ADAPTIVE VIRTUAL REALITY MUSEUMS ON THE WEB

George Lepouras and Costas Vassilakis

ABSTRACT

This chapter presents an architecture for supporting the creation of adaptive virtual reality museums on the web. It argues whether the task of developing adaptive virtual reality museums is a complex one, presenting key challenges, and should thus be facilitated by means of a supporting architecture and relevant tools. The proposed architecture is flexible enough to cater for a variety of user needs, and modular promoting extensibility, maintainability and tailorability. Adoption of this architecture will greatly simplify the development of adaptive virtual reality museums, reducing the needed effort to exhibit digitisation and user profile specification; user profiles are further refined dynamically through the user data recorder and the user modelling engine, which provide input for the virtual environment generator.

Keywords: Virtual reality, Interactive technology, Multimedia database, Web architecture

INTRODUCTION

Museums have long been regarded as keepers and preservers of artefacts and cultural products. However, the notion of a museum with a primary goal of preserving has been changing during the past years and is being replaced by one that couples education and entertainment. Under this view a museum is an institution open to the world with an objective of aiding visitors learn while they are kept content. To this end, the Internet and especially the web has offered museums the medium to open up to the public and reach a wider, ever-increasing audience.

The technologies underlying the web provide a strong background for building online multimedia applications that can help museums attract Internet visitors. Towards this objective, one key aspect in developing a successful web application is the ability to supply the proper information for the targeted user group. This is especially true for web-based applications where the users visiting may come from a variety of cultures, educational backgrounds, ages and have different preferences and objectives. This requirement can be met by implementing an application that will adapt to the user's profile and each time an Internet user visits the web site, it will provide the appropriate set of information in the most efficient way.

Virtual reality technologies on the other hand have been evolving during the past years, to leave from research laboratories and to find application in a number of areas. Virtual reality promises the creation of environments that are vivid, life-like and highly interactive and where the user will be able to emerge in a synthetic world that may non-existing or it may be too difficult or too dangerous to visit in a real world situation. In this respect, virtual reality technologies may find direct applications in the museums, providing memorable experiences by helping users visualise and interact with exhibits.

The objective of this chapter is to specify the architecture of a system that combines the benefits of the aforementioned technologies in delivering adaptive virtual museums on the web. Such a system may be used for both edutainment and research activities for a variety of potential target groups.

BACKGROUND

The term Virtual Museum was coined by Tsichritzis and Gibbs (1991), where they describe the concept of a virtual museum and the technologies needed to realise it. However, in the past years the term virtual museum has come to denote anything from a simple multimedia presentation of selected museum content to a high-end, state of the art installation, with 3D projection facilities where the user immerses in a virtual environment. In the context of this work we will use the term Virtual Reality Museum to refer to a virtual environment build with 3D technologies, not necessarily immersive, but one where the user is able to navigate in a three dimensional exhibition. To this end, an adaptive Virtual Reality Museum, denotes a dynamically, custom built environment that fits the user's preferences as well as her cultural and educational background. An adaptive virtual environment provides each user with a different view of itself, taking into account the user's profile. A virtual reality museum comprises of the structure of the virtual museum building, the objects that are placed and exhibited within as well as the interaction methods by which the user can navigate and interact with the objects. All of the aspects of the environment may be adjusted to better meet the user's preferences and profile:

1. The structure of the 3D space, i.e. the halls comprising the virtual museum, their interconnection paths (corridors, teleports etc) and the exhibit placement in each hall.
2. The exhibits that are available for viewing and the resources used for their presentation (audio, video, documents). For example, if the user's profile is one of a primary education student, the environment should change the presentation of exhibits, documenting texts, and even hide some exhibits if they are not suitable for display to children.
3. The interaction methods available to the user, both for navigation within the virtual environment and for the manipulation of the exhibits. For instance, users with little computer experience may only be allowed to walk around the exhibits, while more experienced users can be presented with the option to "grab" and rotate or move the exhibits. Alternatively, the same options may be available with varying degrees of complexity: for example an elementary school student may be allowed to disassemble some complex machinery by simply selecting the appropriate spot, while a trainee mechanic should perform the same task by precisely following the "real-world" procedure.

Information that will enable the system to appropriately adjust the content can either be directly provided by the user (e.g. "I am using a 56.6K modem", "I am a researcher"), or be inferred by the system from the user's interaction pattern (e.g. resources requested insofar, time needed for the user to download content, etc). If during her interaction, the user shows more interest to specific exhibits than others, the system should take this preference into consideration when constructing the next museum hall.

