Multiscale acquisition and presentation of very large artifacts: the case of Portalada

Marco Callieri [(1)], Antoni Chica [(2)], Matteo Dellepiane [(1)], Isaac Besora [(2)], Massimiliano Corsini [(1)], Jordi Moyés [(2)], Guido Ranzuglia [(1)], Roberto Scopigno [(1)], Pere Brunet [(2)]

The dichotomy between full detail representation and the efficient management of data digitization is still a big issue in the context of the acquisition and visualization of 3D objects, especially in the field of the Cultural Heritage. Modern scanning devices enable very detailed geometry to be acquired, but it is usually quite hard to apply these technologies to large artifacts. In this paper we present a project aimed at virtually reconstructing the impressive (7x11 m.) portal of the Ripoll Monastery, Spain. The monument was acquired using triangulation laser scanning technology, producing a dataset of 2212 range maps for a total of more than 1 billion triangles. All the steps of the entire project are described, from the acquisition planning to the final setup for dissemination to the public. We show how time-of-flight laser scanning data can be used to speed-up the alignment process. In addition we show how, after creating a model and repairing imperfections, an interactive and immersive setup enables the public to navigate and display a fully detailed representation of the portal. This paper shows that, after careful planning and with the aid of state-of-the-art algorithms, it is now possible to preserve and visualize highly detailed information, even for very large surfaces.

Categories and Subject Descriptors: I.4.1 [**Digitization and Image Capture**]: Scanning; I.3.7 [**Three-Dimensional Graphics and Realism**]: Virtual reality; I.3.5 [**Computational Geometry and Object Modeling**]: Geometric algorithms, languages, and systems

General Terms: Scanning, Multiresolution visualization, Immersive environments
Additional Key Words and Phrases: Geometry acquisition and processing, out-of-core data management

1. INTRODUCTION

The Benedictine Monastery of Ripoll was founded by Count Guifré el Pilós in 879. The main work of art from the monastery, which is also the main Romanic sculpture from Catalonia, is the entrance (known as **Portalada**, see Figure 1), which dates

Fig. 1. The Portalada: Romanesque entrance of the Ripoll Monastery. Its size is clearly impressive as is the richness of the sculpted relief. Left: general view Middle: the 7 meter high scaffolding used for acquisition. Right: a smaller scaffolding used for the acquisition of the internal part of the portal

back to the 12th century. This entrance has been defined as the "Stone Bible". It is a masterpiece of cultural,historical, social and scientific interest. Ripoll is located 90 Km North of Barcelona, quite a long way from the main tourist routes. New approaches were needed to display this impressive monument to visitors in Barcelona, and disseminate its contents to a wider audience. This factor, together with the gradual deterioration of the stone, was a strong argument for Romanic experts and museum curators at the Museu Nacional d'Art de Catalunya (MNAC) in Barcelona to require a high-accuracy virtual model of this important cultural masterpiece.

This paper presents a project involving the virtual reconstruction and presentation of the Portalada, the interactive installation which was open to the public in the MNAC exhibition named "The Romanic Art and the Mediterranean. Catalonia, Toulouse and Pisa" from February to May 2008.

Visitors to the exhibition were able interact with the virtual reproduction in two immersive setups (VR kiosks). Using a touch-screen and a back-projection display screen with passive stereo, visitors could simply navigate and zoom-in different parts of the entrance, or they could get further information just by touching different "hotspots", which made 3D information boards appear in front of important components of the facade (see Section 6). They were able, for instance, to discover a representation of Moses leading his people or see how that little stone carvings represented the typical work schedule associated with every month of the year. The use of a stereo display greatly helped in creating a more realistic perception of the size and of the true appearance of the stone figures provided by the geometrical detail and colors.

The challenges of the project derived from the requirements of the museum curators. The goal of the project was three-fold: to have a high-fidelity virtual reproduction for the exhibition, to create an analytical tool for experts, and to enable the

current status of the Monastery entrance to be archived. The main contributions of this paper are strictly related to the processing and real-time management of the huge dataset representing the Portalada, including:

—The acquisition and construction of the model which, given the very large sculpted surface (7 x 11 meters), produced a nearly complete 173M faces high resolution model, with a sampling density of the order of one millimeter.

—The derivation of specific scalable algorithms for model repair and simplification.

—The design of a hierarchical data structure for data management and view-dependent navigation.

—The setup of a user-friendly and immersive system that gave the visitors the feeling that they were actually inside the Portalada.

The project relied on several hot topics in computer graphics, such as model acquisition in Cultural Heritage applications, geometry processing algorithms for model repair and simplification, view-dependent visualization and algorithms for gigantic models handling. Previous technical works on these topics are presented in each corresponding section, while an overview of major projects focused on 3D acquisition and visualization for Cultural Heritage is discussed in the next section. Details on the acquisition process are described in Section 3, while Section 4 presents the strategy for creating the model, including alignment, initial repair and color assignment. Section 5 presents the main repair process and the algorithms and data structures designed for view-dependent visualization. Section 6 describes the interface, the layout and the structure of the virtual reality kiosks. Finally, Section 7 discusses the overall process and draws some conclusions.

2. 3D AND CULTURAL HERITAGE: INTEGRATING SIZE AND DETAIL

The Digital Michelangelo Project [Levoy et al. 2000] was one of the milestones in the field of 3D Scanning. This was because the usefulness of 3D Scanning was highlighted to the public, and the Cultural Heritage was presented as one of the main subjects of interest.

In the last few years, geometry acquisition and visualization techniques have improved considerably [Cignoni and Scopigno 2008]. The Cultural Heritage is a very interesting field of research. It involves a great number of possible applications, and a variety of objects that can be different in size, detail, and interest. The project presented in this paper shows interesting solutions for two aspects of Cultural heritage digitization: firstly, the high detail acquisition of very large surfaces, and secondly the real-time visualization and immersive inspection in a museum context.

Regarding the first issue, several previous projects in the field of the Cultural Heritage have presented a mix of lower detail models (obtained with time-of-flight scanning, 3D modeling or photogrammetry) of large areas and triangulation-based scanning of more detailed areas [Guidi et al. 2006; Stumpfel et al. 2003; Valzano et al. 2005; Beraldin 2004; Balzani et al. 2005]. In our specific case, the different resolution data was mixed in the processing phase, in order to solve the general issues associated with triangulation and time-of-flight acquisition.

