The 9th International Symposium on Virtual Reality, Archaeology and Cultural Heritage
VAST (2008)
M. Ashley, S. Hermon, A. Proenca, and K. Rodriguez-Echavarria (Editors)

Real-time visualisation of multiple time dependent reconstruction hypotheses for a cultural heritage site

S. D. Laycock†, R. G. Laycock, D. Drinkwater and A. M. Day

School Of Computing Sciences, University of East Anglia
Norwich, UK, NR4 7TJ

Abstract

Virtual environments have commonly been used for the dissemination, education and public awareness of cultural heritage, owing to the medium's ability to provide an engaging and interactive exhibit. However, virtual environments have yet to be fully utilised as a tool to enhance the work flow for the archaeologist. When attempting to understand a particular cultural heritage site, a variety of data sources and technologies are employed, resulting in potentially conflicting hypotheses regarding the development of a site through time. It is particularly challenging to quantify the validity of a hypothesis without viewing it within the spatial and temporal context. This paper presents new interactive techniques for the exploration of such alternative interpretations for a large cultural heritage site. To illustrate the utility of the presented approach interpretations of St. Andrew's Monastic complex, Norwich, UK, are considered. The data set comprises of more than two hundred individual components, totalling in excess of five million triangles, ranging from 1258AD to present day. Techniques are presented to process the three dimensional models and utilise consumer level hardware to visualise it in real-time. All navigation is undertaken with a haptic interface, which aside from scene traversal permits different portions of space to be selected and for the objects occupying that space to be regressed or alternative reconstructions to be presented.

Categories and Subject Descriptors (according to ACM CCS): I.3.6 [Computer Graphics]: Interaction techniques

1. Introduction

As computer graphics techniques advance the ability to visualise realistic representations of objects and large environments becomes more easily obtained, particularly in an offline rendering methodology. Despite this, the fundamental concern of the reliability of a virtual reconstruction is seldom realised [Rou02]. Archaeologists and historians are often faced with multiple interpretations of the evidence before them, leading to many potential hypotheses for an environment's appearance and functionality. By visualising a single virtual reconstruction the varied nature of the hypotheses for an area is lost, which prevents detailed interpretation and thus failure to adhere to one of the defining characteristics of computer graphics for cultural heritage results [SS02] [SR02]. In this paper techniques are devised to enable multiple virtual reconstructions of a site to be visualised in a single framework.

The novel visualisation framework provides a flexible and extendable environment in which to compare and contrast changes occurring to an individual artefact, group of artefacts or buildings. Over time these artefacts evolve such as a building being transformed subject to a change of ownership or primary function. The input to the framework consists of a collection of three dimensional models, representing the artefacts, that must be sorted by the range of time periods in which they were believed to exist. However, unlike other temporal based systems a real-time geometry-based solution is presented which is capable of storing multiple versions of a given building or artefact to facilitate the free navigation and visualisation of many differing interpretations of a cultural heritage site over time. Since a geometry-based approach is taken, an individual building existing in the past can easily be visualised in a modern day context and the location of this geometry can be manipulated or queried for

† Email: sdl@cmp.uea.ac.uk

information utilising a three dimensional haptic interface device. These devices are capable of inputting three dimensional coordinates and outputting forces to the user to enhance user interaction. Figure 1 illustrates a two dimensional layout of potential data sources and hypotheses for reconstructions for a given environment. The bottom of the figure illustrates supporting evidence for a virtual reconstruction from which multiple hypotheses can be derived. Each hypothesis for a reconstruction results in a different version and these are laid out along the vertical axis of the diagram, such as the differing reconstructions of the tower at 1712. Each version can contain multiple models depicting changes in time period, which are arranged along the horizontal time axis. The user is able to seamlessly navigate through space, time and version resulting in a real-time five dimensional system.

