Challenges in Heritage Documentation with Terrestrial Laser Scanning

Heinz Rüther, Christoph Held, Roshan Bhurtha, Ralph Schröder, Stephen Wessels

Zamani Research Group

Geomatics Division, University of Cape Town

heinz.ruther@uct.ac.za

Abstract

The paper discusses data acquisition and processing of laser scan point clouds based on experiences gained while documenting heritage sites for the African Cultural Heritage Sites and Landscapes database. More than 6000 individual scans of buildings, rock shelters and landscapes have been recorded and have shown the advantages of laser scanning over other methods, but also exposed challenges which need to be addressed to develop the full potential of laser scanning. Experiences with field work, scan registration, hole-filling, data cleaning, modelling and texturing as well as display options are reported and briefly discussed.

A sometimes unrealistic view of the potential of terrestrial laser scanning as the answer to all spatial documentation is emerging among potential users of laser scan data without technical experience of 3D data acquisition and the complexity of the processing. An attempt will be made to correct some of the unrealistic expectations and to make the user community of laser scan data aware of the complexity and limitations of present data processing methods and technologies while not taking away from the unquestionable value and relevance of laser scanning for heritage documentation.

Keyword: Terrestrial Laser Scanning, Cultural Heritage, Digital Preservation, Surface Reconstruction, Texturing, Visualization, Registration, Field Procedures

Introduction

Laser scanning has found its way into heritage documentation rather rapidly and it has partly replaced some of the conventional methods for the spatial documentation of heritage sites. While the technique was exceptionally well received by technical and non-technical users, it is often misunderstood and its capabilities are over- or underrated. The principle of using laser scan data as a means to record sites as a documentation for the future is excellent. Laser-scan generated 3D models when displayed and manipulated on the computer screen are generally well received and praised by archaeologists, conservation experts and architects, however, when it comes to the planned practical applications of such models, an element of bewilderment often replaces the original enthusiasm. It is therefore important to understand the capabilities and limitations of the technology and relate these to the needs of the 3D data user community.

The authors, all members of the Zamani research group at the University of Cape Town, have been exposed to the realities of laser scanning during the ongoing development of the African Cultural Heritage Sites and Landscapes Database, a project funded by the Andrew W. Mellon Foundation and based at the University of Cape Town. The paper briefly describes the database and its concepts and reports on experiences gained with processing close to 100 structures recorded with six different scanners.

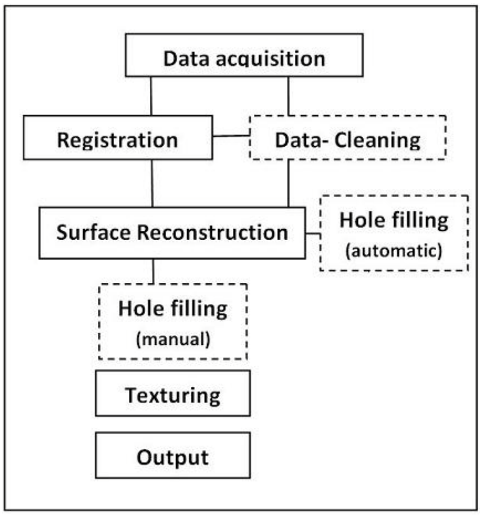

Figure 1: The 3D modelling pipeline. Dotted lines mark optional steps

Data processing

The processing pipeline (Figure 1) for laser scans can be subdivided into data acquisition, registration, data cleaning, modelling, hole filling, texturing and output. Each step of the pipeline will be briefly discussed and some of the challenges associated with each step will be mentioned.

Data acquisition

Laser scanning field work has changed significantly since the first instruments were employed for heritage documentation. In its earliest missions, some five years ago, the Zamani team completed between five and ten full-dome medium resolution scans (2 cm at 20 m distance) per field day and a typical scanning mission returned with 40 to 50 scans per site. Today's field campaigns with phase based scanners can acquire more than 1000 individual scans per site and point clouds have increased from between 20 to 50 million points per site to up to 7 billion points. The dramatic increase in the number of set ups over time was a consequence of the development from early time-of-flight scanners with scan times of two to three hours for a full- dome scan to phase based scanners with scan times of three to six minutes for the same resolution. The highest number of scans with this resolution acquired by a single operator in one day with a Leica HDS 6100 scanner was 140. This was possible not only because of the high scan rate of the phase based scanner, but also because of the one-button operation and the built in data storage and batteries, which did away with the time consuming transport of laptop, batteries and cables from station to station.

