The Digital Cathedral of Siena - Innovative Concepts for Interactive and Immersive Presentation of Cultural Heritage Sites

Johannes Behr (*), *Torsten Fröhlich* (#), *Christian Knöpfle* (#), *Wolfram Kresse* (#), *Bernd Lutz* (#), *Dirk Reiners* (#), *Frank Schöffel* (#)

(*) Center for Computer Graphics (ZGDV e.V.), Rundeturmstr. 6, D-64283 Darmstadt
jbehr@zgdv.de

(#) Fraunhofer Institute for Computer Graphics, Rundeturmstr. 6, D-64283 Darmstadt
{torsten, knoepfle, wkresse, blutz, reiners, schoeffe}@igd.fhg.de

ABSTRACT

In this paper, innovative technologies for the presentation of complex cultural heritage sites with immersive 3D computer graphics are introduced, which are based on new concepts, partly having been adopted from other computer graphics areas to the specific needs for cultural heritage presentations, partly having been developed especially for this purpose. The main focus of this project was to develop new concepts for the integration of historical, architectural, and cultural data related to a cultural heritage site, into an immersive Virtual Reality (VR) Environment suited for presentation of the digital model. These concepts include digital storytelling, virtual 3D avatars and 3D agents. The 3D avatar was implemented as a virtual tourist guide, who accompanies the visitors on their tour and shows them the places of interest. Furthermore, specifics of the very detailed and complex digital models of cultural heritage sites are addressed, in order to meet the high quality requirements of an interactive 3D visualisation. Special focus has been put on both the frame rate and the visual quality. Beside Level-of-detail mechanisms, texture-paging techniques have been developed to allow the use of very high resolution textures. Lighting simulation based on radiosity simulation, which ensure a high visual quality, have been extended for the efficient simulation of global illumination effects within very complex models. These new concepts and techniques have been successfully applied in the immersive VR presentation of a digital model of the cathedral of Siena, which has been shown to the public during the EXPO 2000 in Hannover.

KEYWORDS: Digital Model, Virtual Reality, Digital Storytelling, Interaction, Light Simulation, Avatar

INTRODUCTION

In the recent years, the need for innovative computer graphics technologies to present cultural heritages in a digital way is steadily growing due to of various reasons:

Historical places and buildings are subject of a continuous dilapidation, which is intensified by the ever increasing tourism. To preserve them from total destruction, many historic places are restricted or even closed to the public. Through the use of digital models and accompanying presentations, these places and buildings can be "reopened" again. Furthermore they allow to visualise the change of these places and buildings through time (e.g. [2]).

These novel technologies are not limited to simple non-interactive presentations of the cultural assets (e.g. walk-through), but are also capable to offer the users an unforgettable experience because of the possibility to interact with the virtual scenery. Especially interactivity is a key aspect for comprehensive understanding of the presented topics ("learning by doing") and furthermore heavily improves the attractiveness of the exhibit. Since many museums and similar institutions are using mostly non-interactive media, e.g. text, video and audio, the use of this kind of technology does improve the attractiveness of the museum as a whole.

These new technologies are able to support researchers and experts in the reconstruction process of partly or fully destroyed assets, e.g. to compare different alternatives or preview the effect of the restoration process. One approach was presented by [5], where a framework for cooperative working on the excavation and reconstruction of historic assets was developed.

In our work we focused on presentations for the lay audience, where attractiveness, ease of use and comprehensibility of the presented topics are the most important issues. Thus we developed a virtual reality (VR) based system ([8])), because in the past years, this technology heavily matured and showed its didactic potential in various demonstrations (e.g. [4] or [3]). Furthermore the basic building blocks of VR are realistic and real-time rendering and interactivity, and thus represent the ideal technology for applications in the area of cultural heritage. In this paper we will present the results of our research on VR technology-based information systems for complex cultural heritage sites. With the Siena Cathedral as an example, an immersive virtual environment was created. The Siena Cathedral was chosen because of two main reasons:

- *Historical Background:* The Siena Cathedral is regarded as one of the most beautiful and picturesque cathedrals in Europe. The history of the construction of this Gothic cathedral is highly exceptional. The building was started in 1136. In 1284, the Sienese sculptor and architect Giovanni Pisano started raising the western facade, but only completed the lower story. Later, the three-nave building with the dome and the bell tower were completed. In the first half of the 14th century, the wealthy Sienese had the idea to add a new grand building to the southeast side of the already existing church: the "Duomo Nuovo". In 1348, the work was stopped because of the Black Plague and unsolved statics problems; later on, it was given up completely. Until today, the remaining facade dominates the panorama of Siena. During the centuries, the cathedral was completed in partly Romanesque, partly Gothic style. Today, the splendid marble edifice appears to be a magnificent, integral whole.

