DLNA Interworking for Virtualized Set-top-boxes

Bert Vankeirsbilck*, Chris Develder*, Tom Van Leeuwen$, Dieter Verslype*, Bart Dhoedt*

*Ghent University – IBBT
Gaston Crommenlaan 8 bus 201
B-9050 Ghent, Belgium
{Firstname.Lastname}@intec.UGent.be

$Alcatel - Lucent
Copernicuslaan 50
B-2018 Antwerpen, Belgium
Tom.Van_Leeuwen@alcatel-lucent.com

ABSTRACT

This paper proposes and evaluates a server centered approach to integrating personal content, residing in the user's home network, in the Electronic Program Guide (EPG) of digital broadcast television. With this approach, a server located outside the home of the user takes the integration work for its account, leaving the set-top-box to only handle the interaction with the user. We compare this approach to the situation in which the set-top-box executes the EPG. The benefits of the server centered approach are mainly simplification of the set-top-box leading to less maintenance, reduced energy consumption and lower cost while the functionality is boosted as it could connect in the same way to other services in-home and on the Internet. Another key benefit is that the central service can be easily upgraded without changing anything to the STB itself. We show that the overhead of sending the descriptions of the in-home content to the server is limited. The delay perceived by the user is limited to a typical value of 100 ms. The upstream bandwidth used for sending descriptions of in-home content to the integration server takes on average 100 Bytes per item, thanks to fairly simple compression techniques.

Categories and Subject Descriptors

H.3.4 [**Information Storage and Retrieval**]: Systems and Software – *distributed systems, information networks, performance evaluation (efficiency and effectiveness).*

General Terms

Measurement, Performance, Design, Experimentation.

Keywords

Personal content integration, thin client, Unified Programming Guide, DLNA interworking, Digital TV

1. INTRODUCTION

Digital Television is getting well integrated in the common households nowadays, and step by step, the appliances in the home environment start to communicate with each other. This evolution is fueled by the Digital Living Network Alliance (DLNA), having an explicit mission statement to facilitate *"a global collaboration of 245 of your most trusted brands, all working together to help you create the home entertainment environment you've always imagined"* [2]. In this paper the focus is on the opportunity to control media services available in the user's home from within the Digital Television Electronic Program Guide (EPG).

In this case, the Digital Television provider is an important stakeholder having control of key components in the system. The Digital Television providers usually supply the set-top-boxes (STBs) that are installed in the home environment. For obvious cost reasons (both production and maintenance related), these mass produced user devices should be simple, maintenance-free, cheap and robust, while keeping the functional possibilities as experienced by the end-user as broad as possible.

To this end, we investigate whether it might be feasible and beneficial to apply a server centered approach to control in-home media content from within the EPG. In this approach, the STB is kept as lightweight as possible: it only intercepts user actions (such as remote control button presses) and shows content on the television screen. The remaining functionality needed to integrate both broadcast content and personal, locally stored private content is shifted into the network, i.e. away from the STB. In this way, important software components that would normally execute on the STB are now running on a remote server. We can expect that, in order to optimize the STB lifetime and to keep the device as simple and robust as possible, the Digital Television Service Provider will indeed offer such a server to control this functionality located in its own controlled network rather than in the home environment.

The advantages of this approach are manifold. The STB will be very simple which means a decreased need for maintenance, reduced energy consumption, less hardware faults and lower cost of STB (concerning both hardware and service cost). Furthermore, since the functionality resides on a provider controlled server, the ease of upgrading and changing this functionality is increased because it does not involve the STB itself. There is also an opportunity to make the STB more generic so that it can connect to other services in both the home network and the Internet. For the user, this could mean that he has one single piece of equipment to control all his devices, present in-home and even beyond.

In this paper, we compare the proposed server centered approach to the other extreme where all functionality for merging broadcast and in-home video content resides at the STB.

This paper is structured as follows. In the next section, an architecture is proposed to integrate personal content into an EPG using a server centered approach. In section 3, the feasibility and characteristics of the solution are discussed using experimental results. Section 4 presents the conclusions drawn from the proposed and tested solution.