Essential for a scanning field team is not only field experience, but, more importantly, experience in all aspects of data processing. Only a clear understanding of the data pipeline and the complexities of each step will guarantee that the principal criteria for a successful field campaign can be met. These include:

- complete point cloud coverage of the structure/site, or at least as complete as physically possible. One has to accept that complete coverage of a monument or site is in most cases unachievable.
- appropriate choice of resolution depending on the complexity of the surface, the level of detail required and on the variation in distance and angle to the scanned surfaces, bearing in mind that resolution changes with distance from the scanner.
- sufficient overlap for ICP based registration, where overlap areas must be chosen to contain sufficient surface detail to allow registration algorithms to find a unique solution. Overlaps should also be chosen to allow for similar angles of incidence from the scan stations.
- choice of scanner positions which avoid flat angles of incident for all surfaces, but especially for overlap areas.
- choice of an economic target distribution where targets are visible from a maximum number of instrument setups, without however, sacrificing geometric requirements for stable 3D transformations.

The Zamani project has adopted a field procedure in which the scanner is always levelled, and oriented in the same, arbitrarily chosen, direction. This information, optionally combined with the approximate position and height of the scanner, significantly speeds up the registration process but is not essential. Buildings are typically scanned with 1- 2 cm point spacing (or higher where need be), while terrain is captured with point intervals varying from 10 to 50 cm.

Experience showed that preplanning scanner setups is impractical and, in most cases of heritage documentation, impossible. This is because one cannot assess, without being on site, optimal field-of-view, overlap areas, resolution and necessary additional scans to cover vital occluded detail. Therefore the decision on suitable scanner positions are always done on-the-fly.

Figure 2: Scanner positions in Songo Mnara, Tanzania

Registration

The first step in producing a 3D model is the registration of the individual scans acquired in the field. This involves the transformation of all scans into a single uniform coordinate system. Registration is a vital element in the data processing pipeline, as the accuracy achieved in this stage influences all subsequent processing steps. There are two approaches to combining the individual scans into a single point cloud. One option relies on targets common to some, or all, of the scans while the other approach aligns scans based on overlapping surfaces using well established variants of the Iterated Closest Point (ICP) algorithm [1,2].

Targets are generally easy to identify and locate in point clouds during the registration process, thus reducing processing time especially with software allowing for automatic target detection. Targets also provide high registration accuracies and reduce the danger of mis-registration. The disadvantage of the use of targets is the requirement of high resolution sub-scans of the target area which in turn increases scan times in the field, unless the entire scan is executed at high resolution. It is possible to locate targets outside the scanned object, provided that they are visible from more than one scan position. In heritage documentation, however, one encounters a number of practical problems when using targets. It is often difficult, if not impossible, to physically place targets on fragile or high walls and in difficult-to-access areas. The option of placing targets outside the object is obviously not possible when scanning the inside of buildings. A further argument against targets is the complexity of heritage buildings. For example, the so-called ‘palace’ at the ruined Swahili town of Songo Mnara in Tanzania comprises of more than 100 rooms. In this case it was the large number of targets required which made this approach impractical. This would have required a lot of field time to set up the targets which would also involve planning all the scan positions beforehand, as at least three common targets with other scans need to be visible from every position.

The policy adopted for the scanning of sites for the African Heritage documentation project was therefore to rely on surface based registration and to use only limited numbers of full-sphere targets for very large sites which are placed in exposed positions where they can be viewed from a maximum number of scans. The coordinates of these targets are determined by post-processed GPS surveys wherever possible and they serve to check and position the final model. For some sites point clouds were registered with both, target and surface based methods, but no significant difference between the accuracies of the resulting registered point clouds could be detected.

The number of scanner setups for sites is a function of the complexity of the site and, to a lesser extent, the size of the documented structures. In case of the African Heritage project, scan numbers have increased from an average of 40 per site for scans during the first project year to more than 1000 per site for subsequent documentations. The increase in the number of setups over time was mostly a consequence of technical development of laser scanners, with scan times of two to three hours for a full-dome scan in the early stages, compared to scan times of three to six minutes for the same scan resolution at present. Experience also shows that minimising scan set ups in the field can lead to difficulties when processing the data. Registration of scans is more difficult and more prone to errors when there are less scans and thus less overlap area, which can lead to an unstable registration.