The second issue regards both the management [Hoppe 1999; Cignoni et al. 2003] and display [Cignoni et al. 2005; Rusinkiewicz and Levoy 2000; Botsch et al. 2002] of complex datasets. Multiresolution approaches enable very detailed geometries to be explored interactively. Another key issue is to present the data in the form of appealing applications for the general public. To do this, user-friendly and customizable software applications [Callieri et al. 2008] have been proposed along with more immersive relightable environments [Callieri et al. 2006], the use of Virtual Reality [Gaitatzes et al. 2001; Farella et al. 2005], and custom solutions [Peral et al. 2005]. More generally, one of the main problems is to find a good compromise between ease of use, cost and visual fidelity. In this paper we present a low-cost and easy to assemble solution, where the user can explore the very high detail of the digital model using a very intuitive and simple interface.

3. ACQUISITION

The Portalada of Ripoll is an incredibly rich monument: every single section tells a story through a mixture of the sacred and profane. For this reason, acquisition of the portal needed particular care in in order to preserve the details. Producers of Time-Of-Flight (TOF) scanners have recently provided new devices which enable very large surfaces to be acquired in a short time and with sufficient precision for Cultural Heritage applications. However, triangulation laser scanners still provide sub-millimetric sampling density acquisition and a much lower uncertainty than Pulsed Wave (PW) TOF scanners.
In the case of the Portalada, the study of an individual statue or bas-relief may be studied by an art historian for several years. Moreover, most statues and bas-reliefs present small details (either carved portions or deterioration traces), which can be very important from an artistic and historical point of view, or for conservation purposes. For these reasons, it was decided to digitize the monument using triangulation-based scanning, so as to have a final sampling density of at least 1mm. However, since there are difficulties in scanning such a huge carved surface (7x11 meters) and in processing all the resulting data, we decided to run a double acquisition, using both (TOF and triangulation) technologies, and then to integrate the two data sources.

Various large-scale acquisition projects have used multiple scanning technologies. In most cases this approach is aimed at covering objects of a different scale with the most appropriate instrument, such as the Parthenon and all its statues [Stumpfel et al. 2003]. In other situations more similar to this one, where a single large object had to be completely acquired at medium resolution but at a higher resolution in some interesting areas (such as the acquisition of the Grotta dei Cervi [Beraldin et al. 2006]), 3D scanning devices have been customized in order to be able to work at different resolutions and with different coverage areas. Very large objects, such as the Michelangelo's David [Levoy et al. 2000], have been covered using custom-made moving rigs. In a paper describing the acquisition of the Great Inscription of Gortyna [Remondino et al. 2009], a procedure similar to the one used in this work is described: a high-resolution acquisition of the inscription using laser triangulation

Time Of Flight scanning Campaign	
Scanner model	Leica ScanStation
Acquisition Time	5h
No. of range maps	3
Acquisition Resolution	3 - 6 mm
No. of points acquired	36.2M

Triangulation scanning Campaign	
Scanner model	Minolta Vivid 910
Acquisition Time	5 days
No. of sections	62
No. of range maps acquired	2212
Acquisition Resolution	0.5 - 1 mm
No. of points acquired	500M

Table I.   Time Of Flight and Triangulation Acquisition Data

coupled with a TOF scanning of the entire structure, which was used to provide the context of the inscription but also as a support for processing of the triangulation data.

In our case, the two different technologies were used to cover the same object, with the aim of using the TOF data during the processing. Specifically:

—The area covered by a single range map acquired with a triangulation scanner is usually not bigger than 50x50 cm. In addition, we needed to sample this region from many different directions to produce a complete sampling of the bas-relief detail. Consequently, several hundreds range maps were needed to cover the entire surface. Processing all these range maps is not easy, especially in terms of the alignment phase. The accumulation of alignment errors could have made it impossible align all the range maps. Moreover, the accumulation of residuals could have led to a severe deformation of the final model. The use of a time-of-flight scanned model as a reference for alignment (see Sect. 4.2) solved both these issues.

—The technical constraints of the triangulation scanner (acquisition distance between 50 and 120 cm, small area covered by each shot, difficult in reaching all scanning positions required) made it difficult to have a complete sampling of such a large object. The different point of view and more consistent coverage of a TOF scanning could provide more information and cover parts which could not be acquired by the triangulation scanning (see Sect. 4.3).

Table I present some numerical data regarding the acquisition campaigns. The Time of Flight scanning (Table I) was performed by a team of two technicians using a Leica ScanStation. The Portalada was scanned from three different positions, each scan with a sampling rate of 0.5 cm, providing a data set of 36.2 million points. The acquisition of the geometry and the setup of the markers for the TOF data alignment took a few hours.

The triangulation-based scanning was performed in five days by a team of four people using two Minolta Vivid 910 scanners. During the acquisition planning the monument was subdivided into 62 overlapping sections (the acquisition planning scheme for the frontal part of the portal is shown in Figure 2). Each section was

acquired with a number of range maps varying from 20 to 90, depending on the size and complexity of the shape of the depicted relief. In order to scan the upper part of the portal, a movable 7 m high scaffolding plus a smaller one were used (Figure 1). The final data set was made up of 2212 range maps with a total of more than a billion triangles.


Fig. 2. The planning scheme used to scan the Portalada with the triangulation-based technology: each box corresponds to a section acquired

More than 200 digital photos were taken to sample the apparent color of the entire surface of the Portalada, to be used for color projection on the reconstructed digital surface (see Section 4.4).

4. MODEL GENERATION

In this section we describe the post-processing pipeline setup used to generate the 3D model from the acquired data. Due to the geometric complexity, the size of the object and the sheer amount of data, the processing posed several issues, especially during the alignment operation - these are described in Section 4.2.

The processing stage took place almost entirely in the lab, after completion of the on-site scanning campaign. Due to time constrains when working in the field, it was only possible to perform a check of the data to assess sampling quality, and a preliminary alignment of some areas, in order to verify the data completeness. The main concern in a complex scanning campaign is to accidentally forget to sample sections of the artifact (due to complex topology, scarce planning, over-estimation of already sampled data), so that the problem is discovered only after leaving the scanning site, thus forcing the setup of a second acquisition campaign. Thanks to

the data checks performed during the scanning, these problems were successfully avoided.

4.1   Time of Flight model

The first processed data were the TOF scans, which should serve as the basis for the rest of the data. Thanks to the alignment provided by a few reference targets used during the scanning, the three scans were already perfectly aligned: only some filtering and cleaning of the data was needed. The three scans were merged at different resolutions (6mm and 3mm), in order to obtain models with a good level of detail, yet still easy to manipulate.

The TOF models were used several times during the post-processing phase: for the initial assessment of the data, both for visual reference in the alignment process and to organize the work. They were also used to perform the initial alignment of the photographic dataset, in order to assess the completeness and quality of the photographic information, while the high-resolution model was still being processed. However, the most important use of the TOF model was to represent a geometrical reference during the alignment of the triangulation range maps, as explained in the next section.

