Emmanuel Durand

1. CONTEXT AND RELATED WORKS

This paper describes the work done to improve the visitor's experience while visiting the historical site of Cluny abbey, in Burgundy (France). The church of this abbey, founded in 910 and which was the biggest church of Christendom until the 15th century, has been almost completly destroyed after the French Revolution and so is hard to figure out for the visitors. The usage of virtual reality is one of the way to improve the understanding of the architecture of the site.

Of all cultural heritage subjects, buildings present the difficulty of their size, making it less evident to confront the real site to any virtual representation of it. The obvious solution to this is the installation of on-site devices, or the use of mobile devices.

On-site augmented reality applications dedicated to cultural heritage have appeared more than ten years ago. A well known application which pioneered this research field is the Archeoguide project [Vassilios et al. 2001] (figure 1), which aimed at giving visitors a view of what the famous building of Ancient Greece looked like, on-site. The user, who had to wear a see-through head-mounted display (HMD), was geo-localized and 3D models of the buildings were displayed on them. Computation was taken care of by a laptop the user had to carry in a backpack. Later on, Zoellner et al. [Zoellner et al. 2007] used an ultra mobile PC (UMPC) to handle the computation and display in the iTACITUS project. Tracking was done through GPS localization and improved by video tracking. The goal of this project was to overlay old prints related to what was capture in the live video feed.

Fixed applications also arose, of which we can cite the Virtual Museum Guide from the Fraunhofer Institute. This takes the shape of a dedicated device constituted of a screen mobile around the vertical axis, with a camera targeting a wall covered by a poster. The device added contextual elements (be it texts or images) on top of the captured video feed. On another angle, the CITU showed a prototype of an "augmented reality telescope" which was installed on top of the Arc de Triomphe in Paris for a week. This device gave textual information about Paris from this point of view (figure 2).
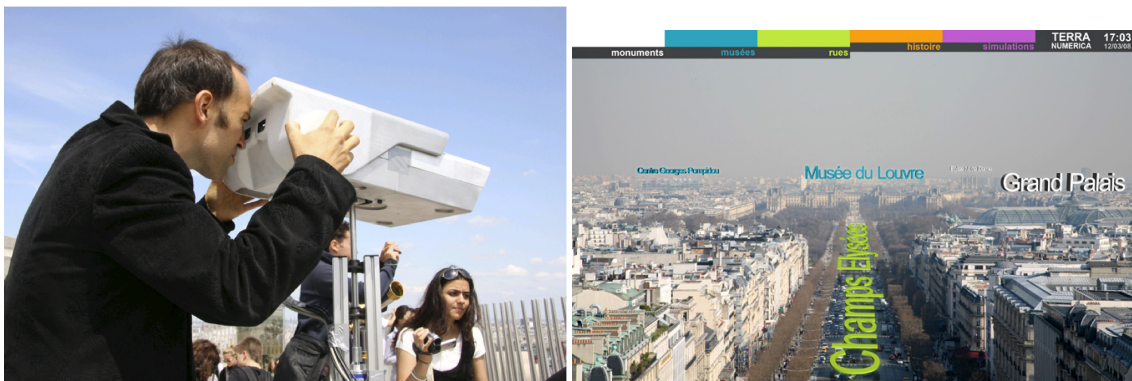
Fig. 1. The Archeoguide project

Fig. 2. The AR Telescope prototype

Finally, the quick expansion of the smartphone market led to the apparition of applications dedicated to these portable devices, of which we can site the Streetmuseum application developed by the Museum of London (figure 3) which gave a similar experience to that of the iTACITUS project, less the specific hardware (UMPC never reached the same market share as smartphones). Also, middleware dedicated to augmented reality on smartphones exist, like the one from Layar.

Fig. 3. The Streetmuseum application

2. SCIENTIFIC ISSUES

Cultural heritage virtual reality applications have the same goal as all the other various techniques applied to cultural heritage: improve the understanding and the knowledge of the visitor about the heritage he is confronted to. The hard part is to allow the visitor to easily find the information he needs, without distracting his attention from the subject.