The project's field team adopted a field technique which significantly reduces the time required for registration, and especially surface based registration. In this approach the scanner is levelled and aligned in the same direction, either visually or with the help of a compass, and scanner positions are surveyed by means of a total station or RTK-GPS or estimated (Fig. 2). This way scans introduced into the registration software are approximately pre-oriented and positioned, thus avoiding the need for a time-consuming, manual pre-alignment, the prerequisite for the more precise alignment via the ICP-Algorithm.

Soft-and hardware limitations become obvious when registering large numbers of scans, which require long processing times and can cause system crashes. The Zamani teams has adopted an approach in which a skeleton model of the entire site or structure is created by first registering a minimum number of individual, usually, longer distance scans. This is followed by a global registration of the skeleton model, provided the data volume allows this. The skeleton model is then filled in with the bulk of short range scans.

It should be noted that the registration phase in the modelling process is crucial and any mis-registrations can prove to be very costly later in the modelling process.

Cleaning of point clouds

The cleaning of the data, i.e. the removal of objects not relevant to the documentation, is an important step in the 3D modelling pipeline. As a rule the Zamani team cleans individual scans of unwanted objects such as trees, people, cables, cars, doors, animals and random objects prior to the final modelling process. Pre-model cleaning results in smaller models, especially on sites with dense vegetation (Fig. 3). Only minor remaining flaws are dealt with on the model.

Cleaning of scans can be performed on individual scans or on the complete point cloud. The latter reduces cleaning time as objects appearing in multiple scans can be removed in one operation (Fig. 4). The disadvantage of this approach is the loss of the individual scan-files: When all points from all individual scans are merged into a single point cloud, they will lose their grid formation and thus the information about their closest point neighbours, which is necessary for the project's current workflow towards a surface model.

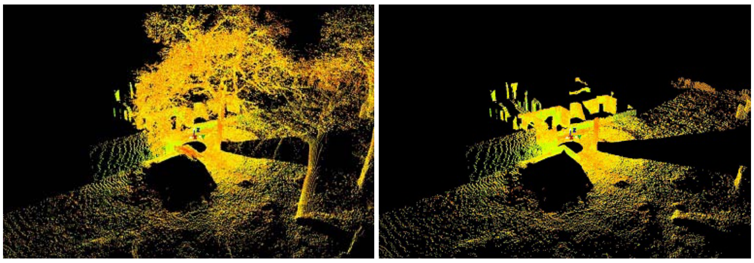

Figure 3: Uncleaned point cloud with 6.7 million points (left), cleaned point cloud with 4 million points (right)

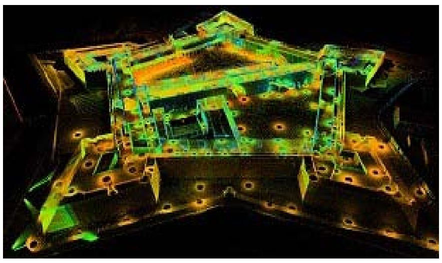

Figure 4: Screenshot of the point cloud of the Castle, Cape Town, South Africa

While some artefacts in scans can be removed automatically most cleaning of scans is still a manual, tedious and time consuming task. A typical scanning project in the African heritage project comprises of 500 – 1000 scans each of which typically requires 0.5 to 2 hours of cleaning time by an experienced operator, while complex and very high resolution scenes can take a full day of cleaning per scan. Here, further development of automatic or semi-automatic cleaning tools, capable of working with large datasets could greatly improve the workflow.

Creating a Surface

Meshing, the conversion of the discrete point set to a continuous surface, represented by triangles, is the next step in the pipeline. Surfaces can either be formed by directly connecting points which then become part of the surface [3,4] or by creating a best-fit surface through the points [5,6,7]. The former approach usually employs the Voronoi diagram and the corresponding Delaunay triangulation to find point neighbours and create connections. This method allows the creation of the model from a complete, unstructured point cloud, but it has the disadvantage of being very susceptible to instrument noise, outliers and alignment errors. Noise turns a surface, which in theory should be infinitely thin, into a narrow 'band' of points. Connecting the closest neighbours in this 'band' without filtering can cause intersecting triangles and faces which are non-consistent in their orientation to their neighbours. Thus a noise-filtering pre-step becomes necessary. Depending on the dataset, this can be a lengthy and complicated process. Noise reduction is generally indiscriminate and unfortunately accompanied by a loss of detail. Noise typically increases in overlapping areas. Mis-registrations also create multiple surfaces and this must be avoided at all cost, regardless of the meshing algorithm used.