4.2   Alignment

During the acquisition, the surface of the Portalada was divided into sections trying to follow the original structure of teh artifact, which is divided in several "scenes" (Figure 2). Each of these sections were acquired independently, mainly for two reasons: since moving the scaffolding was quite cumbersome, it was only possible to work on a small area at a time; moreover, having smaller but self-contained and meaningful sections greatly helped the organization of the data processing.

Each section of the Portalada was acquired following a regular coverage pattern, with a consistent overlap between adjacent range maps. This regularity enabled us to employ automatic approaches [Pingi et al. 2005] for the initial alignment of the range maps. This helped to speed up the alignment process during the data processing.

After this first automatic pair-wise registration, a standard ICP-based fine alignment was applied to each sections in order to obtain a good global alignment for the given section [Callieri et al. 2003]. The ICP (Iterative Closest Points) approach uses the redundancy between adjacent scans to compute a precise geometric alignment. The idea is to find a small set of corresponding points on the overlapping area of the two range maps and move the range maps in order to minimize the distance between the selected corresponding points. This step is repeated until the two surfaces have a satisfactory level of closeness. The overall alignment of the sections needed the most human intervention, and took nearly 10 days. At this point, each section of the Portalada was ready to be merged. However, we were interested in building one model for the entire portal. It was thus necessary to find a way to integrate all the sections (i.e. the entire set of range maps) into a single reference space.

Each section contained a "border" of about 20-30 cm that it shared with its neighboring sections. In principle, by using the overlap it should have been possible to use the standard alignment procedure between adjacent sections to obtain

a global positioning for all the range maps. However, it was not possible to compute global alignment on a group of more than 2000 range maps using a standard ICP-based aligner, due to the sheer size of data. Additionally, we were concerned about the deformation that might occur while aligning such a large object using range maps which, in proportion, were really small. Indeed, the global alignment step [Pulli 1999] helps to distribute the alignment error among all alignment arcs, but it only works effectively if the object is a "closed" surface. It does not guarantee to eliminate that all the error accumulation will be eliminated if the surface has large, open borders. The Portalada was a clear example of this situation. Given the lack of the back part of the artifact and the relative small size of range maps with respect to the entire object, we were sure that a standard alignment process would lead to severe deformations.

We thus needed a rigid reference for the alignment (to prevent deformation) and to apply the alignment process hierarchically (to cope with the size of the dataset). We therefore decided to use the model obtained using TOF scanning as a reference: since each TOF scan contained the entire extent of the Portalada, the resulting model would be rigid enough. To apply the hierarchical alignment, we again exploited the data subdivision in sections, which enabled us to perform ICP at multiple *levels*: inside a section, across sections and toward the rigid reference. We thus built an alignment process that was viable in terms of the time and human intervention required but, more importantly, provided accurate results. The idea of using "external" references to obtain a distortionless, precise alignment is quite common. When working with TOF scanners, the use of markers and of a total station is a kind of *assisted* alignment. Photogrammetry is another tool that can be used for this task: by measuring the position of control points, it is possible to use them as a rigid reference for the alignment. Various on-the-field projects, such as the acquisition of the Donatello's Maddalena [Guidi et al. 2004], have proven that this techniques is quite effective. Photogrammetry tools are also part of the management software of 3D scanning devices (for example, Minolta and Breuckmann), and are used to provide additional strength to the alignment process. In our case, photogrammetry was discarded because of the difficulties in producing a photographic coverage that was good enough for photogrammetric usage (given the size of the object and the difficulty in placing the camera in a position to cover the entire object). Also, the great number of range maps would have required the photogrammetric determination of too many points. TOF scanning was a more feasible way to produce a rigid alignment reference, and also proved helpful in other phases of the processing.

The first step was to place all range maps in the common reference space in a correct enough position to be able to use the ICP between the various range maps and between the range map and the rigid reference (i.e. with an initial misalignment of few millimeters). Each merged section was manually positioned on the TOF model using a feature-based alignment tool and then an ICP step was performed to obtain a better positioning. By applying the obtained roto-translation for each section to all its range maps, we automatically placed all the range maps in the global reference space. This placement was good however, as we did not consider the cross-alignment between range maps of different sections, there were

still misalignments in the border areas between different sections.

With this placement as the starting point, the hierarchical alignment process was made up of three different steps, which were iterated several times to ensure a stable and accurate final alignment:

—alignment **inside each section**: this step was similar to the one mentioned at the beginning of this subsection. All range maps of a single section were processed using the ICP. This step enforced local alignment.

—alignment **across adjacent sections**: the range maps in each overlapping "border" between sections, group by group, were processed using ICP. This step enforced coherence across sections.

—alignment **toward the TOF model**: each range map in the dataset was aligned to the TOF model, which was kept unmovable. This step prevented any deformation of the dataset.

At the end of each iteration, the alignment result were checked to decide if another iteration was necessary. Beside the final error level (measured as the mean and maximum residual error across all ICP arcs) the decision to do another run was also based on *how much* the range maps moved in the various stages of the process. Since this was the first time we employed this strategy, the thresholds for the convergence were arbitrarily chosen. The alignment process converged in four iterations.

What we did can basically be considered as a *manual, out-of-core, multiscale* global alignment step. Standard alignment tools, both commercial and academic, are designed to work on datasets that are smaller than the one produced in this project. Since in our case it was impossible to directly manage the entire range map dataset in our alignment tool, it was necessary to define an alignment pipeline with a lot of data swapping and a manual organization of the work. Each step consisted in executing the alignment tool with different parameters on different subsets of the dataset. Various batch files and "project" files (containing lists of range maps and roto-translation matrices) helped considerably in managing the whole process. Manual steps included the initial placement of the range maps, the definition of the range map groups for the overlapping borders between sections and the decision to continue or stop the iterations. The computational part of the alignments (the neighborhood adjacency graph, creation of arcs, local and global ICP steps, error evaluation) was completely automatic.

Experience shows that each large-scale project has several specific constraints and characteristics that make it unique. For this reason, writing a specific program to cope with such singularities may be not a good idea (it was also not possible for the Portalada scanning due to time pressures: acquisition and processing were performed in just seven weeks). A better solution to avoid manual intervention would be to extend the existing alignment tools to be script-ready. In this way it would be possible to use the basic alignment operations (pair/group alignment, global optimization) *parametrically*, by adapting the process for the specific needs of the project.