Virtual environments enable the user's immersion in a synthetic world and can provide a vivid, life-like experience. Virtual environments have found applications in a number of different areas: Strickland et al. (1997) describe a virtual environment for the cure of phobias, Alison et al. (1997) a virtual environment that lets students assume the persona of an adolescent gorilla and interact as part of a gorilla family unit, while Charitos et al. (2000) describe an environment used for aiding the organisation of autistic children behaviour in everyday tasks. In the case of museums, virtual reality systems can help visitors visualise sites that may have been destroyed or they cannot visit, interact with exhibits they would not be allowed to view closely in a real world museum or simply provide a more vivid, animated and entertaining presentation. Locally executed virtual reality systems can be found in museums and institutions such as the Cave at the Foundation of Hellenic World (2002) and the Tokyo National Museum (2000). Other museums provide through their web-site simple, static 3D representations of spaces or of artefacts. For example, the Getty museum (1999) presents a model of the forum built by the Emperor Trajan, while the Natural History Museum (2002), presents a virtual environment created on the plans of James Cook's ship, the Endeavour. On the other hand the Canadian Museum of Civilization and the Institute for Information Technology at the National Research Council of Canada (NRC) (1999) produced a virtual museum of Inuit culture, while the Natural History Museum (2003) displays a gallery of virtual objects.

In all these cases a variety of technologies has been employed to implement the virtual environment and the 3D artefacts. Technologies and languages such as QTVR (Apple Inc., 2000) Java 3D (Sun Microsystems, 2001) or VRML (Web 3D Consortium, 1997), have been extensively used to this end. Of them, VRML, which stands for Virtual Reality Mark-up Language, has become the standard for presenting virtual reality content on the web. Using VRML one can create static as well as animated dynamic 3D and multimedia objects, and even link them to media such as text, sounds, movies, and images, creating a rich information space.

A web-based, adaptive, virtual reality museum offers a number of benefits and at the same time creates a number of challenges that have to be tackled during development time. Both benefits and challenges stem from the combination of the involved technologies. The web offers the potential for museums to communicate their message to a wide audience, which they could not reach before. The adaptability provides the means for creating user-dependent, personalised applications, while virtual reality technologies promise a more vivid, life-like experience. Web technologies support the requirements for the creation of adaptive applications. Such systems are dynamic in their nature because it is not efficient to a-priory create multiple instances of the same system, one for each user or user profile that may visit the web site.

O'Donell et al. (2001) present ILEX a dynamic hypertext system that arose from considering the problem of providing labels for exhibits in a museum gallery. Although the authors do not directly address the issue of a VR museum, their work addresses the generation of contextually relevant descriptions for museums exhibits. Perteli et al. (1999) describe aspects of monitoring visitor’s behaviour to detect emotions in an augmented reality museum. Marcucci and Paternò (2002) introduce an approach based on user modelling techniques to providing intelligent support through different devices to organise and make a generic sightseeing visit involving indoor (a Museum) and outdoor areas (the city historical centre). Finally, Brusilovsky (2003) describes the notion of adaptive navigation and attempts to build a case for systems that are able to adapt the very adaptation technology to the given user and context.

CHALLENGES

Although the benefits outlined in the previous section support the adoption of adaptive, web-based virtual environments for the presentation needs of museums, so far there have not been any examples of such systems. This reluctance can be attributed to a number of reasons.

Creation of Content

Developing an adaptive virtual environment is a far more complex task than that of implementing an equivalent hypermedia application. To some extent, this can be anticipated, since the creation of both the content and the interaction methods can be more complex in 3D than in 2D. Creation of the content for a static virtual environment is already a cumbersome and time-consuming process (Lepouras et al. 2001). The virtual space has to be designed and developed, objects to populate the space have to be modelled or digitised and inserted, and presentation issues such as lighting have to be taken care of. The designer has also to specify and design the actions a user will be able to perform in the virtual environment, for example whether which objects will be interactive, how the user will be able to perceive what actions can be performed on specific objects, etc.

Complexity of the system's architecture

The requirement for an adaptive environment poses extra complexity in this undertaking. The environment has to be created dynamically based on the users preferences and it has also to monitor and interpret the user's navigation and interaction patterns. To this end, a system has to be built that will create the virtual environment according to the user profile and preferences, transmit the virtual environment to the user, monitor and interpret the user's actions and adapt the virtual environment to it.

Since this environment will be transmitted to the user through the web, a few more restrictions have to be addressed by the system. The generated environment has to be optimised to minimise transmission time. Furthermore, the requirements set for the user's computer specifications should not be too demanding, otherwise there may exist users who will not be able to view the virtual environment properly or even at all. This implies that the processing and interpreting of the user's actions should not take place in the client system but on the server side.