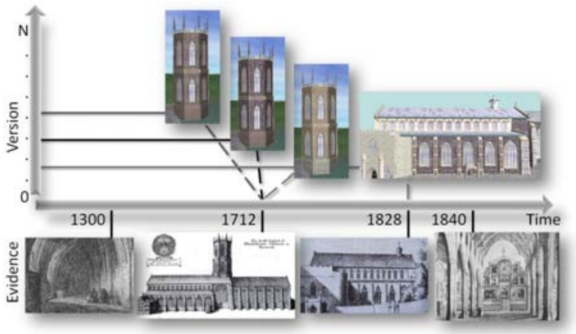

Figure 1: *The reconstructions of a site vary across time and version, as new hypotheses for an object's appearance or purpose are considered.*

1.1. Main Contributions

The main contribution of this work is the conception and development of techniques enabling parts or the entirety of a large scale cultural heritage site to be adjusted interactively in real-time. A user can select a volume of space in which both the time period and the particular hypothetical version can be modified and displayed. An archaeologist is then able to answer queries such as what a pre-existing building may have looked like in the context of modern surroundings. In addition it is possible to verify the validity of a potential reconstruction through a range of time periods. The paper discusses the processes of model construction, rendering and navigating it in real-time, adjusting the time period and version.

2. Previous Work

Extending a three dimensional reconstruction of a site by incorporating time as a fourth dimension is a natural progression for cultural heritage reconstructions. However, unlike the three-dimensional geometry of an artefact the exact time period in which an artefact resides is often not known as precisely. There are many image-based strategies utilising Flash or QuickTime VR which aim to allow the user to blend between time periods to gain an understanding of how an area altered throughout the passage of time. These systems sometimes employ rising and sinking animations or image compositing techniques to depict the temporal existence of a historic entity [ZCG05]. The Provincial Archaeological Museum of Ename utilised QuickTime VR to present the changes occurring to the entire village over the centuries [PSC01]. Vergauwen et al. furthered this work to enable virtual reconstructions of the past to be visualised on present day imagery [VPW*04]. Image-based systems provide a good way of visualising the effects of changes to time but do not promote the ability to visualise what part of a building would look like if visualised in a different time period without the addition of some tedious arrangement of travelling mattes, for example.

Often in cultural heritage the quest for enhanced visual realism has drawn research away from other important factors such as interactivity, sound or touch [RD03]. The previously discussed image-based solutions exhibit some user control but do not facilitate the user in gaining a good understanding of the site under investigation in terms of the spatial proximity of objects. Meeting the requirements of high levels of user interaction and high quality virtual reconstructions has always been a challenging endeavour. To attempt to incorporate user interaction whilst maintaining high quality imagery El-Hakim et al. incorporated both an X3D derived viewer and a high quality rendered animation of the artefacts under investigation [EHML*06]. Haptic feedback has also been used to interact with pre-rendered virtual reconstructions, however, only limited navigation control was available [LLD06]. With current consumer level rendering hardware a real-time implementation can achieve relatively high visual quality whilst benefiting significantly from a high level of interaction. Wüst et al. developed a real-time system for the efficient generation, management and visualisation of large three dimensional environments for cultural heritage projects [WNL04]. Callieri et al. recently developed a visualisation technique for laser scanned objects [CPCS08]. However, no research has considered techniques capable of visualising varied hypotheses, despite this being a property desired by archaeologists and historians. The next section presents information related to St Andrew's Monastic complex, Norwich, UK, which is used as the case study to demonstrate the utility of the techniques presented in this paper.

3. Case Study: St. Andrew's Monastic complex, Norwich, UK

The St. Andrew's Monastic complex in Norwich has a long and varied history ranging from its early beginnings as an English friary through to it being the venue for conferences and concerts today. It is 'the only English friary to survive intact from the medieval period', since most were destroyed or abandoned after Henry VIII's break with the Roman Catholic Church, which makes the St. Andrew's monastic complex an important case for study and conservation, [Nas26]. The site is one of five large monastic complexes that existed in Norwich before the reformation and had originally been granted in 1258AD to the Friars de Sacco. They built a small friary on the site around a simple single storey brick chapel. When they were ordered to leave in 1307, the site was given to the Dominican Black Friars who had established themselves in Norwich in the 13th century. Excavations have revealed that the Dominicans altered this chapel by adding a rib vaulted ceiling, external buttresses, larger windows and an upper storey to the older chapels. These alterations, illustrated in Figure 2 showing Hodgson's 1856 Oil painting of Becket's Chapel, resulted in creating a priory for sixty friars.