2. ARCHITECTURE

In general, a high level architecture for digital television is built using components located in the Home Network and components residing in the Access Network. The Home Network comprises all devices that can be interconnected in the user's home environment. These include personal devices such as multimedia players, telephones, personal computers, laptops, handhelds or next generation microwave ovens and dishwashers, as well as devices rented from an external party, such as a set-top box or a home gateway. The Access Network is the connection to services outside the home environment.

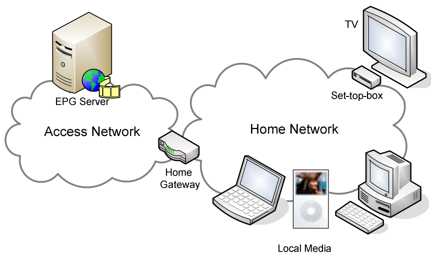

Figure 1: Architecture for thin set-top boxes

Figure 1 presents the architecture for the proposed server centered approach for integrating in-home media devices within a unified EPG. The local media reside on devices that are connected to the Home Network, such as a laptop, PC or a personal media player. The Home Gateway (HGW) is the link between the outside world and the devices in the Home Network. It is able to discover the in-home devices and communicating descriptions of their content to external services. In the Home Network, there is a STB that only takes care of interaction with the user. The unified EPG is constructed on a server in the Access Network, and is able to merge broadcast television content with the description of the user's personal content acquired through the Home Gateway. The concept of a unified EPG has been introduced in [6]. Remark that only descriptions of the in-home content are transmitted to the EPG server, and that when the content needs to be played, the STB acquires it from within the Home Network. The communication between the STB and the EPG server adheres to the principles of Thin Client Computing: the user interaction is sent over the network to the EPG server, and the graphical updates of the EPG are transmitted the other way around.

3. EXPERIMENTAL RESULTS

We have deployed a realistic Digital TV testbed according to the architecture presented in the previous section. On the Home Gateway, we have implemented support for UPnP AV [8] allowing dynamic discovery and use of multimedia devices in the Home Network, allowing the use of Windows Media Player on a laptop as an in-home media server. The EPG Server and Home Gateway communicate through webservices. As thin client protocol, we adapted a VNC server and viewer pair [7], in order to merge client side rendered video with server rendered content. This adaptation required changes to both client and server code. The VNC viewer utilizes a recent FFmpeg library for decoding video [3]. We have implemented a custom EPG in Java in conformity with XML-TV [9], a standard format for describing programming guides. In order to simulate the uplink bandwidth in the Access Network, we used a Click Router [5] that applied simple bandwidth shaping. The details of the experimental testbed are summarized in Table 1.

Table 1: Testbed hardware and software details

Set Top Box	AMD Athlon 64 2000 MHz, 2GB RAM DDR1 Ubuntu 9.10 Adapted tightVNC v1.3.9 FFmpeg v. SVN-r20647 (2009-11-06)
EPG Server	AMD Athlon 64 2100 MHz, 2.48 GB RAM Ubuntu 9.10 Adapted tight VNC v1.3.9 Custom EPG using Java and XML-TV
Media Server	Intel Core Duo P8400 2.26 GHz, 3.45 GB RAM Windows Media Player 11 (UPnP AV compliant)
Home Gateway	Athlon 64 2100 MHz, 1GB RAM UPnP AV compliant WebService interface to EPG Server
Click Router	AMD Athlon 64 2000Mhz, 960 MB RAM Click Modular Router
Network cards	All networks cards used: 100Mbps

Figure 2: Screenshot of the custom made unified EPG

Figure 2 shows how our unified EPG integrates broadcasted video and personal content. The top right video is the broadcast channel or personal content the user is currently watching. Meanwhile, the bottom menu can contain the controls to change broadcast channel or to browse the in-home devices and their content. The figure shows the latter option. In the left banner, the thumbnail picture of the selected personal video is shown.

Through an experimental evaluation, we will compare our server centered approach to a set-top-box centered approach. For this comparison, important metrics are discussed in the following subsections. First we tackle the uplink bandwidth that is consumed for transferring the description of the in-home media, to integrate its content into the EPG. Then, we evaluate the delay a user can expect between invoking a browse operation through the remote control and having the results visible on his screen.

