A Collaborative Virtual Environment for Public Consultation in the Urban Planning Process

Nicolas Gaborit
University of Manchester
gaboritn@cs.man.ac.uk

Toby Howard
University of Manchester
toby@cs.man.ac.uk

Abstract

This paper extends research in Collaborative Virtual Environments (CVEs) by exploring their use in a purposeful interactive public application. Though often discussed, few such studies have been undertaken as the mass availability of 3D platforms is a recent phenomenon. The application selected is that of Urban Design and review. The experiments use a virtual cityscape in which planning changes are presented to the public at large. Respondents inhabit the shared environment and can engage in dialogue with other participants, modify and comment upon proposed designs.

We want to prove that the ready accessibility of the 3D presentation, together with the communication stimulus of the shared environment, will facilitate useful communication and engagement by the public in the planning processes. The research involves investigating the mechanisms of interaction, cooperation and design commentary within this process.

1. Introduction

The idea of Virtual Environments has been around for a long time now. A significant part of the proposed novelty that VEs offer is the 'natural' 3D interface which, ex-hypothesi, will be easier for untrained users to work with; the ability to share an environment with other users, and to experience and exchange information with them directly; and of course a natural vehicle for tasks that work with inherently three dimensional data. Evidence for the improved communication availed by VEs includes major commercial successes of large screen projection based VEs for communicating ideas in design review, and in sales and marketing. Firms as Shell or BMW are using this kind of technology.

However, comparatively few applications have been built that involve large numbers of people in a concrete task amenable to external evaluation, and so it is hard to gauge how effective VEs really are for these purposes. Examples of work that goes some way to this goal are Diamond Park [19], ITW (Inhabiting The Web) [21], Inhabited TV [20], alphaworlds [5], military simulations, and the work in multi-media art such as that of Knowbotic [22], and distributed computer games. However, there are no examples where Virtual Environments have been used as a natural forum for public concerns.

That so few examples exist is largely due to the immaturity of the software technology, together with the cost of 3D hardware capable of working with non-trivial environments. In the last years however, such hardware has gone from costing tens of thousands of pounds, to less than two hundred. Serious applications that exploit the widespread availability of such capabilities have yet to mature.

The targeted environments are large-scale shared virtual environment accessible to many users. As we want public to be consulted from their home computer, the environment has to be efficient on standard computers and user friendly.

In this paper, we will firstly review existing work in this area and described CVEs as tools for Urban Planning. Then we will describe the environment specifications and interactions within it, illustrated with a case study. And finally show accomplished work results with the implementation of a 3D cityscape and conclude.

2. Background

In this section we will review firstly technologies involved in Cityscape development using VR, and then focus on the main application of CVEs to public participation, virtual communities.

Many city models have been constructed using the VRML modeling language, and can readily be found on the web. Some of them are very accurate, as the city of Tokyo model [8] made by Planet 9 studios, one of the main companies of VRML cities design. Though often constructed to a high degree of accuracy, these models are quite limited. Firstly, they cover small parts of the city - this is a limitation of what can be viewed using a VRML equipped browser at acceptable frame rates. Secondly, interaction is extremely limited -- usually to passive navigation for a single user. The achievement of these models is the model itself; not what can be done with it. Finally, current VRML specifications appear to become limited when dealing with large scale cityscapes [9].

On the other hand, city models have been developed using specific VR software. For example, a cityscape has been developed in the University of Manchester for the citizen-oriented Escape Project [1]. UEA Norwich is also working in the area of building and rendering urban environments, in the areas of high--level parametric specification of buildings [10]. Other works concern the use of VR optimization techniques in order to display large cityscapes [11][12]. Finally, Image base rendering techniques has been used to create virtual cityscapes [13][14].

Finally, an important class of CVEs is that of the ``virtual communities'', interactive city-like VEs. A number of them use the Active World [5] technology, such as Alpha World. These are interesting as they actively seek to build a community of users, albeit without the pretension to represent a real cityscape, or a serious interaction with an urban planning process. Nonetheless, they are perhaps the best examples of large-scale communities of VE users engaged in a shared city-like landscape. The Cornell University of New York used the active world technology to build a 3D virtual museum shared environment. It is said [6]: "We are confident that it is possible to present data within the virtual world dynamically and to allow visitors to contribute content to the space". Thus, our idea to make users contribute to the model by adding information or propose new design ideas seems promising. However they experienced a "difficult learning curve". This means that an important part of the evaluation will focus on how people manage to interact with our environment.