4.3 Reconstruction

After generating a correct alignment for all the range maps, the next step was to create a single surface for the entire object. We used a merging tool which adopts a

Fig. 3. Snapshots of the geometry of the 3D model obtained. Note the variety of the decorative elements and, in close-up, their intricate details

standard volumetric algorithm, based on distance fields. This merging tool manages huge datasets by dividing the object bounding box into various parts and working on each component separately. The resulting pieces are then put together, and the replicated border vertices are eliminated, thanks to a specific feature of the out-of-core reconstruction code [Callieri et al. 2003].

Although this dataset was the biggest we had ever worked on, it was not one of the most difficult to merge. This was due to the "flat" nature of the surface and the small extent of each range map: each part of the object was covered at the most by six or seven range maps. Since the reconstruction is mostly carried out using local data, all the necessary information easily fitted in memory. In order to speedup this step, the merging was simultaneously performed on multiple machines. The total merging time required was 26 hours. The merging produced 406 separate blocks, for a total size of 170M triangles and 3.5 gigabytes. The level of detail obtained can be seen in Figure 3.

The merging process was performed several times, in order to spot local misalignments and uncovered areas. Some small misalignments still remained since, given the dataset size and the multi-step alignment process, it was impossible to completely check every area of the object. These alignment problems became apparent in the merging phase and were corrected by manually finding the out-of-place range maps and recalculating the alignment in that specific area.

The inability to reach some parts of the surface with the triangulation scanner resulted in small, uncovered areas. This lack of data was clearly visible in some low detailed areas of the middle arch and in some difficult-to-reach parts on the high-reliefs. We were happy to find that no major holes were present in the more important and detailed areas. The missing parts were filled by carefully cutting small areas from the TOF range scans. The merging tool was also configured to use this kind of data only when no other data were available. An example of the effects of this process can be seen in Figure 4: the area filled with TOF data appear much smoother than the surroundings generated by triangulation, but the integration is seamless and the lower frequency structures are represented with sufficient detail.


Fig. 4. An example of an area of the model where TOF data have been used to fill non-sampled areas with the triangulation device. Because of the lower detail of TOF data the resulting geometry is smoother, but well integrated and with enough detail for the larger decorations.

4.4 Color management

Once the geometry of the portal was reconstructed, it was necessary to add the color information, in order to enhance the realism of the visualization. This is usually done by taking a set of images which covers the entire surface of an object, registering them on the corresponding 3D model, and projecting the color values on it, storing the data by means of a texture parametrization, or by encoding color per vertex.

Many projects in the Cultural Heritage field have to deal with the acquisition of color information, so as to increase realism of the 3D model but also for study purposes. In some cases, such as the Byzantine Crypt of Santa Cristina [Valzano et al. 2004], a very detailed pictorial detail had to be carefully captured (with particular attention to the color accuracy) and reassembled in very large texture datasets. In other cases, such as the digitalization of the Parthenon [Stumpfel et al. 2003], special care was devoted to measuring the exact optical response of the sampled surfaces.

In our case, the color was aimed at conveying a sense of realism by presenting to the user the appearance of the stones used in the construction of this artefact. However, given the lack of a pictorial detail and of a relationship between the grain

of the stone and the sculpted geometry, the level of detail of the color information may as well have been the same as the acquired geometry, presenting fewer constrains with respect to the aforementioned projects. On the other hand, due to the size of the portal, the difficulty in placing the camera in all the positions needed and the particular lighting conditions (we could not shield the sunlight coming from outside, nor provide a diffuse enough artificial lighting), the photographic coverage proved to be quite complex.

As with the geometry acquisition, in the photographic campaign there was coverage of the surface of the portal from a short range, but several images from greater distances were also taken. In Figure 5(left) we show some examples of the images captured: the portion of depicted space was also determined by the scaffolding position from which each part of the surface was acquired; the resulting dataset was composed of more than 200 photos. Figure5(right) is a snapshot of the 3D model, where the camera viewpoint of the images used for color projection is visualized as a string in the surrounding space: as can be noted, camera viewpoints are grouped in correspondence to the scaffolding positions chosen to cover each part of the surface of the Portalada.


Fig. 5. Left: some examples of images used for color projection: to provide maximum coverage, the images framed the various parts of the portal from different directions and different distances. Right: camera position associated with the images used for color management.

From the whole set of images, we manually selected a subset of 163 images which could cover the entire surface, providing sufficiently coherent lighting. We registered the images towards the 3D model using a tool for image registration [Franken et al. 2005], finding the viewpoint and intrinsic camera parameters associated with each image.

At this point it was necessary to mix all the information from the various photos to apply color to the digital 3D surface. Given the size of the object, its topological complexity and the required color resolution, it was impossible to use a texture-map based approach: there was simply no way to automatically build a parametrization for the entire object. In any case, given the amount of color information and the hardware size limit for texture maps, this would have required a very large number of textures.

We resorted to per-vertex color mapping: a color value is stored in each vertex of the

3D model. This representation is quite compact, easy to use and compatible with many tools for 3D modeling and rendering. In addition, given the detail of the 3D model we were using, it was possible to convey enough color information. For the per-vertex mapping we used a color mapping tool [Callieri et al. 2008]: this software implements a weighted blending strategy to compute the most appropriate color to map in every vertex of the surface. Basically, the quality of each pixel in each image is evaluated using different metrics (such as closeness to camera, spatial resolution, focus) in order to build a weight mask that can be used when mapping to effectively mix data from different photos. Thanks to its out-of-core strategy, it is possible to work on very large datasets, such as the one presented here. The unavoidable differences (illumination, exposure, white balancing) between the photos were thus efficiently blended, obtaining a generally satisfying result.

The images taken from longer distances had a role similar to the one played by the time-of-flight scans in the model reconstruction i.e. to cover any parts that were not taken by short range images. The color projection mechanism of [Callieri et al. 2008] automatically assigned a lower weight to these low quality images in the zones where other images covered the surface with a better resolution/quality data.

The color coverage provided by the photographic campaign was almost complete; snapshots of details of the final colored model are shown in Figure 6.


Fig. 6. Left and Middle: Two snapshots of details of the colored mesh of the Portalada. Right: Corresponding close-ups rendered with no texture and wireframe to reveal the underlying triangle mesh.

5.   MODEL PROCESSING AND REALTIME VISUALIZATION

The geometric processing of the final model of the Portalada involved an initial repair step, in order to eliminate holes and topological errors, followed by the creation of the multiresolution model and the in-core octree data structure used for the realtime visualization.

