Enhancing Semantic and Social Navigation in Information-Rich Virtual Worlds

Spyros Vosinakis[1] and Ioannis Papadakis[2]

[1]Department of Product and Systems Design Engineering, University of the Aegean, Greece

[2]Department of Archives and Library Science, Ionian University, Greece

_____

Abstract

Virtual Worlds are a popular medium for communication and collaboration in 3D and they are being utilized as shared information spaces in various application areas, such as education, culture and entertainment. Being presented as 3D visualization spaces they adopt a spatial model of navigation, in which information artifacts have a certain position in space and users can access them through their avatar. However, compared to traditional hypermedia and web-based information systems, Virtual Worlds offer little support for discovering semantically related artifacts and for following the preferences and actions of other users. Furthermore, there is a lack of interoperability with other information systems that would enhance the world's objects with semantically related content. In this paper we present a framework for extending virtual worlds with interactions and visualizations that support improved semantic and social browsing in the 3D environment, and propose an integrated architecture for their interface with external information sources through the employment of Linked Data. Following the proposed framework, we have implemented a prototype environment that presents a gallery of famous paintings and we have performed a user evaluation in order to assess its usability, focusing especially on its novel features for semantic and social browsing.

Keywords: Virtual Worlds, linked data, social navigation, semantic navigation.

_____

Introduction

Virtual Worlds (VWs) are 3D simulation spaces presenting realistic or imaginary environments, in which users are co-present as animated characters (avatars) and interact with each other and with the world's contents (Bartle, 2003). An important characteristic of these environments is the sense of presence that they generate to their users, i.e. the feeling that they are part of the artificial environment they are interacting with (Zahorik and Jenison, 1998; Witmer and Singer, 1998), as well as the relevant sense of co-presence, i.e. the awareness of other users collocated in the 3D space (Slater et al, 2000). There has been a noticeable rise in popularity of VWs as a medium for communication, collaboration and entertainment in the last few years (e.g. the world of Second Life has over 20 million registered users[i]). This is not difficult to explain, as VWs have some unique affordances to offer. The freedom given to users to express themselves, to experiment, to configure their representation and to develop a kind of social life in the artificial environment have shown to be highly engaging (Herman, Coombe, and Lewis, 2006).

A prospective application area of VWs beyond being a place for user communication and social interaction is to serve as an information space, i.e. to contain artifacts that convey related information, either explicitly or on demand, about specific subjects. Typical examples of such environments include virtual museums and exhibitions (Lepouras et al, 2005), digital libraries (Cubaud and Topol,

2001), information visualization environments (Andrews, 1995) and educational places (Prasolova-Forland et al, 2006). A common characteristic of such cases is the existence of a large collection of artifacts placed in the environment that have some 3D representation and are possibly associated with additional information presented in various forms (e.g. text, video, animation) within or outside the virtual space. Following the terminology of Bowman et al (2003), we call these environments information-rich Virtual Worlds. Typically, visitors of these worlds are aiming to explore the information associated with its contents, to search for specific objects or content that match their needs, and to meet with users with similar interests. Being places that allow for synchronous communication and collaboration, information-rich Virtual Worlds may be used for the emergence and support of virtual communities of interest.

Virtual Worlds inherently adopt a spatial model of navigation, as they employ metaphors that mimic (or sometimes enhance) traveling in physical space. However, in the case of information-rich Virtual Worlds, spatial navigation alone is not enough; users may wish to follow semantic links and social trails during their search for information. Dourish and Chalmers (1994) have defined three modes of navigation in information spaces: spatial, semantic and social. Spatial navigation is based on the arrangement of information elements (or, artifacts) in their presentation space, semantic navigation is the following of links to other artifacts with semantic relevance, and social navigation is driven by the actions and preferences of other users in the information space. Jeffrey and Mark (2003) have argued that navigation within a VW could be considered social navigation, due to the fact that the actions of others in the shared space influence the users' actions. This is true, but it is highly restricting compared to Web applications that support social navigation. The reason is that the only perceived user actions in the VW are of those users who are directly visible at that time, whilst on the Web visitors can indirectly get information about the actions of other users. Moreover, visitors of a VW are restrained to an isolated environment without any direct access to other worlds or information spaces. This derives from the fact that VWs' creators are more concerned in attracting as many visitors as possible than agreeing upon the necessary standards and communication protocols that will result in the creation of interoperable worlds, capable of providing their users the opportunity to become members of an integrated environment consisting of diverse yet interoperable worlds or information spaces. Following from the above observations is the fact that although today's VWs may be used as shared information spaces in various application areas, they offer little support for semantic and social navigation compared to other means of accessing information.

In this paper we investigate possible visualization metaphors and interaction techniques that may be used in information-rich VWs in order to allow for improved semantic and social navigation of their inhabitants. In this context, we propose the concept of Thematic Virtual Worlds, i.e. shared 3D environments enhanced with social and semantic navigation features, and we present an integrated framework for designing and implementing them. Moreover, we facilitate the creation of interoperable, open VWs that also provide the opportunity to their visitors to interact with external information spaces. This is achieved through the employment of "linked data", which is described as a method of exposing, sharing, and connecting data via dereferenceable URIs on the Web. As a case study of the proposed framework, we have set up a virtual gallery and performed a user study in order to evaluate its usability and to gain empirical observations about the usage of the social and semantic features of the environment.

The rest of the paper is structured as follows: in the next section we further explain the need to enhance VWs with semantic and social aspects and we present the background work towards this goal. In Section 3 we introduce our proposed

framework, emphasizing in the novel features of the environment and describing its architecture. Section 4 presents an implemented prototype of the proposed framework and Section 5 describes the user evaluation process and displays the results. Our conclusions are presented in Section 6.

Adding Semantic and Social Navigation Support to Virtual Worlds

Virtual Reality applications are based on the metaphor of an artificial space, in which users can navigate from an egocentric point-of-view and can interact with its contents in real time in a physical and intuitive manner. Therefore, a fundamental element of all such environments is the existence of a 3D space that contains a collection of objects and a set of rules that govern their behavior. In today's VWs this metaphor is extended by allowing the coexistence of multiple users, embodied as avatars in a shared space. In this case, the existence and actions of others can be immediately perceived, thus social interactions such as communication and collaboration can be supported within the environment. VWs, can be, therefore, accounted as social spaces, in the sense that they contain a user community and support mutual awareness and interaction. Finally, in the case of information-rich VWs, one may distinguish a third kind of space, the semantic space, which contains the attributes, concepts and relations that refer to a subset of the 3D objects, the information artifacts. The goal of our research is to propose new representations for the seamless integration of these three spaces into an enhanced environment that combines the advantages of VWs, information spaces and social networks.