Figure 2: *David Hodgson's,(1798-1864), 1856 Oil painting of Becket's chapel alongside a virtual reconstruction.*

4. Model Acquisition

Obtaining a complete and accurate picture of the evolution of the monastic complex is a challenging problem and one which when solved would offer further insight into the heritage of the site and the time periods the monastic complex lived through. The route of the challenge, for many cultural heritage reconstructions, lies in piecing together the diverse and often incomplete data that supports the layout and function of the site over time. In the proposed approach multiple reconstruction hypotheses are constructed where each reconstruction, realised in the form of a textured 3D model, is valid for a given time period. The 3D models may originate from a number of sources including laser scanning or 3D modelling packages such as 3DS Max or Maya. However, to ensure the validity of the latter approach the modelling work must be conducted closely with experts in the field and archived data resources.

The reconstruction of the 3D content for St. Andrew's monastic complex was undertaken through consultation with both archaeologists and historians. Furthermore, since this was a public building there are a series of architectural drawings in existence and a large number of other images, which were created over the last 400 years, including prints, drawings, and paintings. It is the quality of these artefacts that enables high fidelity reconstructions to be undertaken, such as those illustrated in Figure 3 for St Andrew's Hall at 1780 and 1840. It is fortunate that this information is available, but such a rich source of information is not always available, nor is its coverage complete or its life span fully determined.

Figure 3: *Top: Artistic impression of St Andrew's Hall in 1780 alongside the virtual reconstruction, Bottom: C. Greenwood's impression of St. Andrew's Hall in 1840 alongside the virtual reconstruction.*

Additional information, to guide the reconstruction process and to help to determine the life span for a model, is obtained from archaeological evidence originating from observations undertaken whilst building work was taking place in 1863. Furthermore, a series of archaeological excavations were carried out at various times over the last century. Through discussions with the experts and by referring to the archived resources a collection of over two hundred three dimensional components were constructed, requiring over five million triangles and 291 texture maps in total. Each component was created and saved in 3DS Max in separate files, with each representing a group of objects, e.g. the south wall, crypt or ceiling beams. Not all components could be accurately reconstructed in their entirety, since some of them are not illustrated in any of the archived resources. For instance, whilst the main arches of the cloisters remain, the inner detail has long since disappeared. Consequently the inner cloisters were modelled according to the general pattern of the period, examples of which exist in the nearby Norwich Cathedral. The lack of complete data meant that possible alternatives should be considered leading to a collection of components that were stamped based on time and version.

All the models for a particular version are located in a single folder and the file names of each model are assigned two separate times, which indicate the limits of the object's possible times of existence. This enables a model, such as the south wall, to be dated based on the supporting evidence. A high quality visualisation of the cultural heritage site, at a given time period and depicting a particular hypothesis, can be created by combining all the separate model files for the chosen time period from a particular folder into the same scene. Each component is positioned within the British National Grid coordinate system, enabling the alignment of the objects and the integration of georeferenced data. To ensure that the task of collecting objects into a single scene is not a laborious one a Maxscript file is used to extract those files whose time period limits encompass the desired time period. Once the models are imported a render can be obtained, such as those presented in Figures 2 and 3. These offer a single snapshot of the site at a given time period and may be combined using post production effects as discussed in the previous work section. The high quality results are visually appealing, although the offline rendering paradigm is not the most desirable when attempting to validate many possible hypotheses quickly, nor does it permit an intuitive approach to view different components within the context of an alternative interpretation or time period. It is the introduction of interactivity that enables the full exploration and analysis of the multidimensional space of models.

5. Rendering large scale cultural heritage sites in real-time

Laing et al. illustrated through experiments that users navigating virtual environments feel more motivated and attentive when actively navigating a desktop virtual model of an environment than by simple passive observation [LCC*07]. Furthermore, Laing et al. determined that human scale navigation is important to obtain a realistic impression of the environment. Drawing from their conclusions this system utilises a real-time rendering approach to provide full control to the users with the aim of increasing the fidelity and level of immersion. The increased fidelity of the scene, is often attributed to the inclusion of millions of triangles to accurately depict the geometric structure of 3D artefacts. This amount of triangles is not prohibitive, if the geometry is mapped to the graphics accelerator card appropriately.

