Workshop Proceedings of the 9th International Conference on Intelligent Environments
J.A. Botía and D. Charitos (Eds.)

doi:10.3233/978-1-61499-286-8-246

A Content-Aware Cloud Platform for Virtual Reality Web Advertising

Markos ZAMPOGLOU [a,1], Athanasios G. MALAMOS [a], Emmanuel SARDIS [b], Anastasios DOULAMIS [c], Kostas KAPETANAKIS [a], Konstantinos KONTAKIS [a], Vrettos MOULOS [b], and George VAFIADIS [b],

[a] *Multimedia Content Laboratory, School of Applied Informatics and Multimedia, TEI of Crete, Heraklion, Greece, www: http://www.medialab.teicrete.gr*
[b] *Institute of Communication and Computer Systems, ICCS, 9, Iroon Polytechniou Str. Athens, Greece www: http://www.iccs.gr/eng/*
[c] *Dept. of Production Engineering and Management, Technical University of Crete, Chania, Greece*

Abstract. In this paper we present our work towards a remote, cloud-based, large-scale platform aiming at offering Virtual Reality advertisements to end-users on a wide range of terminal devices. While the remote and distributed nature of the platform allows us to display complex content on lightweight devices, the system architecture can also adapt to high-quality content for high-end displays. More importantly, we implement an extended version of the MPEG-7 standard for VR scenes, which allows us to manage the 3D worlds stored in the cloud in a content-aware manner. The system also integrates the dominant standards in the fields of Web 3D and multimedia broadcasting, in order to achieve maximum extensibility and interoperability, towards a platform that will be able to offer lightweight yet immersive VR marketing experiences within Web-based environments.

Keywords. VR advertising, cloud-based 3D repositories, remote rendering, iPromotion, Web 3D

Introduction

While web advertising has been an important commercial activity for more than a decade, it fundamentally still relies predominantly on the traditional mediums of text, image and sound data -however, during the same time period, VR technologies have shown impressive progress. While the potential of VR marketing has been evaluated in the past, with research demonstrating its potential benefits in terms of informativeness and user enjoyment [1,2], the concept has so far remained theoretical. We have undertaken the effort of building a platform, which we have named iPromotion, over which VR advertisements will be distributed to users over the web, allowing them to interact with virtual representations of the advertised goods, thus giving them a fuller, more informative and more pleasurable experience.

[1] Corresponding Author: mzampog@epp.teicrete.gr

1. Project aims and description

As a goal, the project has set two distinct usage scenarios for our platform: the first, which we call the ”Web & Mobile - Based Scenario” allows a user to retrieve and interact with a complex 3D scene using a mobile phone, while the second, called ”Large-Scale Demo,” concerns the remote reproduction of a high-quality 3D scene on a touch table or wall projector. The fact that these two tasks appear on the surface to have radically different specifications is misleading: an integrated framework can implement both, by adapting the volume of data sent and the computational load assigned to the client on each occasion. We are in the process of building this framework, and have it comply with the dominant standards in the industry, to increase extensibility and interoperability.

Currently, there is a growing trend of integrating 3D graphics in web pages, under the HTML5 standard. The graphics are taken up and rendered by the browser and the device’s graphics card. However, local processing of the 3D data is problematic in a range of scenarios, where bandwidth and computational resources are limited. This can be the case, for example, in our ”Web & Mobile - Based Scenario”.

To tackle this, we have taken a remote and distributed approach: a cloud storage repository, a remote rendering grid which provides the end devices with only the renderer output, and a number of subsystems for organizing and coordinating the flow of information, to achieve the desired adaptability. Furthermore, the platform is supplemented by a 3D scene authoring tool and an application for automatically extracting MPEG-7 descriptions from 3D scenes, to allow for intelligent content-based search over the 3D scene repository. Finally, the rights of authors over the scenes they have created are maintained and propagated through the entire life cycle of a scene, using the MPEG-21 framework.

The rest of the paper is organized as follows: Section 2 presents the platform architecture. Section 3 presents the more innovative aspects of the platform in greater detail, and finally Section 4 deals with our further steps for the completion of the platform.