- *Technical Challenge:* The Siena cathedral is a building with an enormous complexity in respect of its size, ornaments, sculptures and paintings. Thus a nearly photo-realistic visualisation in real-time of the Siena cathedral is a challenging task but absolutely necessary for an attractive presentation. Consequently a major part of our work focused on that topic. This will be explained in more detail in chapter Rendering.

Because of the rich historic background this cathedral is also well suited as an information system. Thus the developed system allows the visitors to virtually explore the cathedral, to experience its architecture, dimensions, and atmosphere, and obtain architectural, cultural, and historical information as well. Navigation and interaction in the scenario are realized by an innovative man–machine interface that can be operated intuitively. That interface has two aspects: A touchscreen based input interface (see chapter The ancient book) and a virtual avatar as output interface (see chapter The virtual tourist guide). The visualisation output is done in realtime on a stereoscopic large-screen projection. Thus, even large groups can participate in the presentation as seen on EXPO 2000 (see chapter Presentation at EXPO 2000).

MODELLING

Data acquisition

A very important point in creating a Virtual Reality presentation of a cultural object is to have a sufficient information base. This information is required to build the virtual models of the objects and to generate the textures, but they are also necessary for the development of the storyboard and as information visitors will get from the system.

For the modelling and texturing process we used mostly plans, maps and photographs. If it is not possible to get all the needed photographs of the objects, which normally is the case for ancient objects, old paintings or prints can be used. Textural information can help during the modelling process, e.g. to clarify ambiguous measurements or to obtain background information. For the cathedral of Siena, we used the maps in the book [7], which we also received in electronic form from the land register bureau. These maps were the results of a complete photogrammetric measuring of the cathedral and included all necessary measurements and details. This allowed us to generate a very accurate and detailed model of the cathedral. To get all photographs required for the generation of textures we visited the cathedral twice and took about 1000 pictures of it. This was possible with the support of the "Opera della Metropolitana (Siena)".

In addition to texts and pictures, multimedial data can be used to give visitors information about the cultural objects. Depending on the kind of presentation texts, pictures, videos and sound can be used.

Generation of the virtual models

For cultural objects a special focus has to be put on the quality of the virtual model. It is used as a stage for the presentation and it should attract the people to attend the presentation. The virtual model will also provide visual information to the visitors. One of the problems while creating virtual models of cultural objects is the conflict between an exact and detailed model and the demands of realtime rendering. Cultural objects mostly consist of very detailed and ornamented geometries. Due to their age and the ancient methods of construction, most surfaces are not flat and even objects like pillars never look the same. This makes the modeling process very difficult and raises the number of faces required to model the objects.

To insure that the virtual models fit the original maps exactly, we created a texture of the ground plan of the cathedral and mapped it on a face, in this way we could compare the virtual model with the ground plan in the virtual environment (see Fig. 1)

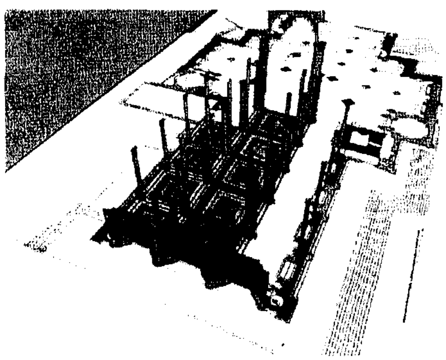

Fig. 1: Construction of the virtual cathedral

All the models where generated with commercial 3D modelling system. These models then had to be converted to FHS, the data format used by Virtual Design 2, the VR system used for the presentation.

The converted models were checked to ensure that the conversion worked well, no parts were missing and that no geometric errors (such as holes, double faces, ...) occur. After the verification, the models had to be optimised for real-time rendering. Therefore the hierarchy was adapted so that techniques like "Culling" or "Level of Detail" could be used by the renderer.