The introduction of vertex and texture buffer objects to the NVIDIA Geforce range now permits much more geometry to be efficiently loaded and rendered. Therefore all the models are loaded into the vertex buffer objects, VBOs, to aid their rendering in real-time. Whilst a notable improvement is achieved using VBOs, this technique should augment the existing real-time techniques. Implemented in the current rendering engine, written using C++ and OpenGL, are view frustum culling and occlusion culling. The view frustum culling is performed by the CPU and entails the determination of the existence of a model's axis aligned bounding box within the view frustum. The occlusion culling is performed on the graphics processing unit, GPU, using hierarchical occlusion culling exposed by the Occlusion Queries extension in OpenGL. During scene traversal, to exploit the occlusion queries to their full potential the objects in the scene are sorted based on their distance to the camera. By sorting the objects, they may be rendered in front to back order to ensure that the objects nearer the camera have maximum effect as occluders for the rest of the scene.

To further reduce the rendering computation all the triangles within each object are sorted based on their applied material. This minimises the state changes required by OpenGL and is performed as a preprocessing stage. In order to maintain a visually accurate scene all 300 textures from the original reconstructions are employed, but are first automatically converted into DirectDraw surface, DDS, images. The size of the image is constrained to reduce the memory overhead; although the DDS format provides a 6:1 compression ratio. These techniques ensure that the cultural heritage site can be traversed smoothly and at real-time frame rates.

A real-time visualisation offers a number of advantages via the local adjustments of objects independently from the rest of the site. Once in this framework additional information such as the time period or version of an object can be visualised using the OpenGL Shading Language, GLSL. Section 7 details how the rendering of objects is altered efficiently. Furthermore, the flexibility of a real-time system can be exploited by employing clip planes to enable the internal structure and cross section of a building to be observed whilst visualising parts of the exterior as well. This is achieved by a simple adjustment to the vertex shader and rendering implementation.

6. Scene Navigation

Once the site is visualised in a real-time framework an enhanced understanding of the virtual reconstruction can be obtained provided that the user is able to interact freely with their environment [LCC*07]. Two separate modes, labelled study and fly, have been developed to facilitate free navigation of the virtual environment. The study camera allows the user to rotate and translate with respect to a fixed point in the scene. The fly mode adopts the usual flight simulator style camera movements to enable free flight around a virtual reconstruction. An additional passive mode is implemented to play back paths that have either been pre-recorded or created in a modelling package. This facilitates the possibility for a user to record a particular path of interest.

A novel haptic rendering approach is developed to provide force feedback to the user whilst navigating and exploring the three dimensional scene. Adjustments to the user's camera position and look direction can be manipulated utilising a unique force feedback box. The haptic feedback device, SensAble's Phantom Omni, employed has two buttons one of which must be pressed to adjust the view. The position of the haptic device, HIP, is recorded at the instant a button is depressed. The position at the button event, P_E, defines the centre of a cube with dimension, $2DIM$. Figure 4 illustrates the cube, HIP and P_E that results during a typical viewing adjustment. Equation 2 details the force returned to the haptic device and Equation 3 can be used to determine how much of a viewing adjustment should be undertaken. The force box paradigm has the property that the user must move by a distance greater than *DIM* before an adjustment in that axis takes place. For instance, if the box is used to control the study operation the user may be permitted to move a small distance in the y-axis with no effect, whilst rotating the view based on a more significant movement in the x-axis. This property alleviates issues with imprecise human movement whilst facilitating the ability to easily perform perfect adjustments to one parameter of a viewing equation. (The threshold used for *DIM* may be altered by the user or turned off, if they prefer.) Furthermore, the higher sampling rate of the haptic device of 1000Hz permits smooth camera control and the force feedback leads to a more immersive user experience.

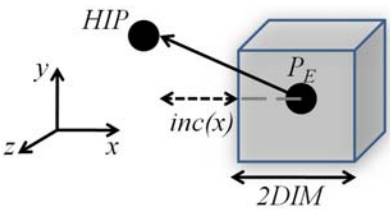

Figure 4: *The force feedback box centred at P_E aids the manipulation of three independent variables.*

$$
\begin{aligned}
d &= HIP - P_E && (1)\\
force &= -dk && (2)\\
inc(v) &= \begin{cases} (d_v - DIM) & \text{if } d_v > DIM \\ 0 & \text{otherwise} \end{cases} && (3)\\
\text{where} \quad k &= spring\ constant\\
P_E &= position\ of\ the\ HIP\ at\ a\ button\ event\\
DIM &= half\ the\ width\ of\ the\ force\ box\\
v &\in \{x,y,z\}
\end{aligned}
$$