Conveying the museum's message

One distinguishing factor between virtual reality museums and other virtual environments is the fact that a museum conveys a message. The message a museum conveys to its visitors can be greatly affected not only by the selected exhibits, but also by their presentation. Changing the presentation order of the exhibits, the lighting or just the positioning of an artefact can completely alter the message perceived by the visitors. If the virtual museum is a static one, the designer can decide during development time the optimum presentation method and layout. However, in an adaptive virtual museum, this role has to be carried out by the system that generates the virtual environment, a task that can be difficult (Sarini, M. & Strapparava, C., 1998). For example, if one hundred or more exhibits satisfy the user's preferences and profile, the system will have to select only a limited number of them to construct the virtual museum, to facilitate transmission of the virtual space to the end user. The selected artefacts will have to have a meaningful grouping for the user and will have to be presented in such a way.

To illustrate the difference the layout of exhibits can make in the overall message communicated to the user one may think of a collection of ancient tools. By presenting tools of the same purpose in chronological order a museum's visitor may view the alterations of the same tool through time. By grouping the tools by purpose the visitor may view how differences in shapes altered the tool's usage. By grouping equivalent purpose tools of the same era by their origin a visitor can understand the differences in technology between the status of technology in different civilizations and even their influence to each other.

These challenges have to be successfully confronted in order to built a system that will carry out the process of constructing and dynamically updating the virtual environment.

RELATED WORK

Chittaro and Ranon (2002) present a generic architecture for implementing a system that can generate adaptive 3D web sites. This architecture is based on five basic components as illustrated in the next figure.

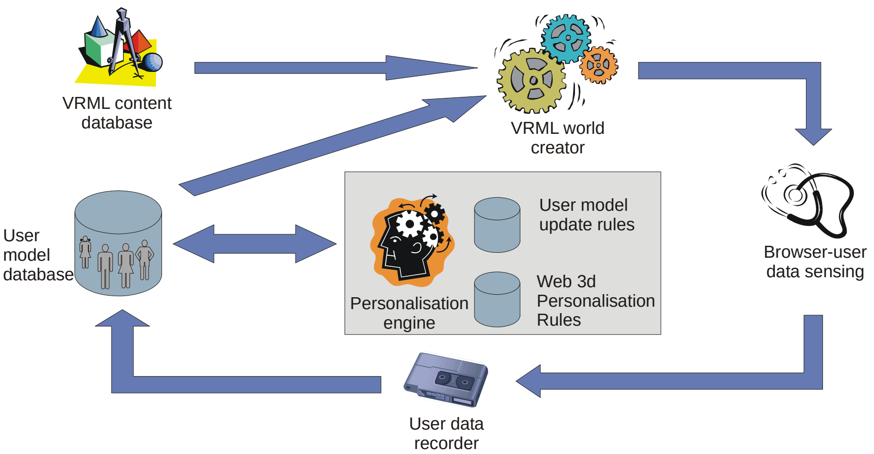

Figure 1 Generic architecture for an adaptive virtual environment creation system

Key component in this generic architecture is the *Personalization component*, which updates the user models and chooses the personalisation that should be made. The personalisation component is comprised of the *Personalisation Engine*, the *User Model Update Rules* and the *Web3D Personalisation Rules*. The personalisation engine is based on the other two sub-components. The User Model Update Rules perform deductions in order to update the user model, while the Web 3D Personalisation Rules choose what personalisation should be carried out, depending on the model of the user. Choices made by the personalisation component are consequently stored in the *User Model Database* and can be updated at predetermined time intervals or after a number of visits to the web site. The user model database contains models of users and updates these models from the data received by the personalisation component and by the *Usage Data Recorder*. The user data recorder has a *User Data Sensing* component that resides on the user's browser, monitors users actions, does some initial processing and sends the data first to the user data recorder and then to the user model database. The *user model database* is queried by the *VRML World Creator* to retrieve personalisation choices along with VRML objects from the *VRML Content database*, in order to formulate the virtual reality environment and transmit it to the user.

In the aforementioned publication this generic architecture is examined in the context of a 3D e-commerce application. The case study discusses a 3D virtual store where some of the environment's features such as store size, style, product display, etc. are adapted to the customer's profile. However, the construction of a virtual museum is a more complex task. As previously noted, the number of potential exhibits can be very large, and the virtual environment has to accommodate these exhibits in a meaningful distribution. These exhibits can be complemented with extra resources such as text, audio or video. Furthermore, digitised exhibits tend to be large in size and their presentation in 3D usually requires more resources than the display of a simple box on a self. This necessitates the partitioning of space in smaller spaces to reduce download time and to ease the navigation of users. Additionally, if the creation of a new virtual museum hall depends on the user's interaction patterns in halls the user visited during the same session, the virtual environment creation engine has to update the user profile during the user's visit and this demands the adoption of an efficient algorithm for the selection of exhibits. These considerations force the adoption of a differentiated architecture that will fully support the creation online, adaptive, Virtual Reality museums.