In the context of using mixed reality in a historical site, this means that the application should immerse the user in its virtual content, without breaking the connection with the real site. Also if the site is displayed in one of its previous form, it should be possible to compare easily the differences between now and the virtual "then", so that to help the visitor to comprehend the history of the site.

This leads to the immediate question of how to immerse in a virtual representation of the site sufficiently precise to show details in common with the real site, while not drawing attention from it.

3. PROPOSED APPROACH

From the context of this work (described in the introduction), we chose to create an uchronic window which would show the user a view of the site as if it was still as it existed at a specific time in the past. Through such an imaginary window, the user would have a full view on the virtual site while still being conscious of the surrounding real site. The result is, from the user's point of view, a merging of the real and virtual site, which may be enhanced by the use of mixed reality techniques on the virtual side.

To achieve this result, we underlined to main principles. The first is the geometric consistency between the real and virtuals worlds. This principle implies that whenever a virtual object is the representation of a still existing real couterpart, both should merge at the edge of the display device used. As an effect, this would reduce the distance between both sides of the display and ease the correspondence between them.

The second principle is the visual consistency, which suggests that both realities are easier to merge if they are visually close. This observation has already been explored in the case of augmented reality where multiple video sources, real and virtual, would be processed so as to look identical in the end: in addition to the obvious choice of rendering virtual content photo-realistically, works from Fischer et al. and Chen et al. showed how reducing the realism of real view can help in the context of augmented reality [Fischer et al. 2005; Chen et al. 2008]. In the current context it is obviously not possible to modify the look of the real site, so the way to go is to render the virtual site in a photo-realistic way, to get close to the real view.

The concept of a window implies a sufficiently wide display to act as such. Though portable tablets could suffice to a certain extent, positioning is not precise enough to match the first principle of geometric consistency. Thus it was decided to use a fixed display, orientable around two axis so that it would be possible to target any area surrounding the device.

4. GEOMETRIC CONSISTENCY

4.1 Original design

For our purpose, geometric consistency implies that virtual models and real objects are aligned. To meet this objective, the point of view of the user must be known and reproduced accordingly in the virtual rendering. Once it is known, all the virtual camera parameters are immediately accessible as soon as we know the dimensions and position of the display.

The user's position is not measured but imposed. The device is free to move around two axis, as the user handles it by grabbing it from both side and moving it around. Thus being in front of the device, at arms length, is the most confortable position to adopt when using it. By chosing mean values for the
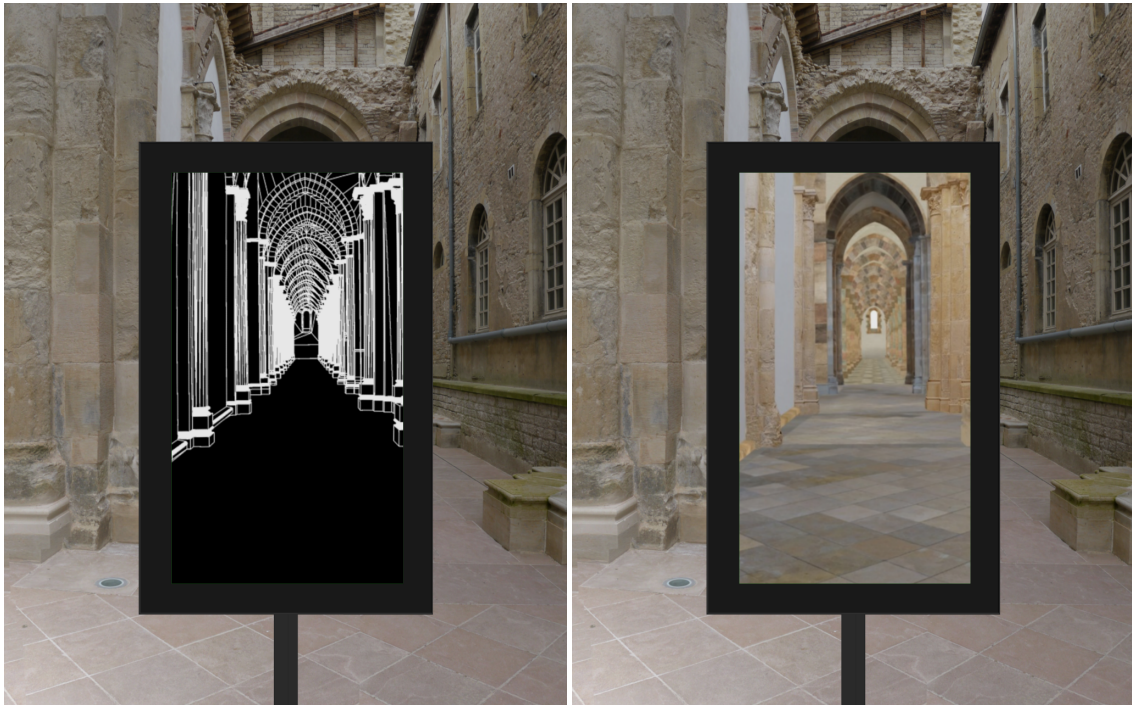
Fig. 4. Left: geometric consistency; Right: Visual consistency