7. Time Period and Version Adjustment

Navigating a cultural heritage site in real-time enables the exploration and consequently further understanding of the spatial relationships between objects. However, techniques are required to reflect the multiple reconstructions based on varied hypotheses for an environment. As stated in Section 4 many different models are constructed for a particular heritage site, each being assigned a candidate valid time period and version number. There are often inconsistencies, which have to be resolved, with both the time periods and the interpretations. In the case of St. Andrew's Hall inconsistencies in the association of time periods result in areas where little or no information is available. For example, there have been several locations of the main entrance, and at least two different main entrance porticoes have existed on the south side following the removal of the medieval houses. The exact time periods for the existence of these porticoes is unclear. Whereas records for the current porticoed entrances are more concrete confirming that they were both constructed at the time of the Victorian renovation in 1863. The lack of information for an area results in an inevitable uncertainty with regard to the reconstruction hypotheses. For instance, the Dominicans friary had a tall polygonal tower, which fell down in 1712. The only reliable surviving contemporary image of the tower comes from 'Monasticum Anglicanum' of the 17th century illustrated in the bottom of Figure 1. In particular, there is conjecture as to the materials that were used for the facing of the tower.

With the lack of supporting archived data, available materials of the time were considered for the tower walls. Norwich is in an area dominated by chalk where there is little stone but an abundance of flint. This was the main building material for medieval churches and larger buildings in the city, where flint was used to face rubble cores. In Norwich bricks were used where normally stone might be expected to be used. For example, on window edgings and the main arches on the exterior walls of the friary cloisters. The different uses of material during the period in question motivated discussion and resulted in multiple reconstructions for the tower being modelled, as illustrated in the top of Figure 1.

In order that each component's life span and its visual appearance can be visualised and cross referenced with other time periods, all models are combined into a single visualisation framework. This is achieved by recording both the local and the global time and version information. The global information concerns the current time period and version number being displayed to the user, for the entire site, and is shown in the top right hand corner of the screen. Figure 6 illustrates a screen shot of the program presenting the monastic complex at the global time of 2006. By holding the blue button of the haptic device and dragging horizontally the global time period is increased or decreased. By dragging vertically the global version number is altered. The alteration of the version number permits alternative reconstructions to be observed within the current time period. The local time period and local version number refer to the time period and version number being visualised in a selected volume of space. By holding the white button of the haptic device a three-dimensional volume of space is easily constructed and any objects within that space are immediately assigned the local time period and version, which initially is set equal to the global values. Once an area of space has been defined the act of dragging with the blue button held down will alter the time and version numbers of only those objects that intersect the selected volume. This process was undertaken in Figure 6 to select the volume of space encompassing the tower and local time period and version are altered from the global values to 1696 and 1 respectively.

The separation into local and global information facilitates the composition of objects from different hypotheses and time periods. Since the start and end times for an object's time period are not always completely determined it is desirable to visualise an object within a specified threshold value, ε. (The value for ε is initially set to 20 years but it can be modified at any time.) Therefore to facilitate this, transparency is used to enable the visualisation of many potentially overlapping time periods for the objects in the scene. The following equation states how the transparency value, tA, for a time period, t, can be calculated based on an object's start time, sT, end time, eT, and the user specified threshold, ε.

$$tA(t) = \begin{cases} (t-(sT-\varepsilon))/\varepsilon & \text{if sT-}\varepsilon \leq \text{t} \leq \text{sT} \\ 1-((t-eT)/\varepsilon) & \text{if eT} \leq \text{t} \leq \text{eT+}\varepsilon \\ 1.0 & \text{if sT} \leq \text{t} \leq \text{eT} \\ 0.0 & \text{otherwise} \end{cases}$$

In addition to the transparency for the time period, an alpha value is used to ensure the smooth transition between different versions. All the objects within a single folder have the same version number. Any object which has a version number within plus or minus O of the global version number will be completely visible. O is the opaque width, which is set to 0.3. A blend width, B, of 0.4 is defined where objects will blend between fully visible and transparent. The following equation states how the version transparency, vA, is determined based on a version number, v, where W is equal to the version width, $W = B + O$.