Best-fit algorithms automatically perform noise reduction by approximating the surface. Secondary information about points, such as the relation to their neighbours, might lead to additional smoothing. For example, the indirect determination of point normals by averaging between the normals of their neighbours, could have a smoothing effect. Other approaches employ pre-triangulated scans to extract a surface by averaging between overlapping areas and thus smoothing the input data.

Depending on the source data, calculating vertex normals can be very trivial or very complex. It is trivial when the data is provided in grid form, which is the case for individual scans. Here the normals can be retrieved by using the scanning grid information which indicates nearest neighbours. If there is no such information available, or if two or more overlapping scans are aligned and merged into one cloud, the grid structure is lost and nearest neighbours have to be found in time consuming computing processes.

In general, a compromise has to be made between reliability and completeness of the model, as 'completeness' might require the inclusion of uncertain data, which could, for example, be created by an inappropriately set threshold, which decides whether a point's neighbour is close enough to be used. Setting this threshold too large will create artefacts and connect objects which should not be connected, or, if set too small, will create holes, even though data is available.

This is especially problematic for individual scans where sampling density decreases with distance from the scanner, resulting in increasingly unreliable point connections and normals. This problem is reduced when working with complete point clouds and thus considering points from other scans. In this way, sampling density is increased in previously sparse areas resulting in a more complete and still reliable model.

It would seem that none of the currently available modelling algorithms produce a perfect model. At present, a combination of multiple versions of the same model, created with different algorithms, appears to lead to the most complete and detailed model, although thorough completeness cannot be achieved. This merging process is not trivial as the various models can differ significantly in resolution and detail (Figure 5a, 5b).

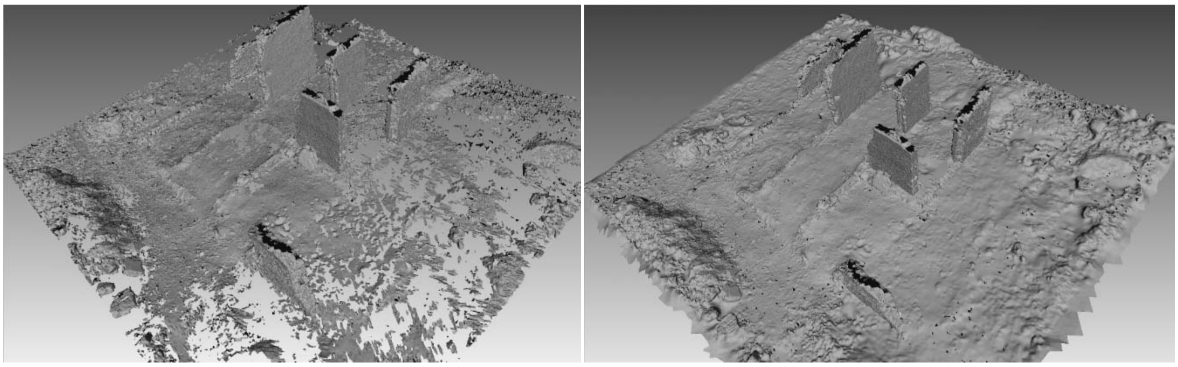

Figure 5a(left):Model meshed only with the volumetric merging method, Figure 5b(right):Combination of the Poisson algorithm for the ground and a volumetric merging of the walls

Another possible cause for artefacts in the modelling process is the nature of the scanned surface. For example grass, leaves and branches of bushes or trees or moving objects like curtains, cars and people should be cleaned out before creating the model. While the latter are generally easy to identify and remove, grass and leaves are more complex, since they are typically only represented by a few points. They might be modelled by the meshing algorithm to produce only a single sided surface with a complex boundary, since a blade of grass or a leaf of a tree or bush is hardly ever scanned from two sides, which is difficult to repair at a later stage. Current vertex-normal-based implementations [8,9], tested by the project, have a good and reasonable smoothing effect on grass without smoothing out too much important detail on the objects of interest, but unfortunately they are not very robust when applied to such irregular surfaces and often crash.

Meshing the large models encountered in the African Heritage project cannot be done with standard, in-core software and thus out-of core tools, and /or streaming techniques with the automatic ability to split up data, are becoming essential. The present approach of the Zamani team is to split the data into subsets to be used for parallel processes with out-of-core tools. Overlaps at subset boundaries are introduced to allow smooth connections of neighbouring sets.