height of the device and arms length, a average virtual camera is calculated and used for the rendering of the virtual scene.

Finally, the orientation of the device is measured through absolute sensors, ensuring a non drifting input of the actual posture of the device and the user. Although it is worth noting that the center of rotation is situated close to the display to reduce point of view variations that could be otherwise visible, due to the rendering method used and explained in the next section.

4.2 Alternative design

The previous proposition is especially effective when real objects are close to the visitor / device. Looking at real ancient columns going down from the ceiling to the floor, passing behind the display while still being visible in their virtual, historically reproduced form is an interesting view. But outdoor setups are also necessary to picture the whole building or site. Such point of views can be too far away from the building so that geometric consistency is not as useful.

For these situations, we imagined a two steps alternative. While immobile, the device would still act as a window, showing the historical site as it was once. Then, when moved by a visitor, the field of view of the virtual camera would tighten to zoom and show details invisible at this distance, even more if they desappeared with time. While the user changes the orientation of the display, the virtual camera follows the movement and so acts as a spotting scope.

5. VISUAL CONSISTENCY

As explained earlier, visual consistency is the second condition to attain the goal of simulating a uchronic window. Though multiple methods have been proposed to enhance the coherency between

real and virtual video feeds, our specific case where one the visual sources is the real surrounding of the user, so there is no choice but reproducing correctly this look.

5.1 Lighting reproduction

Towards this goal we used an image-based rendering method (IBR) as described by Callieri and Debevec [Callieri et al. 2005] and later Pessoa et al. [Pessoa et al. 2008; Pessoa et al. 2010]. This choice is dictated by the complexity of the virtual mockup of the historical site, which reduces the possiblity to render the scene in real-time with the lighting effects needed to obtain a photo-realistic look. This complexity is mostly hidden with an IBR technique, by pre-computing a large part of the render leaving only the lighting to transpose to the virtual result.

The IBR technique used in this work is applied to the rendering of a full panorama from the point of view of the device. From the observation that a given lighting environment is equivalent (from a given point of view) to a sum of directional lights, it is possible to a certain extent reproduce any lighting conditions by summing up multiple rendering of the same panorama with the only difference between them being the position of a single directional light. Then, reproducing a lighting conditions to a virtual panorama is a three steps procedure:

—render the panorama for a number n of different, unit directional lights,

—capture a high dynamic range (HDR) image probe of the lighting environment, and decompose it to the same directions as used for rendering,

—accumulate the renderings according to the weights calculated for each direction in the previous step.