5.1   3D model surface repair

Model repair is a hot research topic (see for instance [Bischoff et al. 2005]) for which a good number of recent algorithms have been proposed. The main problem in our case was, once again, the huge amount of data. Apart from detecting and repairing small imperfections (badly-oriented triangles, noise, spurious polygons), the main challenge was related to the complexity of the model: 173 Mtriangles and 14622 holes of different sizes and shapes. Holes and cracks were automatically detected through a search for cycles of border edges. Simple holes can be repaired using known techniques, but these algorithms fail when applied to complex holes. A novel algorithm that solves and repairs these complex configurations was specially designed for this purpose [P. Brunet and Vinacua 2009].


Fig. 7. The model repairing process. top-left: the original geometry; top-center: voxelization to be used to compute the approximating surface; top-right: the hole-filling approximating surface; lower-left: the original surface plus a patch of the closing surface; lower-right: the closing surface was integrated into the full resolution model

As a first step in our algorithm, holes and cracks are automatically detected by searching for cycles of border edges. An extended bounding box is computed for

each cycle in order to extract a section of the mesh surrounding the hole. Hereafter, the term *hole submesh* indicates the result of clipping the original mesh around a hole with this extended bounding box. For every cycle, the hole submesh information together with the extended box coordinates is inserted into a hole list.

For the Portalada model, this initial hole detecting process ended up with a list of 14622 holes and their corresponding submeshes. A first classification gave a total of 8485 holes with less than 20 border edges, and 6137 bigger and more complex holes (Figure 7). For the first holes, our algorithm uses a simple constrained Delaunay triangulation of the border vertices (without Steiner points). For small holes, this simple technique met our requirements.

For non-small holes, our first idea was to project the border edges onto an auxiliary plane, compute Steiner points on this plane, triangulate them, and project them back onto an approximating surface [Esteve et al. 2008] in order to create the repairing patch. This procedure was able to repair 5149 out of the remaining 6137 holes, but failed in the remaining 988 holes. These holes had a rather complex shape, and projecting the border cycle into a plane without foldovers was not possible. For these complex situations, and for each submesh, the algorithm first computes an implicit surface that approximates the region inside the border cycle even in the most complex cases. This approximating surface is then used to compute Steiner points (3D points inside the hole) with a density similar to the mesh vertices around the hole. These Steiner points are triangulated, thus producing the necessary new triangles inside the hole. The CGAL triangulation algorithm is used in this step [Boissonnat and Oudot 2006]. To connect this new patch to the initial submesh, a standard sewing algorithm is used, based on an advancing front that connects closest pairs of vertices. This algorithm is iterated over every hole submesh in the initial list.

The algorithm used to compute the approximating surface was introduced by Esteve et al [Esteve et al. 2008]. The resulting implicit surface is defined as the zero-set of a functional uniform B-Spline (in 4D) defined on the 3D domain of the discrete voxel representation of the hole surroundings [P. Brunet and Vinacua 2009]. A multigrid representation helps accelerate the computation of the Spline surface, which approximates the geometry around the hole. The algorithm iterates a smoothing and a fitting step for each resolution. The criterion used for smoothing is based on trying to increase the surface continuity from C2 to C3 at the boundary faces of the voxels, the resulting filter being a constrained least-squares approximation of the 4 x 4 x 5 weights of two adjacent voxels. The fitting step attempts to displace the surface so that it interpolates the center points of the constraint cells, while minimizing the deviation from the previous control points. As a result, the algorithm guarantees that the surface stabs the cells intersecting the original mesh.

A final post-process adds a white noise to the new mesh vertices, assigns them an average color computed from the colors around the hole and uses a Laplacian filter to remove color discontinuities. The variance of the white noise is computed from the vertex deviations in the submesh, with the goal of having the new patch as similar as possible to the surrounding appearance of the carved stone. We observed that not having such a noise addition also resulted in smooth filling patches.

This process was iterated over every hole of the original geometry. The final

repaired mesh had around 173 millions triangles, with an average edge length of 1.4 millimeters.

5.2   Multiresolution encoding for efficient visualization

The main goal of the next step was to structure the 3D data to ensure a real-time, interactive visualization of this gigantic mesh on a standard PC. The requirement was that visitors should be able to move freely along the monument and zoom in on any place or any detail. The system should then react in real-time and present the digital model in the zoomed-in region, with the maximum resolution of the digital mesh. Our approach was based on a preprocessing step and on the concurrent use of a number of recent techniques:

—View-dependent, out-of-core algorithms for real-time interactive rendering of gigantic meshes. The progressive buffers in [Sander and Mitchell 2005] and the far voxels algorithm in [Gobbetti and Marton 2005] are good references of algorithms for the interactive inspection of huge models, using textures, levels of detail and hybrid geometric/impostor representations.

—Hierarchical data structures of large granularity, in order to avoid too many changes and updates in the view-dependent front, in a similar way to the algorithms in [Gobbetti and Marton 2004] that use layered point clouds with hierarchical data structures.

—Rendering of a low-resolution mesh with high-quality textures, for distant observer positions. Our approach uses the results from [Tarini et al. 2003].

In our current implementation, we observed that it is possible to obtain a high visual quality using a three-level multiresolution model. At observer locations that are further away than a certain distance from the Portalada (2 m in our implementation), we simply rendered a textured base mesh which is permanently located in the GPU memory through a number of vertex buffer objects. We decided (after experimenting with relief impostors, [Andujar et al. 2007]) to use a low resolution version of these patches made up of 10% of the original triangles.

In the preprocessing step, several data structures and geometric models are computed:

—A textured base mesh, with 200 K triangles. A 200 K triangle simplification of the initial mesh was first obtained using classical vertex clustering techniques. Textures were then obtained by projecting points of the high resolution mesh onto the 3D points of the low-resolution triangles that correspond to each of their texels. Color textures and normal maps were encoded on a texture atlas, which was 550 MBytes in size.

—A set of disjoint high-resolution mesh patches Mk. The overall mesh is divided into 3000 small mesh patches M1,M2, ...,M3000 with around 60000 triangles each. Each triangle of the global 173 Mtriangles mesh belongs to one and only one of the Mk patches, while vertices are duplicated along boundaries between neighbor meshes. The result of rendering all Mk mesh patches is obviously exactly the same as rendering the huge high-resolution mesh. Managing small meshes is much easier, and the memory overhead for duplicate vertices is very small (less than 2 %). These small patches enable the performance of the CPU and GPU

memory caches and of the CPU-GPU traffic to be optimized. Apart from the 3000 high-resolution patches (with an average size of 20 x 20 cm), we grouped them and simplified them to the 10% of triangles, obtaining a total of 300 low-resolution patches. The whole model was therefore available at three LOD levels: high-resolution patches, low-resolution patches and textured base mesh.

—An octree data structure that manages the camera movements and the view dependent visualization at each frame.

