Contents

1

Workflow for Creating and Rendering Huge Point Models

In this chapter, we describe a specific workflow for digitally documenting large monuments and cultural heritage sites. We argue that for large sites where the amount of gathered data is significant, a *point-based* processing workflow from the beginning to the end is advantageous. In contrast to the traditional approach of mesh-based cultural heritage documentation, point-based processing allows easy visualization, exploration, and interaction with huge models. We base our description on two concrete practical example surveys.

The point-based processing workflow has several steps. Before the actual data survey can be executed, planning the surveying campaign such as settling the scan positions for laser scanning and organizing the necessary equipment like lamps or power sources has to be done. Then the data survey is carried out, which is done by point-sampling the geometry of the objects with a laser scanner. During the survey, digital photos are taken by a camera, which are then used to color the point samples. Additional photos can be taken of specific objects of interest, which are later used to generate high-quality textured mesh models. This technique is used in areas with interesting textures, like wall paintings. Next, we describe our approach regarding postprocessing, like cleaning the scanned point data, registering the scans and photographs in a common coordinate system, coloring the point samples from photos, and creating floor plans and cross-sections. Finally, visualization and editing methods are described, as well as creating a camera path and rendering a movie following this path through the point and mesh models.

1.1 Data Survey - Scanning

The surveying process is done on the site, and includes all work done at the location of the monument. The preparation is important, as the survey is done only once, and the amount and quality of the data determines the success of the scanning campaign. If the purpose of the survey is solely the recording of the geometry of a monument, it is important to plan scan positions so that the whole monument is covered. When also photographs of the monument shall be taken, it is important to plan for a good lighting situation during the survey, because using only flash light from a camera might result in changing shadows or unwanted reflections from surfaces. Daylight or generally the weather may affect the data collection and scanning process, so ideal circumstances of season (temperature, natural vegetable cover) should be considered. Before the surveying process can be started, the owner of the property has to endorse it and also the exploitation of the recorded data has to be negotiated. Regarding this point, we had various experiences. Permissions of access, working licenses, appointments of ownership for the data output, and the terms of use and publication should be prepared in advance, especially if the place of work is not identical to the place and nationality of the working team.

1.1.1 Motivation of Data Sets

The basis of every archaeological work is the documentation – in our case of monuments and their setting in the environment – to allow an optimal "reading" of the monument regarding its use, meaning, dating, and function. Very often, not only the main answers but also the most important questions for a correct interpretation of dating and function of a monument can be given only after a detailed study of the documentation – the higher the quality of documentation is, the better the interpretation will be. In this situation, laser scanning typically offers the best answers to the following problems: the extents of the site, the different aspects such as indoor or outdoor spaces at various levels, varying building materials like bricks or stone, changing surfaces (for example mosaics or wall paintings), different lighting conditions from direct sunlight to complete darkness, and last but not least the usually very difficult and limited accessibility of the site for reasons of security or property. The specific characteristics of each monument have to be respected in the planning process of the documentation to guarantee best legibility for the interpretation.

Generally, the aim of the digital documentation influences the choice of hard- and software, and the way a team of collaborators and their project work flow should be planned. We had the best experience in a mixed team with both sides, archaeologists and technicians, to communicate onsite the specific questions what the documentation should offer, and to develop directly the strategy to give answers with the digital data set. Every further question can afterward be answered only in the limits of the already collected data set. In this sense, more data are, at least in the beginning of the process, better.

(a)

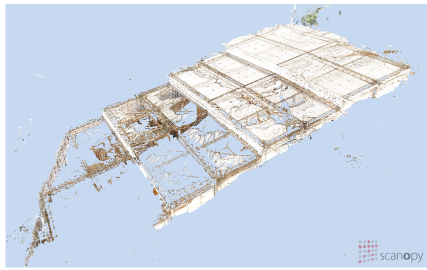

(b)

(c)

(d)

Figure 1.1. Image (a) shows a part of the mosaic in the cupola of Centcelles. Image (b) shows the complete data set of the terrace house 2 in Ephesos. The scans were done from beneath the protective roof. Image (c) shows the lifting ramp used for some outside scans. Image (d) shows the scan positions outside the building visualized in the point cloud of the outer areas.