TOWARDS AN ADAPTIVE VR MUSEUM

The architecture proposed in this paper elaborates and extends the generic architecture presented in the previous section. In our approach we adopt a holistic view of the system that generates adaptive virtual worlds; the proposed system architecture is illustrated in Figure 2. In such a system, the two main groups involved are designers and end users. Designers can be domain experts or they work closely with domain experts in order to prepare usage scenarios for the virtual environment. They define target groups, assign usage scenarios to them and define the usage data to be collected. An important task performed by designers is the assignment of *semantic information* to each resource within the system. Semantic information is structured in the form of *properties* and their corresponding *values* and may be used in querying the resource database to extract elements meeting specific criteria.
Designers are also given the capability to provide multiple versions of the same resource, which are again associated with semantic information. This information is read by the virtual environment generator and combined with the user preferences extracted from the user model database, in order to select the most appropriate resource version to include in the virtual world, depending on the current user profile. In order to perform these tasks, the designers are equipped with appropriate administration tools that are used to populate, query and maintain the VRML content database and the user model database.
Another extension in the proposed architecture is that the user modelling engine consults not only the user model database, but also the VRML content database. This allows the user modelling engine to use the semantic information stored together with the resources in the VRML content database, enabling the derivation of user preferences in a higher level of abstraction; for example, in an archaeological museum the semantic information would

enable the system to deduce that the visitor shows interest in Egyptian antiquities dated within the reign of Ramses, provided that the designers have associated to the resources the properties “origin” and “chronology”. It should be noted here that the user modelling engine automatically takes into account all semantic information associated with the exhibits, and there is no need for designers to perform any extra work in order to support this functionality. The user modelling was based on the paradigm described in (Oberlander, J., et al. 1998) and (Not E. et al. 1998). The Veron and Levasseur (1983) categorisation was adopted in order to help classify users’ behaviours and adapt the virtual museum to their need. This categorisation was used to identify users and select not only the most appropriate resources for the exhibits presented but also the most adequate interaction technique. For example, for visitors who prefer to move in the centre of the room and do not look at details of artwork but prefer to have a more holistic observation, there is not a need to include scripts for the rotation and manipulation of certain exhibits. On the other hand, if a visitor prefers to follow the path proposed by the curator, the virtual environment generator may select spaces with clearly predetermined paths that afford this.

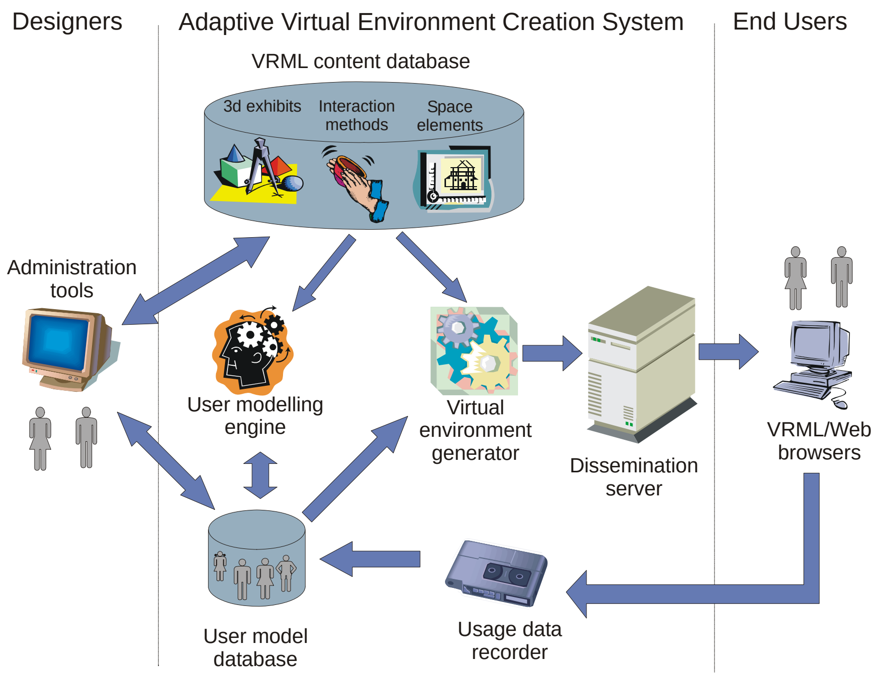

Figure 2. System architecture

Finally, the proposed approach enables the designers to provide predefined templates for virtual spaces, objects and interaction techniques, extending thus the supported dimensions of adaptability to include not only the presented items, but also the surrounding environment and the methods that are available to the users for their interaction with the virtual world. The *virtual environment generator* component encompasses a parser for these templates, which according to the production rules and the user's profile and preferences, fills in appropriate placeholders in the templates with data from the *VRML content database*. A brief description of the different modules of the proposed architecture is provided in the following paragraphs.