Under the hypothesis that the renderings have been calculated with a physically correct rendering engine, and that the camera used for grabbing the light probe is colorimetrically and photometrically calibrated, the accumulation step will produce a correct, high dynamic range reproduction of the lighting conditions [see annexe?].

5.2 Augmented reality

Taking advantage of the lighting reproduction, it is possible to ease the merging of virtual and real live video feeds. One of the difficulties of this step is the consistency between the various merged sources, which needs tuning the color balances and exposure / brightness settings for them to match. With changing lighting conditions, these parameters are subject to change too, so on-site installations need automatic methods to update the parameters.

Let's now consider the previous calibrated lighting reproduction method, combined with a calibrated camera grabbing a live video feed of the scene we want to augment with virtual elements (see figure 5). As both sources are calibrated, merging them becomes a simple matter from a color and exposure perspective and is a two step process:

—merge geometrically both sources (with any cropping, masking and geometrical transformation needed),

—apply a tone-mapping on the resulting HDR image to obtain a displayable image.

The main benefit of this approach is its robustness to lighting changes, keeping at any time an identical look for all merged sources.

6. THE PILOTE SITE OF CLUNY ABBEY (BURGUNDY, FRANCE)

As described in the introduction, the development of this device had from the beginning a target in the site of the old and mostly destroyed church of Cluny. A small part of the original building remains

Fig. 5.   Merging of the virtual church into the live video-feed (top: real view, bottom: augmented view)

(around 10%), and a recurring comment from visitors was that it was difficult to imagine the site in its previous state as most of the traces in the ground have been replaced with newer buildings. Four devices are currently installed, giving very different points of views of the church (figure 6).
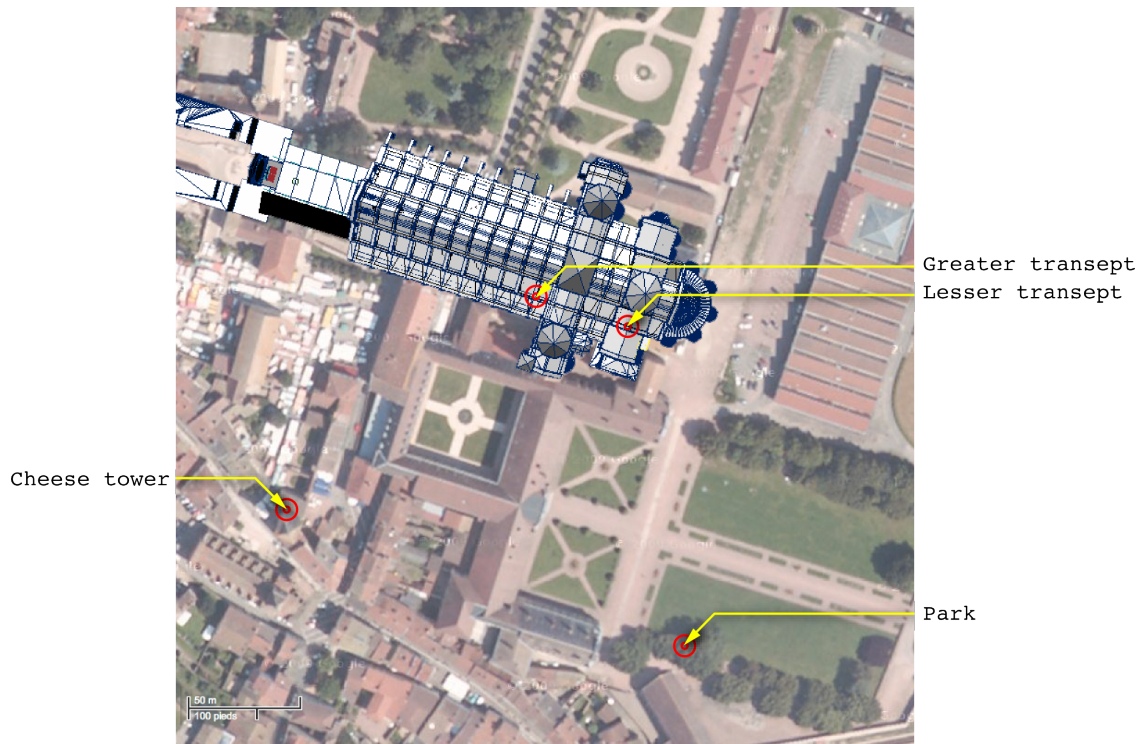
Fig. 6.   Location of the four devices around the church of Cluny

