Standards in Virtual Worlds
Virtual Travel Use Case Metaverse1 Project

José Manuel Cabello, Innovalia Association

José María Franco, CBT Communication & Engineering

Antonio Collado, Avantalia Solutions

Jordi Janer, Universitat Pompeu Fabra

Samuel Cruz-Lara, University of Nancy 2 (University Institute of Technology, Computer Science Department)

David Oyarzun, Vicomtech

Albert Armisen, Information & Image Management System

Roland Geraerts, Utrecht University, department of Information and Computing Sciences

Abstract:

Nowadays, tourism has become a very important industry in the international economy. Information and communication technologies are in constant development; they progress worldwide and across sectors. Their applications in tourism and tourist resources is rapidly increasing, reaching new, innovative and sometimes amazing results in terms of effectiveness, productivity, quality, and customer satisfaction. Exploring the interaction between technologies and tourism is difficult and challenging. Specifically, using virtual world technologies as a new means of information for potential tourists is a big challenge where the actual methods, goals and needs still need to be exactly identified.

Within the Metaverse1 project, Innovalia Spain (in association with other European partners) has created a "technologies laboratory" for the tourism sector within a virtual world; this laboratory is called the Virtual Travel Use Case. A new island has been implemented in Second Life, as the playfield of the Virtual Travel Use Case. This island represents the hot spots of Gran Canaria (one of the seven islands making up the archipelago of Canary Islands, Spain). This is the starting point for analysing, testing and validating developments of tourism in virtual worlds, and the application part of the related research activity.

The main technologies implemented in the Virtual Travel Use Case are listed below. They aim at making the tourists' experience in a virtual world even more interesting.

UNIVERSITAT POMPEU FABRA has developed a platform for authoring and

rendering Soundscapes (sound ambiances) in virtual worlds.

INRIA Talaris has developed multilingual tools for tourists within virtual worlds, the MLIF: The Multi Lingual Information Framework for virtual worlds.

UTRECHT UNIVERSITY has researched and developed the Path and Camara Planning tools in virtual worlds.

In addition to these four technologies, we contributed to the MPEG-V Standard. In this respect, the work of VICOMTECH focused on defining the attributes that can identify each avatar as unique, and afterwards, a high-level XML-based specification of these attributes has been created. The specification has been called ADML (Avatar Definition Markup Language) and it has been included into the new MPEG-V standard for interoperability between virtual worlds.

IMS has designed and implemented the Video streaming model within virtual worlds.

This paper aims at analyzing why and how virtual worlds can become an important platform for tourism-oriented areas to promote a destination in general, and their local heritage and tourist added-value services in particular. The document will also introduce the design of the first prototypes and the validation results of the four specific technologies tested at the Virtual Travel Use Case (Soundscape generation, Multilinguality, Video streaming and Path and Camera Planning). Finally, the contribution to the MPEG-V standard will also be detailed in the paper.

1. Introduction

There are some factors that make tourism an attractive sector to develop different kinds of technologies, such as information and virtual technologies. On the one hand, tourism promotes and commercializes activities offered far from the location of the client. On the other hand, tourism needs specific technologies such as promotional tools, which have to look interesting and attractive for the potential tourist [Buhalis, 1998].

Nowadays, virtual technologies offer a great deal of features, especially for business models. While we assume that the future of navigation will be in a 3D space, its web interface will be 2D. Furthermore, the virtual worlds developed within an open source policy will offer new kinds of possibilities, through interaction, communication, training, etc. But above all, these virtual worlds will offer opportunities to develop creative skills (http://www.hispagrid.com/).

The use of virtual technologies is becoming increasingly important, especially in the tourism sector. They offer multiple possibilities, not only for the potential tourist, but also for the destinations which implement this kind of technologies. While visiting some important virtual platforms such as Second Life, we can notice that important tourist destinations are creating their own virtual environment. Thus, the virtual tourist can access all kinds of information, such as pictures, videos, general and specific tourist information in a virtual environment. They can even walk around in different parts of these places. Furthermore, the visitor can contact all kinds of tourist businesses or the tourist information office to get help if he or she is interested in organizing a non-virtual trip to the destination. In fact, the main goal is to stimulate the traveller to make a real visit. It is thus obvious that virtual technologies are going to be a strong marketing tool. This new initiative shows a clear change in the promotion of tourist destinations.

We should note that new virtual platforms are created every day, with a specific goal or purpose. The tourism sector is commonly playing an important role in these virtual platforms. Thus, virtual environments offer different possibilities to the user. That is to say, the new virtual platforms are working on common areas, focusing on the enhancement of collaborative environments and training, realistic simulations, and the implementation of Web 2.0 technologies with social network tools (like wikis or blogs with immersive 3D tools).

2. Challenge and objective of the use case

In today's modern society, an immense amount of travel and tourism information is given using a wide variety of formats, such as web-sites, pamphlets or commercial advertisements. All this information, routed through ICTs, can also be centralized, filtered and made accessible within a virtual world.

Tourist destinations are experiencing increasing competitiveness when they market their offer and strive to successfully attract travellers. This increased competiveness, combined with one of the most important needs of travellers when choosing their next destinations as well as with the ability to access information on-demand, has spurred interest in the use of virtual technologies as a new marketing tool. It is suggested that a well-designed virtual world, which is able to satisfy this on-demand need, should be easily accessible, navigable, and well-structured. It should also have the capacity to spur travel and re-visitation to the physical destination.

The first goal of the Use Case is to develop a wide-ranging virtual laboratory for testing and validating different technologies. As a test case, it acts as a virtual platform to provide tourism information to virtual travellers as potential, real life visitors to Gran Canaria.

Figure 1: Metaverse1 project - Use Case 1: Virtual Travel

The conceptual basis for this virtual world is to test innovative technologies linked to the delivery of tourist information that is generally found at the physical destination of the Tourist Information Centre. Obviously, the Use Case adapted all the information to the new virtual platform. Examples of the information included in the Use Case are:

- General information – island features, geography, weather patterns, population and culture.

- Historical/specific information - museums and historical events/locations.

A virtual travel takes place in a technological environment where multimedia elements play a major role in the way information is accessed and disseminated. Careful consideration must be given to the ways this static information is best provided through various multimedia file formats.

The second goal of this Use Case is to tackle the growing division between the older and more tech-savvy younger generation of travellers with respect to the technologies they use as information resources. We accomplish this goal by keeping in mind that virtual travellers themselves must have attained a modest level of technical skills to enter and navigate the virtual world from the onset.

a. Why Gran Canaria?

Gran Canaria Island belongs to the Canary Islands in Spain, it is the largest island together with Tenerife, and its capital city has been designated as the place with the best climate in the world. Apart from the so-called "tourism of sun and beach", the cultural offer and events have led the city to be an attractive destination to visit, with theatres, opera, music and water sports.

Gran Canaria has been the location selected to validate the Use Case, as it is one of the most important tourist destinations of Spain, receiving more than 12million tourists every year. To get a good representation of the island and have a consistent version for the virtual laboratory, we decided to divide the island in five virtual tourist zones, as a representation of the sightseeing place of Gran Canaria.

Let us have a closer look at the Use Case in Second Life. Within the virtual island there is the Central Information Point, recreating a tourist information board, where the user will be able to access real information about the island and interact with some of the technologies implemented, such as video streaming. The famous Santa Ana Square and the Vegueta Cathedral, the Maspalomas Beach and the Nestor Museum can also be found there. These places represent the main virtual spaces of Gran Canaria Island in the virtual platform.

The pursued objective is to obtain an interesting virtual tourist scenario for both tourists and tourism professionals, but with an innovative stroke to help gain potential tourists. Using the appropriate tools, it is possible to capitalize on the interactivity and immersion presented on the whole virtual world. With this purpose, the consortium

developed important advanced technologies, offering new possibilities to the potential tourists. In addition, the laboratory made it possible to prepare and organise a trip from home, getting information in advance, e.g. about the hot spots in Gran Canaria.