Virtual Worlds as Information Spaces

Virtual Worlds are quite commonly perceived as digital, artificial environments, represented as 3D spaces where avatars and artifacts coexist and interact. Today's popular VWs such as World of Warcraft and Second Life focus their efforts in providing various means of interaction between their visitors and/or

artifacts. However, VWs are not just a way to graphically represent walls, people and artifacts. From another point of view, VW may be seen as information-rich spaces where information not only is statically presented, but also dynamically produced and disseminated.

As compared to the dominant websites, VWs suffer certain drawbacks as far as information representation is concerned. After all those years of web prevalence, users are accustomed in scrolling long pages in order to explore information that is conveniently arranged within the screen's boundaries. This is certainly not the case for VWs, where information has to share a common space with walls, possibly moving objects and avatars. Exploring for information is certainly not as straightforward as sliding a scrollbar, since a visitor of a VW has to move around a 3D space in order to overcome walls and other obstacles that are in sight. Moreover, websites are governed by the powerful hyperlink metaphor, which provides an easy way to move from one website to another. The only equivalent of hyperlinking in 3D environments is the ability to "teleport" the user to distant places. However, similarly to the hyperlink metaphor, extensive use of this kind of navigation may distort the users' sense of presence and cause disorientation (Bowman et al, 1997). Finally, perhaps the most important issue that VWs have to solve is the lack of interoperability in VWs design and implementation. There is a noticeable lack of standards for interconnecting different VW platforms with each other and/or with the Web. This constitutes today's VWs as isolated environments, where visitors are 'trapped' inside without any option of retrieving and sharing information with other possibly interesting VWs and/or information spaces.

Semantic Navigation in Virtual Worlds

Supporting semantic navigation in VWs can be achieved by letting users explore the information space based on semantic relationships between its elements. This ability is therefore based on two distinctive

features: a) the representation of the semantic space, i.e. the assignment of semantic properties to the objects of the environment and the search for similarities between them and b) the metaphors and interaction techniques for the visual representation of semantic links and the user navigation between them. The semantic representation of artifacts (or, resources) is an issue that has been extensively researched in Web-based information systems. On the other hand, the visualization of semantic associations between elements is still an open issue. The spatial arrangement of items based on their semantic proximity is a simple solution that is, however, restricted by the dimensionality of the 3D environment and the inherent need for realistic space representations in VWs. Therefore, additional visualizations that associate spatially distant elements based on one or more common properties must be employed. Characteristics such as user customization and adaptivity of semantic links may be also preferable, as these features are expected to simplify the user interface and reduce information overload.

The use of annotations, i.e. the association of VW elements with additional semantic information is a metaphor that has been utilized in various application areas (e.g. in the Immersive Redliner presented by Jung and Do (2000)). Polys and Bowman (2004) present a review of design approaches for the presentation of abstract information in virtual environments. They focus on the visualization and layout of annotations and on displaying their association with 3D objects. Bazargan and Falquet (2009) also present and classify a number of techniques for representing non-geometric information in virtual environments. A generic environment for adding semantic annotations in virtual environments is presented by Kleinermann et al (2007).

Techniques to visualize associations between information artifacts have been used at the field of information visualization. In this case, researchers have proposed 3D representations of linked documents in which users can navigate and search for information. A commonly used metaphor is a 3D graph, where nodes represented as 2D or 3D objects correspond to information elements and links correspond to lines connecting them. Andrews (1995) presents a tool for information visualization, in which a hierarchical collection of documents is mapped as an "information landscape", i.e. a 3D graph of interconnected documents. DocuWorld (Einsfeld et al, 2006) is a prototype for the 3D visualization and navigation of a document collection. It presents metadata and semantic associations between documents and dynamically rearranges them depending on the user task. StarWalker (Chen et al, 1999) is a multi-user virtual environment that presents semantically organized information. It visualizes documents as spheres and their semantic relations as links connecting them, and has been used as a testbed to study the behavior and search strategy of users in collaborative information retrieval tasks.

Social Navigation in Virtual Worlds

A basic prerequisite for social navigation is the awareness of the presence and actions of others. In the case of VWs, the perception of other avatars is conveying information about where they are, where they are looking at and what they are doing. However, this information is spatially and temporally restricted; only the avatars that are online and stand within the user's field of view are perceivable and, thus, it is only the behavior of this limited number of visible users that drives social navigation. The problem of spatial restriction can be overcome by displaying the position of other online users in a mini-map, i.e. a small top-down view of the environment that is being used as a navigational aid in VWs (Draken and Silbert, 1993), and/or by adding textual descriptions of critical user actions. Concerning past activities, Grammenos et al (2006) have proposed the visualization of user trails as a means to support social navigation. Movements and actions of users leave trails that appear in the environment and can be perceived later on by others. They propose the metaphors of footprints for user navigation, fingerprints for user

actions and fossils as generic landmarks that can be left in places of special interests.

The Scope of Our Research

The aim of our research is to apply the functionality and search strategies found in hypermedia and Web information spaces in the context of 3D multi-user environments in order to improve the effectiveness of information-rich VWs in searching for information and supporting virtual communities. We do not aim to propose novel ways of rendering information spaces in 3D or to discover more efficient means of visualizing abstract information attached to virtual objects and user activities and preferences. Rather, we are interested in proposing a generic solution that can extend existing VWs utilizing some of the previously-mentioned solutions, where applicable, towards interoperability and improved support for navigating through the information artifacts. In this context, we focus on selecting appropriate open protocols and architectures to ensure connectivity with external sources and on proposing metaphors and interaction techniques that allow successful

navigational aids found in hypermedia and social networks to be applied in VWs.