This way, the models for the African Heritage project are usually extracted at a resolution of 2 cm with a software variant of the Volumetric Integration approach [10], but much higher resolution models from the raw data could be generated, if the need arises (Fig. 6).

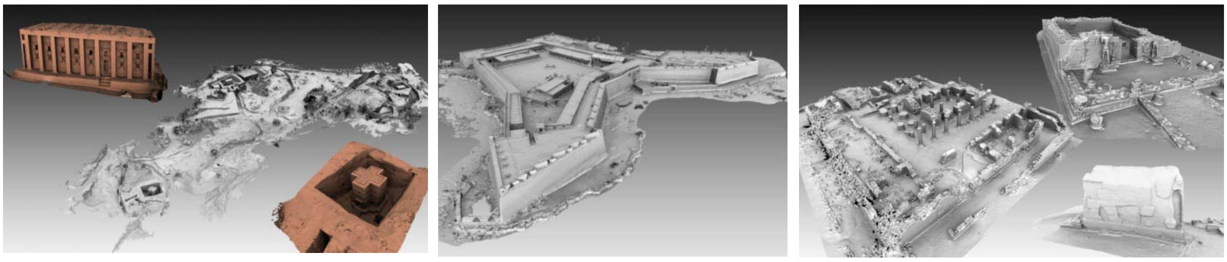

Figure 6: 3D model screenshot of two churches in Lalibela and the full 3D landscape of the Lalibela church complex, Ethiopia (left), screenshot of the 3D Model of San Sebastian Fortress on Mozambique Island, Mozambique (centre), 3D model screenshot of Musawwarat es-Sufra, Sudan (right)

Hole filling or surface augmentation

Only in very exceptional cases, if ever, is a laser scan model free of scan holes. Holes occur especially when scanning very complex and irregular objects, or wherever else a surface is invisible to the scanner. Typical examples are ornamental building facades or upward facing surfaces where no scan positions can be found above the surface, such as window ledges and roofs. The same problem arises when acquiring photography for texturing the model.

If no scan data is available, current software fills in small holes automatically and offers conventional modelling or cloning options for larger patches (Fig 7). This semi-automatic approach is a time intensive task, since it is heavily user-based.

Methods for automated hole-filling or surface augmentation have been developed [11] but their use in heritage documentation is questionable. Surface augmentation algorithms can fill holes plausibly, which makes it difficult and even impossible to distinguish between real surface data and artificially introduced patches. This is not acceptable for the scientific documentation of cultural heritage sites, where data might be used for research or restoration. On the other hand, hole free surfaces are aesthetically more appealing and necessary to produce interactive 3D walkthroughs. Watertight models are also required when producing physical to-scale models with a 3D printer. It would seem desirable for the heritage scanning community if software performing hole-filling and model-viewing could make use of a standard display format clearly indicating augmented surface portions, on request.

Photorealistic texturing

Colour can enhance a 3D-Model not only visually, but also assist with interpretation and diagnosis. Photorealistic colour of a surface can, for example, assist with the detection and monitoring of eroded, chemically changed or restored surface areas. While textured models might be desirable, but not essential, for many applications, they are critical when modelling rock art shelters (Fig. 8)**.**

Laser-scanner manufacturers increasingly equip their instruments with photo- or video cameras to achieve the colourisation of the scanned surface. These built-in systems are convenient to use, but so far they do not seem to reach the quality of independent, external cameras. Images tend to have a blurry or milky appearance. Some manufactures thus offer adapters for external cameras which take a panoramic image from the position of the scanner. The biggest disadvantage of these approaches is the inconsistent lighting between scanning positions. A scanning campaign of a heritage sites usually spans several days with the scanner operating throughout the day and in some cases at night. This necessarily leads to scanning under very different light conditions and if the images are taken concurrently with the scans, they can differ significantly throughout the model. This suggests that the optimal way to acquire texturing photography for a model is independent photography, taken over a minimal time period, or on different days with approximately the same lighting conditions.

To place an image onto a 3D-Model, the camera parameters, internal and external, have to be known. An accurate projection can only be accomplished with the correct position and orientation of the camera in space, the correct focal length and distortion free images. The best approach is to use images acquired with a calibrated camera and lens. But this method also has its disadvantage as it requires a preset fixed focus, which might result in some of the captured images being blurred. Theoretically the camera can be calibrated individually for each photo but this is extremely time consuming, whatever technique one uses for the calibration.