It should also be noted that the Use Case provides full interactivity with the environment and with other visitors, creating a real interactive tourist destination community.

b. Demand from the Gran Canaria Tourist Board for the Virtual Travel Technological Solution.

Recognizing the ability to recreate real world places in virtual environments, several tourist boards and organisations have taken steps into Second Life. Bearing in mind the global target audience of Second Life, their "higher than average" disposable incomes and knowing that these organisations broadly use the Internet for marketing purposes, the deployment of Metaverse initiatives supports in many ways other activities. Furthermore, since many people use the Internet to conduct research about places before actually visiting them, the visual presentation of these venues provides valuable information.

Our Use Case has been a clear example of such practice. The Gran Canaria Tourist Office is an autonomous public institution dependent on the Gran Canaria Island Government. Its ultimate objective is to guarantee both the development and management of the tourist industry on the island of Gran Canaria. The office has decided to strongly collaborate with our Use Case at Metaverse1. The cooperation with this institution was a key reason for the Use Case development. In this respect, an official agreement was signed between this institution and the project consortium, granting access to official information and multimedia material to be used within the Use Case.

The final results of the virtual travel Use Case has persuaded Gran Canaria to become the first tourist destination in Spain that uses virtual technologies as a marketing tool to attract more tourists. In addition, it was an excellent opportunity to present the main tourist resources of Gran Canaria, and to get a better position on the tourism market, through an innovative research and tools.

As a direct result of its participation in Metaverse1, this important tourism institution of the Canary Islands is going to develop and virtualize more tourist spaces of the island. This pilot experience brought remarkable results.

c. Why Second Life?

To successfully implement this Use Case, several different technologies have been assessed and weighed on their individual characteristics and capacities to meet the needs and technical requirements for this particular virtual travel scenario. The outcome of this evaluation has pinpointed the use of three main technologies, a Virtual World

Metaverse platform, known as Second Life (SL) and the use of virtual sound technology combined with navigation and realistic building modelling.

One of the features of Second Life is the ability to create and handle scripts in a specific language (Linden Script), which in turn allows users to develop, animate and program several aspects within the virtual world. These features range from a canyon that launches people (like in circus) to a telephone that uses a real life messaging system for mobile devices all around the world. The operational programming of Second Life provides users with the possibility to edit any of these different features or items, keeping the intellectual property rights with the user who created them.

Following the background idea of the Metaverse1 project, which is enabling users to develop a global standard among real and virtual worlds, it was necessary to implement specific technologies, making the virtual travel a real experience. The user visiting the Gran Canaria Island should experience a unique sensation that should be close to the experience in the real environment of the real tourist destination, real information, real image, and finally real needs when they organize and prepare a trip.

The technologies presented here are not the only ones developed during the project, but these technologies were the most appropriated to be implemented according to the philosophy of the Use Case: travel to a real destination through a virtual situation.

In the following section, the global picture developed in the Use Case will be explained, and general information will be given on the technologies implemented in the Virtual Travel Use Case. We will explain how they have contributed to the main goal.
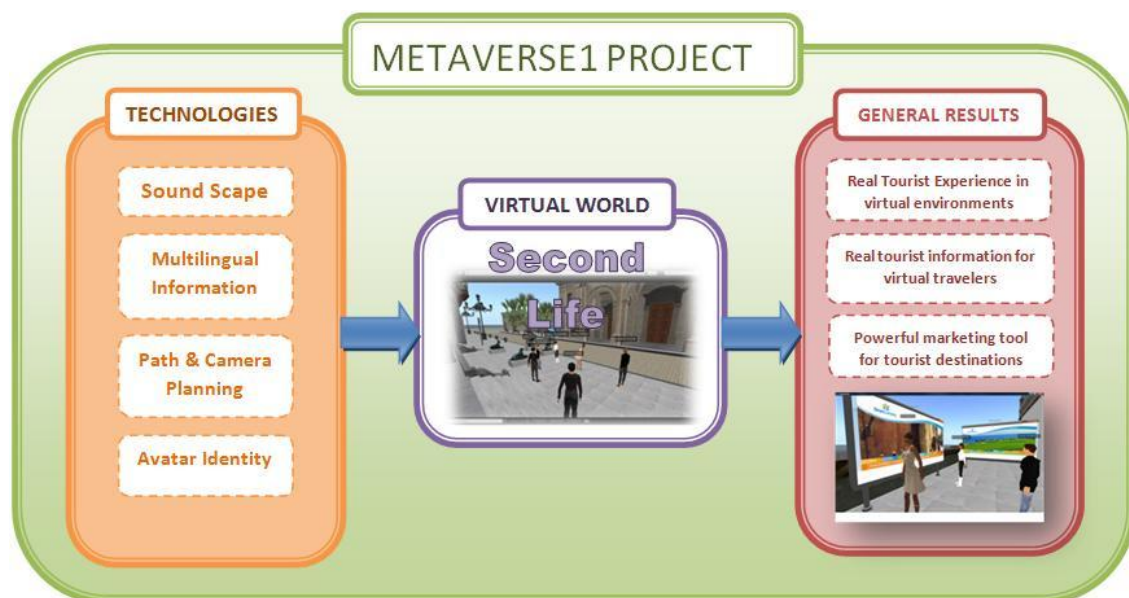

Figure 2: Metaverse1 Project: Overview of the Use Case

3. Technologies implemented

a. Introduction

As it has been mentioned in the previous sections, the Virtual Travel Use Case has served as an authentic test bed, giving the opportunity to validate the technologies identified in the Metaverse1 project. According to the main goal of Metaverse1 project: "Global Standard among Real and Virtual Worlds", some specific technologies have been validated in this use case. In the following section, a short overview on some of them will be given, especially on those that have been implemented, despite the technological constraints and limitations of the Second Life virtual platform.

The following section provides a global overview about the concepts related to these technologies and the contribution to the state-of-the-art.

b. Soundscape generation

With respect to the audio elements, a platform that allows recreating a realistic soundscape has been introduced. This was not possible in Second Life, due to its closed architecture and limited audio generation capabilities.

The audio is a crucial element in immersive virtual environments. Its main role is the creation of a sound ambiance or *soundscape*. One of the technologies that were needed in the Virtual Travel Use Case was a Soundscape Generation system that overcomes the audio limitations of Second Life.

Soundscape design is generally a slow manual process, requiring expert skills and access to large sound effect libraries, or to complex synthesis algorithms. Current technological scenarios (e.g. online communities, convergence of web and mobile technologies) can foster new paradigms for a collaborative design of soundscapes.