The Thematic Virtual Worlds Framework

The proposed framework consists of a set of functionality enhancements of Virtual Worlds embedded in an integrated architecture that aims to improve the semantic and social searches of users during their visits. First, we define the concept of a 'Thematic Virtual World' as a Virtual World related to a certain subject (or theme), that aims to bring together people interested in it and lets them acquire and exchange information from their interaction with the environment and with each other. Therefore, the goal is to disseminate and extend a knowledge corpus within a Virtual World and to support the emergence of virtual communities and special interest groups. Possible application areas that could benefit from the concept of Thematic Virtual Worlds are: virtual museums, digital libraries, virtual educational environments, knowledge communities and serious games.
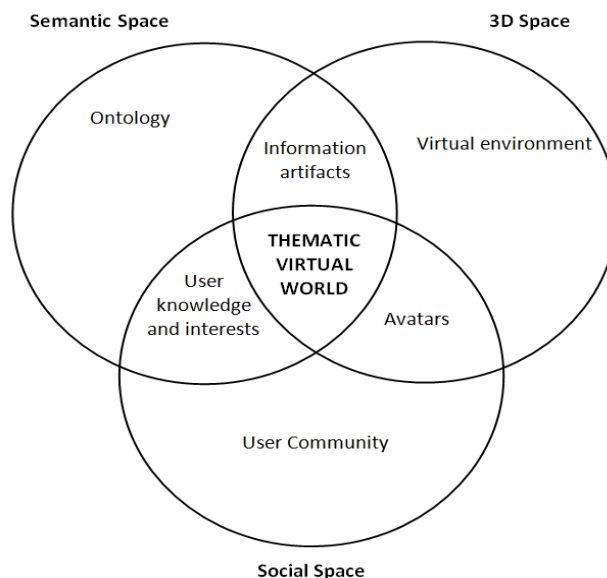

Fig 1. The Concept of 'Thematic Virtual Worlds'

A Thematic Virtual World consists of:

1. **An *Interactive 3D Environment*** possibly containing a set of rules that govern the behavior of its elements.

2. **A *Formal Ontology*** that contains all the concepts related to the theme.

3. **A Set of *Information Artifacts***, i.e. elements of the environment that are related to concepts of the ontology and may possibly be linked to external sources of information.

4. **A *User Community*** that may enter the 3D environment as avatars. The members of the community are related to concepts of the ontology in terms of their knowledge and interests.

Fig. 1 depicts the concept of Thematic Virtual Worlds as the intersection of semantic, social and 3D space. The requirements in terms of functionality towards this goal, as identified in the previous sections, can be summarized as follows:

- *Spatial Navigation*: users should be able to easily navigate to desired places or objects of the environment. Therefore, designers should consider usability guidelines for navigation in 3D environments (e.g. Gabbard, 1997) and should include navigational aids, such as mini-maps and landmarks.

- *Social Navigation*: users should be able to chat with each other, to participate in group discussions, to see where other users are located and to be informed about current and past actions of other users.

- *Semantic Navigation*: users should be able to search for artifacts, other users or

group activities based on desired concepts. Furthermore, the concepts should be dynamically linked with external information sources from the internet, in order to allow for further searches.

The proposed framework attempts to address these needs by introducing a number of novel features that extend VWs' functionality. These features are outlined below.

Thematic Discussions

Erickson and Kellog (2000) claim that knowledge communities can be supported by providing the necessary means to support "long running, deep and coherent conversations". They believe that users should be able to "search, navigate and visualize their conversations" and they suggest that knowledge production and use "will proceed most easily in a semiprivate environment". The authors agree with this approach and propose an application in the context of VWs through the metaphor of "thematic discussions". A thematic discussion may be initiated by any user and can be associated with one or more concepts of the domain ontology. It is visualized as a region drawn on the ground and only users that place themselves within this region are participating in the discussion; all other users are not reading the discussion messages. Thus, thematic discussions take place in the semi-private environment defined by the drawn place. A discussion becomes inactive if there is no chat activity for a long period of time and it no longer exists in the environment. Furthermore, all activity that takes place within the discussion is logged and may be recalled later even if it is deactivated. Users may reactivate a past discussion if they wish. Figure 2 shows a thematic discussion as implemented in our prototype.
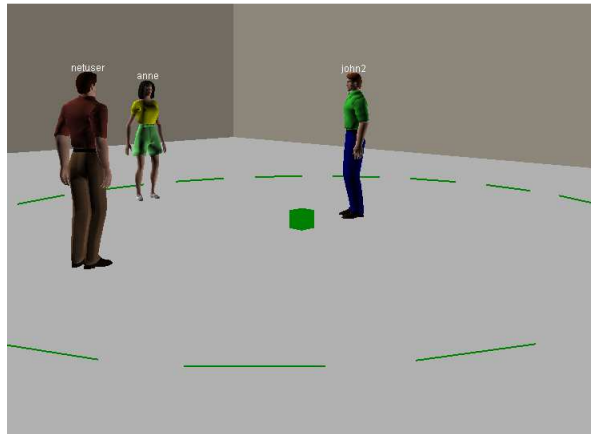
Fig 2. Screenshot of the Prototype Implementation with Users Participating in a Thematic Discussion

The differences between the proposed concept of thematic discussions and of emergent group discussions that may take place in any VW are the following:

1. ***Thematic Discussions are Semi-Private*:** any user can be aware of an ongoing discussion and any user may enter or leave a discussion, but only users that are within the discussion area are reading each other's messages.

2. ***Thematic Discussions are Logged*:** users can read the contents of past discussions and may re-activate them on demand.

3. ***Thematic Discussions are Associated with Domain Concepts*:** users are informed about the topics of the discussion and so they can decide whether to join or not, based on their interests. Furthermore, they can search for ongoing or past discussions based on domain concepts as we shall describe later.

User Trails and Tags

A basic prerequisite for social navigation is to have an awareness of other users' actions and opinions. Ideally, this awareness should include not only directly perceivable events, but also past user actions. In order to satisfy these requirements, we adopt two metaphors that are quite popular in today's web pages and apply them in the context of Thematic Virtual Worlds: a) users may optionally add tags in places, information artifacts or discussions, and b) they leave trails of their motion in the environment. Figure 3 shows user trails in the prototype implementation.

The use of tags lets users know what others are thinking or suggesting about the environment. Tags have been used in collaborative virtual environments in the form of 'annotations' to support asynchronous communication between collaborators (e.g. Jung and Do, 2000). Given the fact that if a lot of tagged elements exist concurrently in the users' field of view their labels will take up much space blocking the view of the 3D environment, designers have to use rules to limit the number of tags being displayed. In our implemented prototype tags are rendered as labels associated to the element they refer to, and they are being displayed only if the user clicks on an artifact or discussion.

Fig 3. Screenshot of the Prototype Implementation Showing User Trails

The visualization of trails is a way to understand the places the majority of users have been moving to and the elements they have been looking at. The environment keeps log of users' motion and their paths are visualized as continuous lines drawn on the ground. The concurrent presentation of a number of user trails provides rich visual information about the popular places of the environment as well as the navigation routes that users have been using in the environment. User trails are visualized upon user request and may be filtered based on specific subjects. E.g. one may see only the trails of users interested in a given subset of the domain ontology.

Semantic Filters

The traditional metaphor of semantic navigation in the Web is the following of hyperlinks from one page to another with relevant content. In the case of VWs, however, where the elements have certain location in the 3D space and user navigational model is spatial, new means of representation have to be adopted. The use of teleportation links from one place to another could be one possible approach. However, frequent use of teleportation as a means of navigating in a VW has been accused as causing disorientation and distortion of the users' sense of presence.