3. CVEs as tools for Urban Planning

Having reviewed the existing urban planning practices, and relevant VE technologies, we now consider the role of CVEs in the Urban Planning process. The aim is to suggest specific application strategies for CVEs in Urban planning consultation with large scale public involvement.

The use of computer science technology has been proven to enhance public participation in the Urban Planning process [18]. Furthermore, G. Schmitt et al. say [15]: "The improved evaluation and visualization component of VR will lead to a wealth of new discoveries in architecture that could not occur otherwise." Thus, using computer science, and specifically VE technologies in the Urban Planning process seems relevant.

Urban designers work largely with 2D design. There are good reasons for this. It would be presumptuous to assert that 3D tools should replace existing practices. In the public mind however, cities are three dimensional edifices and ex-hypothesi, the public would expect a planning proposal to be presented in 3D and find such a presentation easier to comprehend and work with. Further evidence here is the common use of physical models in public planning consultations. One difficulty in this process is that such models are often presented as clean architects models, ``artists impressions'' whereas the existing situation is presented as unpleasant photographs of reality, with all that that entails (see the presentation of the Piccadilly scheme in Manchester as evidence of this). One advantage of the VE technique here is that of a neutral presentation for both the existing, and the proposed situations. Alternatives can be presented side-by-side, or overlaid upon the same model under user control. Members of the public could then more readily compare their direct experience of reality with the VE model of the existing, and the VE model of the proposed. Thus it is necessary to make available the contrast between the existing and proposed within the same VE presentation. Additional detailed information about proposals can be presented via auxiliary links, or embedded information spaces, as in the escape project [1].

Further to this is the idea of feedback and engagement from the public participants. A number of possibilities present themselves for exploration in the project. At a minimum level, members of the public would be able to offer commentaries - graffiti if you will - on the proposals. This should allow interaction with other participants. More advanced would be interaction with the model to suggest design alternatives. Beyond this there is the question of how all this information, the numerous feedback and modifications, are to be presented in a useful way to the designers, and how this fits into their work process. But all these are interesting problems to be addressed in the research.

4. Describing the virtual environment

The targeted environments are large-scale shared virtual environments. The environments we are studying here are model-based ones, on which people are consulted on design choices on an object (the city for example) which will be represented by a model. The model itself is not the only component of the environment. Indeed, in a public consultation process, the model is always shown with some information to guide the public (for example on an urban planning public consultation process, a map or a small scale model of the city will be shown with a key). So, this kind of environment is dual, as it contains a model and some associated information. We will describe now each of them.

4.1. The model

The model is subdivided into three layers, as shown by Figure 1. The logical layer contains the data structures, providing a logical description of the model structure. These structures form a hierarchy tree, to subdivide the model into, let us define, areas. The subdivision into areas is useful for public consultation, as urban projects can be bounded to a specific area. The physical layer deals with the description of the objects of model, and finally the 3D layer contains the geometrical representations of the objects. The distinction of the 3D layer from the physical layer allows the use of different 3D layers easily.

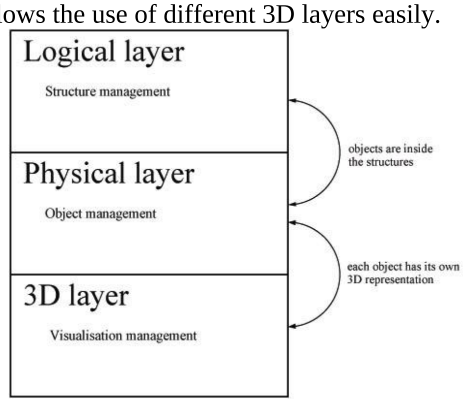

Figure 1. The model

To the point of view of a city model, the logical and physical layer can be seen as GIS (Geographic Information System) data, and the 3D layer can be seen as a 3D extension to GIS, as work from the Center for Advanced Spatial Analysis group of the University College London [2][3].

We could object about the utility of the logical layer. What is the interest in subdividing the model into "areas"? The answer comes from urban planning. Indeed, as urban projects can be bounded to a specific area, the utility of the subdivision becomes obvious. Thus, a consultation will depend on an area of the logical layer of the model, and separating this area from the others, it becomes possible to manage different interaction and display rules (for an example, restricting interaction with the model to the consultation area and use less complex 3D models to display the other areas).