2. System architecture

The system (Figure 1) is a service-oriented platform integrating multiple subunits. The sub-units include, besides the core cloud-based STaaS (STorage as a Service) repositories and the grid rendering servers, an authoring tool and a subunit for the automatic extraction of MPEG-7 descriptors and MPEG-21 encapsulation, plus the subunits that manage the content adaptation to accommodate the entire range of possible scenarios.

The STaaS Model Repositories hold all available 3D models and scenes, stored in X3D format, as well as any meta-information available on them, plus any supplementary texture files for the 3D scenes. X3D is an open, XML-based language for 3D graphics, specifically designed for display over the web. We consider X3D to be a good option for its syntactic simplicity and openness, combined with the fact that it’s based on XML and is thus easier to integrate with the MPEG-7/21 standards.

The meta-information kept in the cloud repository is not limited to text describing the scene in high-level terms, but also includes low-level content descriptors and semantic metadata, and is stored in MPEG-7 format. Furthermore, the X3D files and their corresponding MPEG-7 descriptors are integrated within the MPEG-21 framework, which also allows the integration of digital authorship and usage rights for each scene. The use

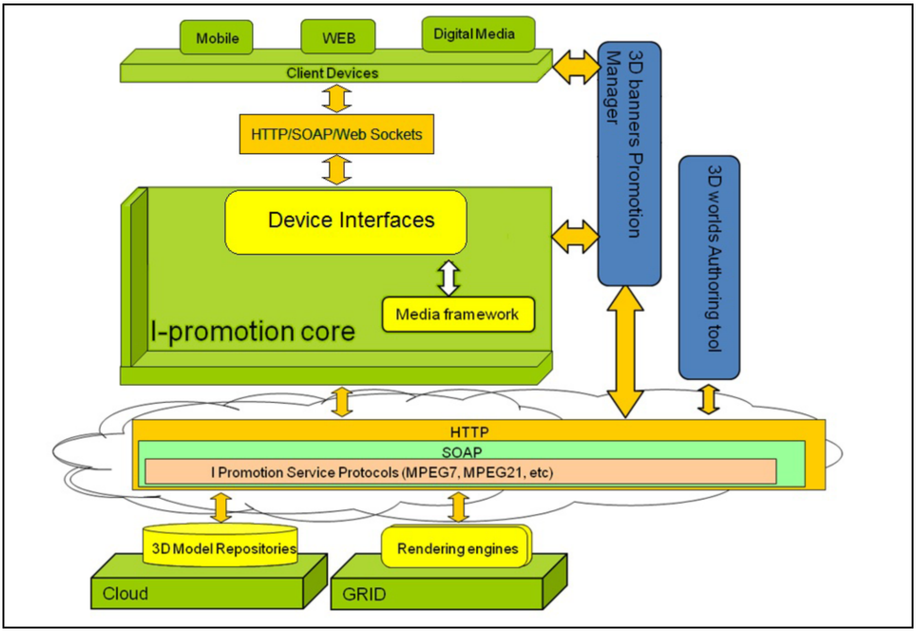

Figure 1. The iPromotion platform architecture

of content description allows for advanced search capabilities, both from visitors, who can chose -or be assigned- a 3D world according to high-level search criteria, and authors using the 3D Worlds Authoring Tool. The latter option is provided so as to encourage content sharing and re-use on the part of the authors -provided the digital rights status of the content allows it.

When a user requests a scene with particular high-level characteristics, the repository is traversed and a search is performed on the MPEG-7 descriptors of the various scenes stored therein. The relevant scenes are selected, and the corresponding MPEG-21 files are read. The corresponding X3D files are then extracted from the repository, along with their supplementary files (e.g. texture images), and forwarded to the Grid Rendering Engine, which renders the 3D model and outputs a sequence of images. Before forwarding the image stream to the end user, the system needs to identify the current user specifications and adapt the output parameters (such as format, resolution and frame rate) to the capabilities of the end device and the connection. The Device Interfaces subunit uses the MPEG-21 Multimedia Adaptation framework for this and then gives the relevant instructions to the Media Framework subsystem, which takes care of any necessary conversions. The User Interface for all use cases always takes the form of an HTML5 web page.