The proposed solution is the use of semantic filters as a metaphor for semantic navigation assistance. The concept of semantic filters is the following: at any time the user may select a subset of concepts from the domain ontology and set them as filters. These filters are then used to search one or more of the following elements: information artifacts, users currently online, active and past discussions, user trails. Figure 4 shows a screenshot of our prototype in which the search results for users, artifacts and discussions related to 'Baroque' are visualized.
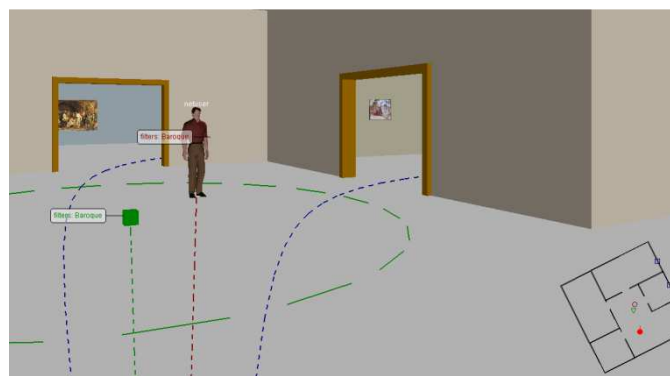

Fig 4. Screenshot of the Prototype Implementation Showing a Visualization of Semantic Search Results for Users, Discussions and Artifacts Related to 'Baroque'

In the implemented prototype we used color coding for the visualization of the search results to distinguish between the first three types of elements (artifacts, users and discussions). For each element that matches the given filters:

a) A safe (collision-free) path is drawn on the ground from the user's point of view to the element.

b) The element is being marked on the mini-map.

c) A label is attached to the element displaying the concepts it is associated with.

Finally, in the case of user trails, only the trails of users whose interests match the semantic filters are being drawn.

Interface with Linked Data

Thematic Virtual Worlds should be capable of communicating with other information spaces of possibly interesting content through their participation to the Web of Data (Bizer et al, (2009)). The vision of the Web of Data refers to an information environment where information is published in a machine-understandable fashion from diverse information providers. The publication process involves enriching published data with semantic information that adds meaning to the data both for humans and computers. Moreover, data would be capable of referencing other data whenever such cross-referencing seems appropriate. Currently, the vision of the Web of Data is facilitated through the employment of linked data (Berners-Lee et al, 2001), which refers to a number of principles and the corresponding technologies to make such principles happen. More specifically, the idea of linked data is governed by four 'rules'[ii]:

1. Use Uniform Resource Identifiers – URI as names for things.

2. Use active URI so that people can look up those names.

3. When someone looks up a URI, provide useful information, using semantic web standards.

4. Include links to other URIs, so that users can discover more things.

The above rules actually imply that the cell of linked data is the 'resource', which refers to anything that can be referenced through a URI. In the case of virtual worlds, resources could be artifacts, avatars, actions between them, relations, etc. Such URIs are 'active' in the sense that people and/or computers can find interesting information about them through the employment of adequate semantic web technologies, such as the Resource Description Framework – RDF[iii] and SPARQL[iv]. More specifically, the RDF standard provides a way to represent information as a graph where the nodes refer to resources and arcs refer to the relations between resources. In order to store an RDF graph in a machine-understandable format, it may be serialized in a number of ways. A popular RDF serialization syntax is the N3 syntax[v], where information about resources is expressed as triples of the form:

```
<subject> <predicate> <object> .
```

According to a trivial scenario, subjects are resources, predicates are the relations between resources and objects can be either other resources or literals. For example, the triple

```
<http://www.dbpedia.org/resource/Paris>
<http://www.dbpedia.org/property/islocated>
<http://www.dbpedia.org/resource/France> .
```

Says that Paris is located in France, whereas the triple

```
<http://www.dbpedia.org/resource/Paris>
<http://www.dbpedia.org/property/hasPopulation> "2.000.000" .
```

Says that Paris has 2,000,000 inhabitants. A repository of triples is called a triplestore and can be stored in various ways, ranging from flat files to specialized RDMS. Quite commonly, specialized RDMS such as Openlink Virtuoso (Idehen, 2011) or D2R(Bizer and Cyganiak, 2006) also provide a queryable access point to the underlying triplestore, capable of addressing queries complying to the SPARQL language.

Evidently, the realization of the notion of the Web of Data depends on the amount of linked data that is available online as well as the degree of interlinking between resources existing within the various trplestores of the linked data – LOD cloud[vi]. Along these lines, DBpedia[vii] emerged as an effort to realize the vision of linked data. More specifically, Auer et al (2007) transformed the semi-structured information contained within Wikipedia to RDF, stored it in accordingly designed triplestores and made it publicly available as data dumps[viii]. Moreover, they provided a SPARQL endpoint in accordance to the aforementioned linked data rules. The emergence of DBpedia inspired many other parties to follow similar actions in order to make their own data publicly available.

In order to apply the above principles in a VW, we make the assumption that a VW is yet another information space, where users interact with each other and with various items of the surrounding environment. The outcome of such interactions is the production of information. Semantic web standards such as URI, RDF and SPARQL should be employed in order to store such information in a way that it is usable from other information spaces outside the VW. At the same time, a 'linked data ready' VW should be capable of providing its visitors the opportunity to look for 'interesting' information outside the strictly defined borders of the VW per se. In order to accomplish such a task, the Graphical User Interface – GUI should be able to translate the visitors' requests for information

during their presence within the VW, to appropriate queries (e.g. SPARQL queries). Such queries should be accordingly addressed to suitable triplestores that act as external information spaces. The corresponding responses should be tunneled back to the GUI, which is responsible for the visualization of the information in a way that is easily comprehensible from the visitors that addressed the initial query. It should be noted that the complexity of the information exchange with the external information spaces should be hidden from the end users.

Architecture

The proposed framework is based on an accordingly designed, 3-tier architecture (Fig. 5). User clients are connected to the Virtual World – VW server, which is responsible for the synchronization of the multi-user 3D environment, and all data management processes take place at the Linked Open Data – LOD server, which provides the interface between the Virtual World and the linked data.

VW Clients

VW clients serve as the presentation layer of the architecture. They render the 3D environment in real-time, handle the user interactions, notify the VW server of any changes in the environment and update the environment according to any messages received from the VW server.

VW Server

The VW server's role is to transmit the appropriate messages, in order to ensure consistency in all client environments.