1.1.2 Description of Data Sets

We applied 3D-laser scanning for the documentation of the late Roman villa complex of Centcelles/ Tarragona, Spain [Adenstedt et al. 14], which is on the UNESCO's world heritage list for its wonderful mosaics in the main entrance hall (vestibulum), a circular space with an inner diameter of about 10 m and a 13 m-high cupola (see Figure 1.1). As a second example, we show some interesting views of the interior of terrace house 2, at Ephesus, Turkey (see also Section 1.2.3). This Roman peristyle house conserved large parts of its original mosaic floor and wall paintings. Again, images and architecture were planned together to impress the patron's guests during his dinner invitations. For reasons of conservations, tourists are not allowed to enter the rooms and step on the original surfaces.

For the scanning process, we had full access to and could cover the entire monument by systematically scanning it from different positions. The type of scanner we used makes use of full wave-form processing [Ullrich and Reichert 05]. This makes scanning of difficult surface materials like glass mosaics or marble possible, as bad point samples can often be identified during post-processing (see also Section 1.2.1), although some reflections cannot be avoided. But, other scanners may have more problems with certain surfaces. In both cases (Centcelles and Ephesus), the indoor lighting conditions were very good and we did not have to add artificial light. In Ephesos, a semi-transparent white canvas is spanned over the monument as roof, which causes a diffuse light at the complete monument. Therefore, the light was relative independent of the time-of-day. Similarly, in Centcelles, the light could enter the room with the cupola only through windows built into thick walls, causing a diffuse light inside the room. This significantly accelerated the working process. In the outside areas, the time-of-day has a significant impact when photos are taken during the scanning process, as the lighting direction (and therefore the shadows cast by the objects) moves with the sun. If consistent lighting is required, possible solutions are to cover the scanned areas by some canvas (as in Ephesos) to create diffuse lighting, or to scan at the same time-of-day over several days. Since our focus was on the wall paintings in the cupola of Centcelles, we did not take special care about the lighting situation for the outdoor scans. The only required preparation was the distribution of reflectors (reflecting targets) for the automatized registration of various scan positions in one point cloud.

1.1.3 Workflow at the Monument

The equipment to survey the entire site at Centcelles with its environment and the remaining buildings from in- and outside were two hydraulic lifting ramps (one for in- and one for outside), a total station, and a 3D-image-Laser scanner (Riegl VZ-400 [Riegl Laser Measurement Systems 14a]). The complete documentation was performed in only three days of field work: First, on the entire site, a grid of about 180 reflector points was distributed and measured with a total station. The layout of the grid was determined beforehand during a site inspection, relying on the experience of the scanning team and not on an algorithmic solution. Second, inside that grid of reflector points, a total of 114 overlapping scans were taken, each with a 360° scan and a set of six HD photos for the coloring of the points. For a reasonable resolution of the sampled geometry (e.g., 0.04 degrees horizontal and vertical angular step width of the laser beam), about 20 million points per scan were acquired, but this can also be increased if necessary, up to 10^9 points per scan.

Sets of scans were taken from all around the building (24 scans, see Figure 1.1), from above with the taller platform (10 scans, see also Fig-

(a) (b)

Figure 1.2. Image (a) shows the laser scanner on the lifting ramp. Image (b) shows the operator cowering beneath the scanner. This image was taken from the camera mounted on top of the scanner when taking images for the previously laser-scanned area. The scanner was tilted 90° to capture the top of the cupola.

ure 1.1), from inside the hall, and below in a double crypt (50 scans). Special attention was paid to the mosaic dome with the most interesting iconographic program from a smaller platform, allowing closer view (20 scans). In this case, the amount of scans also allowed for more photos and photos taken from a closer distance. When adding up all scans, the complete point data set consists of about $2.3*10^9$ points. The registration of all single scans into one huge point cloud is done semi-automatically: once the same reflector points in overlapping areas are identified, the registration runs automatically. Reflector identification is user-assisted: the software proposes a number of potential reflectors in each scan based on maxima of reflectivity, and the user can fine-tune the selection and adjust the exact reflector positions based on high-resolution close-up scans of the reflectors.