$$vA(v) = \begin{cases} (v-(N-W))/B & \text{if N-W} \leq \text{v} \leq \text{N-O} \\ 1-((v-(N+O)/B) & \text{if N+O} \leq \text{v} \leq \text{N+W} \\ 1.0 & \text{if N-O} \leq \text{v} \leq \text{N+O} \\ 0.0 & \text{otherwise} \end{cases}$$

For a given time period, t, and version, v an object's alpha is equal to $tA(t) * vA(v)$. Figure 5 illustrates the blend regions for an object.

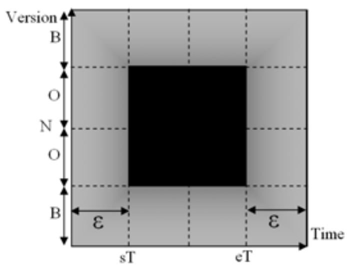

Figure 5: *The alpha value for an object is based upon both time and version.*

As stated in Section 5 the objects in the scene are depth sorted to enhance the benefits obtained from occlusion culling. However, to ensure that the transparency is also correctly visualised two sorted lists of objects are obtained, whenever the camera moves or the scene is modified. The first list represents the opaque objects, which may be rendered as stated earlier, and the second represents the transparent objects. During real-time traversal these lists of objects are compiled from those that are completely interior to the view frustum based upon the calculated alpha values. For those objects, in the transparent list, the distance from the camera is determined and for any object with the same distance as another object in the list the order is adjusted such that objects with a larger opacity value are pushed to the front of the list. The renderer subsequently renders all the opaque geometry first, followed by the remaining objects ordered according to their distance from the camera and their alpha value.

The described technique enables the site to be visualised at multiple time periods and reconstruction hypotheses. To improve the level of immersion and therefore increase the understanding of the cultural heritage site, the user is given the ability to interact directly with the objects, as opposed to just observing them. The haptic device's ability to input three dimensional information significantly improves upon a 2D interface and therefore enables the user to easily select any object in the scene and to modify the object's transformation.

8. Selecting and Modifying objects

To select objects the intersections between the haptic device position, *HIP*, and the triangles forming the objects in the scene must be calculated. To ensure this can be achieved efficiently, given that reconstructions can comprise of millions of triangles, spatial data structures are employed. Initially a regular grid of boxes, labelled the scene_grid, is constructed around the entire virtual reconstruction site. A linked list of objects is stored with each grid box containing those objects which pass through it. A regular grid is also constructed around each of the individual objects, labelled the object_grid. An octree is utilised to determine which triangles pass through which boxes in the object_grid. With these structures in place determining exactly if any triangles forming an object are close to the *HIP*, becomes a very efficient task. Firstly a constant time operation takes place to determine which box from the scene_grid contains the *HIP*, assuming the *HIP* is initially transformed by the inverse of the viewing matrix. The *HIP* can then be tested with each object intersecting the previously identified scene_grid box to determine the closest triangles. Since each object can be transformed to any location in the scene the *HIP* is transformed into the local coordinate frame of each object before testing for triangles in close proximity. By storing the viewing matrix for the scene and a separate transformation matrix for each object a fully dynamic and navigable environment can be realised.

When the user moves the *HIP* close to an object the *HIP* position is recorded, P_E, and the system currently displays just the name and duration of existence for the selected object. However, there is great potential to link to a relational database to allow object specific information such as supporting evidence or other comments regarding the reconstruction process to be visualised. The object remains in selection until the user pulls the haptic device a sufficient distance away. A spring force is calculated based on the vector between the HIP and P_E. The haptic device provides a means to easily select objects within the environment. In addition to this by depressing one of the two buttons on the device the system supports the functionality of rotating and translating selected objects within the reconstruction. It is envisaged that this facility is mostly to be used for fine adjustment since the proposed system's main function is that of a comparison tool as opposed to an interactive modeller. The force box paradigm, from Section 6, is employed to give the user the ability to displace an object in the x-axis even if small movements in other axes occur. This strategy provides significant benefits over the employment of transformation gizmos utilised in modelling packages, which are tedious to use, particularly for novice users.