<u>An online platform for virtual soundscapes</u>

An online platform has been developed [Schirosa, 2010; Janer, 2011], aiming at simplifying the process of soundscape authoring, able to generate a realistic and interactive soundscape. Moreover, our synthesis engine and server architecture support several independent listeners simultaneously. The client application, that is Second Life in this use case, is responsible for sending position updates, and for receiving the soundscape as a web stream.
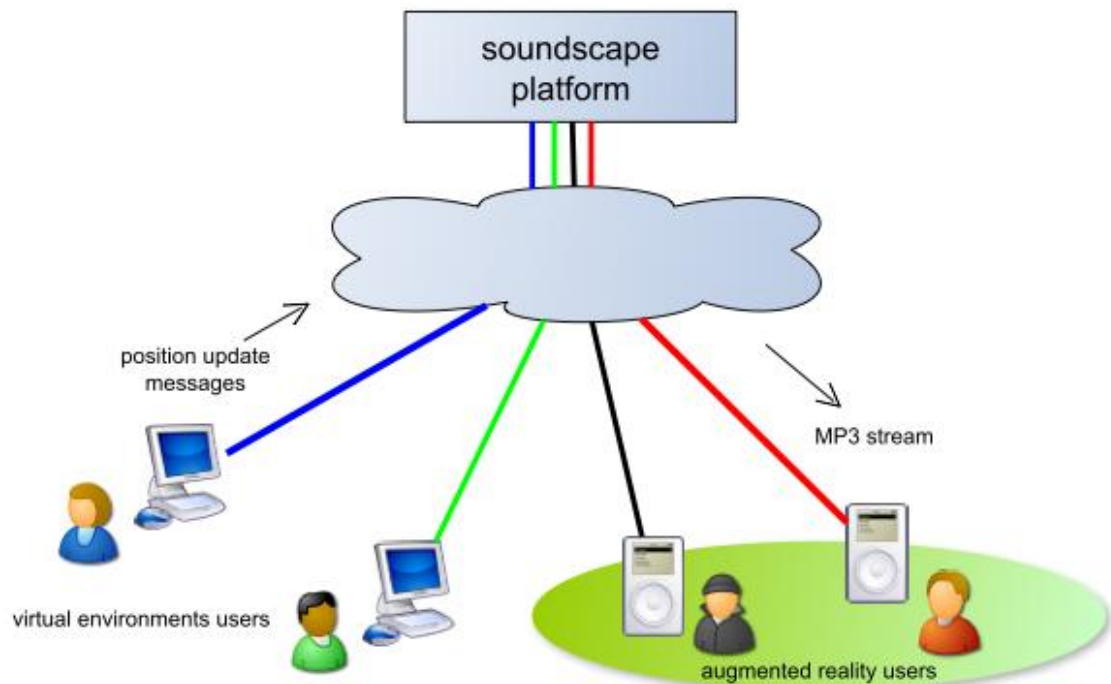
Figure 3: The soundscape system architecture can support virtual or augmented reality applications

When recreating a real place in a virtual environment, the visual content is an oversimplified version of the reality. This often gives a sense of a static and empty space. The developed soundscape improves the sense of presence by providing acoustic cues from the original real place (e.g. children playing on a beach, people chatting, etc).

In the context of the Virtual Travel Use Case, one objective is to recreate the soundscape of a real location into a virtual world running on Second Life. It involves the sonification of three real locations in the city of Las Palmas (Canary Islands, Spain): a beach, a square and a museum, both indoors and outdoors.

In order to bring realism in the recreation of a sonic environment, static or repetitive loops should be avoided. Also, to increase the sense of presence, the quality of synthesized sounds should be as faithful to the real location as possible.

Therefore, we address the soundscape design from a preliminary study of the real locations and on-site field-recordings. An acoustic schema of the environment categorizes sounds as "events" or "ambiance" types, as found in [Valle, 2009; Schirosa 2010].

The built-in sound design tools of Second Life are rather limited. Basically, an audio sample is assigned to a virtual object. When an avatar approaches the virtual object or performs a predefined action, the audio sample is triggered. Hence, the sound sample will sound always the same way. As a unique alternative, Second Life offers to play a sound-effects track as background for a whole area, but the content can be only chosen among predefined factory presets. Moreover, Second Life is a closed system,

where the content creator has to pay for uploading sound files to a proprietary server (SL grid). The audio system of Second Life seems too inflexible for our purposes, offering reduced aesthetic and interaction possibilities.

To overcome these constraints, we decided to implement an external soundscape generation system, which can be integrated through a web streaming in Second Life, as well as in other virtual environments. We believe that the design (or authoring) process should be collaborative by encompassing user-generated content.
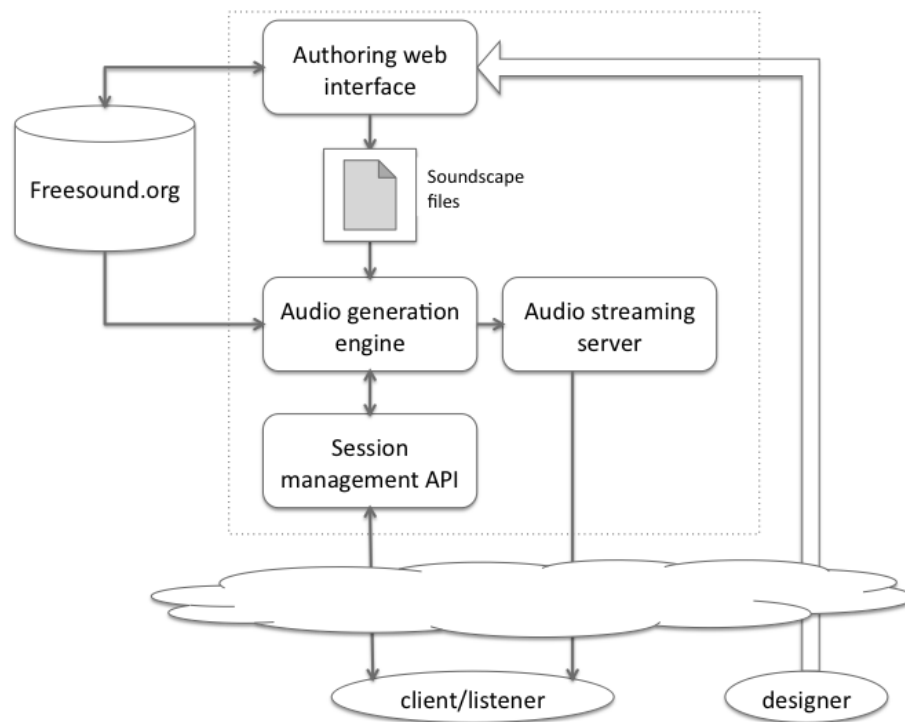

Figure 4: Diagram of the soundscape generation system

Authoring a new Soundscape

A virtual soundscape is here structured in four levels [Valle 2009; Schirosa 2010]: global, zone, concept and events. A zone is populated with multiple sound concepts. Additionally, a zone can have an associated background sound, referred to as ambiance. The synthesis of a sound concept is based on sample concatenation of events and their temporal evolution.

Our system aims at facilitating the authoring process, which can be summarized as:

    1. Localization: the designer places and labels all sound concepts in a 2D coordinates system. Alternatively, the system can import and convert a standard Keyhole Markup Language (KML)[1] [3][Cruz-Lara et al, 2008]

---

[1] http://code.google.com/apis/kml/

file with placemarks (in longitude and latitude) generated in world browsers such as Google Earth. Each sound concept is defined as either: a) a fixed point source, b) a point source randomly located in an area (e.g. bird calls), or c) as an area source that is present in the whole area (e.g. wind).

2. Sample search: each sound concept has a number of associated sound files that are retrieved from Freesound, an online repository of user-contributed sounds [Freesound, 2005]. To facilitate the sample search, the editor application provides direct links to the Freesound API. In addition to the common text query (folksonomy-based), the search engine also features content-based ranking, based on an ecological acoustics taxonomy [Roma, 2010]. Finally, a soundscape design is stored in two separate XML files.

To synthesize a sound concept, several parameters can be configured (e.g. concept probability, regular vs. arrhythmic triggering, randomness, or number of simultaneous samples). A sound concept is composed of various events (samples). To ensure an autonomous and endless generation, a graph model drives the temporal sequencing of events, triggering a sample after a given time specified on each edge. A detailed description of the method behind the graph model is found in [Schirosa, 2010].

The synthesis algorithm is programmed in SuperCollider [Winson, 2011]. After downloading the necessary files from Freesound using a cache mechanism, the algorithm streams audio directly from the disk, therefore reducing the RAM usage.

A spatialization module manages all active listeners, rendering an individualized soundscape stereo mix depending on its position and orientation. Finally, the audio outputs of the synthesis engine are routed to a streaming server.

Soundscapes in the Virtual Travel Use Case

To integrate the soundscape system into the Virtual Travel prototype in Second Life, an additional module has to be implemented. A client command-line application acts as a proxy between the SL application, the SL grid server and our audio streaming server. When an avatar enters the virtual island, the proxy sends a "new listener" message through an HTTP API, and then it receives the listener id and the streaming URL, which are used for further communication with the streaming server. Avatar position and rotation messages are sent from the SL client to update the listener position in the soundscape generation server. A customized streaming URL is passed on to the SL client that renders the audio stream as part of the "music URL" associated with a region in the virtual world. Various simultaneous listeners can navigate through the same sound space.