Two separate tools we have developed are the 3D Worlds Authoring Tool and the X3D to MPEG-7 Tool. The former aims at introducing new 3D scenes into the repository, by allowing authors to edit and combine already existing scenes. The Tool is a standalone application, which connects with the repository and offers semantic search capabilities based on MPEG-7 descriptions. When a user locates the models they are interested in, they can import them to the Tool, perform minor modifications (such as altering their spatial or parameters animation parameters) and store them back into the repository cloud as a new 3D scene. Furthermore, the Tool allows authors to add textual meta-tags to a scene which are consecutively incorporated into its MPEG-7 description. The X3D to MPEG-7 Tool draws from our previous work on MPEG-7 extensions for the description of X3D scenes [3], which we have further streamlined by developing an XSLT-based service for the automatic extraction of descriptors from the files placed in the cloud repository [4]. Further details are offered in subsection 3.3.

3. Novel contributions

3.1. Remote rendering

To achieve independence from the computational limitations of the various end devices, we use a grid rendering system located in our remote servers. Deferring a portion of the computational cost or 3D rendering to a remote server is a well-studied practice [5,6]. For our needs, we have modified the open Xj3D Browser [7] by converting it into an Axis2 Java Web service running over a server and communicating via SOAP messages. When the system wants to provide a user with a scene, the client page sends a message that initiates a session. The renderer loads the scene from the cloud, and begins transmitting (through the appropriate subsystems) the rendered output images. The available options that the User Interface provides include, besides scene load/unload commands, navigation and interactions within the scene: as the user interacts with the client-side HTML interface, appropriate SOAP messages are sent to the platform, and instruct the renderer to adjust the scene accordingly.

3.2. Cloud Storage

One of the main objectives of the proposed cloud storage system is to be able to handle any kind of multimedia files or metadata services that could be sent from the multimedia provider. With this in mind, our data model schema is not strictly structured as a standard RDBMs schema, but can instead include any data/metadata that are required for the application functionality. The file structure hierarchy is described as a JSON file that includes all the necessary files into field brackets, where common fields in a collection documents may hold different types of data.

The objective of this architecture is twofold: (a) define and enable an interface of several heterogeneous information and service entities, (b) create an active pool of technologies, Operating Systems, devices and policies where the cloud storage system could work. To fulfill the requirements, an interoperable structure framework was created, inspired by the NoSQL notion and the data-intensive constrains. The storage system is based on HDFS (The blue box that can be seen in Figure 2) and uses the OPTIMIS data management service [8] to achieve better federation results. A final modification we have performed on the Hadoop framework is the installation of Corona [9], an alternative job-scheduling system that is more efficient for small tasks, but also tasks that are not Map-reduce at all.

The next two layers are responsible for the interoperability, simplicity and compatibility of the framework with a plethora of document multimedia types and applications.

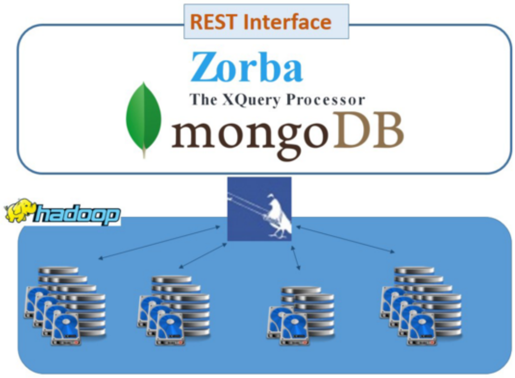

Figure 2. The file storage architecture

The white box in Figure 2 presents the technologies that are used in order to overcome these issues. MongoDB was the best candidate for the data storage as it gives us the flexibility to store data without worrying about following the exact schema. The JSON file is saved as it is in the database and is treated as object. Over that layer we use the Zorba framework [10] as most of the multimedia applications communicating with the database are based on xquery. Zorba undertakes that transformation and provides a unified layer for the framework. Finally the interconnection with the applications and the iPromotion service protocol is provided by a REST service (using the HTTP protocol).