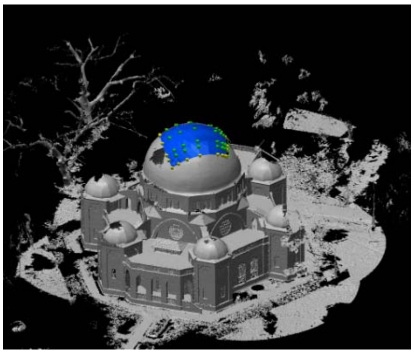

Figure 7: 3D model showing filling techniques of the Peace Memorial Museum in Zanzibar, Tanzania.

Figure 8: 3D model screenshot of the textured Keurbos cave, South Africa.

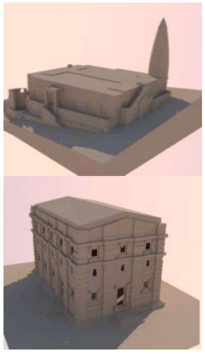

Figure 9: Line models

The internal as well as the external parameters can also be estimated or post-calibrated. If the parameters are not known at all, current software asks the user to find corresponding points on the model and on the image. But even if the amount of correlated points is reduced to a bare minimum, these points need to be chosen accurately and thus, covering a large structure completely in detail with hundreds or even thousands of images, will take a significant amount of time. Full-dome panoramic images, converted to cube maps, can be used to reduce the amount of images significantly. Per panorama, only one sub-image needs to be registered, while the remaining images differ only in their orientation. Panoramic images are nearly free of lens distortions, since the camera parameters are estimated when stitching the panorama image. Computational tools to refine initial camera parameter estimations and thus increase the accuracy of the image-model-fit have been successfully demonstrated [12, 13].

Current texturing techniques encourage the user to minimise the number of photos to reduce processing time. However, advances in structure-from-motion algorithms [14] suggest the acquisition and use of large numbers of images. These algorithms require large overlaps and small separations between camera positions and can almost automatically estimate all camera parameters required for texture mapping. However, software using these methods to colour the model, commercial, freeware or academic is still rare.

If the task of finding the camera parameters is solved, it has to be decided how to project the images onto the surface. There are two ways, either by projective texture mapping or by assigning colours to the vertices. The first approach requires that each vertex (3D model coordinates) on the model is associated with a corresponding point (2D image coordinates) on one or more of the images. It is not trivial to handle these extremely large datasets of extensive models with hundreds of associated images.

It is much easier to assign a colour directly to each vertex, and then interpolate the colour on the surface between the vertices. This requires that the original point cloud resolution is high enough to represent the required detail. If this is not the case one can artificially densify the point cloud by subdividing the original triangles. In this case each new point needs to be located in the image set to obtain its corresponding 'real' colour. Practically, this requires tools which allow the loading of large models and to then further increase their size. At present such tools appear elusive.

Data Visualisation

The virtual 3D model of a cultural heritage site should be highly detailed, dimensionally accurate, complete and ideally it should be textured. A complete, detailed and textured model of a typical heritage site is likely to be of a data volume well above the capability of today's PC graphics cards and viewing software. Current viewing solutions either make use of data decimation, subdividing the model into blocks, streaming the data, or building it in a progressive way.

Line or low resolution representations (Fig. 9) can be created to approximate the surfaces of the structure. These will show only the most important features and are commonly used by architects and engineers. Methods of extracting features automatically do exist, such as using surface curvature [15], but often involve manual guidance which is a tedious process. Although these models are visually appealing and, due to their small file size, easy to view, they lack the detail that is recorded by laser scanning, and could as well be produced by simpler, cheaper methods.

Based on the experiences gained during this project, a visualization tool for 3D models should have the following specifications: optimization for large models (both point and triangle-based), measuring capabilities (distance, volume, area), drawing capabilities (marking areas and tracing features, like cracks or rock art etc.), generation of ortho-images, interactive slicing, walkthroughs, embedding hyperlinks into the 3D-surface, changeable lighting, and layered texturing or colouring to incorporate different information such as infrared, x-ray, false colours, tracings, historic imagery etc. Currently, the project makes extensive use of Meshlab as a viewer and editor [16].