<u>Conclusion</u>

In terms of scalability, the implemented architecture has two main limitations. On the one hand, the system introduces latency due to audio compression in the streaming server, plus the network delay. On the other hand, processing audio data in real-time for one soundscape, with a single listener requires 5% of the CPU capacity. A solution for the scalability problems is to implement the sound generation engine on the client. For example, an implementation in HTML5 could run on web browsers.

With respect to the use of standards, the soundscape data is stored in an extended KML format. KML (Keyhole Markup Language) is an open format based on XML, used to describe geographic data developed by Google. Hence, our sounsdcape system can also be used in Mixed and Augmented Reality applications. All developed software is available on the project website2.

c. MLIF: The Multi Lingual Information Framework

As with TMF[3] (ISO 16642:2003) in terminology, MLIF[4] (ISO FDIS 24616) provides a metamodel and a set of generic data categories for various application domains. MLIF describes not only the basic linguistic elements (such as sentence, syntactic component, word, and part-of-speech), but can also be used to represent the structure of the document (such as title, paragraph, and section).

---

[2] http://mtg.upf.edu/technologies/soundscapes
[3] TMF. ISO 16642 (2003). Computer Application in Terminology – Terminological Markup Framework, Geneva, International Organization for Standardization.
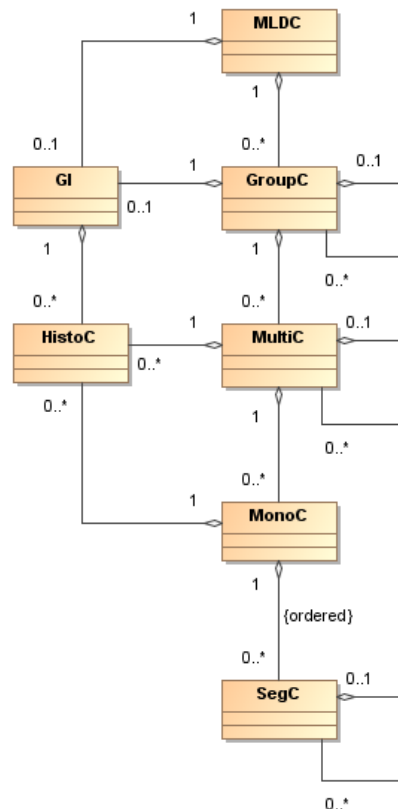
Figure 5: The MLIF metamodel

What is a metamodel?

A metamodel does not describe one specific format, but acts as a high level mechanism based on the following elementary notions: structure, information and methodology. The metamodel can be defined as a generic structure shared by all other formats. It decomposes the organization of a specific standard into basic components. A metamodel should be a generic mechanism for representing content within a specific context. A metamodel actually summarizes the organization of data. The structuring elements of the metamodel are called "components" and they may be "decorated" with information units. A metamodel should also comprise a flexible specification platform for elementary units. This specification platform should be coupled with a reference set of descriptors that should be used to parameterize specific applications dealing with content.

A metamodel contains several information units related to a given format, which are referred to as "Data Categories". A selection of data categories can be derived as a subset of a Data Category Registry (DCR) [ISO 12620:2009; DCR TC37]. The DCR defines a set of data categories accepted by an ISO committee. The overall goal of the DCR is not to impose a specific set of data categories, but rather to ensure that the semantics of these data categories is well defined and understood. A data category is the generic term that references a concept. There is one and only one identifier for a data category in a DCR. All data categories are represented by a unique set of descriptors.

For example, the data category /languageIdentifier/ indicates the name of a language which is described by 2 [ISO 639-1] or 3 [ISO 639-2] digits. A Data Category Selection (DCS) is needed in order to define, in combination with a metamodel, the various constraints that apply to a given domain-specific information structure or interchange format. A DCS and a metamodel can represent the organization of an individual application, the organization of a specific domain.

Due to the genericity of its metamodel and the facility to adorn it with data categories, MLIF can be used in several types of applications. The latest version of MLIF provides strategies for the interoperability and/or linking of models including, but not limited to: LISA TMX, OASIS XLIFF, W3C SMILText and W3C ITS.

MLIF in the Virtual Travel Use Case

Within the Virtual Travel Use Case, the representation and the management of textual multilingual information relies on MLIF. The Multilingual-Assisted Chat Interface is a tool that offers new features to chat users in virtual worlds. It is directly embedded in some viewers for virtual worlds, including Second Life and Solipsis.

Figure 7 shows the general high-level architecture of the main components of the Multilingual Assisted Chat Interface. Three colours are used in this scheme. Each one represents a certain category of components:

- **orange:** components belonging to the virtual world (especially the viewer);

- **blue:** the web service components that we developed, mainly business components;

- **green:** external web services and corpora, as well as data storage.

The circled numbers represent the chronological order of the interactions between the components, when a message is sent or when a word is clicked on. The corresponding explanations are written below:

1. Every message sent by a user is first sent to the virtual world server. When the client (i.e. the viewer) of the user we are writing to receives a message, it is forwarded to a Message Manager, i.e. the component set dealing with the chat messages. We needed to modify these components both in Snowglobe and in Solipsis.

2. Before it displays the message on the Chat Interface, the Message Manager sends an HTTP Request to the web service (the Grammatical Analyzer) in order to get the MLIF data representing the sentence and its several grammatical components.

3. The Grammatical Analyzer connects to external Grammatical Corpora in order to get the grammatical label for each word of the message.

4. When the Message Manager receives the MLIF data structure, representing the original message with a grammatical label (as an HTTP response), the MLIF data is parsed and turned into a format enabling the colouring and hyperlinking of each word. The colour depends on the settings of the user interface.

5. An action is performed: when the user whom the text was sent to reads the message, they click on a word that they do not understand (in the Chat Interface) in order to retrieve synonyms, definitions and translations for this word.

6. After clicking on the word, they are redirected to a Web Interface, which is going to display all the desired information.

7. Loading the web page involves calling the Word Request Manager so that it retrieves information from external web services.

8. The Word Request Manager retrieves definitions and synonyms from WordNet.

9. The Word Request Manager connects to Google Translate in order to retrieve translations.

10. Once the Word Request Manager has received all the desired information, it puts it together in an MLIF data structure. This MLIF data is stored in a database, so that it can be retrieved again later if required.

11. The MLIF data is then sent to an MLIF to HTML Parser, which is going to transform the MLIF data into a user-friendly HTML code.

12. The HTML code obtained previously is finally displayed by the Web Interface to the user, who can easily read the information that they needed.
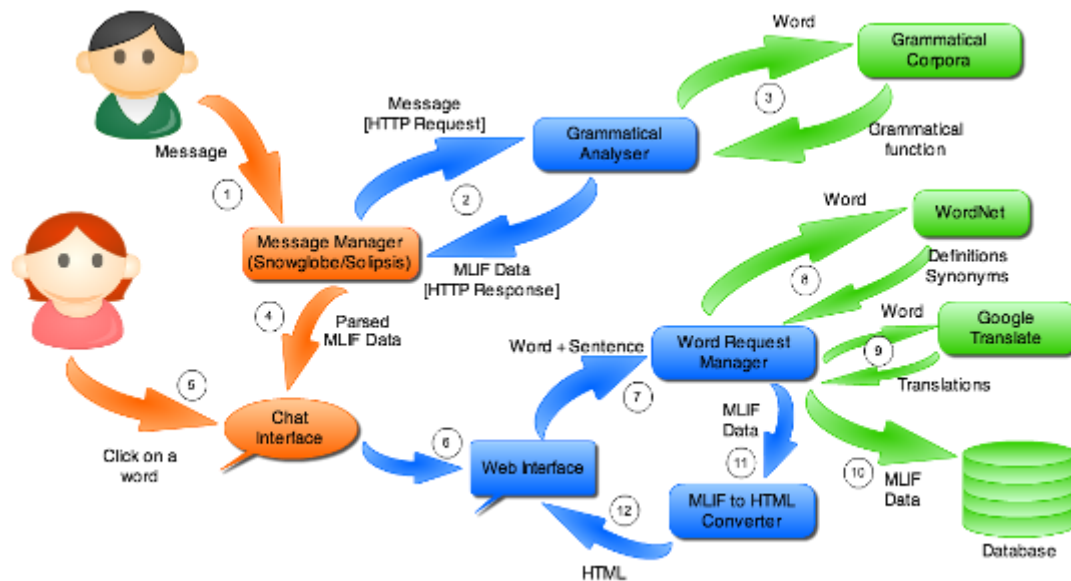
Figure 6: MLIF in the ITEA2 METAVERSE1 Project

Conclusion

The multilingual needs are increasing every day and virtual worlds are a good example of fields of development where applications supporting multilinguality are becoming absolutely necessary.