4.2. Information

Information contains model information, which is brought by the ideas men of the environment, and consultation information, which is given by the people who are consulted. In this section we decide to merge consultation and model information, as they are technically the same information. Indeed, a piece of information brought by the public could have been considered as a piece of model information if the comment was left by an idea man.

Where is located information on the environment? It is obvious that a piece of information will be located on the model. For example, a description of an object can be considered as attached to this object. However some pieces of information are not necessarily directly linked to objects of the model. For instance a comment based on a particular view of the model (for example "in this landscape there is a lack of trees") will not be linked to a specific object of the model.

Thus, Figure 2 shows pieces of information which are directly located on the model, and other pieces of information which are independent from it. We define information/consultation boards the objects that contain these data. These objects can be considered just as information containers. They have two layers; the physical layer describing the object himself and the 3D layer its representation, as well as its localization in the 3D space.

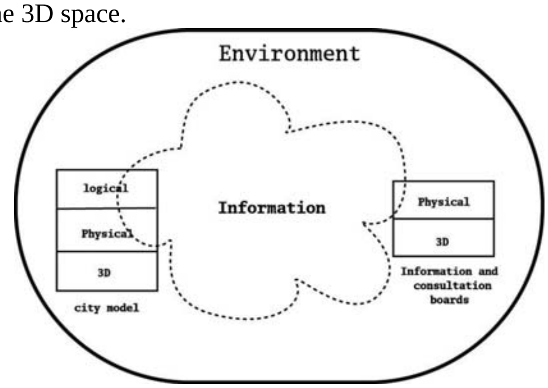

Figure 2. Information on the environment

Looking now at a lower level, shown by Figure 3, we can now classify information. Firstly, a piece of information is linked to the logical layer of the model. We can classify this as "general" information, as it does not refer to the details of the model provided by the physical layer. We call it logical information.

A second piece of information is linked to the physical layer. As opposite to the previous kind of information, we classify it as "local" information, as each piece of this information is related to a single object of the model. We call it physical information.

Finally, the last piece of information is linked to the boards we defined before. This piece of information can be considered as "spatial", as, not directly related to the model, depends mostly on the coordinates of these boards in the 3D space. We call it spatial information.

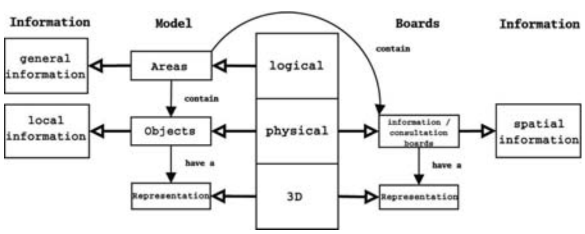

Figure 3. Information classification

We can notice a link between the boards and the areas. In fact it is necessary, as we saw that a consultation is dependant of an area in section 4.1. Thus each piece of information must be related to an area to be used during a consultation.

Spatial information would not be possible without expanding our environment to 3D. Thus we can see that VR permits to introduce a new kind of information, and so enrich the environment. Indeed, this information can depend on its localization. Here is the novelty aspect of this information, as, on the contrary to real small scale models to the city it becomes possible to have information about a group of objects based on their visualization from a specific viewpoint and thus public comments are now local and so more precise.

We observe now logical information. Firstly it cannot depend on the physical objects of the data structure on which the information is about. For instance, it can be for an area its name. We can call this information ex nihilo logical information. But logical information can depend on the physical objects contained by the data structure. As an example, it can be the number of objects of a specific type, or a mean of a numerical attribute of some objects (it can be for a city the average height of the buildings of an area). Thus this piece of information is generated by the physical information of these objects. We will call it generated logical information.

5. Populating the virtual environment

Having described the environment, we must study how people will experience it. Firstly we need to define which kind of user will populate the environment, and then describe the possible interactions.

5.1 Users

The different kind of users who interact with the environment can be subdivided into three groups.

Firstly, we have the people managing the environment. They can be compared to the "root" user on an Unix-like system. Their role is to make sure the system is working properly. During a consultation they may for example ban users who disturb the consultation, or allow some people to do some specific interactions. We will call these users managers.

Secondly, we have the people modeling the environment. They stand for programmers on a computer system and are the urban designers of the virtual cityscape. Their role is to modify the environment. They also can access feedback from consultations. We will call these users designers.