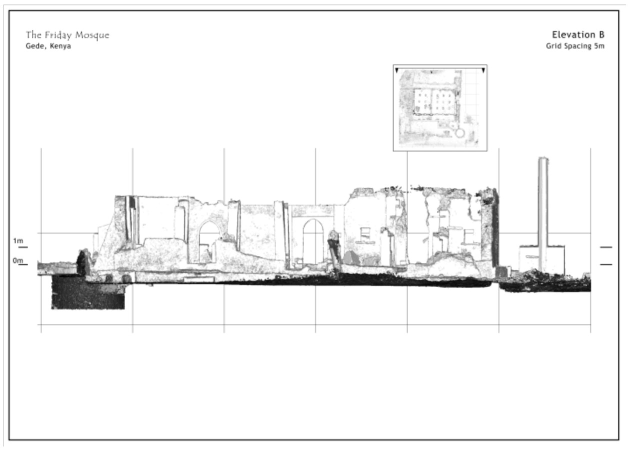

Figure 10: Elevation of the Friday Mosque in Gede, Kenya

The African Cultural Heritage Sites and Landscapes Database

The African Cultural Heritage Sites and Landscapes Database project is dedicated to recording, and thus contributing to the protecting of, cultural heritage monuments and landscapes throughout the African continent. The project recognizes the urgent need to:

- create metrically correct digital (3D) documentations of African Heritage Sites for future generations
- provide spatial information of heritage sites for conservation and restoration
- provide spatial information of heritage sites for education and research
- promote awareness of Architectural African Heritage (historical sites, buildings and towns) within Africa and worldwide
- provide management tools for site management at local and regional level

The project was initiated with a grant from the Andrew W. Mellon foundation and is executed by the Zamani Research Group, based at the University of Cape Town (UCT). The project uses state-of-the-art technologies, such as Laser Scanning, GIS, close-range and aerial Photogrammetry, Remote Sensing, as well as virtual reality technology, to create an integrated database of important African cultural heritage sites. The output of the project is made available by the JStor digital library, New York (http://JStor.org and http://aluka.org). Full data sets are available for research, education, conservation and restoration projects. A small subset of the data can be viewed on http://www.zamaniproject.org.

The Zamani group has to date documented some 27 African heritage sites in Ghana, Mali, Mozambique, Tanzania, Kenya, Ethiopia, Sudan, Egypt, Algeria, Zimbabwe, Cameroon and South Africa. All projects are executed by members of the unit, with the support of staff members of Antiquities or equivalent Government Departments. A total of approximately ten terabytes of data on African sites has been generated by the Zamani group over the past six years. These data are augmented by relevant contextual information data selected by JStor/Aluka.

All projects are carried out in close cooperation with local museum and heritage authorities as well as staff and students of local universities and the University of Cape Town. The generated information and database has been disseminated through seven workshops, the JStor/Aluka website and a much smaller UCT website.

Figure 11: Gede Palace near Mombasa, Kenya

The deliverables of the project are:

- an integrated database consisting of a Spatial/Geographic Information System (GIS) for each of the sites
- 3D computer models of structures and parts of towns and landscapes (Fig. 11)
- elevations, sections, ground plans and roof plans (Fig. 10)
- computer visualisations with walk-through and other inspection capabilities (where feasible)
- spherical panorama tours
- contextual photographs
- site related digitized documents, scientific papers, excavation reports and similar material

During the process of creating the database, a methodology for the documentation of African heritage sites has been developed and optimal ways are explored in which the data can be used by African heritage authorities and by museum officials and researchers in Africa and worldwide.

Conclusion

Laser scanning has established a prominent position in the spatial documentation of heritage sites and rightfully so. It is a powerful tool which makes it possible to collect valuable and accurate spatial information in relatively short field time. It played a prominent role in the acquisition of data for the African Heritage Database and it is unlikely that other techniques would have been able to acquire similar data volumes. However, during the course of the African Heritage Project it also became obvious that laser scanning alone does not hold all the answers to heritage documentation. Photogrammetry still has an important role to play, both, using cameras directly linked to the scanner or using traditional close-range photogrammetry, with independent cameras. Both approaches can contribute to texturing and feature extraction, while close-range photogrammetry is also essential for the capturing of surfaces which are difficult to access by scanners. It is also obvious that there is still room for development in the many complex steps in the laser scanning pipeline, from data collection and scanner specifications, registration, cleaning, meshing and hole-filling, and texturing.