To enhance interoperability between virtual worlds, applications and corpora, it is obvious that standards should come into place. This is the main goal of MLIF, and, on a more global perspective, of the Metaverse1 project.

d. Path and Camera Planning

Navigation within Second Life can be hard, especially for novice users. Therefore, a support solution has been set up to assist the user. This assistance can be used in any virtual world.

A mechanism has been designed that allows the user to easily navigate through the virtual world (i.e. Second Life) without directly controlling the motion of the avatar. Users should be led to and along interesting places in a smooth way.

The user can travel through the virtual world by using a *transportation device*, such as a *virtual air-scooter* (or *magic carpet* in the Use Case). This device contains a dashboard with a map of the world. By clicking on a certain point on the map, the device flies to the indicated position, using smooth motion, avoiding collisions with elements from the environment and leading the user along interesting locations on the route. Together with the device path, a smooth camera motion is provided, giving the user a clear view of the environment. This type of motion consists of camera placements and camera aim positions. Hence, three (corresponding) paths are specified for a tour. In

addition to a personal tour, the user can also choose a more general guided tour.

Virtual tour guides

General technical details

A planning algorithm needs to have a simplified geometrical representation of the world. Initially, we use the 2D footprint of the world, which is made up of a collection of geometric primitives, such as points, lines, and polygons, all lying in the ground plane. While the planning is performed in 2D, the motions can be projected on an elevated terrain. Next, starting from the footprint, a data structure is built, i.e. the *Explicit Corridor Map* (ECM), which is a navigation mesh representing the walkable areas in the environment [Geraerts, 2010]. The ECM consists of the *medial axis* of the walkable spaces in the environment, where some points on the medial axis are connected to their closest obstacle points. We refer the reader to Figure 7 for an example. This data structure is generated automatically using the graphics card, allowing for real-time computations. Its size is optimal, i.e. linear in the number of vertices describing the obstacles. Recently, this data structure has been extended for multi-layered environments, such as an airport or a multi-storey building [van Toll et al, 2011].

Given an ECM, the radius of a character, and a query (consisting of a start and goal position), we extract a two-dimensional corridor which provides the global route for the character (or transportation device). Next, we extract a shortest path that has a user-specified amount of minimum clearance to the obstacles. These calculations are *optimal*, i.e. linear in the number of cells describing the corridor. This path acts as an *Indicative route* for the character, providing the input for the *Indicative route method* (IRM)]Karamouzas et al, 2009]. The IRM uses the corridor, indicative route, and query to generate a smooth path (while avoiding other moving characters). This path is the final path of the character, or transportation device in our case.

Next, a camera path is extracted, taking into account the ECM and the character's final path, see Figure 8. More information about this procedure can be found in [Geraerts, 2009].

Demo

Since the Second Life scripting language is too limited for implementing our algorithms, we revert to an external program which can be queried using http-requests within Second Life (see below).

Figure 7: Footprint of an island and its corresponding navigation mesh

Figure 8: Camera and character path, displayed inside a collision-free corridor

Implementation on our server

Pre-processing phase

As input, we need a set of polygons (or a black-and-white bitmap) which denotes the non-traversable (i.e. forbidden) areas. If the user provides such a bitmap, where the black pixels denote the obstacles and the white ones the walkable space, it is converted into a set of polygons by an algorithm that traces the outline of connected black pixels.

Input: A set of polygons

An Explicit Corridor Map can be constructed by uploading a primitives (PRI) file through our server. A PRI file must contain at least the following two lines, where `xMin`, `yMin`, `xMax`, `yMax`, and `value` are variables:

- `bounding_box xMin yMin xMax yMax`

- `max_dist 0.72` *(internal parameter)*

Next, the PRI file may contain instances of the following primitives:

- `point x y`

- `line $x_1$ $y_1$ $x_2$ $y_2$`

- `polygon nrPoints $x_1$ $y_1$ $x_2$ $y_2$ $x_i$ $y_i$`

- `border` *(places a bounding box around the environment)*

After uploading the PRI file, an ECM (Explicit Corridor Map) file is generated and stored on the server.

Query phase: obtaining a path

The ECM can now be queried through our server, where 10 parameters must be set. The first three ones are strings and the remainders are numbers:

- *environment*. This is the name of the environment for which an ECM has been generated.

- *algorithm*.   Currently,   three   algorithms   are   provided: `shortest_path`, `smooth_path` and `camera_path`.

- *output*. The output can either be a set of coordinates or a picture.

- *startx*, *starty*, *goalx*, *goaly*: the start and goal coordinates of the tour.

- *radius*: the radius of the character (a large radius will cause the character to avoid narrow passages when it does not fit).

- *clearance*: the minimum distance to the obstacles along the character's path.

- *timestep*: the sampling density of the resulting path.

For integration purposes with Second Life, the user needs to choose for the coordinates as output. For the `shortest_path` and `smooth_path` algorithms, a sequence of the elements `character.x   character.y` is provided. For the `camera_path` algorithm, we have the following additional output: `character.x character.y camera_pos.x camera_pos.y camera_aim.x camera_aim.y`. Currently, no z-coordinates are computed. This coordinate can be retrieved by querying the current height within Second Life. The camera is placed two meters above this z-coordinate to avoid bumping into other avatars as much as possible. It is though possible to add a method for avoiding collisions with these avatars [Karamouzas and Overmars, 2010], but this might be inefficient due to Second Life's architecture.

Integration in the Virtual Travel Use Case

The coordinate points that are returned by the query are then used to create the tour within Second Life. This can be achieved by using an http-request. We scale the coordinates to the area used in Second Life and then use the `llSetPos()` function to move our vehicle from one waypoint to the next along the path of the tour.

SL-scripts, or: how to integrate a virtual tour into other Second Life locations

First, the footprint of your environment or island is created (as described above) and uploaded to our server. Next, an SL-script that retrieves the coordinates of the path(s) is run. Finally, the script that moves a device is attached to a vehicle. By clicking on the object, a vehicle will be created and initiated.

To select a destination from a clickable map, we put several small objects on a map, each with their own script. When a selection is made, the map object

communicates with the vehicle object, which then moves the avatar to the destination, while providing a pleasant view of the environment.

Demo

Demonstrations can be viewed on our Sentona Island[5], see Figure 9. The user is presented with a clickable graphical representation of the island to choose a destination (see Figure 10,

Figure **11** and

Figure **12**). It should be noted that while the user has a discrete number of options here, the planning algorithm supports arbitrary start and goal positions. A tour vehicle then takes the user to the destination along a path that is computed in real time on our server (see Figure 13).

Our algorithm has also been integrated into the virtual tour of Gran Canaria[6], where a magic carpet is used to explore the island.

Figure 9: Sentona Island

---

[5] See http://slurl.com/secondlife/Sentona%20Island.

[6] See http://slurl.com/secondlife/INNOVALIA%20Virtual%20Travel/128/38/49.

Figure 10: Virtual tour starter (red globe) and map interface with clickable destinations

Figure 11: A starting point of the virtual tour

Figure 12: A user selects a destination from the interface

Figure 13: Avatar on virtual tour cushion moving along a path. The text below shows communication between map, vehicle and path planning server.