Our octree supports the hierarchical management of the data structure of the multiresolution meshes (see Figure 8). The octree data structure is permanently in-core during the navigation. In our current implementation, the number of leaf nodes containing parts of the mesh (surface terminal nodes) is 283617, and the number of terminal void nodes is 389945. The octree contains 96286 non-terminal (grey) nodes. The depth of the octree is 10 levels, given a node edge size of 3.51 centimeters for a Universe size of 18 x 18 x 18 meters. Surface terminal nodes and grey nodes contain a mask, which points to the mesh patches Mk partially or totally contained in the node's cube. Grey nodes also represent a hierarchical linear distance field which approximates the distance from the observer to the Portalada. Each octree grey node stores a 3D point, its distance value and the gradient of the distance field. During octree construction, the maximum deviation between the linear approximation of the distance and the exact distance to the Portalada is computed for each node being processed. If this deviation exceeds a predefined tolerance, the node is subdivided and the distance gradient is computed for each of its child nodes. Octree nodes also contain pointers to the Portalada multi-resolution meshes.
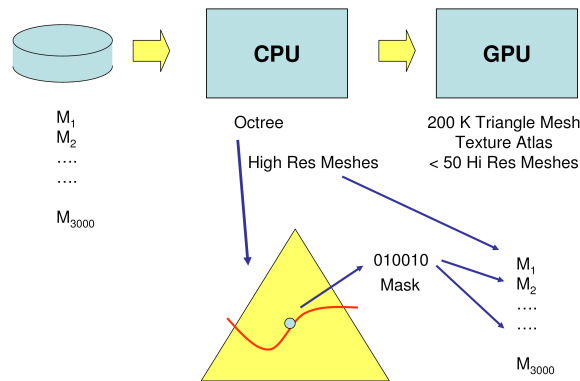

Fig. 8. Data flow among disks, the CPU and the GPU, and patch management based on masks and the octree data structure

The interactive visualization uses a view-dependent algorithm which is derived from the methods in [Gobbetti and Marton 2004] and [Gobbetti and Marton 2005]. An octree traversal is performed at each frame to first detect if the movement of

the observer is valid or not. If the new observer location is valid (he/she is not exiting the octree Universe Cube and their distance to the surface of the Portalada is greater than 30 centimeters), the view-dependent front is obtained and the list of models that should be rendered is computed. The octree traversal is also used for hierarchical frustum culling and generating the list of visible patches for the present frame. Nodes of the active octree front are computed at each frame on the basis of the size of the projection of their corresponding cube in pixel coordinates. If this size is larger than a predefined tolerance in pixels, the front is searched for in the descendants of the node. If it is not larger, the node itself is added to the list of front nodes for this frame. The projection test is very fast, it is implemented as a linear search in a precomputed array, the size of the array being the number of octree levels. This array stores the expected squared distance to the observer for front nodes of the different levels in the octree. Distances between the observer and the surface of the Portalada are also computed. In this case, the hierarchical linear distance field is used, requiring a single test per octree level. The observer position is pushed down along the corresponding branch of the tree, while the distance is computed by the linear distance field approximation in the node. The branch descent finishes as soon as the distance error is below a prescribed tolerance threshold. Once the list of front nodes has been obtained, a list of high-resolution mesh patches that have to be rendered in this frame is computed. Two-level node masks are used to select the right set of high-resolution mesh patches. Each octree node at a depth greater than four, has a 11 byte mask that represents which of the high-resolution meshes Mk intersect the node, within the set of meshes Ma intersecting its ancestor at the fourth level of the tree. The 11 bytes mask allows a maximum cardinality of 88 for the sets Ma. The mask works like a bit mask, the Mk mesh in Ma being intersected by the octree node iff the k-th bit in the mask has been set to one. We observed that the number of elements in the sets Ma is never greater than 88. When a new front node is detected, the list of high resolution mesh patches that should be rendered in this frame is updated by simple bit operations: look-up in a mask dictionary corresponding to Ma in order to obtain a global mask, and a bitwise or operation.

Each visible patch inside the frustum is assigned a priority which is related to the distance in pixels between the projection of a representative mesh point and the center of the viewport. The rendering algorithm uses CPU and GPU storages. We found that a good performance is achieved when the GPU cache size is equivalent to 50 patches (either high or low resolution). Patches with a high priority will be rendered at a high resolution, whereas the low-resolution version of the patches is used for the ones with a lower priority in this frame. As already mentioned, the GPU cache permanently stores the textured base mesh apart from the 50 high or low resolution patches. In this way, we ensure that the system will always be able to render the model, no matter how fast the interaction movements are.

A lazy CPU-GPU communication algorithm is used, where a single patch per frame Mk is sent from the CPU to the GPU: the one that has the highest priority in the CPU list, from the ones still not being sent, is chosen. We use a fragment program that implements shading and shadow algorithms. Progressive alpha blending is performed so that every new patch sent to the GPU is almost transparent at

the first frame and becomes increasingly opaque along the next ten frames.

6. INTERFACE AND IMMERSIVE INSPECTION

Since the goal of the visualization system was to provide the extreme detail fidelity of the digital model, the presentation system needed to enable visitors to navigate and zoom-in very close on any part of the monument. Visitors were thus able to look at hidden details that cannot be seen by people going to the actual monument in Ripoll, since such details might be located 6-7 meters from the floor. In addition, visitors were able to retrieve pictures, texts and views of destroyed parts of some of the sculptures that cover the whole portico.

Visitors could interact with the digital model of the Portalada through two VR kiosks, which were identical. Visitors were directed to one or to the other depending on the queues and the demand. Visitors were provided with passive stereo glasses at the entrance to the kiosks and asked to return them at the exit. These two kiosks were designed and constructed during the project, by following the specifications given by the Museum curators.


Fig. 9. Left: the interface of the touch-screen that controls the kiosk. Right: interaction with the hotspot to access the multimedia information

One of the important features of the VR kiosks was that they were low-cost and easy to assembly. The components of each kiosk included one dual-core PC for the visualization tasks, a second PC to manage the interface, two DLP projectors with circular polarizing filters in a custom-designed calibration frame, a 2 x 1.5 meters back-projection screen and a touch screen for the users interaction, see Figure 10. All the above-mentioned parts were off the shelf components. A touch-screen interface was selected because of anti-vandalic requirements, typical in public exhibitions.

The user interface in the touch-screen (Figure 9) presented a view of the Portalada and a few buttons and widgets. The representation of the Portalada without the

surrounding walls of the Monastery was intended to highlight its similarity with the Arcs of the Roman Empire. Visitors were able to approach by pressing on the perspective arrows, to perform a *pan* operation to go to specific parts of the monument by simply moving the finger along the Portalada to the place they were interested in, or to rotate the camera by virtually rotating the *trackball* at the lower right corner.