VW Clients – Server Interaction

The message exchange between any client and the VW server in our implementation is the following:

- **User Entrance:** when a registered user enters the environment its client sends a user entrance message to the VW server, which in respect transmits the message to all clients enhanced with information such as the user's real name and user id, its initial position and the name of its avatar.

- **User Navigation:** a change in the position or orientation of a user causes the client to send a user navigation message to the VW server, which transmits it to all other clients that make the appropriate changes to the avatar that moved.

- **User Communication:** A communication message may be global, personal (whisper) or a part of a discussion. Clients transmit any messages to the VW server, which delivers them to the appropriate clients depending on the message type.

- **Creation or Re-Activation of a Discussion:** when a user creates a new discussion or re-activates a past discussion, the client sends the appropriate message to the VW server, which transmits it to all other clients,

which, in turn, render the discussion area on screen and attach the respective information.

- **Tag:** users may tag information artifacts and discussions. The tag is being sent by the client to the VW server, which informs the LOD server.

- **Semantic Search:** when a user defines one or more semantic filters and wishes to search for related exhibits, users, discussions or trails, its client sends a semantic search message to the VW server, which performs the appropriate queries to the LOD server and sends back the results to the client in order to visualize them.

- **Request for Information:** when the user requests to view the information attached to an artifact, a message is being sent by the client to the VW server, which queries the LOD server in order to find relevant information sources and add them as hyperlinks. The resulting document that aggregates artifact information with related linked open data sources is being sent back to the client.
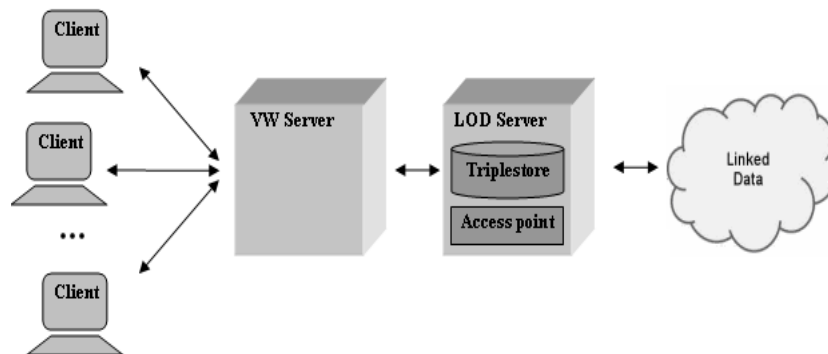

Fig 5. Proposed Architecture

LOD Server

As discussed earlier in this paper, the proposed framework is based on linked data technologies. More specifically, the LOD Server consists of a triplestore containing all the information that is available to the Virtual World. Such

information is encoded in Notation 3 – N3 format and stored in flat file system.

VW Server – LOD Server Interaction

Information exchange between the triplestore within the LOD server and the VW server is facilitated through the employment of an accordingly designed

protocol. The LOD server accepts requests from the VW Server and replies with suitable responses. There are two types of requests that can be addressed to the LOD Server: SET requests and GET requests. SET requests add data to the triplestore in N3 format and GET requests query the triplestore.

LOD Server – Linked Data Interaction

Depending on the type of a request, it is possible for the LOD server to forward a request to an external information space. For example, a request for related information on a given subject could be addressed to DBpedia's SPARQL endpoint (http://dbpedia.org/sparql), which, in turn replies with related hyperinks coming from various data providers. The hyperlinks are tunneled back to the VW server which, in turn, transmits them to the clients for visualization.

The proposed framework is also capable of satisfying requests for information about the underlying VW that derive from other information spaces. Such external systems should be linked data compliant and could vary from other linked data – ready VWs to any member of the LOD cloud. More specifically, the LOD server maintains a publicly available SPARQL endpoint capable of being employed from any other external system.

Prototype Implementation

We have created a prototype implementation of the proposed Thematic Virtual Worlds framework and used it to set up a case study in order to perform an initial user evaluation of its usability and performance. The environment is a virtual gallery containing famous paintings of the 19th and 20th century. Visitors are able to browse the gallery, discuss with each other, perform semantic searches, and lookup information about the artistic movement or style of each painting in related information sources through the employment of linked data. Moreover, external information spaces are able to make requests about the information that is stored within the virtual gallery in the form of linked data (i.e. triplestore) by taking advantage of suitable semantic web technologies (i.e. SPARQL endpoint). Figure 6 presents a screenshot of the case study.

Fig 6. Screenshot of the Prototype Implementation

In order to setup the case study, we have selected 40 images of well known paintings with public domain license and modeled an equal number of VRML models presenting them in a canvas with their original dimensions. We associated each painting with a number of categories concerning the artistic movement it belongs to, its subject, and the technique being used. The set of all these categories formed our domain ontology; visitors of the virtual gallery could declare one or more of them as their interests, use them in semantic searches, or add new discussions about them. Finally, we modeled a large interior space as exhibition hall and arranged the exhibits in its rooms based on the artistic movement.

The implemented prototype is integrating a number of technologies. The client has been implemented as a Java application running over the Web using the Java Web Start technology. All 3D models are stored in VRML format and the visualization process on the client applications is using the Java3D library and the java3d-vrml97 loaders to import and visualize the models. The VW server is a java application that communicates with the clients using standard TCP/IP sockets and a dedicated communication protocol that supports the message exchange presented in 3.5. Furthermore, it posts XML Requests to the LOD Server and accordingly receives XML Responses using the Apache Commons libraries.

The triplestore within the LOD server stores all information in Notation 3 - N3 format. More specifically, the triplestore contains facts represented as triples containing 3 entities: subject, predicate and object. Each entity is expressed as a URI belonging to an accordingly defined namespace (i.e. http://vw.org) or a literal.

For example, the triple:

```
<http://vw.org/resource#impressionism>
<http://vw.org/property#type>
"topic" .
```

Indicates that the resource 'impressionism' is a topic, whereas the triple:

```
<http://vw.org/resource#spyrosv>
<http://vw.org/property#interestedIn>
<http://vw.org/resource#impressionism> .
```

Indicates that the resource 'spyrosv' is interested in the resource 'impressionism'.

Information exchange between the triplestore and the VW server is facilitated through the employment of a protocol encoded in custom XML format. The protocol is implemented in Python and it is based on the RDFLib library[ix]. As discussed earlier, there are two types of XML requests that can be addressed to the LOD Server: SET requests and GET requests. Each request is expressed as an XML message, which is addressed from the VW server to the LOD server.