Generation of the textures

Textures are used in virtual worlds to achieve photo realistic scenes. To generate realistic surfaces like wood or marble or to integrate flat objects like paintings, images of these objects are used. Normally one can retrieve these images from texture libraries. In most cases, this is not possible for cultural objects, since here it is important to get original images of most of these surfaces to get an authentic impression of the objects.

The problem with these objects is, that it is often very difficult to take photographs due to several reasons:

- The surfaces/objects don't exist anymore.
- It is forbidden to use a flash or artificial lights.
- It is impossible to reach the objects or to take photographs without perspective distortions.
- It is impossible to take pictures without interfering objects or persons.

Therefore the pictures have to be retouched to achieve good results. This process includes the removal of occluding objects, the correction of perspective distortions, the correction of different lighting conditions, and a color correction.

LIGHTING SIMULATION

Increased realism

To significantly increase the realism and visual quality of the VR model, a light simulation has to be performed. Standard hardware lighting algorithms are inadequate for scenes such as the cathedral. A large room, lit only by daylight coming from windows high above the ground, and artificial lighting at many places in the ceiling and on the walls, cannot be represented realistically by a simple lightsource provided by the rendering hardware, casting no shadows.

Therefore, state-of-the-art lighting simulation algorithms were applied, namely the radiosity method described in [6]. In a radiosity simulation, the light distribution between diffuse surfaces is computed. Multiple lightsources and light types, shadows, and interreflections between surfaces are taken into account, physical correctness within the energy distribution is acchieved.

Radiosity for complex scenes is a difficult task. In theory, every object can exchange light energy with every other object. In addition, to be able to display the light distribution across a surface, polygons have to be further subdivided into smaller patches, since only for the vertices of such patches the arriving light can be displayed.

For a highly detailed model such as the Cathedral, special radiosity optimizations had to be implemented. To reduce the complexity of the simulation, hierarchical radiosity with clustering has been used [SILL94]. For this algorithm, a hierarchy of clustered objects is created from the geometrical model. A cluster is a set of polygons or other clusters. The idea is that, for example, for a chair modelled in high detail and lit by a distant lightsource not every single polygon has to be processed individually. Instead, the chair is treated as a single entity, a cluster. Only with the notion of clusters it was possible to compute a radiosity simulation for the entire interior of the cathedral.

Five hundred spot and point lightsources were placed to create a realistic environment representing the actual lighting conditions within the cathedral. An additional 300 emissive polygons were placed outside the building as a sky dome to simulate daylight. The daylight enters through the upper windows and the two rosette windows, and propagates its way down to the floor. For all light sources, appropriate color temperatures and intensities needed to be set (see
Fig. 2). In Fig. one can see the a shot taken from the VR demonstration.