The visiting path goes through the remaining southern parts of the transepts, where two devices are installed and configured by the original design described in the section Geometric consistency 4 (see figure 7). These devices show fully virtual panoramas of the inside of the church, with the reproduction of the lighting environment.

Next along the visiting path is the park located right next to the remains. At the opposite side of this park has been installed a device using the alternative design described earlier. From this point of view, a full view of the church is visible and merged with the existing, more recent building, as well as a precise view of the details of the ouside of the choir, the southern transept and most tower (see figure 8).

Finally, a last device is installed outside of the visiting path, in a nearby tower from which visitors have a unique side view of the site of the church. This device has introduced the use of augmented reality and paired with a camera grabbing the view, the virtual and real views being merged as described in the previous section (see figure 9).

7.   DISCUSSION

This paper described the work done to design and build an on-site virtual reality system which complements and integrates with the surrounding remains and helps the understanding of the historical site without hiding or replacing it. From the comments visitors and guides of the Cluny abbey returned, these devices proved to be efficient and met their goal: it made the work of guides more illustrative, and visitors on their own can enjoy full views of the building.

Fig. 7.   Indoor device in the transept of the church

Fig. 8.   Outdoor device installed in the park

Fig. 9. Device located in a distant tower

Among the improvements this device could benefit from, the most immediate are those related to the enhancement of the window analogy. Though it seems difficult to allow multiple user to have a correct view with a single device, integrating head tracking for a single user would greatly improve the feeling of being in front of a window. Also, replacing the actual rendering method using fixed panoramas to a full 3D rendering engine (but keeping lighting reproduction by using IBR on textures) would allow the user to see the parallaxe (either when moving the device or his head) and thus feel the depth in the display.

REFERENCES

M. Callieri, Paul E. Debevec, J. Pair, and R. Scopigno. 2005. A realtime immersive application with realistic lighting: The Parthenon. (2005).

Jiajian Chen, Greg Turk, and Blair MacIntyre. 2008. Watercolor Inspired Non-Photorealistic Rendering for Augmented Reality. *VRST* (2008), 231–234.

Jan Fischer, Dirk Bartz, and Wolfgang Strasser. 2005. Stylized Augmented Reality for Improved Immersion. *Virtual Reality* (March 2005), 195–202.

S. Pessoa, G. Moura, J. Lima, V. Teichried, and J. Kelner. 2010. Photorealistic rendering for Augmented Reality: A global illumination and BRDF solution. In *Virtual Reality Conference*. 3–10.

Saulo A. Pessoa, Eduardo L. Apolinario, Guilherme de S. Moura, and Joo Paulo S. do M. Lima. 2008. Illumination Techniques for Photorealistic Rendering in Augmented Reality. *ISMAR* (2008), 223–232.

Vlahakis Vassilios, Karigiannis John, Tsotros Manolis, Gounaris Michael, Almeida Luis, Stricker Didier, Gleue Tim, Christou Ioannis T., Carlucci Renzo, and Ioannidis Nikos. 2001. ARCHEOGUIDE: First results of an Augmented Reality, Mobile Computing System in Cultural Heritage Sites. *VAST* (2001).

Michael Zoellner, Didier Stricker, Gabi Bleser, and Yulian Pastarmov. 2007. iTacitus - Novel Interaction and Tracking Paradigms for Mobile AR. *VAST* (2007).