SET Requests

The following types of SET requests are supported:

1. Add discussion

2. Set inactive discussion

3. Add person

4. Add topic

5. Add relation

6. Add artifact

7. Add tag

Each request is expressed as an XML message, which is addressed from the VW server to the LOD server. For example, in case of an Ádd tag' request, the following message is created:

```
<request action="set"><query id="q1" type="tag">
<tag kind="resource" id="medieval"/>
<istagged kind="resource" id="ampelonas"/>
</query></request>
```

More specifically, the VW server requests to add the tag: 'medieval' to the artifact 'ampelonas'. The outcome of this request is the creation of the following triple in the triplestore:

```
<http://virtualworlds.org/resource#ampelonas>
<http://virtualworlds.org/property#istagged>
<http://virtualworlds.org/resource#medieval> .
```

The above triple declares that the resource id='ampelonas' is tagged with the resource id='medieval'.

GET Requests

According to the protocol, the following types of GET requests are supported:

1. Get predicates of a given subject

2. Get subjects of a given object/predicate

3. Get objects of a given subject/predicate

4. Get related sites of a given resource

Again, similarly to SET requests, each request is expressed as an XML message, which is addressed from the VW server to the LOD server. For example, in case a request for related sites, the following message is created:

```
<request action="get"><query id="q1" type="relatedsites">
<topic kind="resource" id="Impressionism"/>
</query></request>
```

The above request results to the following response, which essentially fetches related sites to the resource id='Impressionism'.

```
<response action="get"><reply type="success" id="q1">
<resource hostname="umbel.org" type="relatedsite">
http://umbel.org/umbel/ne/wikipedia/Impressionism </resource><resource
hostname="rdf.freebase.com" type="relatedsite">
http://rdf.freebase.com/ns/guid.9202a8c04000641f800000000001f201</resource>
<resource hostname="wikipedia.org" type="relatedsite">
http://wikipedia.org/wiki/Impressionism </resource></reply></response>
```

The request is addressed to DBpedia's SPARQL endpoint (http://dbpedia.org/sparql), which, in turn replies with three related hyperinks coming from umbel.com, rdf.freebase.com and wikipedia.org respectively. The three hyperlinks are tunneled back to the VW server which, in turn, delivers them to the client that initiated the entire message exchange.

We have to note that the technologies employed for our prototype are not the only approach to implementing the proposed framework. Existing VW platforms could be extended with the appropriate tools and visualizations to support the features of Thematic Virtual Worlds introduced in this paper, provided that they have the ability to exchange data with an external LOD server via HTTP. Given that Second Life supports HTTP requests using its embedded scripting language (LSL), it would be possible to implement our proposed architecture in that platform. However, there are

limitations in terms of visualization and interface customization using LSL. One such example is the dynamic creation of a safe (collision-free) path that leads to a specific place, which is not possible to be created using LSL, so one will have to use an alternative visualization technique, e.g. a radar. Therefore, the main reason that led us to our own implementation of the VW client and server instead of being based on an existing platform was to have more freedom in choosing the appropriate visualizations and designing the user interface.

Evaluation

We set up a user evaluation of the prototype implementation in order to test its performance with a number of concurrently connected users, to assess the usability of the proposed metaphors, and to gain empirical observations from the usage of the environment. The main focus of our evaluation was not on the implementation itself, but on the introduced metaphors and concepts, and on the users' response to them. The evaluation took place in the Computer Laboratory of the Department of Product and Systems Design Engineering. The client software was running in PCs with Windows XP professional, Intel Pentium4 3.6GHz CPU, NVIDIA Quadro FX 540 Graphics card and 2GB Ram. We asked a number of users to concurrently connect to the prototype environment and to follow a specified scenario and we collected data from monitoring of students' behaviour, questionnaires and personal interviews.

The evaluation scenario was the following: we asked the users to register to the system and to declare their interests by selecting a subset of the domain categories. Then, we let them concurrently connect to the environment. We spent 5 minutes in an introductory session, in which the functionality of the environment was explained and demonstrated.

After that, we let them experiment on their own with the environment and instructed them to do the following:

1. To browse the collection of paintings, find at least three paintings that match their interests and tag them.

2. To try and find other users with similar interests and/or visualize their trails.

3. To participate in a discussion about a major artistic movement.

The session lasted 30 minutes, during which all users' movements and actions were logged by the VW server. After that, the users were asked to complete a questionnaire in order to assess their opinion about each of the semantic and social navigation features of the environment and of the system in general. Finally, the users were interviewed in order to provide further comments and suggestions.

A group of thirteen users (four male and eight female) participated in the evaluation process, all of which were undergraduate students of the department of Product and Systems Design Engineering, University of the Aegean. Initially, the users were asked to record their experience in first-person navigation (e.g. from playing computer games) in a 5-level Likert scale. Four of them declared that they were highly experienced, whilst the rest of them had medium or little experience. No user reported having no experience at all.

The overall performance of the system during the evaluation session was satisfactory. There were no problems by the concurrent interaction of all users, as the refresh rate of the client environments was not significantly decreased by the introduction of a multitude of animated avatars. However, delays were noted during the login procedure and during the request for additional information about the paintings, caused by the increased traffic between the VW server and the LOD server.

During their interaction session, users seemed to have little difficulty in understanding the new metaphors. They used the semantic search tool quite often (on average 9.3 times per user) in order to find paintings that were related to categories that captured their attention. Also, they easily managed to detect the search results on the mini-map and to follow the drawn paths in order to approach the items. Three new discussions were generated, but the dialogs were short and typical given the limited time of the interaction session. A large number of the paintings (29 out of 40) have been tagged by users. Most of them were comments about their aesthetic quality. All users visited at least three external links related to the paintings, mostly concerning the creator's biography or information about the artistic movement. Finally 10 out of 13 users requested to see the trails of other users with similar interests and attempted to follow them to see what they have visited.
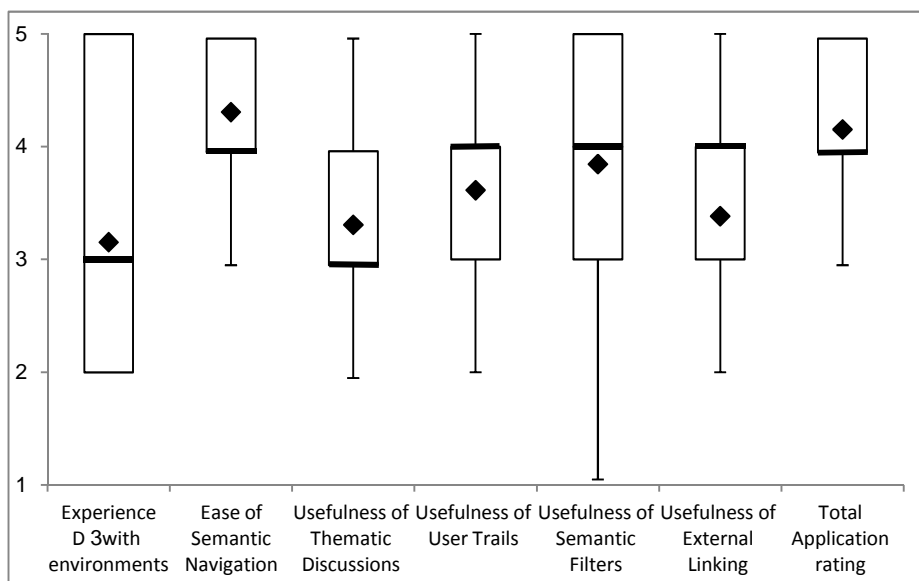