VRML CONTENT DATABASE

The *VRML content database* stores all the resources that are available for the creation of the virtual environments. It holds parts that are domain-dependent such as the digital representations of the exhibits, and some parts that are more or less independent of the domain such as the library of interaction methods and the digital representations of exhibition spaces. Both the library of interaction methods and the digital representations of exhibition spaces are only indirectly dependent on the domain, in the sense that they can be associated to exhibits.

In more detail, the *VRML content database* encompasses:

- *Exhibit digital representations.* Each exhibit may be represented by a variety of methods, such as 3D photographs, 3D models, 2D photographs, videos, animation, audio clips texts etc. The most suitable means for representing each exhibit is determined during the design phase by museum curators, taking into account the exhibit nature, the intended audience and the required functionality. An exhibit representation may include resources of different kinds (e.g. 3D models *and* videos) and multiple resources of the same type (e.g. multiple descriptive texts with varying detail, multiple 3D models with different levels of detail etc). Resource multiplicity is of essence in an adaptive environment, since the virtual world generation engine has the possibility to select the resource that matches best the user profile and preferences and thus present a different world to each visitor. For example, brief descriptive texts may be more appropriate for the ordinary visitor, whereas long, detailed texts are more suitable for researchers.
- *Interaction method library*. The VRML content database contains a number of pre-defined interaction methods that can be dynamically associated with exhibits, in order to provide the visitors with the ability for exhibit manipulation. The interaction methods have been decoupled from the exhibit digital representations, since the same exhibit may be required to be associated with different interaction methods, depending on the user requests, the user profile etc. Since interaction methods are inherently dependent on the type of the digital representation (e.g. a *rotation* method applies to 3D photographs and 3D models but not to audio clips), each interaction

method is tagged with appropriate information that allows it to be combined with the appropriate representation types.

- *Exhibition space element digital representations.* Complementary to the exhibits, the VRML content database includes digital representations of the elements required by the *Virtual Environment Generator* to formulate the overall virtual environment that is presented to the visitor. Exhibition space elements include halls, corridors, foyers, display cases, textures etc. The exhibition space elements may contain *item placeholders*, which are filled-in with appropriate resources when a specific virtual environment is instantiated. For example, an exhibition hall may contain exhibit placeholders where exhibits or showcases will be positioned, whereas a corridor may contain sign placeholders where text providing navigational directions will be placed.

Visitors should be able to either browse and/or search the VRML database to locate exhibit resources that are of interest to them. Browsing may be facilitated via suitable categorisation schemata, which allow visitors to drill down into exhibit groups having common properties. Searching may be enabled by associating each exhibit with a number of *attributes*, assigning a specific *value* to each of these attributes and allowing visitors to request the exhibits that have specific properties set to specific values. For instance, visitors in an archaeological museum may search for the exhibits that have the attribute *civilization* set to *Egyptian* and the property *usage* set to *religious*.

Another requirement for the VRML content database is the ability to store multilingual versions of language-dependent resources, such as descriptive texts and narrative audio clips. The *Virtual Environment Generator* will then select and embed in the generated virtual environments the resource versions that match most closely the user profile. This is of essence in the context of web-accessible environments, where visitors have different national and cultural backgrounds. The VRML content database schema should also be extensible, to allow for the incorporation of new resource types whenever this is required. A schema that meets the requirements listed above is presented in (Charitos et al., 2000) and (Lepouras et al., 2003).

USER MODEL DATABASE

The VRML content database hosts the resources modelling the virtual world content, and would by itself be sufficient to support the generation of the virtual environment. In order, however, to support adaptivity, the virtual world content resources should be complemented with information regarding the user profiles and bindings between the user profiles and specific resources or resource properties. This information is exploited by the Virtual Environment Generator to select the exhibit resources that are most suited to the current profile of the museum visitor and thus support the dimension of adaptivity. In more detail, the user model database contains the following information:

- *Static user profile data,* i.e. data that characterise invariable (in the context of interaction) aspects of the user, such as the user's native language, the educational background, the purpose of the visit (e.g. education, recreation, research), the Internet connection bandwidth etc. Static user profile data may either be expressly declared by the user, or inferred by the system through appropriate measurements (e.g. the time needed to download a specific file provides clues for the connection bandwidth).
- *Dynamic user profile data,* i.e. data that reflect the user's behaviour and preferences that may change during the interaction. These include the history of exhibits presented insofar, the preference towards certain media types or interaction methods, the interest shown for various exhibit categories and so on. Dynamic user profile data are collected through the *usage data recorder* (presented in the next section). The collected data are essentially detailed information regarding the user's interaction with the virtual environment (e.g. time spent viewing each exhibit) thus an additional

module, the *personalization engine* is introduced that retrieves the detailed data and employs rules to construct higher-level information regarding the user's profile. The personalization engine is discussed in more detail in the following paragraphs.

The user model database contains a number of predefined generic user models, in order to facilitate quick visitor configuration (visitors only describe generic aspects of their profile, rather than entering tedious details); the profile selected by each visitor is subsequently refined to reflect individual preferences and needs through the dynamic collection of data pertaining to the user's behaviour. These user models additionally serve the purpose of efficiency, by allowing parameters that apply to multiple visitors to be represented using a single database entry. This effectively reduces the data volume that must be stored and manipulated within the user model database and consequently leads to better performance.

USAGE DATA RECORDER – USER MODELLING ENGINE

The usage data recorder module is responsible for collecting and storing data regarding the interaction of the user with the virtual environment. This data is subsequently exploited by other modules to support system adaptivity. The usage data recorder comprises of two parts, namely the *data collector agent* and the *usage data recorder server.* The data collector agent is a library of *event hooks* that are packed into the virtual worlds sent to the museum visitors. These event hooks are triggered when certain actions are performed by the visitor, such as acquiring or losing visibility to an exhibit, moving close to an exhibit or moving away from it, beginning and ending of interaction with an exhibit, request for a specific resource type etc. When an event hook is triggered a *usage data record* is formulated, which includes the event type, the virtual environment item to which the event pertains (mainly exhibits and exhibit resources), the associated timestamp and the user identifier.

The usage data records are sent to the adaptive virtual environment, where they are intercepted by the *usage data recorder server* module. Transmission of usage data records from the data collector agent to the *usage data recorder server* is performed in the background while the user navigates within the virtual environment (during this time the communication line is usually idle, thus this approach minimises the interference with user actions) and takes place when a certain number of usage data records has been amassed. Upon reception of a new usage data record batch, the usage data recorder extracts the individual records and arranges for their storage into the user model database. After this step, the usage data recorder invokes the user modelling engine, which is responsible for analysing the low-level data collected by the usage data recorder to infer a more detailed profile of the user's behaviour and preferences. The first step in the inference procedure is the combination of "event beginning-event end" records, to determine the duration of the events, such as the viewing of an exhibit or the interaction with it. This step is only performed for events to which duration is applicable; for instantaneous events (e.g. requesting a resource associated with a certain exhibit), only the number of such events is computed. Once the preparatory step has been completed, the following rules are employed to determine the user interests:

1. Resources that have been viewed or examined for a long duration are assigned a high interest rating (in the range 4 to 6).
2. Resources that have been viewed or examined for a medium duration receive an intermediate interest rating (in the range 1 to 3).
3. Resources that have been visible within the interaction but have not drawn the user's attention (the user has not moved close to the respective objects or has passed through them very quickly) receive a negative interest rating (in the range –3 to –1).
4. No interest characterisation is associated with resources that have not come into visibility scope in the current interaction time window, since the user may not even be aware of their presence.

5. Resources explicitly requested for (e.g. descriptive texts, enlarged photographs etc) are assigned a high interest rating (in the range 4 to 6).

When all interest ratings have been assigned, the VRML content database is queried to extract the property and corresponding value combinations of the exhibits to which rated resources pertain. Then, for every distinct property/value combination retrieved, the interest ratings of the resources related to it are summed up (a property/value combination is considered to be related to a resource if the exhibit to which the resource pertains is tagged with the specific property/value combination) to compute the *interest measure* of the property/value combinations. The user model database is finally updated to include the newly derived information regarding the user preferences. Similar algorithms are employed for determining the user preferences towards media types and interaction methods.

VIRTUAL ENVIRONMENT GENERATOR

The virtual environment generator receives requests for the creation of virtual environments that will be sent to the users. Upon reception of a request the virtual environment extracts from the request the user identity and any request parameters denoting specific user preferences. Then, the user model database is queried to retrieve previously stored express or inferred user preferences, which are merged with those extracted from the request. If contradicting preferences exist in the request and the user model database, then the preferences stated in the request take precedence over ones fetched from the user model database, since the former have been explicitly stated by the user and constitute thus "hard" requirements, overriding any system-deduced preferences or even preferences previously stated by the user. For example, if a researcher accessing the virtual museum using a low-speed modem requests to view the high-resolution digital representation of a statue, the request will be honoured, despite the presence of a generic rule in the user profile database stating that low-resolution representations are more suitable for the specific connection type.