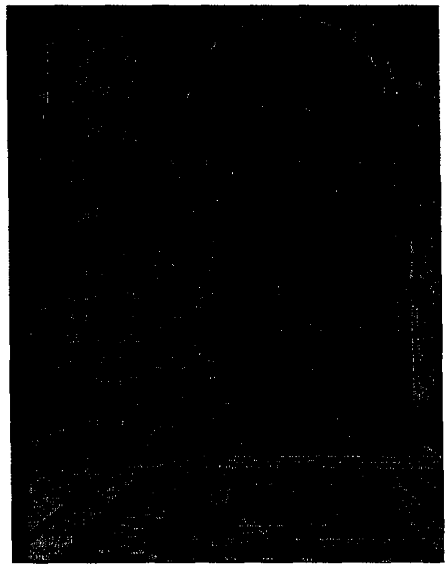

Fig. 2: a view of the cathedral, with standard OpenGL shading

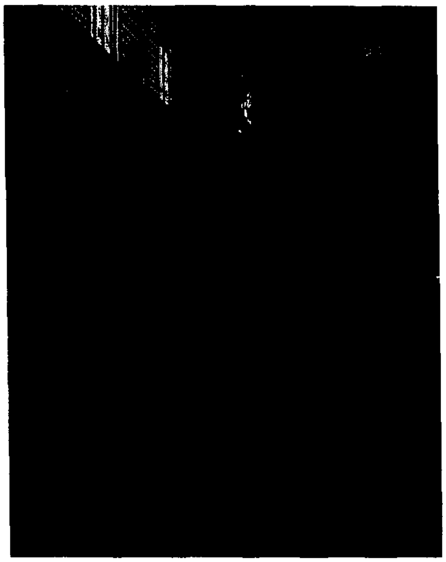

Fig. 3: the result of the radiosity simulation, rendered in real-time

Fig. 4: Libreria Picolomini

Timings

The polygonal input scene consisted of 150000 textured polygons. The simulation of the direct light took 23 hours. However, within that time, 800 individual light sources had to be processed, so the actual time for a single light source is about 1.7 minutes. The radiosity algorithm automatically decided where polygons had to be subdivided into smaller patches to represent light conditions and shadows accurately, so that the final radiosity solution consisted of 280000 polygons.

To complete the radiosity simulation, a converged solution had to be computed that took interreflections and indirect light into account. This converged solution, which theoretically consisted of multiple bounces of light of about 22 billion light interactions, took another 50 hours to compute. In the process, 10 million links were generated. Timings were done on an SGI Onyx2 IR2 on a 200MHz MIPS R10000. Memory consumption for the light simulation alone was about 500 MB.

RENDERING

Even though the model has been approximated by simple planar faces and images, the sheer amount of data that the model represents will overwhelm even the fastest graphics workstations if just used directly. But not all of them are visible or equally important at any given time.

Parts of the model that are behind the viewer or outside his visible area don't have to be displayed and are not sent to the graphics hardware. But from certain vantage points nearly the whole model is visible, so more work has to be done. Level of detail is one useful method of handling this.

Level of Detail

The idea of levels of detail is based on the observation that a screen only has a finite number of pixel. Thus very small details that are far away from the viewer will occupy a small number of pixel on the screen, or even multiple polygons will occupy a single pixel. In these cases it clearly makes no sense to display them all, as only one of them will contribute to the pixel. Thus there has to be a way to reduce the number of polygons for far-away objects and to reduce the load on the graphics hardware to increase rendering speed.

One efficient way to do that is using not one, but a number of differently detailed models for parts of the scene. Depending on the size of the model on the screen or on the distance to the model a more or less detailed version is chosen. This is very efficient for compact model parts, e.g. the columns. They are also a very good candidate for level of detail handling because they are modeled as curves and as such it is easy to reduce their polygon count by reparametrization of their underlying curve.

Alternatively levels of detail can be automatically generated by polygon reduction algorithms, but they tend to work best on highly detailed models of a few objects. For the siena cathedral, which is a very big model, consisting of a very large number of parts they don't perform well. Therefore all the levels of detail we used were hand-made, which gives the best quality, albeit at the high price of having to actually model all the different versions. To reduce the amount of work it was only done for round objects, where it could be done simply by reducing a single parameter in the modelling system.

Another problem is the sheer amount of textures used. Especially the floor's mosaics are very high resolution images; all in all the textures add up to ~300 MB. Graphics hardware needs to constantly access these textures, so they have to be stored in dedicated texture memory that is on the graphics card. But the amount of memory on the card is severely limited, and nowhere near 300 MB. Thus the textures have to be handled more explicitly by texture paging.

Texture Paging

The concepts behind texture paging are similar to the ones behind levels of detail. Textures are usually used in a hierarchy of versions, starting from the original texture with each step in the hierarchy half as large as the previous step. This image hierarchy is called a mip-map and is used to filter the areas where multiple pixel on the texture are mapped to a single pixel on screen. For textures on objects that are far away on the screen only the smaller versions will be used, as the size of the object on screen doesn't allow more pixel to be seen. The graphics hardware can't predict which parts of the texture will be used next, thus it usually has to upload all of them. If not all of them are used, a lot of texture memory is wasted, possibly resulting in exhausted texture memory. This enforces moving older textures out, and thus forcing them to be reloaded for the next frame, if they are visible again. This texture thrashing can severely impact the performance of the system, sometimes making it unusable. The 300MB of textures in the cathedral, when used on our machine with 64MB texture memory, will drop the framerate from several frames per second to ~10 seconds per frame, which is absolutely intolerable for a VR application.