9. Camera collision detection

The spatial data structures, utilised previously, also offer benefits to achieving an enhanced sense of presence when navigating the site using fly mode. This is realised through the calculation of collisions between the camera and the triangles in the environment. On detection of a collision the camera's position is constrained to the triangle it collided with and a noticeable discontinuity in the force feedback is presented to signify an impact. The force discontinuity is created by setting the position at the previous button event, P_E, to be equal to the current HIP and thus reducing the force to zero on impact. To perform the collision detection a line segment is constructed emanating from the camera's eye position and terminating a small distance away in the direction of the camera's look direction. The closest intersection between the segment and the triangles in the scene must then be calculated. Firstly the range of scene_grid boxes the segment overlaps must be determined to enable a list of potentially colliding objects to be determined. The range of object_grid boxes can then be determined easily once the segment has been transformed to the local coordinate system of the object. An optimised segment-triangle intersection approach is utilised to determine the intersection between the segment and each of the triangles located in the range of the object_grid boxes. The intersection algorithm returns a percentage identifying the distance the intersection point is from the start of the segment.

The camera can then be reversed based on the distance between the end of the segment and the point of intersection to ensure that it is left in a free location. Since the segment is very short the range of boxes in both the scene_grid and the object grids are very small and thus this approach is able to reject many potential collisions early in the algorithm ensuring efficient operation in practice. Given that the system is currently not CPU intensive the camera collision detection stage can be performed with no loss in frame rates even when the reconstruction contains several million triangles.

10. Results

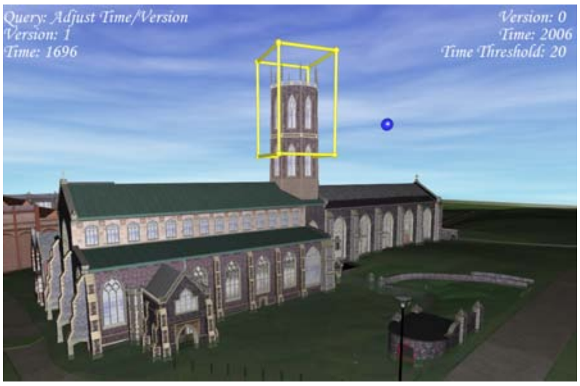

Figure 6: *A real-time screenshot of St. Andrew's Monastic complex in 2006 with a reconstruction of the tower, which fell in 1712.*

Figure 6 presents a screen shot from the real time simulation. It illustrates the presentation of the St. Andrew's Hall monastic complex in 2006. In the screen shot the user has selected a volume of space and adjusted the time period and version numbers to present the tower, which fell in 1712, within a current context.

Whilst the main result of this work is the research into techniques to display multiple time-dependent hypotheses it is important that the navigation of the site can be achieved at interactive rates. In order to show the efficiency of the proposed techniques the 210 components representing St. Andrew's monastic complex were loaded directly from 3DS max into the rendering system. Table 1 illustrates the results for four tests, with each test being measured on the number of objects, #Objs, the number of triangles, #Tris, the number of visible triangles, #Vis and the minimum and maximum frames per second observed. The scene was duplicated and translated 300 metres in the x direction to increase the polygon counts for tests 3 and 4. For tests 1 and 3 a 360 degree spin inside St Andrew's Hall was undertaken and for tests 2 and 4 a view from around the entire model was carried out. Due to the nature of visibility culling the difference between the minimum and maximum frame rates, fps, is inevitable. In fact, when positioned immediately facing a building the frame rate will rise to above several hundred frames per second. However, the results illustrate that an interactive frame rate is achieved when displaying 3.0M triangles pertinent to a particular time period from a scene comprised of 10.5M triangles. (The number of visible triangles will fluctuate during the traversal, therefore the average number of visible tri-

Test	#Objs	#Tris	#Vis	min FPS	max FPS
1	210	5.3M	0.8M	50	75
2	210	5.3M	1.5M	46	50
3	420	10.5M	1.0M	25	30
4	420	10.5M	3.0M	12	14

Table 1: *Real-time rendering results for the cultural heritage site.*

angles is presented.) The haptic interaction is calculated in under 0.003ms on the CPU. All the tests were undertaken on a Pentium 2.4GHz Quad core PC with 3GB of RAM and a Geforce 8800 GTX graphics card with 768MB of video memory.