This design of the navigation interface followed the guidelines described in the work that presented Virtual Inspector [Callieri et al. 2008], the tool for museum-oriented installations: most of the users of these kinds of kiosks are not expert computer users, and need a simple and accessible way to manipulate the 3D object. For this reason, the users should always have a view over the complete object, which may be used as a navigation map. They also need simple and intuitive controls which prevent "wrong" positions and, even more important, a *reset* button. This setup, coupled with the characteristics of the displayed object, helped in creating an even more minimalistic interface.

The interface PC sent the camera parameters to the navigation PC, which detected if the new position was valid or not. If it was'nt, the camera was not updated and the visitor had the perception of colliding with the surface of the Portalada. Updating the camera basically meant updating the observer and viewing reference points, the latter necessary being the intersection of the viewing direction with the surface of the monument. The echo on the touch-screen consisted of the displaying of a semitransparent quadrilateral representing the projection of the results of intersecting the view frustum with a plane parallel to the Z-far plane and containing the view reference point. This resulted in quite intuitive feedback of the part of the Portalada the camera was looking at. The shape of this semitransparent frame changed with the camera rotation, so that it always enclosed the portion of the monument that was displayed on the stereo screen.

The user interface had two more buttons: a reset button ("inici") that moved the camera back to a general view of the Portalada, and a *mode* button. The application supported three modes, which could be reached in a "circular" way by repeatedly pushing the mod button: the navigation mode, the sun illumination mode and the information mode. The navigation mode allowed for standard navigation as already described, the sun illumination mode simulated solar lighting while the user changes the time of the day and the day of the year using two standard sliders, and the information mode displayed additional information on request. Additional information (historical photos or text) were prepared by the museum experts. Special 3D marks popped-up during the interactive inspection at their appropriate locations, and visitors could select them with simple buttons on the touch-screen. Text and photos were then displayed at appropriate camera perspectives (see Figure 9).

Figure 10 shows the final arrangement of one of the two VR kiosks installed in the museum and the interaction with the 3D model. A simple usability test was performed on a group of ten people, to tune various parameters involved, such as the sensibility of the touch screen and the speed and acceleration of the camera changes. Users in the museum were observed during the first few days of the exhibition, and some informal interviews were carried out at the time the glasses were

Fig. 10.   Setup of one of the VR kiosks at the MNAC museum

returned. The feedback, both from this test group and from the exhibition visitors, was very positive. The public was especially impressed by the quality of the details in the close-ups and by the user-friendly interaction and the 3D perception of the reliefs. Visitors were able to inspect tiny details of the upper parts of the monument which are inaccessible in the real site. Visitors were spontaneously grouped into sets of three to eight people, one acting as the group interactor that operated the touch-screen. The option of having additional information was considered very positive by most of the users, whereas other visitors did not use it. After this positive evaluation, the museum decided to include a non-stereo kiosk in its permanent exhibition. The stereo setting was considered quite positive for a temporary exhibition, but not suitable for a permanent one.

7. CONCLUSIONS

We have presented a project concerning the virtual reconstruction of the entrance of the Ripoll monastery in Catalonia, Spain. The project is the result of a successful cooperation among the MNAC museum, UPC in Barcelona and ISTI-CNR in Pisa. The project involved a number of important challenges, from the size of the Portalada, the quality and precision requirements, to the extremely tight schedule: the whole project had to be finished within four months.

The acquisition process and the alignment for the generation of the high resolution triangles meshes involved the use of novel techniques. Specific scalable algorithms were designed for model repair and simplification, and an octree data structure was designed for data management and view-dependent navigation. A user-friendly interface was designed for two immersive, passive stereo VR kiosks. In the whole process, the huge sampled dataset was transformed into a valid mesh of triangles with millimetric resolution, which can now be inspected in real time in the PC-driven kiosks. The informal evaluation by the museum visitors was very positive.

As most of the 3D Scanning campaigns on complex objects, the proposed project was completed after a careful planning and adopting ad-hoc solutions. Nevertheless, the generic procedure of mixing TOF and triangulation laser data, and the iterative process for alignment and merging can be applied to the acquisition of any big structure. The future experience on similar cases will help in decreasing the needed processing time, and possibly defining a generic procedure for acquisition

and processing of data.

REFERENCES

ANDUJAR, C., BOO, J., BRUNET, P., FAIREN, M., NAVAZO, I., VAZQUEZ, P., AND VINACUA, A. 2007. Omni-directional relief impostors. *Computer Graphics Forum 26,* 3 (Sept.), 553–560.

BALZANI, M., CALLIERI, M., CAPUTO, G., CIGNONI, P., DELLEPIANE, M., MONTANI, C., PINGI, P., PONCHIO, F., SCOPIGNO, R., TOMASI, A., AND UCCELLI, F. 2005. Using multiple scanning technologies for the 3d acquisition of torcello's basilica. In *Int. Workshop 3D-ARCH'2005 - 3D Virtual Reconstruction and Visualization of Complex Architectures, August 22-24, 2005, Mestre-Venice, Italy.*

BERALDIN, J.-A. 2004. Integration of laser scanning and close-range photogrammetry the last decade and beyond. In *Proc. IAPRS 35(5) Istanbul, Turkey.* 972–983.

BERALDIN, J.-A., BLAIS, F., COURNOYER, L., PICARD, M., GAMACHE, D., VALZANO, V., BANDIERA, A., AND GORGOGLIONE, M. 2006. Multiresolution digital 3d imaging system applied to the recording of grotto sites: The case of the grotta dei cervi. In *7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST).*

BISCHOFF, S., PAVIC, D., AND KOBBELT, L. 2005. Automatic restoration of polygon models. *ACM Transactions on Graphics 24,* 4 (October), 1332–1352.

BOISSONNAT, J.-D. AND OUDOT, S. 2006. Provably good sampling and meshing of lipschitz surfaces. In *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry.* ACM Press, New York, NY, USA, 337–346.

BOTSCH, M., WIRATANAYA, A., AND KOBBELT, L. 2002. Efficient high quality rendering of point sampled geometry. In *Proceedings of the 13th Eurographics Workshop on Rendering (RENDERING TECHNIQUES-02),* S. Gibson and P. Debevec, Eds. Eurographics Association, Aire-la-Ville, Switzerland, 53–64.

CALLIERI, M., CIGNONI, P., CORSINI, M., AND SCOPIGNO, R. 2008. Masked photo blending: mapping dense photographic dataset on dense 3d models. *Computer & Graphics (In press).*

CALLIERI, M., CIGNONI, P., GANOVELLI, F., MONTANI, C., PINGI, P., AND SCOPIGNO, R. 2003. VCLab's tools for 3D range data processing. In *VAST 2003.* Brighton, UK, 13–22.