It should be noted here that the virtual environment generator always obeys restrictions placed by content administrators, regarding the suitability of certain resources for each user or user category. Thus, if a visitor profiled as “high school student” explicitly requests to view a specific forensic science museum exhibit that has been characterised as “available only to adults”, the request will be rejected.

Once the virtual environment generator has collected all the preferences and restrictions that apply to the request, it queries the VRML content database to retrieve the elements that will be used to create the virtual environment that will be sent to the user. Firstly, the exhibit resources that best match the preferences and restrictions are retrieved. The virtual world generator uses upper bounds both in terms of the download size of the virtual world and of the number of artefacts that will be presented, thus the matching resource list is appropriately pruned to comply with these limits. During the resource selection the history of exhibits viewed insofar is also consulted, in order to avoid the repetitive presentation of the same resources. After the resources that will be presented have been determined, for each one of them the most appropriate interaction methods are selected, based on the resource type and the user profile. For a more technical description of an algorithm for a virtual environment generator the reader may view (Lepouras, 2004). Then, the exhibit physical dimensions are examined, so as to calculate the respective dimensions of the showcases and exhibition halls and, finally, the exhibition hall templates are retrieved from the VRML content database, scaled to the dimensions computed in the previous step and the placeholders are filled in with the pertinent resources. Once this step is complete, the virtual environment has been fully constructed and is sent to the visitor.

FUTURE TRENDS

The architectural approach described in this chapter although focussed on the museums application area, is generic enough to be used as a basis for implementing systems in other application areas such as e-commerce and e-training. The vividness of the virtual reality coupled with the personalisation capabilities of adaptivity technologies can aid in delivering e-training environments that are easy to use, and will hold the trainees' interest throughout the session. In the context of the e-commerce application area, users will benefit from the adaptivity since they will be able to easily locate objects of interest, while the virtual reality will provide a more accurate and tangible reproduction of the original item.

The adaptability described herein can be characterised as *lazy*, in the sense that there exists a certain latency between the user's interaction and the system's reaction. In the best case the system will take into account the user's preferences when the user requests a new virtual environment, therefore while the user navigates in a virtual world no further adaptability to the user's interaction can be anticipated. In order to provide a more *eager* form of adaptability one has to investigate which parts of the system's adaptability functions can migrate to the client side, allowing for immediate response to the user's preferences without the need for creating and transmitting a new environment. The behavioural enhancements to VRML introduced by the emerging X3D standard (Web3D Consortium, 2002) will be considered in relation to this task.

Further research opportunities include more elaborate algorithms for resource selection, with the possible use of a knowledge-based system. These algorithms should be easily configurable by designers and domain experts, and should produce meaningful groupings of objects without sacrificing performance.

CONCLUSION

New technologies create possibilities for the appreciation of museum content to a larger audience. While the number of challenges in the application of these technologies is still considerable, their potential justifies the effort required for their application. These technologies as previously described can formulate a generic framework, applicable not only to museums or cultural institutions, but in any thematic area where a stimulating environment for learning, training, conducting research, collaborating or even shopping.

REFERENCES

Tsichritzis, D., Gibbs, S. (1991) Virtual Museums and Virtual Realities. Proceedings of the International Conference on Hypermedia and Interactivity in Museums, Pittsburgh.

Strickland, D., Larry Hodges, L., North, M., Weghorst, S. (1997) Overcoming phobias by virtual exposure. Communications of the ACM, Vol. 40, No. 8, pp. 34 – 39.

Allison, D., Wills B., Bowman, D., Wineman, J., and Hodges L. F. (1997) The Virtual Reality Gorilla Exhibit. IEEE Computer Graphics and Applications, November/December

Charitos, D., Karadanos, G., Sereti, E., Triantafillou, S., Koukouvinou, S., Martakos, D. (2000). Employing virtual reality for aiding the organisation of autistic children behaviour in everyday tasks. Proceedings of ICDVRAT 2000, Alghero.

