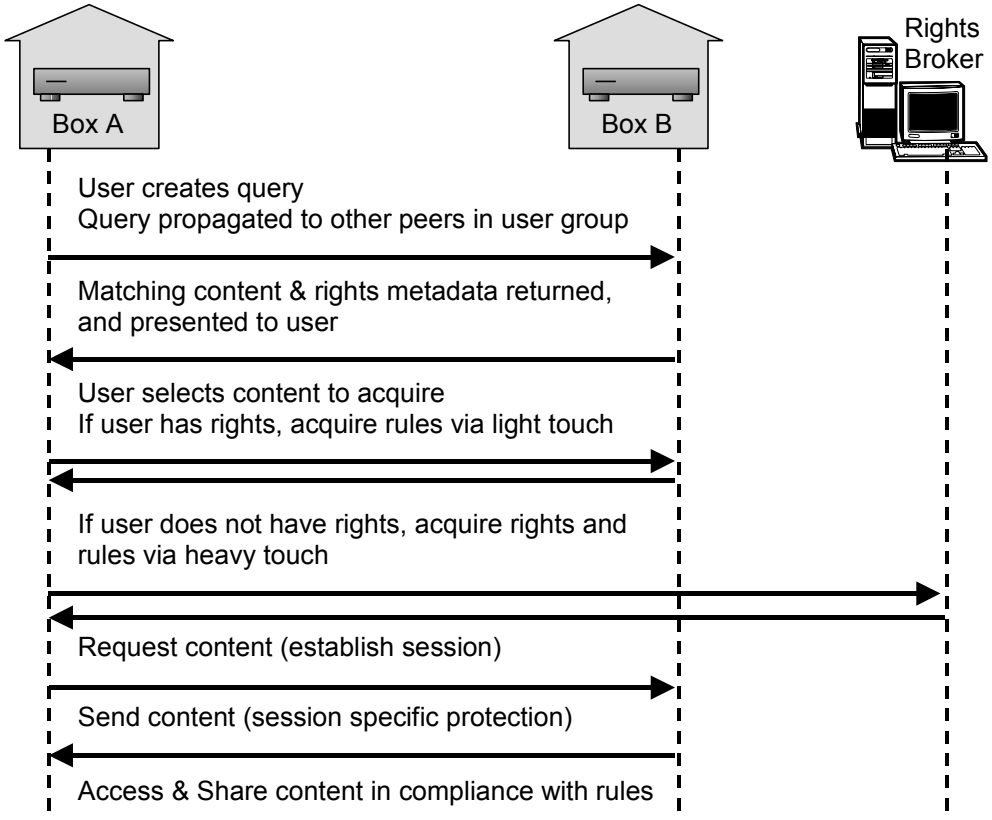

Figure 6 – Rights negotiation

CONCLUSIONS

The project has explored many issues involved in the design and deployment of a wide spread, consumer, TV-centric, peer-to-peer network of connected homes. Starting from a set of scenarios of potential use we have proceeded to refine the usage models, to design and adapt protocols for P2P content sharing, for DRM, and APIs to access this new functionality. We are developing a number of trial applications based on the usage scenarios to evaluate the value of the new service ideas and the underlying technology.

The Share it! system demonstrator is being shown on the IBC Booth 3.141 “IST-Village (ISTV)”. We invite anyone interested in finding out more about the Share it! project or the topics discussed in this paper to visit the booth.

REFERENCES

[1] The Share it! project web page, http://www.extra.research.philips.com/euprojects/share_it/

[2] The DVB project web page, http://www.dvb.org/

[3] The TV-Anytime forum web page, http://www.tv-anytime.org

[4] de Jong, Frans (editor), Deliverable 3: Description of example applications, Public deliverable from the Share It! project. IST-2000-28703.

[5] The TV-Anytime Forum: Content referencing, Specification Series: S-4, http://www.tv-anytime.org

[6] Project JXTA Technology: Creating Connected Communities, http://www.jxta.org/project/www/docs/JXTA-Exec-Brief-032803.pdf

[7] The JXTA project web page, http://www.jxta.org

[8] JXTA v2.0 Protocols Specification, http://spec.jxta.org/v1.0/docbook/JXTAProtocols.html Protocol

[9] Brendon Wilson: JXTA, http://www.brendonwilson.com/projects/jxta/pdf/JXTA.pdf

[10] The TV-Anytime Forum: Metadata Part A, Specification Series: S-3, http://www.tv-anytime.org

[11] The TV-Anytime Forum: Metadata Services over a Bi-directional Network, Specification Series: S-6, http://www.tv-anytime.org

ACKNOWLEDGEMENTS

The authors would like to thank the members of the Share it! consortium for their co-operation and their contributions to this work. This includes the teams at Philips Research Redhill and Eindhoven, BBC, NDS, NOB, TNO-T, University of Linköping, University of Ljubljana, Fraunhofer FOKUS Berlin, and Elisa Corporation. We would also like to thank our respective companies for permission to publish this paper.