3.1 Size of personal content description

Regarding the bandwidth used to transfer a description of the personal audio/video content, we have investigated the possibilities and effect of both lossy and lossless compression of the DIDL_Lite standard. This standard is derived from a subset of DIDL, the Digital Item Description Language, recently developed within ISO/MPEG21 [4, 10]. Concerning the lossless compression we applied the GZIP algorithm from the java.util.zip library. For lossy compression, the Home Gateway parses the DIDL_Lite formatted description retrieved from the in-home Media Server, and strips it down to only retain the description entries that are supported by the EPG. Although the lossy compression approach might seem to be aimed to restrict the functionality, this mechanism could bring an important advantage of increasing privacy about the personal content. The attributes that were kept for the description are: the item title, media type (audio, video, container), resources in the form of URLs and thumbnails in the form of UPnP:albumArt. Measurements with the combination of both kinds of compression were also conducted.

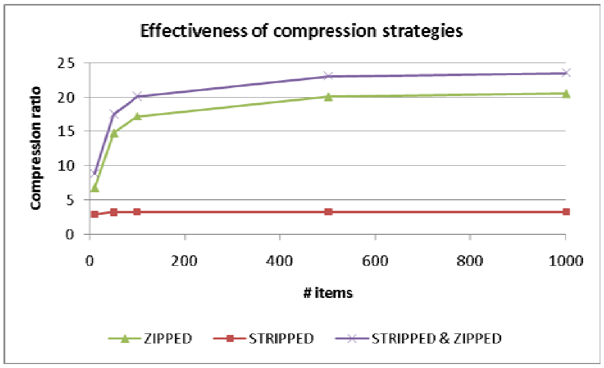

Figure 3: Effectiveness of compression strategies

Figure 3 presents the effectiveness of the proposed compression strategies. Applying lossy stripping of unneeded description attributes reduces the transferred data size with a constant factor of 3.25. The lossless GZIP algorithm is known to be efficient for textual data, and shows to compress better for larger data sizes according to a logarithmic function. In our tests, the GZIP compressor was able to take profit of the repetitive structure of the data and reached a compression factor of 20.5 for 1000 item descriptions. Combining both the lossy and lossless compression techniques, the total size reduction factor further increased to 23.5.

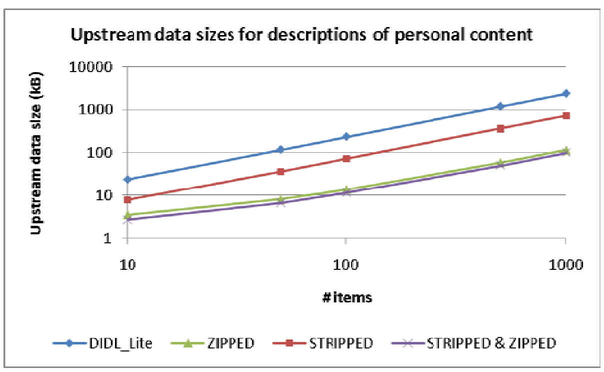

Figure 4: Upstream bandwidth for browsing personal content.

Figure 4 proves the need for the compression strategies. Using the full DIDL_Lite description would be undesirable from an uplink bandwidth perspective. The required bandwidth increases linearly with the number of items of which the description needs to be transferred. The figure shows that the description of 1000 items takes 2.4 MB to be transferred from the Home Gateway to the remote server. Stripping the description attributes compacts the size less than zipping the contents. Combining stripping with zipping slightly increases the compression ratio of zipping the original DIDL_Lite description.

Eventually, transmission of the descriptions of 1000 items from the Home Gateway to the server that builds the EPG requires only 100 kB. This proves that for the user, uplink bandwidth limits will not be an issue for applying a server centered approach. From the service provider perspective, the incoming link capacity will not be an issue either.