The field work was done by two persons in a total of five days (3 days on site, 2 days for traveling). About 4 scans were acquired per working hour. In areas where the scanner had to be moved only a few meters, up to 10 scans per hour were possible (e.g., in the double crypt below the building). Scanning was done over the whole day (from 8.00h to 18.00h), without longer breaks. At most 30 minutes elapsed between two consecutive scans, when having to equipment (e.g., from indoors to outdoors).

During acquisition, two problems were most notable: one was the operation of the scanner on lifting ramps, and the other was the handling of unwanted reflections of the laser beam. When operating a scanner on a lifting ramp, any movement of the lifting ramp during scanning will bias the result. Since the scans were initiated manually, a person had to be on the ramp together with the scanner. During scanning, the person had to cower beneath the scanner and remain motionless, to avoid being scanned.

This situation is shown in Figure 1.2. A complete remote control of the scanner and the lifting ramp would have alleviated this problem. The unwanted reflections were handled during post-processing, and we used the per-point information provided by the full wave-form scanner to identify and eliminate a large part of the unwanted reflections. Reflections that could not be identified this way were removed manually when cleaning the point clouds (see also Section 1.2.1).

The workflow for the scanning of the second site was similar, and the entire peristyle house of unit 1 of terrace house 2 was scanned in only one day of work. This resulted in a complete data set of 45 scan positions with some 875 million points. Approximately 6 scans per hour were taken, with a one-hour break at noon. This relatively fast working process is only possible if one knows the site already quite well, and if one can work under optimal conditions (weather, wind, electricity, etc.). If such circumstances are not guaranteed, more time will be required.

1.2 Postprocessing

In the postprocessing step, the recorded data is pre-screened, then scans are prepared for visualization and further investigation. After post-processing, the point cloud allows a complete 3D view of the colored monument, with the possibility of distance and high-quote measuring, free virtual movement in- and outside the monument, and, with special interest for the interpretation of the mosaic program, a virtual reconstruction of the access of the building and the visibility of the mosaics in the view axes calculated by the ancient owner. The "classical" duties such as ground plan generation and sectioning were done with a much higher accuracy and with the now accessible dimension of the 3D development. The archaeologist's question of why and how the ancient owner chose this form of architecture and images could not only be answered in a complex way, but the monument is also available for other archaeologists or students, for the cultural heritage management and also for the public in an incomparable accuracy and vicinity. In a colored mesh model, a virtual visit was created and a video of it was made available on the Internet [ÖAW and TU Wien 13].

1.2.1 Registration and Cleaning

The two registration methods we used are registration by markers and registration by bundle adjustment. The used markers are highly reflective and can thus be found automatically during the scan process. Registration by reflectors is done by matching manually distributed reflectors within the monument in the scans. This method is not sensitive to outliers elsewhere in the point cloud. Registration by bundle adjustment looks for feature points in several scans and is sensitive to outliers, especially when these outliers appear in clustered form, as they might be mistaken for a point cloud feature.

Depending on the registration method used, it might be advisable to clean the scans before registration. When using registration by markers, additional high-resolution scans of the markers are made, which are then used for the registration. The surveying scan itself is not used. This is in contrast to the registration by bundle adjustment. With this method, features within two or more scans are used to adjust them properly. Therefore, cleaning the scans before registering them by bundle adjustment is necessary if many outliers exist that do not represent the actual geometry of the scanned objects.

If full wave-form processing is available on the scanner [Ullrich and Reichert 05], then it is known whether only one echo or several echos of one laser pulse were received. When scanning monuments, only points which reflect a single echo or which are first echo points are meaningful, all other points are reflections on windows or other highly reflective materials that contain no meaningful information. Therefore, automatic cleaning methods can be used that take advantage of this information. This information is stored by the scanner at every point during recording and can be used to filter the scans accordingly. In the scans we used for the point models, only the points which resulted from a single echo or which were the first echo of a laser pulse were kept.

Another information that can be used to remove outliers is the quality of the reflected laser pulse. It is recorded as deviation at each point by the scanner. The deviation captures the difference of the reflected laser pulse to the emitted laser pulse with respect to the shape of the signal. If the shape of the reflected signal is identical to the emitted signal, the deviation is zero. The more the shape differs, the larger the deviation value becomes. For the laser scans we did, only cleanly reflected laser pulses are of interest, as more distorted ones hint at pulses that are reflected from corners or walls at grazing angles. These reflections can be quite imprecise, with errors up to several centimeters. For the deviation, values from 0 to 65,535 are possible, and we deleted all points with a deviation larger then 17, so all points that were kept had a value of 17 or less. This threshold value was found by visual examination of the scans before and after deleting points according to this threshold.