Fig 7. Boxplots, Medians (Thick Line) and Averages of the Likert Scale Questionnaire Results Concerning Users' Experience with 3D Environments, Ease of Detecting Semantic Navigation Results, Usefulness of the Metaphors Introduced, and Total Application Rating

Users were asked to comment on the usefulness of the introduced features for VWs containing information artifacts (e.g. 3D digital libraries, virtual museums, educational environments, etc). They replied in a 5-level Likert scale, with 1 meaning not useful at all and 5 meaning very useful. The average values of their responses (Fig. 7) suggest that the majority found them quite useful, with semantic filters being the most useful one (3.85) and thematic discussions the less (3.3).

Users were also asked to assess their difficulty in detecting the search results of semantic filters using the same scale (1 meaning impossible and 5 meaning very easy) and the mean value of their answers was 4.3, which implies that they did not have significant difficulties following the navigational cues. All users that reported to be highly experienced in 3D environments found it very easy to navigate to the search results. This finding was expected since the search results were presented as drawn routes and marks on the minimap, which are common navigational cues in 3D environments. Finally, users were asked to rate the system in total in terms of its usefulness as a platform for Thematic Virtual World and their mean rating was 4.15. No correlation was found between users' experience with 3D environments and their satisfaction from the application.

The results of the questionnaire are presented in Figure 7 including boxplots, medians and averages.

Most user comments focused on the general usefulness of the platform, as well as of the problems they encountered during their interaction. Two users commented that they experienced difficulties in navigation. One of them explained that it was difficult to maneuver his/her avatar in order to position himself correctly to have a good view of the paintings. Another user commented that the image quality of the paintings was low, so he/she missed some interesting details. In order to address this issue, an extra option should be added to the user interface, enabling users to view another version of the painting with higher resolution. On the other hand, some users commented that the idea is very interesting, especially if applied in educational environments, and one user made the following comment: "a very useful application, because the user can have a very good initial idea about the exhibition he/she wishes to visit, mark the places with interesting exhibits and any related information that he might like to view, so that he doesn't have to read brochures before entering the physical space of an exhibition" (translated from Greek). Finally, the majority of users were fairly enthusiastic about using the system and some of them asked whether there are other similar 3D environments of online museums and galleries worldwide.

The evaluation results suggest that the users found the proposed metaphors and the concept of Thematic Virtual Worlds useful for navigating and interacting with information-rich 3D environments. They did not have significant difficulties in understanding and using the introduced metaphors and they made a fairly good use of them during the evaluation session. Most of the difficulties observed or reported by the users had to do with the specific implementation, rather than the proposed concepts. The navigational difficulties reported could be overcome by including a variety of navigational methods to choose from (e.g. first person / 3rd person /

examine), and the ability to move sideways (strafe) would improve the browsing of the exhibits. Based on the observations and the recorded user behaviors, users could successfully perform semantic navigation activities, i.e. to find exhibits related to their preferences, to visit related exhibits or related external information sources. On the other hand, the social navigation activities, i.e. discussions and user trails, were considerably fewer. This result was quite expected given the limited time and the number of people using the environment. A long-term evaluation of the environment with a large number of users would probably lead to more revealing results concerning its ability to promote social navigation. However, we decided to perform the evaluation in the laboratory in order to be able to directly observe the users and their reactions while they were interacting with the environment, and to interview them afterwards.

The implemented prototype did not present significant difficulties in terms of performance. However, its ability to handle a large number of concurrent users and/or a much larger number of information artifacts has to be tested in further evaluations. A long term usage of the environment would require the employment of a strategy for removing a percentage of user trails, discussions and tags, based on some criteria. Such an approach would avoid a possible reduction of performance by searching in extensive amounts of data. Furthermore, in the case of a large number of results additional visualization techniques will be required to avoid overloading the user's view with too many visual cues.

Conclusions

In this paper, we have introduced the concept of the 'Thematic Virtual World' as a VW related to a certain subject that aims to bring together people interested in it. During their presence in the VW, visitors acquire and exchange information from their interaction with the environment and with each other. A Thematic Virtual World contains metaphors that underpin the

semantic and social aspect of this particular environment.

The implemented prototype provides the functionality that derives from the intersection of three different spaces, namely the semantic space, the 3D space and the social space. More specifically, the virtual gallery introduces thematic discussions as a special kind of group discussions featuring semi-privacy, logging, and association with domain concepts. Also, user trails are supported, giving this way the opportunity to visitors to follow others that exhibit similar interests. The virtual gallery avoids information overload of the VW by introducing a number of semantic filters that give the opportunity to visitors of the gallery to filter paintings, online users, active and past discussions and user trails. Finally, visitors of the virtual gallery are able to gather information from other, external information systems through the employment of linked data technologies. At the same time, other VWs or information systems are able to query the information that is stored within the virtual gallery.

The evaluation results were rather encouraging, in the sense that the proposed functionality enhancements, as presented in the implemented prototype, were considered useful by the users and opened new possibilities for semantic and social navigation. However, we have to note that the prototype was only one specific way of implementing the framework; there are multiple ways of designing and visualizing the proposed features. Therefore, a number of further studies will be required in order to test, evaluate and re-design more usable solutions for the concepts presented in the Thematic Virtual Worlds framework.

Educational Virtual Worlds, Virtual Museums and Serious Games may significantly benefit by following the proposed approach. Given that VWs are a highly engaging medium for collaboration and recreation and that they have a number of unique affordances, they are considered as prospective future learning environments. Therefore, it is expected 3D environments will be enhanced with educational content and this mass of information will require more efficient navigational strategies and connectivity. It is important for educational environments to be able to update their content from external sources and dynamically generate hyperlinks with related concepts. Even more important is the ability to allow external pages to retrieve data about the activity that takes place within the VW, e.g. a Web page could inform visitors about the contents of the Virtual World, the interests of its inhabitants and the discussions that took place in it.