The Foundation of the Hellenic World (2002). Virtual reconstruction of ancient Miletus. http://www.ime.gr [Last accessed 30th October 2003]

Tokyo National Museum (2000). Recreation of the Toshodai-ji temple as a large-scale, high-resolution virtual environment. http://www.sgi.com/features/2000/april/ toppan/index.html [Last accessed: February 2, 2004]

Getty Museum -Forum of Trajan (1999). A virtual reconstruction of the destroyed forum. Accessible at: http://www.getty.edu/artsednet/Exhibitions/Trajan/ [Last accessed 30th October 2003]

Natural History Museum (2002). 1768: the voyage of the Endeavour. Accessible at http://www.nhm.ac.uk/museum/tempexhib/voyages/endeavour.html [Last accessed 30th October 2003]

Canadian Museum of Civilization (1999). The Inuit 3D museum. Accessible at http://www.civilization.ca/aborig/inuit3d/vmcinuit_e.html [Last accessed 30th October 2003]

Natural History Museum (2003). Gallery of Virtual Objects. Accessible at: http://www.nhm.ac.uk/museum/vr/ [Last accessed 30th October 2003]

Apple Inc. (2000). QuickTime VR environment. For more information: http://www.apple.com/quicktime/qtvr/ [Last accessed 30th October 2003]

Sun Microsystems (2001). Java 3D Application Programming Interface. Accessible at http://java.sun.com/products/java-media/3D/ [Last accessed 30th October 2003]

Web 3D Consortium (1997). VRML97 Functional specification and VRML97 External Authoring Interface (EAI) International Standard ISO/IEC 14772-1:1997, ISO/IEC 14772-2:2002, available through http://www.web3d.org/fs_specifications.htm [Last accessed 30th October 2003]

O'Donnell, M., Mellish, C., Oberlander, J., Knott, A. (2001). ILEX: an architecture for a dynamic hypertext generation system. Natural Language Engineering, Vol. 7 (3), pp. 225-250.

Petrelli,D., Not, E., Zancanaro, M. (1999). Getting Engaged and Getting Tired: What Is in a Museum Experience. UM'99, Workshop on 'Attitude, Personality and Emotions in User-Adapted Interaction', Banff, June 1999, available at http://aos2.uniba.it:8080/pr-ws-um.html, [Last accessed 9 Feb. 2004]

Marucci, L., Paternò F. (2002). Helping users through ubiquitous personalised, interactive support in a sightseeing visit. Proceedings ECCE11 - Eleventh European Conference on Cognitive Ergonomics, Catania, 2002, pp.245-250.

Brusilovsky, P. (2003), Adaptive navigation support in educational hypermedia: the role of student knowledge level and the case for meta-adaptation, *British Journal of Educational Technology*, Vol. 34 (4) pp 487-497.

Lepouras, G., Charitos, D., Vassilakis, C., Charissi, A., Halatsi, L. Building a VR - Museum in a Museum. Third International Virtual Reality Conference, VRIC2001, Laval, France, May 16-20, 2001.

Sarini, M. & Strapparava, C. (1998). Building a User Model for a Museum Exploration and Information-Providing Adaptive System. Proceedings of the 2nd Workshop on Adaptive Hypertext and Hypermedia, HYPERTEXT'98, Pittsburgh, USA, June 20-24, 1998.

[20] Chittaro L. & Ranon R. (2002). Dynamic Generation of Dynamic VRML Content: a General Approach and its Application to 3D E-Commerce. Proceedings of Web3D Conference, pp. 145-154.

Oberlander, J., O'Donnell, M., Mellish, C., and Knott, A. (1998). Conversation in the museum: experiments in dynamic hypermedia with the intelligent labelling explorer. The New Review of Hypermedia and Multimedia, Vol. 4, 1998 pp 11 – 32.

Not, E., Petrelli, D., Sarini, M., Stock, O., Strapparava, C., and Zancanaro, M. (1998). Hypernavigation in the physical space: adapting presentations to the user and to the situational context. The New Review of Hypermedia and Multimedia, Vol. 4, 1998 pp 33 – 46

Veron, E., Levasseur, M. (1983) Ethnographie de l’exposition. Paris, Bibliothèque publique d' Information, Centre Georges Pompidou, 1983.

Charitos, D., Lepouras, G., Vassilakis, C., Katifori, V., Halatsi, L. (2000) An Approach to Designing and Implementing Virtual Museums. Proceedings of the Seventh UK VR-SIG Conference, Glasgow, September 19-20.

Lepouras, G., Katifori, V., Vassilakis, C., Charissi A. (2003) Facilitating VR Museums Web Presence. Proceedings of Human Computer Interaction International 2003 Conference.

Lepouras, G. (2004). Applying Clustering Algorithms to Web-based Adaptive Virtual Environments. To appear in Journal of Computational Methods in Sciences and Engineering

Web3D Consortium (2002). Extensible 3D (X3D) Graphics. Accessible at

http://web3d.org/x3d.html [Last accessed 30th October 2003]