3.2 Delay

From the user perspective, the responsiveness of the EPG is of key importance. We express responsiveness in function of the delay between the registration of a user action on the STB and the presentation of the corresponding graphical reaction. When comparing the server centered approach with the STB centered approach, the uplink bandwidth in the Access Network will have a considerable influence on the extra delay to present the content available in the Home Network, and there will also be some overhead delay to transfer the user interaction between the STB and the remote server. Both kinds of delays are discussed

3.2.1 Transfer of in-home content descriptions

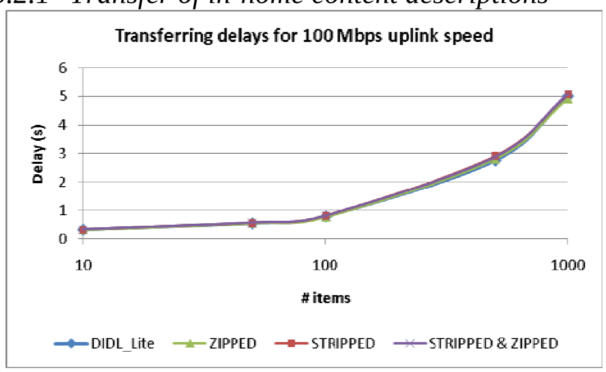

Figure 5: STB centered approach, compression vs. delay

Figure 5 presents the effect of the compression methods for an in-home situation, where a throughput value of 100 Mbps can be assumed. Thus the figure can be interpreted as the case that the EPG is constructed on the STB. It is shown that in this case, the total delay is increasing linearly with the number of items, and that for the compression methods the reduction in networking delay approximately equals the processing delay to perform the compression. Even in this in-home setting, getting the descriptions of 1000 items takes 5 seconds. However, optimizations are possible, such as incrementally loading the descriptions and preloading items of subdirectories.

The explanation above assumes that the EPG that is executed on the STB uses the HGW as an intermediary component to handle the discovery of the in-home devices and the listing of their contents. If the set-top-box would take over this functionality, and no compression is applied, our tests indicated that the time to acquire the descriptions of 1000 items would be 4.67 seconds.

Figure 6 shows the influence of the proposed compression methods on the delay in the case that the descriptions are transferred to the EPG server over a slow uplink network. Included in these delay measurements are the time it takes the HGW to acquire the data over UPnP from the media device in-home, the delay caused by stripping and/or compressing the data on the HGW, the delay for transporting the data in the Access Network, the decompression delay and the parsing time of the description at the EPG server. This figure shows that transmission of the uncompressed data would be undesirable from

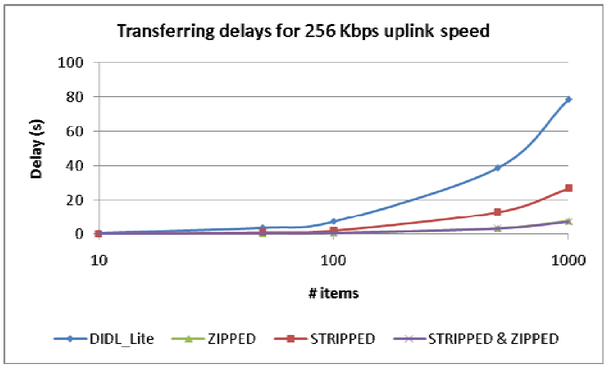

Figure 6: Server centered approach, compression delay

the delay perspective. For 1000 items the delay comes to an unacceptable 78 seconds. Compressing them with the combined method brings this number down to 7.3 seconds. In this case, incrementally loading the descriptions and preloading items of subdirectories can lower this value.