As the OpenGL API demands a complete set of mipmaps, and the graphics hardware is not smart enough to only store the needed ones, the application has to help reducing texture demands. Objects farther away from the viewer will occupy small screen areas and need smaller textures. To allow that, the smaller levels of the image hierarchy are put together to form a new texture, which is used instead of the old one. The open problem is deciding when to use it.

It is possible to do that manually by using the smaller texture on the appropriate level of detail (see chapter Level of Detail). But that would have demanded manual work to apply it to all of the textured objects, which is every object in the scene, and also to create new geometry for objects that are very simple, just textured, like the floor. Instead an automatic mechanism was developed, that analyses the size of the object on screen and selects the appropriate texture for the object. As it is impossible to be close to all of the floor at the same time, this guarantees that most of the floor tiles are only needed in their low-res version, preventing texture memory from being overcommited and having the framerate drop to non-interactive speeds.

USER INTERFACE

The user interface consists of two parts: the virtual tourist guide as an "output device" and an ancient book (realized through a touch-screen) as an input device.

The virtual tourist guide

To give the visitors feedback of their actions, to guide them and to provide historical information, we implemented a virtual tourist guide named "Luigi". Luigi was realized as a virtual avatar with the appearance of a real person vested with a traditional costume. He can move around in the cathedral, guiding the visitors to interesting regions. Furthermore he can talk via audio output to the visitors in order to inform them about architectural or historical facts of the current area. Luigi was modelled and animated in real-time with advanced techniques for facial and body animation, which are described in more detail in chapter Character Animation.

Fig. 5: The virtual tourist guide

The ancient book

To interact with and navigate through a virtual world, special input devices such as a spacemouse or data glove are normally used. These devices can't be used for a virtual reality presentation at a public exhibition. The input device for such an event has to be very robust and intuitive to use. For our presentation we used a touch-screen as input device and a specially designed interface for navigation and interaction. Touch-screens are very common and most people know how to use them, so that no special instructions were needed. To not confuse people by the freedom of movement possible in a virtual world, the possibilities of movement were restricted. The virtual camera couldn't change its height and the tilt angels were limited. To avoid that people move into pillars, statues or walls a collision detection system was integrated. The devised algorithm gave us the possibility to stop the movement of the camera in front of walls or to automatically guide the camera around obstacles.

The user interface was designed like an ancient book (see Fig.), to fit into the atmosphere of the whole presentation. The right page was used to redisplay the information Luigi gave to the visitors. The left side is completely used for interaction and navigation.

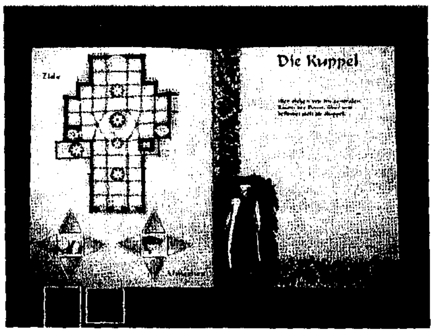

Fig. 6: The interface

The map in the upper part of the page is used for the transition between regions in the cathedral. Visitors will see all the regions they can reach from their current position and ask Luigi to bring them there by touching a region on the map. The arrows in the lower part are used for the free navigation inside the cathedral. The four arrows on the left side move the camera and the arrows on the right side are used to turn around and to look up and down. These elements give the visitors the possibility to move around in the cathedral and to take a closer look at special objects. The three bookmarks on the bottom of the book are used to get more detailed information on architecture, art or history of the region the visitor is standing in.

The user interface was realized using Macromedia director. A Lingo script handles the user input and communicates with the virtual reality system and the audio server.

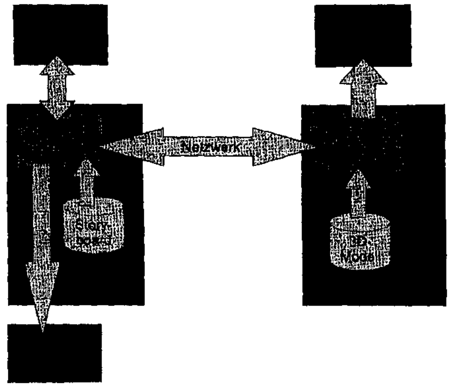

Fig. 7: Schematic view of the system

CHARACTER ANIMATION

Character Animation

In the analysis and design phase at the beginning of the project some strict goals for the animation system where set:

- Must perform in real-time (< 30ms for all animation tasks)
- Must handle high polygon counts (main actor has > 10k triangles)
- Must be able to hold the animation sequences of 1h in system memory
- Must handle seamless surface animation
- Must handle speech-sync facial animation
- Must allow to combine different animation sequences at runtime
- Must allow to model the animation in standard packages (e.g. 3Dmax)

Since the animation sequences were modelled and designed with standard tools (e.g. 3D Studio Max) the simplest technique would be to use linear coordinate interpolation for real-time playback. Since one request was to hold more than one hour for a 10k triangle character, ordinary coordinate key sets would need too much memory resources (~ 20 GByte). Therefore special methods for facial and body animation were developed and implemented.

Body Animation. One very popular structure for body animation is a hierarchy of joints and rigid segments (or bones), which is used by games (e.g. Tomb Raider, see Fig.) and real-time animation systems (e.g. H-Anim). Every joint defines a rotation and holds a geometric segment and the connected (lower) joints. The technique is very time and memory efficient, since the animation is solely defined by the rotation settings per joint, but does not fulfil the requirement for a seamless skin.

Fig. 8: Simple rigid segment character

We used a more advanced method, which is based on the Skeleton-Subspace Deformation. The Skeleton-Subspace Deformation is an algorithm that is used in commercial modelling packages under several rather uninformative names such as skinning, enveloping, etc. The method uses not just rigid geometric segments but a single surface that defines the skin. During the animation process, the skin vertices positions are deformed via the skeletal system (joint transformation). The deformation of the vertices is controlled by weights, which represent the degree of influence on the associated Joints (see Fig. 9).

Fig. 9: Skin deformation using two Joints

Our method allows us to define independent skin elements, which enable the algorithm to handle multiple material settings (e.g. face, clothing) per character. The implementation was highly optimised to allow real-time deformations of characters that have more than 10k triangles.

Facial Animation

The facial animation sequences were modelled using a set of 16 morph-targets. A morph-target is a version of the face geometry with a certain "frozen" facial expression, e.g. eyes open or closed, speaking an "A", etc. The actual shape of the face is then blended from all 16 morph targets.

One goal was to have the lip animation synced with the audio channel. The commercial lipsync system from Lipsync was used to automatically generate the channel settings for the morph-targets.

For the real-time facial calculation we used a simple but very efficient linear blending technique. In our case, an animation is represented by the base geometries (the morph targets) and several keyframes, which are simple weight vectors (see Fig.).

Fig. 10: Lip sync face animation

One set of weight vectors defines the range of influence for every single base shape. The number of base shapes is about one order of magnitude smaller than the number of vertices, resulting in a very compact representation of the animation itself. This is the reason why this scheme is so well suited for the requirement to hold a large number of different animations in system memory.

Character Modelling

As mentioned earlier, a model consists of three major components:

- a polygonal model of the surface of the avatar,
- an underlying hierarchical skeleton model and
- a weight of influence from each bone of the skeleton to each vertex of the geometry

During the modelling process the skeleton is visualized and moved by the animator, eg. performing steps or gestures. The geometry is deformed automatically as defined by the weights.

The polygonal model has been exported in the standard VRML2 format. The skeleton animation and the weights are stored in 3D Studio Max proprietary file formats. In order to combine all three parts of the animation, we have defined an extension of the standard H-Anim file to define geometry, skeleton and weights. A converter has been developed for this purpose to enable the renderin system to load and play back these files. See Fig. for a closer look on Luigi.

Fig. 11: the virtual avatar

Dynamic path planning

While Luigi explains the attractions of a certain place in the cathedral, the users are allowed to move their view arbitrarily. During the transition from one place in the cathedral to the next, the avatar and the virtual camera are animated automatically. When the user chooses to proceed to the next attraction, the camera has to be moved automatically to the starting point of the following animation in order to avoid a sudden jump of the perspective.

The problem to be solved is to find a smooth motion curve for the camera from the actual camera's position and orientation to the starting point and orientation of the animation. These constraints lead us directly to Hermite curves. Hermite curves are parametric cubic curves defined by position and tangent at start and end points and therefore fit to our problem nicely.