1.2.2 Models and Coloring

For making the recorded data easily accessible, we create a point cloud model, and in some specific areas also mesh reconstructions. Consolidating the available data into a virtual model is useful to get an overview of the data, and can be used to reveal connections between different areas of the recorded monument. Registered photographs can be used to create a mesh model by photogrammetric methods (see Chapter 4 previous Chapter). To this end we use the photogrammetric software by Adam Technology [ADAM Technology 14]. If point clouds from laser scans are available, we create a model of all registered laser scans combined with our software Scanopy, which enables us to render point clouds in real time and edit them interactively. Scanopy is developed for internal use, however may be made available on request and individual licensing terms.

Meshes We first describe our approach to capturing detailed representations of specific models using meshes. This is done most efficiently by taking photographs with a camera. These photographs can also be used to create a geometric model of the object. For this purpose, enough photos have to be taken to cover the complete surface of the object of interest [Mayer et al. 11]. This means that for each area of interest, at least one photo should be taken that is oriented orthogonally to this area. For example a room with a 2m x 2m floor space has to be covered with 18 photos or more, as the photos of neighboring areas should overlap. After the photos have been taken, their position and orientation relative to each other can be determined, and for this we used the software CalibCam [ADAM Technology 14]. The software calculates the interior and exterior calibration of the camera as well as the used lens. Using pairs of photographs showing the same area, mesh models of partial areas of the object can be created. This can be done with the software 3DM Analyst [ADAM Technology 14]. To create a single mesh from the previously generated meshes, all vertices of the meshes are imported into the software Geomagic [Geomagic 14], where a clean mesh model can be created. The created mesh has no texture applied yet, but for each taken photograph the position and orientation matrix is available. With this information, a re-texturing of the mesh model can be done. For this we use the same images that were used for creating the mesh model.

After determining the photographs to be used as textures for a mesh, we need to find a method to calculate a consistent texture for the whole mesh from the overlapping photos, which usually also contain lighting variations, leading to seams. For this, we use a two-fold approach. First, we try to find a consistent assignment of photos to mesh triangles, and second, we try to compensate the differences in the lighting of the photographs. Our approach is based on [Lempitsky and Ivanov 07], where in the first phase a photo is assigned to each triangle of the mesh for texturing, and in the second phase the seams are leveled in the gradient domain, i.e., the colors of the photographs along the seams are adjusted to not change abruptly.

We improved the leveling to account for large changes in the brightness of neighboring photos [Birsak 12]. If seams are still visible in the textured model, the remaining artifacts are evened out manually by a professional artist with an image manipulation software like Photoshop [Adobe Photoshop 14]. To easier identify the areas in a photograph that are used for texturing, we developed a tool that creates a mask for the areas of a photo which are used for the texture on the model. This mask can be used as a layer in Photoshop to highlight these areas.

Having (adjusted or original) photographs for texturing the model, they can optionally be stored in a so-called texture atlas. A texture atlas stores all photos that are used for texturing one or more meshes. We implemented a so-called virtual texturing method that uses a texture atlas with a size from 32k x 32k up to 256k x 256k texels and creates a level-of-detail (LOD) representation on top of this texture atlas, i.e., a chain of texture atlases with successively lower resolutions [Mayer et al. 11]. With this virtual texturing approach, the large texture atlases can be partially loaded to graphics memory during rendering. This is accomplished by storing the texture atlases as tiles, which have a size between 64 x 64 and 256 x 256 texels, and these tiles are then loaded to graphics memory. The LOD representations of the texture atlas are used when the full resolution of the texture atlas is not required.

Point Clouds For the overall model, we typically use a point-based approach, avoiding the creation of meshes for the whole monument. During the scanning process, photos have been taken from the scanner's position. These images are registered to the scanner's position and can therefore easily be used to color the point samples. For mapping the photographs to the scanned points, we use the software accompanying the laser scanner, in our case RiScan Pro from Riegl [Riegl Laser Measurement Systems 14b]. The colored point scans are stored in addition to the original point scans on disk.