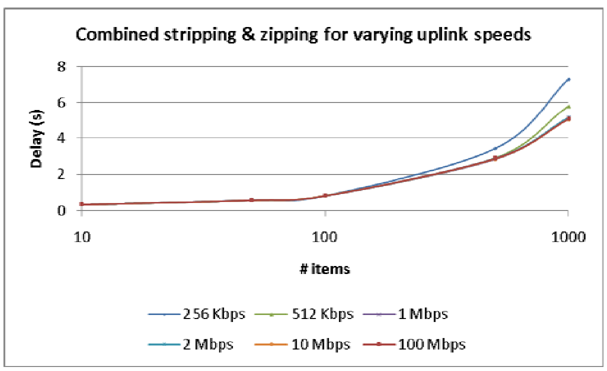

Figure 7: Effect of available uplink bandwidth

As explained before, the data is compressed which will cause a decrease of network delay, but also an increase of the processing delay for performing the compression and decompression. Figure 7 compares the measured total delay for various uplink speeds. It is shown that the compression methods overcome the effect of the slower uplink speeds as compared to the STB centered approach, if the uplink speed is higher than 1 Mbps. Thus, the time to load the personal content descriptions into the EPG server takes only 350 ms longer than when the EPG would be composed on the STB. Also for 512 Kbps, the difference is only noticable when more than 500 item descriptions are to be transferred. For an upstream transmission speed of 256 Kbps, the influence of the network is more apparent and can't be solved by the proposed compression methods. However, current commonly available providers supply 256 Kbps uplink bandwidth for *light* accounts only.

3.2.2 Thin client protocol delay

In order to compare the server centered approach with the set-top-box centered approach, we need to take into account the extra delay that is caused by applying the thin client computing paradigm. In [1], the authors empirically found that the delay of thin client protocols depends on the content to be sent, and state that in reasonable downlink bandwidth networks (1 Mbps) the order of magnitude is 100 ms or less. This is a constant extra delay, independent of the number of items to be displayed.

4. CONCLUSIONS

In this paper we have shown that a server centered approach for merging broadcast and personal content is feasible and beneficial. Experimental evaluation has shown that the extra delay experienced by the user will be limited to 450 ms, and it has been argued that possibilities exist to lower this value to only the delay overhead of the thin client protocol, being about 100 ms. Regarding the bandwidth needed to transfer description of personal media items, the proposed compression mechanisms perform very well so that uplink bandwidth limits are not an issue. Thus, the drawbacks are actually minor, but the benefits of the simple and lightweight set-top-box approach are manifold for both service provider and end user such as decreased cost, energy consumption, and maintenance of the set-top-box, access to many services both in-home and in the Internet with a single device, and easy upgrading and changing of the centralized unified EPG service.

5. ACKNOWLEDGMENTS

This work has been partly funded by the Institute for Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen), through the IWT Vampire project. Bert Vankeirsbilck is funded by a Ph.D grant of the IWT-Vlaanderen. Chris Develder is supported as a post-doctoral fellow of the Research Foundation – Flanders (FWO).

6. REFERENCES

[1] A. Lai, J. Nieh. On the Performance of Thin Client Computing, *ACM Transactions on Computer Systems*, 24(2), pp. 175-209, 2006.

[2] "DLNA for HD Video Streaming in Home Networking Environments", Whitepaper, http://www.dlna.org/about_us/about/DLNA_Whitepaper.pdf

[3] FFmpeg: a complete, cross-platform solution to record, convert and stream audio and video. http://ffmpeg.org/

[4] IEC/IEC CD 21000-2:2001, Information Technology – Multimedia Framework – Part 2: Digital Item Declaration, July 2001.

[5] Morris, R., Kohler, E., Jannotti, J., and Kaashoek, M. F. 1999. The Click modular router. *SIGOPS Oper. Syst. Rev.* 33(5), pp. 217-231, 1999.

[6] Obrist, M., Moser, C., Alliez, D., Holocher, T., and Tscheligi, M. 2009. Connecting TV & PC: an in-situ field evaluation of a unified electronic program guide concept. *In Proc. of the 7th Conference on European interactive Television*, Leuven, Belgium, pp. 91-100, June 2009.

[7] Richardson, Q. Stafford-Fraser, K. R. Wood, A. Hopper. Virtual Network Computing, *IEEE Internet Computing, IEEE Computer Society*, Volume 02, pp. 33-38, 1998.

[8] UPnP AV. http://www.upnp.org/default.asp

[9] XML-TV File format (Dec 2008) http://wiki.xmltv.org/index.php/XMLTVFormat

[10] XML Schema for ContentDirectory:3 Structure and Metadata (DIDL_Lite), UPnP Forum, September 20, 2008. Latest version: http://www.upnp.org/schemas/av/didl-lite-v2.xsd