Conclusion

Our path planning algorithms provide smooth motions for flying devices, while giving the viewer a pleasant view of the environment. A next step is to simulate a whole crowd of non-player characters to make the virtual environment (Second Life in particular) less empty. While there exist techniques for simulating thousands of characters in real-time [van Toll et al, 2012], Second Life's architecture must be changed dramatically to support this extension.

e. Other technologies implemented:

The avatar identity: The contribution of the MPEG-V standard to the Use Case

Avatar identity attributes

One of the aspects that have been studied during the Use Case development is the avatar identity. Many users like to be represented by a unique identity that distinguishes their avatars' appearance and behaviour from others, while it is exchangeable among virtual worlds.

The study has been focused on defining the attributes that can identify each avatar as unique, and, afterwards, a high-level XML-based specification of these attributes has been created. The specification has been called ADML (Avatar Definition Markup Language) and it has been included into the new MPEG-V standard for interoperability between virtual worlds with minor changes.

The starting point for specifying the avatar attributes lies in several studies that define the human identity. Since an avatar can be seen as an imitation of human beings, its identity should follow this simile.

The set of attributes that can computationally model the human identity are usually divided in Body and Mind attributes [Carolis et al, 2002].

- Attributes that define the visual appearance and movements are considered to be body attributes [Moccozet et al, 2006].

- Attributes that define the internal behaviour are considered to be mind attributes [Kasap and Magnenat-Thalmann, 2007; Gebhard, 2005]

It should also be noted that communication skills are part of the identity, but there is no agreement concerning the group in which they should be included. In this study, the Metaverse1 Consortium has concluded that they are affected by both body and mind factors. **Error! Reference source not found.** shows this identity categorization.

Figure 14: Identity structure

This study aims at defining those attributes that allow the definition of an identity which can be exchanged among virtual worlds. The next step is the selection of the attributes that can represent the avatar, independently of a concrete virtual world or events that occur in a virtual world.

- Body attributes. Appearance-related attributes can be easily defined in a way that allows exchange among virtual worlds. However, movements are very dependent on interactions of the virtual world's inhabitants and on context. They are more related to functional capabilities than to the generic identity. Therefore, only appearance attributes will be codified.

- Mind attributes. Emotions and moods are treated according to criteria similar to those used for movements. While personality is inherent to avatar identity and defines its way of perceiving a world, emotions and moods are dependent on the virtual world context. Therefore, only personality will be codified.

Communication skills describe an avatar's capacities to communicate with the virtual world and to interact with other avatars. Therefore, due to their importance, the main communication channels will be codified.

High-level specification

Error! Reference source not found. shows the ADML structure. The next sections explain and analyse the main tags (`appearance`, `personality` and `communication`) that compose it.

This ADML specification, with some minor changes, was included in the new MPEG-V international standard [Preda and Han, 2010]. In-depth descriptions of concrete sections can be found in [Oyarzun et al, 2009; Oyarzun et al, 2010; Oyarzun, 2010].
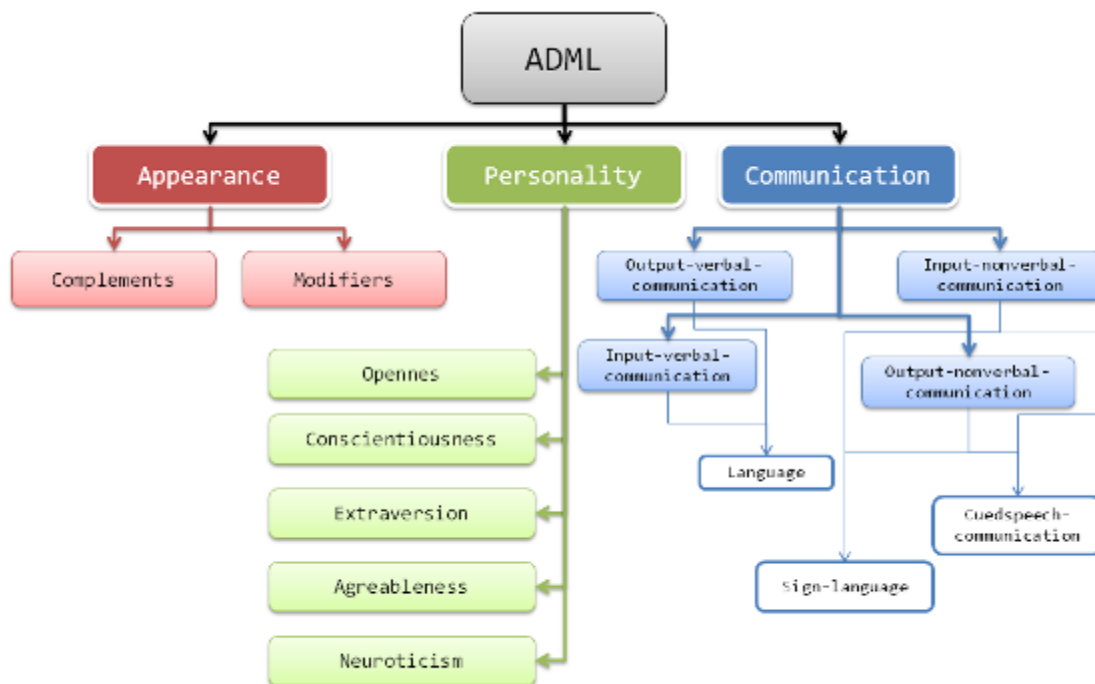

Figure 15: ADML structure

The goals of this part of the standard are:

- Independence from concrete geometric formats. The avatar appearance should not be dependent on geometric formats such as, 3DS OBJ or other open or proprietary formats.

- Unique representation of the avatar in any virtual world. All language-compliant virtual worlds will be able to represent a concrete avatar in the same way, with the same appearance features.

- Extensibility. Language will give the possibility to include new proprietary complements in a concrete virtual world. This will be done through text identifiers; therefore, it will not affect the exchange capabilities of language.

A set of basic avatar models, i.e. man, woman, boy and girl, has been defined to achieve these goals. These models have a hierarchical associated identifier. Following the example, they will be `generic.humans.male`, `generic.humans.female`, `generic.humans.boy` and `generic.humans.girl`. These identifiers will imply a concrete appearance that consist only in the human (or cartoon) shape, without clothes or complements. Any ADML-compliant virtual world will contain the concrete geometric representation for these identifiers. The geometric format is not important; the key requirement is that the representation is correctly associated to the identifier.

There are a set of modifiers that act over the base models, i.e. `eyes-color`, or `height`. Similarly to base models, these modifiers are a set of generic identifiers. The goal of these identifiers is to provide a high level of configurability for the avatar. Additionally, a set of complements are defined, i.e. clothes, jewellery, ornaments, etc. through the tag `complement`. Any virtual world can extend these base models and complements, and add concrete proprietary objects, following the hierarchical structure.

The most extended model for representing the personality in a computational way is the Five Factors model [McCrae and John 1992]. This model has been accepted both by computer scientists and psychologists as a pseudo-standard for representing it. The Five Factors model defines five traits that are the foundation for any personality (openness, conscientiousness, extraversion, agreeableness and neuroticism). Higher or lower values for each trait define a concrete personality.

ADML defines one tag for each of personality trait and an associated value between -1 and 1 that represents the weight of each trait in the personality. The goal of the `personality` tag is to give the possibility to define a personality for the avatar, which can be interpreted by the virtual world architecture. Verbal and non-verbal communication can thus be adapted to each personality, according to this interpretation. ADML also defines the `communication` tag to specify communication skills and preferences of the avatar in the virtual world. The goal of the tag is that the virtual world and the other avatars have knowledge about these preferences and are able (or at least they have the potential capacity) to adapt their inputs and outputs to them, keeping a balance with their own preferences. In this way, all the inputs and outputs will be adapted to each avatar, individually.