From the previous postprocessing steps, we have a registered, cleaned, and colored point cloud from each scan position available. With the current generation of computer hardware it is not possible to simply load all scans into memory for visualizing them together. There are two bottlenecks here, first the video memory of the graphics card is too small, and second the number of points that can be rendered per frame is not sufficient for this simple approach. Other, more subtle problems like pixel overdraw also make this approach infeasible.

Therefore, we build a single point cloud from all scan positions by inserting the points into a hierarchical data structure. This allows us to access the different areas of the complete point data set faster, and furthermore enables a LOD representation during rendering. For the hierarchical

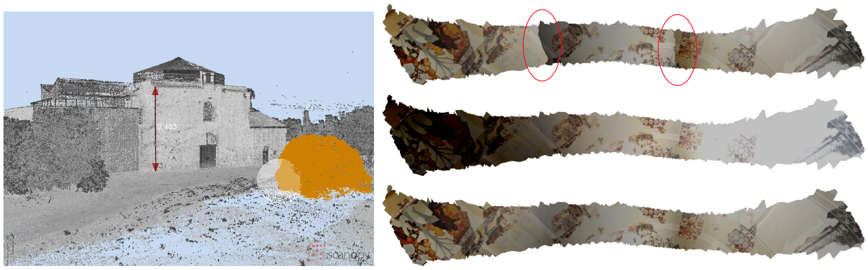

(a) (b)

Figure 1.3. The architecture visible in Image (a) is the central entrance building of an ancient villa, and from outside we want to see the original structure without the modern vegetation. A tape measure for measuring the height of the monument is applied. The point model is colored by the reflectance measured by the laser scanner at each point. The reflectance is determined by the amount of light returned at a measurement point, i.e., the less light reflected the darker the point is colored. The cubic building has a huge dome that conserves inside a unique mosaic program that we wanted to show in best lighting conditions. In Image (b), one part of this mosaic is shown 3 times. The upper row shows 2 seams (circled red) where neighboring photos exhibit different lighting conditions. The middle row is already leveled according to [Lempitsky and Ivanov 07], but the brightness levels are not satisfying. In the lower row, the leveling is done according to [Birsak 12], where brightness levels are balanced.

data structure we use an octree, and in each octree node we inscribe a 3D grid [Scheiblauer and Wimmer 11]. One side length of the grid has 128 cells, and each grid cell can store one point. The points from the original data set are stored in leaf nodes and inner nodes, and the memory consumption of the data structure is the same as for the original point data. We have implemented this data structure in our software Scanopy, where we use it to visualize and edit point clouds. We access the original data directly as it is stored by the laser scanner in a point scan data base, and do not convert it before using it in Scanopy. This way, we do not have to export the point data to a separate file format. The resulting point cloud is stored in a proprietary file format, which enables fast access to the data for rendering and editing.

1.2.3 Interactions to Edit Point Clouds for a Specific Task

Having a point data set stored in our hierarchical data structure, we are able to delete parts of it in an interactive manner. To this end, we extended our point renderer with a selection mechanism, to allow the user to select points from a point cloud [Scheiblauer et al. 09]. The user can move a so-

Figure 1.4. The original structure of a Roman peristyle house. The large parts of its floor mosaics and the wall painting program on the walls allows a direct impression of the ancient housing and furnishing fashions. It is possible to separately show (after marking and editing) the ancient housing and modern protecting structures. In this image, a cutaway view of the theater room in the point model of the terrace house 2 in Ephesos is shown. The walls are made temporarily invisible. The protective roof that is spanned over the monument has been deleted from the point model.

called selection volume, i.e., a sphere or a cuboid, along the points that are currently visible on screen. When pressing the left mouse button, he can mark the points he wants to add to the set of selected points. This way, larger selections can be composed from points of different areas in the point cloud (see Figure 1.3). When pressing the right mouse button, points can be un-marked and removed from the selection again. After having selected a set of points, the user can either hide the selected points temporarily (i.e., making them invisible), or delete them altogether from the point cloud. Making points temporarily invisible can be used to emphasize certain areas of a model, for example by hiding walls that hinder a free sight onto this areas.