Still they are not the complete solution. First, Hermite curves tend to "wiggle" at certain combinations of start and end directions, which would lead to unnatural motion of the camera. Second, the path generated has to avoid collisions with the walls and obstacles. To achieve a more natural motion, we start with turning the camera into the general direction of our destination, the animation starting point. In the next step, we check the direct linear path to the goal. If it is collision free, we are done. Otherwise we insert more waypoints until we have a collision free polygonal path, which is then interpolated smoothly by Hermite curves.

PRESENTATION AT EXPO 2000

The first public presentation of the virtual Cathedral of Siena took place at EXPO 2000 in Hanover/Germany, which lasted from the 1st of June until the 31st of October. The installation was located in the theme park "Basic Needs", which was created by Rajeev Sethi, a famous Indian artist and showed how people from different countries are defining their basic needs. The Cathedral of Siena was placed between the areas "Shelter and Living" and "Immaterial Needs" and thus could be seen as a bridge between architecture and spirituality.

The presentation was shown in a darkened room with a base size of 5m by 4m and 5m height. The projection screen was 4m tall and 2,5m wide and thus presented the virtual cathedral in a "portrait format", which improved the overall impression of the dimensions of the building. The demonstration ran on an ONYX2-InfiniteReality3 and an Electrohome 8500 projector was used to generate the images. See Fig.

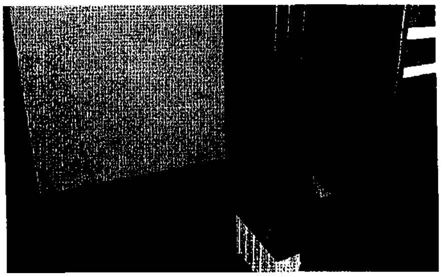

Fig. 12: EXPO installation

The touch-screen was located in front of the screen, where the users could interact with the system. audio speakers were attached to the front wall to support sound feedback, e.g. the voice of Luigi, the virtual avatar. Many people were impressed by the quality of the demonstration, but only a few really tried to use the touch-screen for controlling the demonstration (mostly younger people). We guess that people feared that they could damage the system in some way or make a fool of themselves, in case they don't understand the operation of the presentation. To avoid that problem the virtual avatar could explain the system as soon as the touch-screen is not operated for a specific time.

CONCLUSION AND OUTLOOK

Currently, only the models of the interior of the cathedral are completely finished. The next step will be to include the surrounding area outside the cathedral. The modelling process, as well as the generation of the textures and animations is nearly finished. An alternative would be to include objects, which are already destroyed, which have changed over time, or which were planned but never built. It would be interesting to include the unfinished ``duomo novo", the new cathedral into the model to give visitors the possibility to see how the cathedral could have looked like.

During our tests on EXPO it turned out, that the user interface is already easy to use, but nevertheless people were swamped with the operation of it. Therefore an important task for the future will be the improvement of the screen design of the user interface as well as an avatar who is able to react on the users behaviour in a more flexible way.

From a technical point of view it would be desirable to have an avatar able to act more autonomously. Actually we have to animate every step and gesture by manipulating the skeleton model, export the data using the extended H-Anim format and run the VR-system. The system would be much more flexible and easier to extend if we could simply define the motion paths and actions of the avatar and let the software perform the animation automatically. This is actually work in progress.

ACKNOWLEDGEMENT

Most 3D models are based on plans, which were provided by the "Kunsthistorisches Institut Florenz". The creation of the plans was supported by the bmbf. The main responsible people were Prof. Walter Haas, Prof. Peter A. Riedel, Prof. Max Seidel und Prof. Dethard von Winterfeld.

Furthermore we would like to thank the Opera della Metropolitana Siena and the Messbildstelle Dresden.

Last but not least we would like to thank the numerous people, who were involved in that project. Without their support and work, this project wouldn't have been possible.

The work described herein was funded by the Bundesministerium für Bildung und Forschung (bmb+f) in Germany, project number 01 IR 906 2.