11. Conclusions and Future Work

In this paper a set of techniques have been presented capable of visualising multiple reconstruction hypotheses for the development of a cultural heritage site through time. The inconsistent and partial evidence for some time periods leads to conjecture with regard to the appearance, structure and functionality of cultural artefacts and these techniques provide a means to verify and visualise these hypotheses. The visualisation is achieved on consumer level hardware, which with efficient rendering techniques provides real-time navigation. A haptic device is utilised to obtain 3d input from the user for navigation and modifications to the scene, in order to progress towards a more immersive user experience. Furthermore, the 3d positioning afforded by the haptic device enables an easy method for selecting objects, which may be useful for linking the visualisation techniques to a database containing the supporting evidence for the object's reconstruction.

Currently the techniques are bound to the graphics processing unit, leaving valuable computational power on the CPU to be utilised. This resource could be employed for software scene optimisations such as the computation of improved lighting models. An alternative extension of this work could involve the incorporation of physical models with the objects, facilitating the ability to test an object's functional aspects or to visualise these aspects for an increased awareness of the cultural site.

12. Acknowledgments

The authors wish to thank HEART (http://www.heritagecity.org/), for funding the development of the three dimensional reconstructions.

References

[CPCS08] CALLIERI M., PONCHIO F., CIGNONI P., SCOPIGNO R.: Virtual inspector: A flexible visualizer for dense 3d scanned models. *IEEE CG&A 28*, 1 (2008), 44–54.

[EHML*06] EL-HAKIM S., MACDONALD G., LAPOINTE J., GONZO L., JEMTRUD M.: On the digital reconstruction and interactive presentation of heritage sites through time. In *VAST* (2006), pp. 243–250.

[LCC*07] LAING R., CONNIFF A., CRAIG T., DIAZ C., SCOTT S.: Design and use of a virtual heritage model to enable a comparison of active navigation of buildings and spaces with passive observation. *Automation in Construction 16(6)* (2007), 830–841.

[LLD06] LAYCOCK R., LAYCOCK S., DAY A.: Haptic navigation and exploration of high quality pre-rendered environments. In *VAST* (2006), pp. 17–24.

[Nas26] NASH P.: The sackfriars and blackfriars conventual buildings in the parishes of st andrew and st peter hungate, norwich. *Norfolk Archaeology XXII* (1926).

[PSC01] PLETINCKX D., SILBERMAN N., CALLEBAUT D.: The saint laurentius church in ename and its role in the francia media heritage initiative. In *VAST* (2001), pp. 197–204.

[RD03] ROUSSOU M., DRETTAKIS G.: Photorealism and non-photorealism in virtual heritage representation. In *1st EG Workshop on Graphics and Cultural Heritage, VAST* (2003).

[Rou02] ROUSSOU M.: Virtual heritage: From the research lab to the broad public. In *In Franco Niccolucci (ed.), Virtual Archaeology: Proc. of the VAST 2000 Euroconference Arezzo, Italy, published as BAR International Series 1075, Oxford, Archaeopress* (2002), pp. 93–100.

[SR02] SIDERIS A., ROUSSOU M.: Making a new world out of an old one: In search of a common language for archaeological immersive vr representation. In *8th Int. Conf. on Virtual Systems and Multimedia* (2002), pp. 31–42.

[SS02] SIDIROPOULOS G., SIDERIS A.: Requirements and assumptions in visualization process of urban and surrounding areas. In *Computers Applications in Archaeology* (2002).

[VPW*04] VERGAUWEN M., PLETINCKX D., WILLEMS G., VERBIEST L., VAN GOOL L., HELSEN T.: As time flies by: Mixed image and model-based rendering of an historic landscape from helicopter images. In *VAST* (2004), pp. 241–250.

[WNL04] W*ü*ST T., NEBIKER S., LANDOLT R.: Applying the 3d gis dilas to archaeology and cultural heritage projects - requirements and first results. In *Geo-Imaging Bridging Continents XXth ISPRS Congress* (2004).

[ZCG05] ZUK T., CARPENDALE S., GLANZMAN W.: Visualising temporal uncertainty in 3d virtual reconstructions. In *VAST* (2005), pp. 99–106.