We show an example for selection and deletion in the case of the terrace house 2 in Ephesos: There is a protective roof, which is stretched above the monument, and this roof is also present in the point model, but it is hindering the exploration of the data set, especially from the bird's eye view. This roof was completely deleted from the point cloud, and only the ancient building parts were used for further examinations, but before ultimately deleting the roof, we stored a copy of the original point cloud, in case some errors were made during the editing operation. We also prepared a virtual visit to the most beautiful space, the so-called theater room. For a comprehensive view of the context of mosaic floor and wall painting we selected the rear wall, made it temporarily invisible, and chose a bird-eye view on the monument (see Figure 1.4). Again, the archaeological basics like measuring or leveling are now tasks easy to accomplish.

Another editing task involving deletion is the preparation of the point cloud for recording a camera path. Here every single point that is not part of the monument but rather an artifact due to particles like dust “floating” in midair, becomes an annoyance. Before recording the camera path, it is retraced, so that outliers along the path become visible and can be deleted. Also the layout of the camera path through the model can be enhanced this way, and waypoints can be added or deleted accordingly.

A further editing interaction offered by our system is the annotation of points. As in Google maps, we can create markers annotating any point in the data set, and store them in a data base, or link them with an existing data base. Any finding (for the archaeologist) but, of course, also any kind of information (e.g., for tourists), can be stored within the point cloud, expanding the point cloud or model of a monument to a virtual 3D data base.

1.3 Visualization

The models we created from our data, either textured meshes or point clouds, are huge in terms of memory consumption. Therefore, we cannot simply load all data into graphics memory and render it to the screen. Instead, we select chunks of data that are needed for the current view and load them into the graphics memory asynchronously, so that after some short time the model is visible in full detail. For camera paths, this works as well, but it might be helpful to move the camera slowly along the path, so there is enough time to load the data from disk.

1.3.1 Rendering

When rendering the models, we aim for an interactive visualization of the models, where the user can move around freely inside the model. Therefore we have to load the data currently visible in real time. When rendering meshes, we can access the texture in an out-of-core manner, while the geometry of a mesh is loaded completely into graphics memory at once.

When rendering point clouds, we load the points partially with an out-of-core strategy into the graphics memory.

Meshes As described in Section 1.2.2, we store textures with up to $65*10^9$ texels in size. During rendering, we determine which tiles of the texture are currently visible and how far away the viewpoint is from the mesh model. Depending on this information, we choose an area and a LOD of the texture that is required for the current viewpoint. The tiles covering this area are then loaded to graphics memory [Mayer et al. 11]. A single mesh is about 200k vertices in size, so we can store about 100 of them on the graphics card with 2GB video memory. The rest of the memory is used for storing the tiles used by the virtual texturing algorithm.

Point Clouds A point cloud stored in our hierarchical data structure is rendered by searching for the most important nodes of the hierarchy and rendering them on screen. The importance of the nodes depends on the current viewpoint, and if some nodes are not available in the video memory of the graphics card, they are loaded from disk as one chunk of points per node. This is done asynchronously, so the rendering is not interrupted while the points are being loaded. If the viewpoint is further away from the model then a lower LOD is chosen. A lower LOD requires less points to be rendered to completely represent the model. As the point density throughout the whole model is often not equal, we use a weighted point size for rendering the points. This means that the size of a point on screen is dependent on the density of the points currently rendered [Scheiblauer and Wimmer 11]. With this method we can alleviate the occurrence of holes in undersampled regions.

During coloring the point cloud, we store one color value per point (see also Section 1.2.2). Rendering each point with a single color can lead to artifacts and noise in the color domain, if points are rendered as screen-aligned splats with a side length of more than 1 pixel. When using such splats, the color information from the original photographs can be reconstructed more adequately by using Gaussian splats [Scheiblauer and Pregesbauer 11]. A Gaussian splat averages the color information of neighboring splats depending on their weighted distance, and the color is calculated for each pixel. This alleviates the color artifacts and makes more details visible compared to when using only one color per splat, as can be seen in Figure 1.5.

1.3.2 Films and Camera Paths

We added the possibility to define and edit camera paths in our point rendering software Scanopy. Once a camera path has been defined, the movement of the camera along it can be interrupted at any time, for example when using the camera path during a lecture or a virtual guided tour inside the monument; one can take a new point of view or show other details of the mesh model/point cloud. The most important advantage of this technique is that the point cloud does not need to be modified or manipulated during post-processing. Being able to show only partly a monument is very helpful sometimes, and in our case we do not lose density of data, but can offer a modified view of the entire data set.