3.3. MPEG-7 extension for 3D scene annotation

For our system, we needed to offer efficient, intelligent and content-aware search capabilities for the 3D scenes located in the cloud repository. In the current state-of-the-art (a survey can be found in [3]), the absence of a unified framework that can describe all important aspects of a 3D scene is striking. To cover our needs, we extended the MPEG-7 standard with a number of novel descriptors for the geometry, lighting, animation, structure, viewpoints and potential interactivity of an entire 3D scene, while also modifying and adapting the original MPEG-7 color, texture and motion descriptors. A full list of the descriptors, their format, and their scope is presented in [3].

However, for a large scale system to work, it is necessary that the descriptors are automatically extracted from the scenes. One of the reasons we have chosen X3D for our scenes is that, as both X3D and MPEG-7 are XML-based standards, we can automatically extract MPEG-7 descriptions from X3D files, using eXtensible Stylesheet Language Transformations (XSLT) between the two XML schemas. While normally an XSL Transformation aims at replacing a source XML tree with the XSLT output tree, what we need to do instead is form a secondary, MPEG-7 XML file to accompany the XML file that contains the original X3D scene. We have thus developed an X3D to MPEG-7 Tool that, given an X3D file, outputs the corresponding MPEG-7 description, validates it (according to our extended standard) and places it in the Repositories alongside the original X3D file [4].

4. Conclusions and future work

We have presented our work towards a large-scale, cloud-based VR advertising platform over the Web. The system is still under development, and, while many aspects (such as the remote rendering, mobile client, cloud repository and X3D to MPEG-7 tool are already functional, we are still progressing towards the final system integration. While certain aspects we have described (such as MPEG-21 compliance or the authoring tool) remain to be completed, the core of the system as presented here offers a powerful market tool for remote, distributed and flexible advertising, while offering a technology with vastly broader implications, in remote VR experiences.

Acknowledgements

The research of this paper is granted by the European Union and the Hellenic General Secretary of Research and Technology under the "COOPERATION 2009 / 09SYN-72-956" Framework.

References

[1] D. Mazursky and G. Vinitzky, "Modifying consumer search processes in enhanced on-line interfaces," *Journal of Business Research*, vol. 58, no. 10, pp. 1299 – 1309, 2005.

[2] K.-S. Suh and Y. E. Lee, "The effects of virtual reality on consumer learning: an empirical investigation," *MIS Quarterly*, vol. 29, no. 4, pp. 673–697, 2005.

[3] P. Spala, A. G. Malamos, A. D. Doulamis, and G. Mamakis, "Extending MPEG-7 for efficient annotation of complex web 3D scenes," *Multimedia Tools Appl*, vol. 59, no. 2, pp. 463–504, 2012.

[4] M. Zampoglou, P. Spala, K. Kontakis, A. G. Malamos, and J. A. Ware, "Direct mapping of x3d scenes to mpeg-7 descriptions," in *Proceeding of the 18th International Conference on 3D Web Technology June 20-22, 2013 San Sebastian, Spain*, accepted for publication.

[5] F. Lamberti and A. Sanna, "A streaming-based solution for remote visualization of 3D graphics on mobile devices," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 2, pp. 247–260, 2007.

[6] T. K. Çapin, K. Pulli, and T. Akenine-Möller, "The state of the art in mobile graphics research," *IEEE Computer Graphics and Applications*, vol. 28, no. 4, pp. 74–84, 2008.

[7] The Xj3D Project website, "http://www.xj3d.org/." last accessed 23/2/2013.

[8] G. Kousiouris, G. Vafiadis, and T. A. Varvarigou, "A front-end, hadoop-based data management service for efficient federated clouds," in *Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on*, pp. 511–516, 2011.

[9] The Corona Project website, "https://github.com/facebook/hadoop-20/tree/master/src/contrib/corona." last accessed 28/2/2013.

[10] The Zorba Project website, "http://www.zorba-xquery.com/." last accessed 28/2/2013.