References

Andrews, K. (1995). "Visualising Cyberspace: Information Visualisation in the Harmony Internet Browser," *Proceedings of IEEE Symposium on Intofmation Visualization.*

Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R. & Ives, Z. (2007). "Dbpedia: A Nucleus for a Web of Open Data," Proceedings of ISWC07.

Bartle, R. (2003). Designing Virtual Worlds, *New Riders, USA.*

Bazargan, K. & Falquet, G. (2009). "Specifying the Representation of Non-Geometric Information in 3D Virtual Environments," Proceedings of the 13th International Conference on Human-Computer Interaction, 773-782.

Berners-Lee, T., Hendler, J. & Lassila, O. (2001). "The Semantic Web," *Scientific American,* [Online], [Retrieved July 15 2011], http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21.

Bizer, C. & Cyganiak, R. (2006). "D2R Server – Publishing Relational Databases on the Semantic Web," Poster at the 5th International Semantic Web Conference (ISWC).

Bizer, C., Heath, T. & Berners-Lee, T. (2009). "Linked Data - The Story So Far," *International Journal on Semantic Web and Information Systems (IJSWIS).* [Online], [Retrieved July 15 2011], http://tomheath.com/papers/bizer-heath-berners-lee-ijswis-linked-data.pdf.

Bowman, D. A., Koller, D. & Hodges, L. F. (1997). "Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques," Proceedings of Virtual Reality Annual International Symposium (VRAIS).

Bowman, D. A., North, C., Chen, J., Polys, N. F., Pyla, P. S. & Yilmaz U. (2003), "Information-Rich Virtual Environments: Theory, Tools, and Research Agenda," *Proceedings of the ACM symposium on Virtual reality software and technology (VRST '03). ACM, New York, NY, USA,* 81-90.

Chen, C., Thomas, L., Cole, J. & Chennawasin, C. (1999). "Representing the Semantics of Virtual Spaces," *IEEE Multimedia,* 6(2), 54-63.

Cubaud, P. & Topol, A. (2001). "A VRML-Based User Interface for an Online Digitalized Antiquarian Collection," Proceedings of the Sixth international Conference on 3D Web Technology (Paderbon, Germany). Web3D '01. ACM, New York, NY, 51-59.

Darken, R. P. & Sibert, J. L. (1993). "A Toolset for Navigation in Virtual Environments," *Proceedings of the 6th Annual ACM Symposium on User Interface Software and Technology,* 157-165.

Dourish, P. & Chalmers, M. (1994). "Running Out of Space: Models of Information Navigation," *Proceedings of HCI 94.*

Einsfeld, K., Agne, S., Deller, M., Ebert, A., Klein, B. & Reuschling, C. (2006). "Dynamic Visualization and Navigation of Semantic Virtual Environments," Proceedings of the Conference on Information Visualization, 569-574.

Erickson, T. & Kellogg, W. A. (2000). "Social Translucence: An Approach to Designing Systems That Support Social Processes," *ACM Transacions on Computer-Human Interaction,* 7(1), 59-83.

Gabbard, J. L. (1997). "A Taxonomy of Usability Characteristics in Virtual Environments," *Msc Thesis, Virginia Polytechnic Institute and State University.*

Grammenos, D., Mourouzis, A. & Stephanidis, C. (2006). "Virtual Prints: Augmenting Virtual Environments with Interactive Personal Marks," *International Journal of Human-Computer Studies,* 64(3), 221-239.

Herman, H., Coombe, R. J. & Kaye, L. (2006). "Your Second Life?," *Cultural Studies* 20(2-3),184- 210.

Idehen, K. (2011). "Introducing Openlink Virtuoso: Universal Data Access without Boundaries," *White paper,* [Online], [Retrieved July 15 2011], http://www.openlinksw.com/.

Jeffrey, P. & Mark, G. (2003). "Navigating the Virtual Landscape: Co-Ordinating the Shared Use of Space," Designing Information Spaces: The Social Navigation Approach, K. Höök, D. Benyon, A. J. Munro (Eds), *Springer,* 105–124.

Jung, T. & Do, E. Y. (2000). "Immersive Redliner: Collaborative Design in CyberSpace," *Proceedings of ACADIA 2000, Associations of Computer Aided Design in Architecture National Conf.,* 185–194.

Kleinermann, F., De Troyer, O., Creelle, C. & Pellens, B. (2007). "Adding Semantic Annotations, Navigation Paths and Tour Guides to Existing Virtual Environments," Proceedings of the 13th International Conference on Virtual Systems and Multimedia (VSMM'07), 50-62.

Lepouras, G. & Vassilakis, C. (2005). "Virtual Museums for All: Employing Game Technology for Edutainment," *Virtual Reality,* 96-106.

Polys, N. F. & Bowman, D. A. (2004). "Design and Display of Enhancing Information in Desktop Information-Rich Virtual Environments: Challenges and Techniques," *Virtual Reality,* 8, 41-54.

Prasolova-Førland, E., Sourin, A. & Sourina, O. (2006). "Cybercampuses: Design Issues and Future Directions," *The Visual Computer,* 22(12), 1015-1028.

Slater, M., Sadagic, A., Usoh, M. & Schroeder, R. (2000). "Small-Group Behavior in a Virtual and Real Environment: A Comparative Study," *Presence: Teleoperators & Virtual Environments* 9(1), 37-51.

Witmer, B. G. & Singer, M. J. (1998). "Measuring Presence in Virtual Environments: A Presence Questionnaire," *Presence: Teleoperators & Virtual Environments* 7(3), 225-240.

Zahorik, P. & Jenison, R. L. (1998). "Presence as Being-in-the-World," *Presence: Teleoperators & Virtual Environments* 7(1), 78-89.

[i] User Metrics for Second Life: http://secondlife.com/xmlhttp/secondlife.php, [accessed: 4/1/2011]

[ii] Linked data: http://www.w3.org/DesignIssues/LinkedData.html, [accessed: 13.7.2011]

[iii] Resource Description Framework – RDF: www.w3.org/RDF/, [accessed: 4/7/2011]

[iv] SPARQL: www.w3.org/TR/rdf-sparql-query/, [accessed: 4/7/2011]

[v] N3: www.w3.org/DesignIssues/Notation3, [accessed: 4/7/2011]

[vi] LOD cloud: http://lod-cloud.net, [accessed: 4/1/2011]

[vii] DBpedia: www.dbpedia.org, [accessed: 4/7/2011]

[viii] DBpedia dumps: http://wiki.dbpedia.org/Downloads36, [accessed: 4/7/2011]

[ix] RDFlib: http://www.rdflib.net/ [accessed: 4/1/2011]