Finally, here are the people visiting the environment. They can be considered as program users on a computer system and are the public of an urban planning consultation. These users will mainly access information from the environment and leave feedback. We will call these users visitors.

Obviously there is a hierarchy between these three groups, and so a manager can as well model the environment, and a designer can also just visit the environment.

5.2. Interaction

Interactions can be characterized into two groups: user-environment interactions and user-user interactions.

User-user interactions deal with communication between users, such as chat for example. Indeed, communication is a vital factor during a collaborative task. Consultation can be considered as a collaborative task, as a feedback from a group of individuals is not the sum of their respective feedback. There are two ways communication can be achieved: Firstly communication can be private between a group of two (or more) users. This can be useful when some people want to discuss about the feedback they want to give. Then sometimes it can be useful to reach all the users by using broadcast messages. For example, a designer who made a modification on the model may want to inform all the participants.

User-environment interactions can be related to the model or related to information. Other interactions will deal with managing ways the model is linked to information. Model-related interactions are about modifying the model, which includes adding and removing parts of the model. Information related interactions are about accessing and letting information. They are critical, as they are deeply involved during the consultation. Finally, Environment-related interactions are about the management of the environment, such as adding or removing information and consultation boards.

Having all these different interaction and user type, it becomes necessary to manage these interactions according to the different kind of users. Instead of using a static allocation, the idea is to apply Unix file-system features to the environment. So we can see an interaction as an action (similar to shell commands on Unix) that some people (similar to Unix users) can perform on some objects (similar to files and programs). During a consultation, the manager users will interact with the other users, giving them "rights" on interactions, as would do the root user on a Unix system. Why use flexibility? Consultations can vary from one to another. For example we could imagine a first consultation where public is only allowed to walk through the environment. Another consultation may involve some minor modifications of the environment for people to choose between different scenarios. So, a dynamic interaction allocation seems to be a good solution. Then, why use real-time interaction allocation during a consultation? The idea is here to enhance the way consultations are made. Someone from the public could have at some point a design idea. A manager could hear him and allow him to express this idea (for example move an object) by giving him the appropriate right on some interactions. For urban planning, this idea seams interesting, as for now it is obviously impossible for someone from the public to modify the city model during a consultation!

6. Case Study

Having described the environment and interaction processes, we now describe a public consultation process during an urban planning process, using our environment (Figure 4).

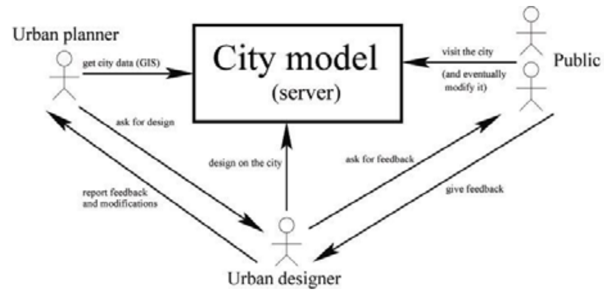

Figure 4. Case study

An urban planner observes city data, and decides that an area of the city needs to be redesigned. This observation could be done getting information from the city model. This is the stage of political decisions. Then he asks an urban designer to design this city area. So he will interact with the city model, altering 3D models of the cityscape, according to the design directions he chose.

After that, the model is opened to the public for consultation. This is the stage of urban planning we are interested in here. The public can interact with the city, by modifying 3D objects, or simply let feedback on the cityscape.

Then the urban designer can collect this feedback by different ways: he can observe how the public modify the model, he can read the feedback they let or simply he can communicate with them during the consultation. After analyzing this information he can alter the city design or report the feedback directly to the urban planner who will make new decisions.

Opposite to the way consultations are undertaken using a small scale model, it is here possible to modify the model within the process. This is a more dynamic approach to consultation (Figure 5).

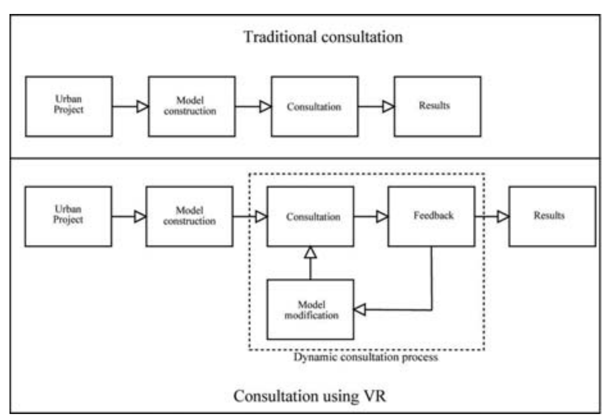

Figure 5. The Consultation process

7. Developing the cityscape (or accomplished work?)

In this section we describe the implementation of the cityscape environment we will use to draw experiments. Firstly we will focus on the cityscape model itself, and then we will describe the navigation and user interface.