In Figure 1.6, a camera path inside the cupola of the virtually textured mesh of Centcelles is shown. As we can load point cloud models and virtual textured mesh models at the same time, we can use the camera paths to show both types of models at the same time. When the camera moves along the path, the out-of-core algorithms stream the required data in real time from disk to the video memory on the graphics card, so that the models appear in their highest resolution from the current viewpoint. The camera path is a great assistance for non-technical users, as they are guided along the most interesting parts of the cupola, without having to learn navigating inside the virtual scene. Even for technical users, a camera path can be helpful, because although they know how to navigate, they might not be aware of the most interesting parts of the point model, so the camera path acts as a guidance tool for them as well.

1.4 Summary

Huge point models resulting from extensive scanning campaigns can be used for the documentation and visualization of vast archaeological monuments. Such point clouds can be used to solve several problems, for example finding the extents of the site, documenting varying building materials, viewing the monument under different lighting conditions, or limited access to a site.

The workflow for creating point clouds by scanning monuments is divided into several parts, starting with a preparation phase. This includes, amongst other things, acquiring the proper equipment, planning the scan positions, and negotiations with the property owner about access to the monument and the exploitation of the recorded data. At the monument, the scanning takes place and also an initial registration of the scans, such that a first impression of the complete data set of the scan campaign is possible. After the work in the field, post-processing steps can be conducted. In these steps, (optionally) cleaning the scans, registering them properly, and coloring them is done. From the colored scans, a compound point cloud of the single scans can be built using specialized data structures, which divide the complete point set into chunks. These enable managing the data set out-of-core, therefore only those chunks of points visible from the current view point have to be loaded to the graphics card. The compound point cloud can still be edited, and so it is possible to prepare it for special view points by making areas temporarily invisible, or deleting outliers that are most noticeable after creating it from the single scans.

Confined areas which exhibit particular surface features, like mosaics or wall paintings, can be captured with photographs, and from these photographs textured meshes can be created. These meshes cannot be edited, but serve for the visualization of landmarks within a point model of a vast monument. The textures can use high-resolution images, as they are also managed out-of-core.

We implemented the out-of-core point rendering and mesh texturing algorithms into our rendering framework, which is developed in-house. A video of the Centcelles cupola rendered with our rendering framework, is available on YouTube.com [ÖAW and TU Wien 13].

Figure 1.5. The figures consist of a group of hunters in the lowest of three round-freezes with best quality mosaics in the dome of the entrance to the Roman villa in Centcelles. The hunting scenes are part of the owners self-representation, the upper freezes show biblical scenes and self-representations of the patron's entire family. In the image, the upper row shows the point model rendered with box splats of 3 pixels side length. Each splat has one color. The lower row shows the same point model rendered with Gaussian Splats, where the color information is calculated per pixel.

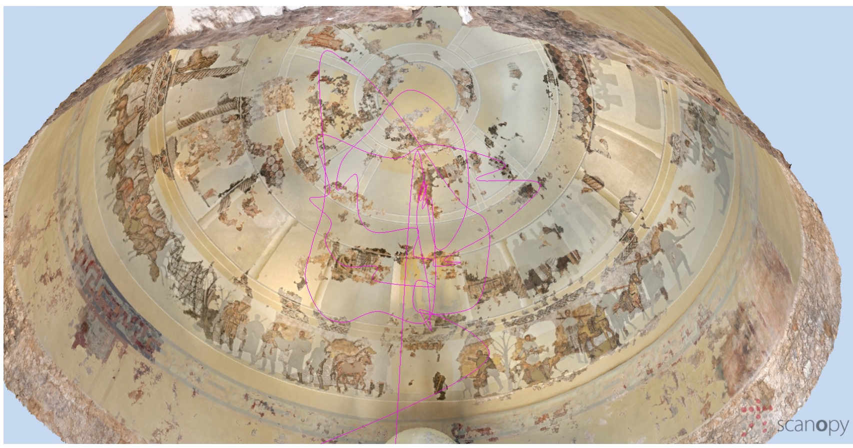

Figure 1.6. The pink line shows a camera path layed inside the cupola of the virtually textured mesh of Centcelles. The camera path has two aims: first, to give a good overview of the general setting of the monument and the dome, and second, to show the main axis of the mosaic program and to pass then every single scene, in all three mosaic freezes and the central circle, at least one time in a relatively slow motion. At any time, the path can be interrupted, turned back, be forwarded, or be left aside for individual explorations, just to ensure the user/viewer has the best understanding of the beautiful and very complex mosaic program.