REFERENCES

1. Marc Alexa and Johannes Behr and Wolfgang Müller,The Morph Node, Web3D - VRML 2000 Proceedings, page 29-34, ISBN 1-58113-211-5

2. Antonicelli, A., Sciscio, G., Rosicarelli, R., Ausiello, G., Catarci, T., Ferrarini, M.: "Exploiting Pompei Cultural Heritage: The Plinius Project", in Short Papers and Demos of Eurographics 1999, pp. 100-103, Milano, Italy

3. Fröhlich, Torsten:"The Virtual Oceanarium - A Visual Computer Simulation of Europe's Largest Aquarium",Communications Of ACM, July 2000

4. Lutz, B., Weintke, M.: "Virtual Dunhuang Art Cave: A cave within a CAVE", in Proceedings of Eurographics, 1999, Milan

5. Li, X., Lu, D., Pan, Y., Hua, Z.: "Virtual Dunhuang Mural Restoration System in Collaborative Network Environments", in Eurographics Forum 19(3), 2000

6. Cindy M. Goral and Kenneth E. Torrance and Donald P. Greenberg and Bennett Battaile: "Modelling the Interaction of Light Between Diffuse Surfaces", Computer Graphics (ACM SIGGRAPH '84 Proceedings), vol.18/3, pp. 212-222

7. Walter Haas, Dethard von Winterfeld: ``Die Kirchen von Siena. Band 3.1.3 Der Dom S. Maria Assunta", 1999, München

8. Pimentel, K, Teixeira, K.: ,Virtual Reality. Through the new looking glass", New York: Windcrest, 1993

9. Francois Sillion. ``A Unified Hierarchical Algorithm for Global Illumination with Scattering Volumes and Object Clusters", IEEE Transactions on Visualization and Computer Graphics, vol.1/3, 1995.

10. N.Magnenat-Thalmann, R.Laperriere, and D.Thalmann: "Joint-Dependent Local Deformations for Hand Animation and Object Grasping", Proc. Graphics Interface, 1988, pp. 26-33.

ABOUT THE AUTHORS

Johannes Behr received his M. Sc. diploma in Advanced Software Engineering form the University of Wolverhampton in 1996. Since 1997, he is working as a research assistant within the Visual Computing department of the Computer Graphics Center (ZGDV). Areas of interests focus on Virtual Reality, Realtime rendering and animation/behaviour descriptions for 3D interactive worlds.
E-mail: *jbehr@zgdv.de*

Torsten Fröhlich works since 1996 as a researcher at the Fraunhofer Institute for Computer Graphics, Dept. Visualization and Virtual Reality. His recent work comprises several Virtual Reality installations for public audiences, among them a Virtual Reality simulation of the marine life inside the Lisbon Oceanarium presented on EXPO'98. His research interests include the interactive simulation of dynamic behavior in Virtual Environments, such as artificial creatures or avatars.
E-mail: *torsten@igd.fhg.de*

Christian Knöpfle worked from 1996 to April 2000 at Fraunhofer-IGD as a researcher in the virtual reality group. During that time he worked mainly on the development of solutions for the automotive industry with the main focus on intuitive interaction. Since April 2000 he is the head of the virtual reality group.
E-mail: *knoepfle@igd.fhg.de*

Wolfram Kresse studied computer science at the Technical University of Darmstadt, Germany. Since 1998 he is working in the field of advanced lighting simulation at the Department for Visualization and Virtual Reality at the Fraunhofer Institute for Computer Graphics in Darmstadt, Germany. In the Eurpoean Union research project ARCADE in collaboration with Gravir/INRIA he developed new radiosity algorithms for the lighting simulation of complex scenes.
E-mail: *wkresse@igd.fhg.de*

Bernd Lutz is working since 1996 in the Fraunhofer Institute for Computer Graphics. He studied electrical engineering at the Technical University of Darmstadt. His research fields are representation of cultural heritage in virtual world and virtual environments for learning.
E-mail: *blutz@igd.fhg.de*

Dirk Reiners has been working in realtime rendering at Fraunhofer IGD since 1994. He has built the rendering components of several systems, including the commercial Virtual Design 2 software. Currently he is the technical director of the OpenSG Forum and the project leader of OpenSG, an open source scenegraph system

(*http://www.opensg.org*).
E-mail: *reiners@igd.fhg.de*

Frank Schöffel was working from 1995 to end of 2000 in the Fraunhofer Institute for Computer Graphics (Fraunhofer-IGD). He studied computer science at the Technical University of Darmstadt. His research field is radiosity based lighting simulation with the focus on dynamic sceneries.
E-mail: *frank.schoeffel@gmx.de*