Communication preferences are defined in ADML by means of two input and two output channels, which ensure multimodality. These channels are voice recognition and gesture recognition in the case of inputs, and verbal and gestural channels as outputs. Basically, channels can be specified as enabled or disabled. When all channels are enabled, they imply that the avatar is able to speak, make gestures and recognize voice and gestures. For generating and recognizing voice, text or voice modes can be specified.

Gestural generation and recognition has three levels: speech related non-verbal language, sign language and cued-speech language. All skills that are related to language (speak, sign language and cued-speech language) have a `language` attribute

to specify the languages that the avatar is able to understand.

Conclusion: ADML and MPEG-V.

The resulting specification has been an input to new MPEG-V standard. Concretely, `personality` and `communication` tags have been directly included with minor naming changes for keeping the standard formalisms. `The appearance tag` concept was included after an in-depth study by [Jovanova and Preda, 2010]. They created a standard skeleton specification, concrete complements and attributes, and additional tags like those related with animation capabilities. Therefore, they went further than the concept itself, creating a specification that allows implementing it in real use cases. Both [Oyarzun, 2010] and [Jovanova and Preda, 2010] work together on the avatar section of the MPEG-V specification.

Video Streaming

 Apart from the technologies mentioned above, there exist current technologies adapted to virtual worlds, offering new possibilities to end users. The following paragraphs will describe a short description of the implementation of Video Streaming in a virtual world, in this case in Second Life.

The application of this technology to the virtual world would pave the way for the streaming-based services previously listed. For example, the possibility to watch a live video feed from the real world, according to the virtual world location, or view the face of the person guiding the virtual visitors. It could be even extended to allow video conferencing between users, thus extending and complementing the virtual world communication language. Although Second Life is able to display a single video stream projected on a surface, it must be shown as a texture; there can only be a single video stream at once, everyone in the parcel can see the video and the user must be the owner of the parcel. These limitations are quite strict if several videos are intended to be displayed or if we want to enable video communication between users, without disclosing it to non-participants.

Real-time video is usually transmitted via UDP transport protocol, instead of TCP. This ensures less overhead is generated by the transport layer, thus more bandwidth is available for data transmission. The most common protocol for this application is RTP (Real Time Protocol). IP cameras already have the capability to stream video using RTP/UDP, although RTP over TCP is also widely available (if no frame loss is a requirement); and control is carried out by RTSP over UDP or RTSP over HTTP (for instance, to avoid a firewall).

Currently, *Second Life* is using *QuickTime* to display video, which might have a big performance impact. We will consider using another decoder, such as *libavcodec* or our own implementation. Although it might seem very complicated to have three protocols just to display a single video stream, the reason behind this is to split data, flow information and control into three different protocols because they serve different

purposes. Part of the RTP specification, it is a protocol commonly used to control flow and detect congestion. It provides statistics and this information might be used to adapt the encoding and to detect transmission errors.

This protocol is typically used to control streaming servers, to establish connections and send commands that provide functionality akin to a VCR, such as play or pause. It is similar to HTTP; nevertheless, RTSP is a stateful protocol. Extra features can be implemented to provide control over the bit rate, display text on the video, set a video codec. These are protocols that will probably be used in our working solution and are widely available in commercial IP cameras.

In order to video stream, we need to focus on two different sides: real-world implementation through IP cameras that record the real world and a virtual world which the user interacts with.

Figure 16: Video streaming protocol

f. Conclusion

Specific information has been given here about technologies used in this particular use case. A general overview is also included, with a short description of the technology, the state of the art and the latest contributions, in order to highlight the innovation developed in the Metaverse1 project, within the virtual travel use case.

Name of technology	Technology description	State-of-the-art of this technology (short description)	New contribution to the Virtual World (state-of-the-art)
Soundscape	Platform for the design and generation of a sound ambiance for virtual environments, based on user-contributed audio databases.	There are two main trends: 1) game engines: offer interactive sound ambiances, with a high degree of complexity. Drawbacks: tedious manual design, fixed content after release. 2) Virtual worlds allow users to customize the sound ambiance. Drawbacks: a reduced functionality in terms of audio, limited to sample playback for	Combination of users: contributed audio database, (e.g. recordings from the real world), with a flexible sound engine for the generation of complex soundscapes in virtual environments. The platform is deployed as a web-service, including: a) analysis and retrieval tools, b) sound design interface, and c) streaming

		virtual objects.	server.
Multi Linguality	MLIF provides a metamodel and a set of generic data categories for modeling and managing multilingual information in various domains such as localization, translation, multimedia. MLIF also provides strategies for the interoperability and/or linking of models including XLIFF, TMX, OAXAL, SMILText, and ITS.	There are many formats for the textual representation such as TMX, XLIFF, SMILText, and ITS). Although they share many identical requirements, there is no interoperability among them. Moreover, their use is mainly limited to the field they are designed for.	A format for representing multilingual textual information and several applications relying on it, such as a chat interface with grammatical colour codes, emotion detection from text, and representation of information about objects in Second Life.
Path and Camera Planing	Based on the walkable surfaces of the environment, we automatically create a navigation mesh (i.e. an Explicit Corridor Map) that can be used for efficient path planning and crowd simulation in virtual worlds.	Current methods usually create non-exact navigation meshes and poorly support dynamic 3D environments.	The method is exact, works in 3D dynamic multi-layered environments, and does not require user intervention. Currently, the method provides high-quality paths for up to 50.000 characters in real-time.
Avatar Modeling	We specify the avatar features in a generic and standard way. The technology allows users to define appearance, personality, communication skills and animations through high level XML languages.	Current state-of-the-art does not allow defining these characteristics as a whole, that is, as an avatar identity. There are some XML languages for defining appearance and animations, but they are not generic.	We define all the features that compose the avatar identity. They are defined in a high level way, through a XML-compliant language. We standardize them by means of new MPEG-V standard.
Video Streaming	This technology consists in developing a prototype that would pave the way for streaming-based services in a virtual world such as Second Life. It has developed a system that makes it possible to watch a	The current trend in video streaming is embracing HD content. The success of standardised CODECs and the increase of computing power have paved the way for MPEG-4 AVC video. These combined with broadband connections allow	The main goal of this technology is to display a video streaming in the Second Life virtual world. This video-stream must be watched by the avatars who visit a specific parcel in Second

	live video feed from the real world, according to the virtual world location, or to view the face of the person guiding the virtual visitors. This is useful not only for the average end user or someone interesting in travelling, but also for new e-learning methods.	end users to watch high quality video; streaming video from the user's location is still a challenge due to very limited upload transfer rates and poor reliability. If you add the resources required for running a virtual world client to the mix, the situation becomes even more difficult. To put it simply, there are no solutions that offer HQ streaming video in a virtual world. We can have streaming video, or we can have virtual worlds; but not both at the same time.	Life in order to give them the possibility to know tourist information or just to watch the media. An object has been developed to implement this technology. This prototype uses the LSL language that allows changing object properties.

4. Conclusions

Many complementary virtual technologies are able to support the tourism sector in general, and marketing tourist destinations in particular. The Virtual Travel Use Case of the Metaverse1 project that is detailed in this document will bring great help to the tourist sector. Other interesting and innovative research projects are paving the way to the tourist-centred virtualization of destinations. But there is still a long and challenging way to go and a number of hindrances to overcome.

The Virtual Travel Use Case has contributed to the creation of a real tourist environment, offering the possibility to experiment the sensation to be in a real destination using a virtual world and improved technologies adapted to the virtual traveller. The results seem to be a good opportunity for tourist destinations such as Gran Canaria Island, as an interactive, innovative and amazing marketing tool for programming the advertising activities to capture tourists, adapting the content to them and using personalized marketing strategies.

From the user's point of view, the results of the validation process were very optimistic, being well received among users who have used the virtual travel scenario as a tool to organize their holidays or decide to visit Gran Canaria, and to experience the sensations of the virtual world, in a real environment. The use case has also served as an authentic test bed place to improve, adapt and develop technologies to be implemented in a virtual world, thus contributing to the creation of a global standard among real and virtual worlds. The technologies developed were implemented in more similar scenarios.