7.1. Developing the cityscape environment

The cityscape model has been developed in C++ using an object-oriented design (Figure 6). Such kind of design has been used in various virtual cityscape development [16][17].

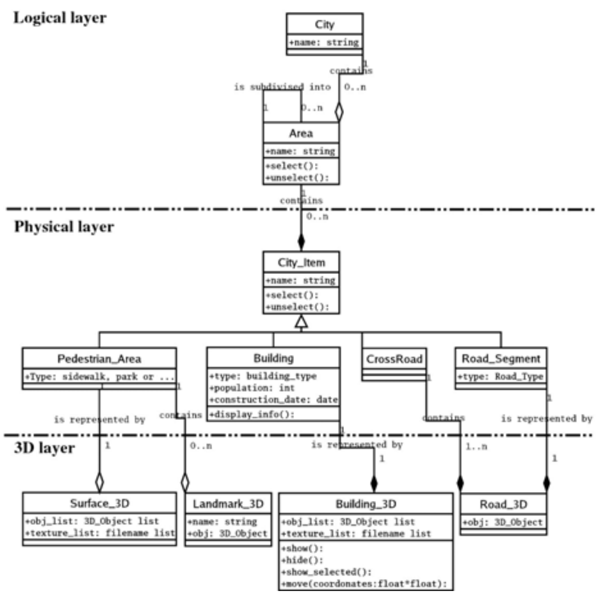

Figure 6. The cityscape model

The city contains an area tree, which define the logical layer of the data structure. Any area can contain city items, which can be pedestrian areas, buildings, crossroads or road segments. This is the physical layer. And then, all city items have an object to describe their 3D representation. This is the 3D layer. There is also the Landmark_3D class, standing for landmark objects which are on pedestrian areas.

The 3D layer has been developed using Maverik [7] (MAnchester Virtual EnviRonment Interface Kernel), a micro-kernel for Virtual Reality, developed at the University of Manchester, using OpenGL. Figure 7 shows some screenshots of the result.

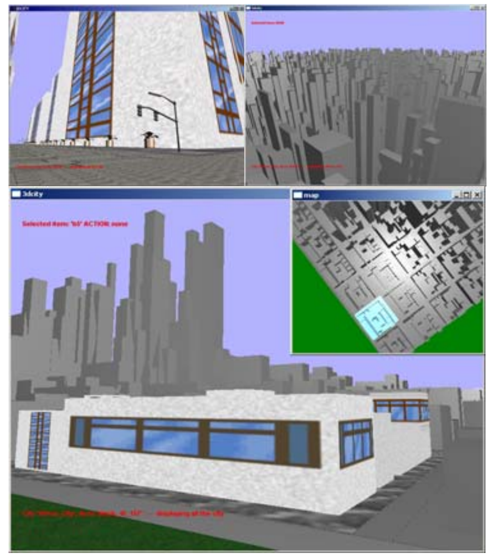

Figure 7. The Cityscape

Some interactions have been developed, as moving and modifying objects, letting comments on Information boards and getting information from the model.

7.2. Navigation and user interface

Intuitive navigation using the mouse has been developed, as many VR systems and computer games have proved their efficiency. It was obviously impossible to include VR-specific input devices, such as VR helmet, as we want people to use the application on their home computers.

An intuitive interface has been developed as well for interaction with objects. When an interaction requires the displacement or modification of an object, this one has to be previously locked, by simply selecting it with the mouse by the user. When an object is selected, any authorized interaction with it is then possible. This locking mechanism allows the user to be the only one to interact with the object. Furthermore it simplifies the interaction process, as it is only possible to interact with one object at once.

To enhance navigation, a map window has been implemented (Figure 8). Using colors, it becomes possible to identify clearly the consultation area from the rest of the city, and within this area the locked object if required.