CALLIERI, M., DEBEVEC, P., PAIR, J., AND SCOPIGNO, R. 2006. A realtime immersive application with realistic lighting: the parthenon. *Computer & Graphics 30,* 3, 368–376.

CALLIERI, M., PONCHIO, F., CIGNONI, P., AND SCOPIGNO, R. 2008. Virtual inspector: a flexible visualizer for dense 3d scanned models. *IEEE Computer Graphics and Applications 28,* 1 (Jan.-Febr.), 44–55.

CIGNONI, P., GANOVELLI, F., GOBBETTI, E., MARTON, F., PONCHIO, F., AND SCOPIGNO, R. 2005. Batched multi triangulation. In *IEEE Visualization 2005.* 27–35.

CIGNONI, P., MONTANI, C., ROCCHINI, C., AND SCOPIGNO, R. 2003. External memory management and simplification of huge meshes. *IEEE Transactions on Visualization and Computer Graphics 9,* 4, 525–537.

CIGNONI, P. AND SCOPIGNO, R. 2008. Sampled 3d models for ch applications: A viable and enabling new medium or just a technological exercise? *J. Comput. Cult. Herit. 1,* 1, 1–23.

ESTEVE, J., VINACUA, A., AND BRUNET, P. 2008. Piecewise algebraic surface computation and smoothing from a discrete model. *Computer Aided Geometric Design 24 (6),* 357–372.

FARELLA, E., BRUNELLI, D., BENINI, L., RICCO, B., AND BONFIGLI, M. E. 2005. Pervasive computing for interactive virtual heritage. *IEEE MultiMedia 12,* 3, 46–58.

FRANKEN, T., DELLEPIANE, M., GANOVELLI, F., CIGNONI, P., MONTANI, C., AND SCOPIGNO, R. 2005. Minimizing user intervention in registering 2D images to 3D models. *The Visual Computer 21,* 8-10 (sep), 619–628.

GAITATZES, A., CHRISTOPOULOS, D., AND ROUSSOU, M. 2001. Reviving the past: cultural heritage meets virtual reality. In *VAST '01: Proceedings of the 2001 conference on Virtual reality, archeology, and cultural heritage.* ACM, New York, NY, USA, 103–110.

GOBBETTI, E. AND MARTON, F. 2004. Layered point clouds: a simple and efficient multiresolution structure for distributing and rendering gigantic point-sampled models. *Computers and Graphics 28,* 815–826.

GOBBETTI, E. AND MARTON, F. 2005. Far Voxels – a multiresolution framework for interactive rendering of huge complex 3d models on commodity graphics platforms. *ACM Transactions on Graphics 24,* 3 (August), 878–885. Proc. SIGGRAPH 2005.

GUIDI, G., BERALDIN, J.-A., AND ATZENI, C. 2004. High-accuracy 3d modeling of cultural heritage: the digitizing of donatello's "maddalena". *Image Processing, IEEE Transactions on 13,* 3, 370 –380.

GUIDI, G., FRISCHER, B., RUSSO, M., SPINETTI, A., CAROSSO, L., AND MICOLI, L. L. 2006. Three-dimensional acquisition of large and detailed cultural heritage objects. *Mach. Vision Appl. 17,* 6, 349–360.

HOPPE, H. 1999. New quadric metric for simplifying meshes with appearance attributes. In *Proc. IEEE Conference on Visualization (VIS99)*. ACM Press, New York, 59–66.

LEVOY, M., PULLI, K., CURLESS, B., RUSINKIEWICZ, S., KOLLER, D., PEREIRA, L., GINZTON, M., ANDERSON, S., DAVIS, J., GINSBERG, J., SHADE, J., AND FULK, D. 2000. The digital michelangelo project: 3D scanning of large statues. In *Proceedings of ACM SIGGRAPH 2000*. 131–144.

P. BRUNET, A. CHICA, I. N. AND VINACUA, A. 2009. Massive mesh hole repair minimizing user intervention. In Dagstuhl Seminar on Geometric Modeling (submitted for publication). *Computing 86*, 101–111.

PERAL, R., SAGASTI, D., AND SILLAURREN, S. 2005. Virtual restoration of cultural heritage through real-time 3d models projection.

PINGI, P., FASANO, A., CIGNONI, P., MONTANI, C., AND SCOPIGNO, R. 2005. Exploiting the scanning sequence for automatic registration of large sets of range maps. *Computer Graphics Forum 24,* 3, 517–526.

PULLI, K. 1999. Multiview registration for large datasets. In *Proc 2nd Int.l Conf. on 3D Digital Imaging and Modeling*. IEEE, 160–168.

REMONDINO, F., GIRARDI, S., RIZZI, A., AND GONZO, L. 2009. 3d modeling of complex and detailed cultural heritage using multi-resolution data. *J. Comput. Cult. Herit. 2,* 1, 1–20.

RUSINKIEWICZ, S. AND LEVOY, M. 2000. QSplat: A multiresolution point rendering system for large meshes. In *Comp. Graph. Proc., Annual Conf. Series (SIGGRAPH 00)*. ACM Press, 343–352.

SANDER, P. V. AND MITCHELL, J. L. 2005. Progressive buffers: view-dependent geometry and texture lod rendering. In *SGP '05: Proceedings of the third Eurographics symposium on Geometry processing*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 129.

STUMPFEL, J., TCHOU, C., HAWKINS, T., MARTINEZ, P., EMERSON, B., BROWNLOW, M., JONES, A., YUN, N., AND DEBEVEC, P. 2003. Digital reunification of the Parthenon and its sculptures. In *4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. VAST.

TARINI, M., CIGNONI, P., AND SCOPIGNO, R. 2003. Visibility based methods and assessment for detail-recovery. In *IEEE Visualization 2003*. IEEE, Seattle, Washington, 457–464.

VALZANO, V., A BERALDIN, J., PICARD, M., EL-HAKIM, S. F., AND GODIN, G. 2004. Virtual heritage: the cases of the byzantine crypt of santa cristina and temple c of selinunte. In *VSMM 2004 : 10th International Conference on Virtual Systems and Multimedia. Hybrid Realities: Digital Partners*. Softopia Japan.

VALZANO, V., BANDIERA, A., BERALDIN, J.-A., PICARD, M., EL-HAKIM, S., GODIN, G., PAQUET, E., AND RIOUX, M. 2005. Fusion of 3d information for efficient modeling of cultural heritage sites with objects. In *Proc. CIPA 2005 XXth Int. Symposium, Torino, Italy*.

...