Obviously, one of the main obstacles concerns dealing with standardization. Virtual worlds need to be standardized to decrease the required development costs for the participant organizations. The Metaverse1 project and its Virtual Travel Use Case have contributed to that standardization process providing inputs to the MPEG-V standard. As a consequence, we may offer a better quality of both the interface and the content to the final user. Moreover, apart from the attractive 3D modelling of the environment, a satisfactory virtual travel needs further and well adapted services for the tourist. This is one of the critical issues that will make a difference between virtual destinations.

The prototypes we have developed and validated through the Use Case are in a continuous development to improve the quality of the service and to adapt to the changing wishes of the traveller.

5. References

1. [Carolis et al., 2002] Carolis, B., Carofiglio, V., Bilvi, M., and Pelachaud, C. (2002). "Apml: A markup language for believable behavior generation. Embodied conversational agents - let's specify and evaluate them!," held in conjunction with AAMAS02.

2. [Chen and Tsai, 2007] Chen, C-F. & Tsai, D.C. (2007). "How destination image and evaluative factors affect behavioral intentions?" *Tourism Management*, 28, 1115–1122.

3. **[Cruz-Lara et al, 2008]** Samuel Cruz-Lara, Nadia Bellalem, Julien Ducret, and Isabelle Krammer. "Standardising the Management and the Representation of Multilingual Data: the Multi Lingual Information Framework," Topics in Language Resources for Translation and Localisation. Editor Elia Yuste. pp. 151-172. John Benjamins Publishers. 2008.

4. [Davis et al, 1989] Davis, F. D.; Bagozzi, R. P.; Warshaw, P. R. (1989), "User acceptance of computer technology: A comparison of two theoretical models", *Management Science* 35: 982–1003

5. [Davis, 1989] Davis, F. D. (1989), "Perceived usefulness, perceived ease of use, and user acceptance of information technology", *MIS Quarterly* 13(3): 319–340.

6. [Freesound, 2005] Universitat Pompeu Fabra. (2005) Freesound.org. Repository of sounds under the Creative Commons license. [Online]. Available: http://www.freesound. Oorg

7. [Gebhard, 2005] Gebhard, P. (2005). "Alma: a layered model of affect". In AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems, pages 29–36, New York, NY, USA. ACM.

8. [Geraerts, 2009] Geraerts, R. "Camera Planning in Virtual Environments using the Corridor Map Method". In The Second International Workshop on Motion in Games, Springer Lecture Notes in Computer Science 5884, pp. 194-209, 2009.

9. [Geraerts, 2010] Geraerts, R. "Planning Short Paths with Clearance using Explicit Corridors". In IEEE International Conference on Robotics and Automation, pp. 1997-2004, 2010.

10. [Janer, 2011] Janer, J., Kersten S., Schirosa M., & Roma G. (2011). "An online platform for interactive soundscapes with user-contributed content". AES 41st International Conference on Audio for Games.

11. [Jovanova and Preda, 2010] Jovanova B., Preda M., "Avatars

interoperability in Virtual Worlds", Multimedia Signal Processing MMSP 2010 IEEE International Workshop on (2010), 263-268.

12. [Karamouzas and Overmars, 2010] Karamouzas, I. and Overmars, M.H. "A velocity-based approach for simulating human collision avoidance". Intelligent Virtual Agents, 6356:180–186, 2010.

13. [Karamouzas et al, 2009] Karamouzas, I., Geraerts, R., and Mark Overmars. "Indicative Routes for Path Planning and Crowd Simulation". In The Fourth International Conference on the Foundations of Digital Games, pp. 113-120, 2009.

14. [Kasap and Magnenat-Thalmann, 2007] Kasap, Z. and Magnenat-Thalmann, N. (2007). "Intelligent virtual humans with autonomy and personality: State of the art". Intelligent Decision Technologies, 1(1-2):3–15.

15. [Kulviwat et al, 2007] Kulviwat, Songpol, Gordon C. Bruner II, Anand Kumar, Suzanne A. Nasco, and Terry Clark (2007), "Toward a Unified Theory of Consumer Acceptance of Technology," Psychology & Marketing, forthcoming.

16. [Kumar et al, 2008] Kumar, S., Chhugani, J., Kim, C. Kim, D., Nguyen, A., Dubey, P., Bienia, C. & Kim, Y. (2008). "Second Life and the New Generation of Virtual Worlds". *Computer*, 41 (9), 46 – 53.

17. [McCrae and John, 1992] McCrae, R. and John, O. (1992). "An introduction to the five-factor model and its applications". Journal of Personality, 60(2):175–215.

18. [Moccozet et al., 2006] Moccozet, L., García-Rojas, A., Vexo, F., Thalmann, D., and Magnenat-Thalmann, N. (2006). "In search for your own virtual individual". In SAMT, pages 26–40.

19. [Nielsen, 1993] Nielsen, J. (1993). "Usability Engineering". Academic Press, Boston, ISBN 0-12-518405-0.

20. [Oyarzun et al. 2010] Oyarzun, D., Carretero, M. P., Mujika, A., Arrieta A. (2010) "Establishing communication channels for Digital Storytelling Applications". Interactive Storytelling, pp. 260-264, Edinburgh, Scotland.

21. [Oyarzun et al., 2009] Oyarzun, D., Ortiz, A., Carretero, M. P., Gelissen, J., García-Alonso, A., Sivan, Y. (2009) "ADML: a framework for representing inhabitants in 3D virtual worlds". Web3D '09: Proceedings of the 14th International Conference on 3D Web Technology, pp. 83-90, Darmstadt, Germany.

22. [Oyarzun, 2010] Oyarzun, D. (2010) "Representation of virtual world inhabitants: a framework for the conceptual definition of avatars". University of the Basque Country (PhD Thesis), Donostia - San Sebastián.

23. [Preda and Han, 2010] Preda, M., Han, J. J. (2010) "Study Text of ISO/IEC

FCD 23005-4 - Virtual world object characteristics", Eds. Marius Preda (Institut TELECOM) Jae Joon Han (Samsung AIT).

24. [Roma, 2010] G. Roma, J. Janer, S. Kersten, M. Schirosa, P. Herrera, and X. Serra, "Ecological acoustics perspective for content-based retrieval of environmental sounds," EURASIP Journal on Audio, Speech, and Music Processing, 2010.

25. [Schirosa, 2010] M. Schirosa, J. Janer, S. Kersten, and G. Roma, "A system for soundscape composition, generation and streaming ," in Proceedings of the CIM - Collo- quium for Musical Informatics, Turin, Italy, 2010.

26. [Valle, 2008] A. Valle, "Geography: a real-time, graph-based composition environment," in NIME 2008: Pro- ceedings, 2008, pp. 253–256, antonio Camurri and Stefania Serafin and Gualtiero Volpe.

27. [Valle, 2009] A. Valle, V. Lombardo, and M. Schirosa, "Simulating the Soundscape through an Analysis/Resynthesis Methodology", Auditory Display 6th International Symposium, CMMR/ICAD 2009, Copenhagen, Denmark, 2009. Springer, 2010, vol. 5954, pp. 330–357.

28. [van Toll et al, 2011] van Toll, W.G., Cook IV, A.F., and Geraerts, R. "Navigation Meshes for Realistic Multi-Layered Environments". In IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3526-3532, 2011.

29. [van Toll et al, 2012] van Toll, W.G., Cook IV, A.F., and Geraerts, R. "Real-Time Density-Based Crowd Simulation". Submitted to Computer Animation and Virtual Worlds, 2012.

30. [Venkatesh, 2003] Venkatesh, V.; Morris, M. G.; Davis, G. B.; Davis, F. D. (2003), "User acceptance of information technology: Toward a unified view", *MIS Quarterly* 27(3): 425–478.

31. [Winson, 2011] S. Wilson, D. Cottle, and N. Collins, Eds., "The SuperCollider Book". Cambridge, Mass.: The MIT Press, In Press.