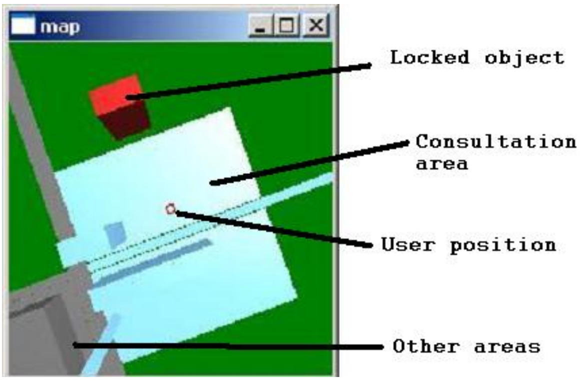

Figure 8. The Map

7.2. First results

Having developed the model, the first experiment focused on display performance of the model.

So we drew measures the frame rate on two different computer systems. The first one, a Pentium III 450Mhz with a Geforce 2 GTS can now be considered as a low-performance range computer system. The other one is an Athlon XP 2000+ with a Geforce 4 Ti4600, which can be considered as a middle-performance range computer system. The first computer was running on Linux and the second one on Windows XP.

We divided the city into blocks of 15 buildings. We selected a block to be the consultation area and made vary the number of the other blocks. As we wanted to evaluate the impact of the size of the cityscape on performance for a consultation scenario, the size of the consultation area did not change. Each block contains also 3D landmarks. The experiment consisted of an automated navigation inside the city near the consultation block, as people would navigate during a consultation. During this movement the frame rate was recorded, as shown in Figure 9.

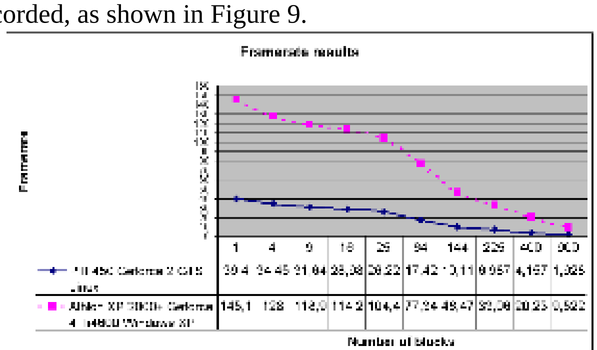

Figure 9. Frame rate results

The results are very encouraging, as if we consider a frame rate of 20 frames per second as an acceptable threshold [4], we can see that the low-performance range computer manages well displaying cities up to about 50 blocks. Thus, the city model will be able to perform well technically on standard computers.

8. Future experiments

The first experiments will assess how easily the users can access the information from the model, and do some interactions. There are two different way to measure this we will use. Firstly we will ask people to do perform tasks in an environment, and measure the time they need to perform them. There will be different tasks, as navigation, accessing information, environment modifications and communication between people. Each task will be assessed independently. Doing this test with people with different experience with 3D will tell us how people with none or few experience with 3D will deal with navigation and interactions. Then we will use a user questionnaire, and so people will answer questions about their experience in the virtual environment. Some questions will be directed, and some will be free.

Then we will simulate a case of public consultation. We will test different levels of interaction from the public (one example: they may be able to change the color of a building and then say on an expression board which color they prefer). Here we will assess how easily urban designers can access the results from a consultation. The result may just be a questionnaire addressed to them. We will also assess how information is dealt in our model, compared to a standard small scale model of a city, in the point of view of people being consulted. Are information and consultation boards really useful? Is it easier to access information this way (using boards rather than paper information)?

9. Conclusion

Having described how could be a CVE oriented towards large scale public consultation, we implemented a virtual 3D cityscape environment prototype. We established that the environment is suitable for public use in terms or performance.

Development of this environment is currently pursued in order to experiment on the way people manage to experience the environment, accessing, adding information and using interactions, as we described on the theoretical model. Finally we want to assess that allowing VR technology to be use by the public at a large scale improves the way public participation is handled. For that we need to prove that the way data are acceded and created gives a great potential to VR as a new tool for urban planning public consultation.

10. References

[1] S.R. Pettifer, S. Gibson and A.J. West, "Visualising the Virtual Cityscape", *Visualisation of Structure and Population with Electronic Landscape*", Mariani, J., Rouncefield, M., O'Brien, J. and Rodden, T. (Editors), Lancaster University Press, Ch. 6 pp. 127-138, 1998.

[2] M. Batty, M. Dodge and B. Jiang, "GIS and Urban Design", *Geographical Information and Planning*, Stillwell, F., and Geertman, S. (Editors), Springer-Verlag, Berlin, 1999.