Bibliography

[ADAM Technology 14] ADAM Technology. "Company website.", 2014. Available online (http://www.adamtech.com.au/).

[Adenstedt et al. 14] Ingrid Adenstedt, Irmengard Mayer, and Norbert Zimmermann. *Der Kuppelbau von Centcelles (Tarragona, Spanien). Neue Forschungen zu einem enigmatischen Denkmal von Weltrang*. Achim Arbeiter, Dieter Korol, Dirce Marzoli, 2014.

[Adobe Photoshop 14] Adobe Photoshop. "Product website.", 2014. Available online (http://www.photoshop.com/).

[Birsak 12] Michael Birsak. "Coloring Meshes of Archaeological Datasets." Master thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria, 2012. Available online (http://www.cg.tuwien.ac.at/research/publications/2012/Birsak-Michael-2012/).

[Geomagic 14] Geomagic. "Company website.", 2014. Available online (http://www.geomagic.com/).

[Lempitsky and Ivanov 07] Victor Lempitsky and Denis Ivanov. "Seamless Mosaicing of Image-Based Texture Maps." In *CVPR*, pp. 1–6, 2007. Available online (http://dx.doi.org/10.1109/CVPR.2007.383078).

[Mayer et al. 11] Irmengard Mayer, Claus Scheiblauer, and Albert Julian Mayer. "Virtual Texturing in the Documentation of Cultural Heritage." *Geoinformatics FCE CTU*. Available online (http://www.cg.tuwien.ac.at/research/publications/2011/mayer_2011_vt/).

[ÖAW and TU Wien 13] ÖAW and TU Wien. "Centcelles - spätantike Mosaikkuppel.", 2013. Available online (https://www.youtube.com/watch?v=pIHYRd_thuM).

[Riegl Laser Measurement Systems 14a] Riegl Laser Measurement Systems. "3D Terrestrial Laser Scanner RIEGL VZ-400.", 2014. Available online (http://www.riegl.com/nc/products/terrestrial-scanning/produktdetail/product/scanner/5/).

[Riegl Laser Measurement Systems 14b] Riegl Laser Measurement Systems. "Company website.", 2014. Available online (http://www.riegl.com/).

[Scheiblauer and Pregesbauer 11] Claus Scheiblauer and Michael Pregesbauer. "Consolidated Visualization of Enormous 3D Scan Point Clouds with Scanopy." In *Proceedings of the 16th International Conference on Cultural Heritage and New Technologies*, pp. 242–247, 2011. Available online (http://www.cg.tuwien.ac.at/research/publications/2011/scheiblauer-2011-chnt/).

[Scheiblauer and Wimmer 11] Claus Scheiblauer and Michael Wimmer. "Out-of-Core Selection and Editing of Huge Point Clouds." *Computers & Graphics* 35:2 (2011), 342–351. Available online (http://www.cg.tuwien.ac.at/research/publications/2011/scheiblauer-2011-oocsehpc/).

[Scheiblauer et al. 09] Claus Scheiblauer, Norbert Zimmermann, and Michael Wimmer. "Interactive Domitilla Catacomb Exploration." In *VAST09: The 10th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, edited by Kurt Debattista, Cinzia Perlingieri, Denis Pitzalis, and Sandro Spina, pp. 65–72. St. Julians, Malta: Eurographics Association, 2009. Available online (http://www.eg.org/EG/DL/WS/VAST/VAST09/065-072.pdf).

[Ullrich and Reichert 05] Andreas Ullrich and Rainer Reichert. "High resolution laser scanner with waveform digitization for subsequent full waveform analysis." In *Proc. SPIE*, 5791, 5791, pp. 82–88, 2005. Available online (http://dx.doi.org/10.1117/12.623847).