[3] M. Dodge, A. Smith, and A. Doyle, "Visualising Urban Environments for Planning and Design*", Proceedings of the Graphics, Visualization and the Social Sciences workshop*, Loughborough, UK., 1997.

[4] W. Jepson, R. Liggett, and S. Friedman, "Virtual Modeling of Urban Environments", *Presence: Teleoperators and Virtual Environments*, vol. 5, no. 1, pp. 72-86, 1996.

[5] Active Worlds website. http://www.activeworlds.com

[6] M. Corbit, "Building Virtual Worlds for Informal Science Learning (SciCentr and SciFair) in the Active Worlds Educational Universe (AWEDU)", *Presence: Teleoperators and Virtual Environments*, vol. 11, no. 1, pp. 55-67, 2002.

[7] R. Hubbold, J. Cook, M. Keates, S. Gibson, T. Howard, A. Murta, A. West, and S. Pettifer. “GNU/MAVERIK : A micro-kernel for large-scale virtual environments”. *Presence: Teleoperators and Virtual Environments*. ISSN 1054-7460. 10:22-34, 2001.

[8] Planet 9 Studios, Virtual Tokyo. http://www.planet9.com/earth/tokyo/index.htm

[9] V. Bourdakis, "The Future of VRML on Large Urban Models", *UKVRSIG97, The Fourth UK VR-SIG Conference*. Brunel University, 1997.

[10] P.J. Birch, V.J. jennings, A.M. Day and D.B. Arnold, "Procedural Modeling of Vernacular Housing for Virtual Heritage Environments", *EGUK 2001*, University College London, 2001.

[11] S.P. Browne, J. Willmott, L.I. Wright, A.M. Day and D.B. Arnold, "Modelling and Rendering Large Urban Environments", *EGUK 2001*, University College London, 2001.

[12] P. Wonka and D. Schmalstieg, "Occluder Shadows for Fast Walkthroughts of Urban Environments", *Eurographics'99*, Milan, 1999.

[13] P.E. Debevec, C.J. Taylor and J. Malik, "Modeling and Rendering Architecture from Photographs: A Hybrid Geometry- and Image-Based Approach", *SIGGRAPH'96*, New Orleans, 1996.

[14] D. Liebowitz, A. Criminisi and A. Zisserman, "Creating Architectural Models from Images", *Eurographics'99*, Milan, 1999.

[15] Schmidtt G. et al., "Toward Virtual Reality in Architecture: Concepts and Scenarios from the Architectural Space Laboratory", *Presence: Teleoperators and Virtual Environments*, vol. 4, no. 3, pp. 267-285, 1995.

[16] G. Thomas and S. Donikian, "Modelling Virtual Cities dedicated to Behavioural Animation", *Eurographics'2000*, Interlaken (Switzerland), 2000.

[17] N. Farenc, R. Boulicand and D. Thalmann, "An Informed Environment didicated to the simulation of virtual humans in urban context", *Eurographics'99*, Milan, 1999.

[18] O. Märker and V. Pipek, "Computer-Supported Participation in Urban Planning from the viewpoint of "Communicative Planning Theory" ", *IFIP 8.5 – Advances in Electronic Government*, Zaragoza (Spain), 2000.

[19] R.C. Waters, D.B. Anderson, J.W. Barrus, D.C. Brogan, M.A. Casey, S.G. McKeown, T. Nitta, I.B. Sterns and W.S. Yerazunis, "Diamond Park and Spline: A Social Virtual Reality System with 3D Animation, Spoken Interaction, and Runtime Modifiability", *Presence: Teleoperators and Virtual Environments*, Vol. 6, no. 4, August 1997, pp. 461-480, 1997.

[20] S. Benford, C. Greenhalgh, C. Brown, G. Walker, T. Regan, J. Morphett, P. Rea and J. Wyver, "Experiments in Inhabited TV", *CHI'98 Late Breaking Results (Conference Summary)*, ACM Press, pp.289-290, 1998.

[21] J. Tromp, and D. Snowdon, "Virtual body language: providing appropriate user interfaces in collaborative virtual environments", *ACM symposium on Virtual reality software and technology (VRST)*, ACM Press, Lausanne (Switzerland), pp. 37-44, 1997.

[22] Knowbotic Research. http://www.krcf.org/krcfhome/