Visually appealing textured 3D computer models, video fly-throughs and interactive presentations appear attractive to the end user and may be preferred to incomplete models with remaining scan holes. Unrealistic expectations with respect to 3D model quality, resolution, size and texturing have been raised through internet publications of seemingly perfect models, thus making it difficult for projects to meet user expectations. Because of this, unintended pressure can be placed on the producer of such data to sacrifice authenticity and accuracy in favour of visual appeal by, for example, filling holes, creating unrealistic lighting effects and hiding low resolution by texture. This is contrary to the ethos of Heritage Documentation, which requires objective, accurate and unmodified data. It would therefore seem imperative to develop specifications which clearly define strengths and limitations of a generated model which can be conveyed to the user. Without such descriptors, the scientific, engineering, and cultural worth of these models may be reduced. However, technical specifications are only then meaningful, if the users are aware of the endemic limitations of this particular methodology. Furthermore, it is important to explore new ways of using laser scanned models, besides for visual presentation, to justify the effort and cost spent on their creation.

Acknowledgement

This project has been made possible with the generous support of the Mellon Foundation, the assistance of Rahim Rajan from Aluka/JStor, and the invaluable advice of the team of experts on African history – Dr. George Abungu, Prof. Susan Macintosh and Prof. Martin Hall. The authors thank Leica Geosystems and Optron Geomatics, who have generously provided additional laser scanning equipment for field campaigns. Our gratitude goes to the Visual Computing Lab of the ISTI-CNR, Pisa, Italy and to InnovMetric Software Inc, Canada.

References:

[1] Besl, P.J. & McKay, N.D.,1992. 'A method for registration of 3D shapes' IEEE Trans. Pattern. Anal. and Mach. Intell, pp.239-256.

[2] Gruen A., Devrim A.* ,2005. 'Least squares 3D surface and curve matching', *ISPRS Journal of Photogrammetry and Remote Sensing 59 (3),* pp.151-174.

[3] Amenta, N., Choi, S. & Kolluri, R. K, 2001.'The Power Crust', *Proceedings of the sixth ACM symposium on Solid modeling and applications, ACM Press,* pp.249-260.

[4] Bernardini, F., Mittleman, J., Rushmeier, H., Silva, C. & Taubin, G., 1999. 'The Ball-Pivoting Algorithm for Surface Reconstruction', *IEEE Transactions on Visualization and Computer Graphics, IEEE Educational Activities Department, 5,* pp.349-359.

[5] Curless, B. & Levoy, M.,1996. ‘A Volumetric Method for Building Complex Models from Range Images’, *SIGGRAPH 96 Conference Proceedings ACM SIGGRAPH, Addison Wesley,August 1996,* pp.303-312.

[6] Hoppe, H., DeRose, T., Duchampt, T., McDonald, J. & Stuetzle, W., 1992. ‘Surface Reconstruction from Unorganized Points’, *ACM SIGGRAPH*, pp.71-78.

[7] Kazhdan, M., Bolitho, M. & Hoppe, H., 2006. ‘Poisson surface reconstruction’, *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing, Eurographics Association*, pp.61-70.

[8] Fiorin, V., Cignoni, P. & Scopigno, R., Gobbetti, E. *(Ed.)* 2007. ‘Out-of-core MLS reconstruction’, *The Ninth IASTED International Conference on Computer Graphics and Imaging - CGIM 2007, ACTA Press*, pp.27-34.

[9] Bolitho, M., Kazhdan, M., Burns, R. & Hoppe, H., 2007. ‘Multilevel streaming for out-of-core surface reconstruction’, *SGP '07: Proceedings of the fifth Eurographics symposium on Geometry processing, Eurographics Association*, pp.69-78.

[10] Callieri, M., Cignoni, P., Ganovelli, F., Montani, C., Pingi, P. & Scopigno, R., 2003. ‘VCLab's Tools for 3D range data processing’, *4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST2003) and First EUROGRAPHICS Workshop on Graphics and Cultural Heritage, Brighton (UK), pp 5-7.*

[11] Sharf, A., Alexa, M. & Cohen-Or, D., 2004, ‘Context-based surface completion’, *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers, ACM*, pp.878-887.

[12] Corsini, M., Dellepiane, M., Ponchio, F. & Scopigno, R. 2009. ‘Image-to-Geometry Registration: a Mutual Information Method exploiting Illumination-related Geometric Properties’, *Computer Graphics Forum*, 28, pp.1755-1764.

[13] Gal, R., Wexler, Y., Ofek, E., Hoppe, H. & Cohen-Or, D., 2010. ‘Seamless Montage for Texturing Models’, *Computer Graphics Forum, Blackwell Publishing*, 29, pp.479-486.

[14] Snavely, N., Seitz, S.M. & Szeliski, R., 2006. ‘Photo tourism: exploring photo collections in 3D’, *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers, ACM*, pp.835-846.

[15] Ho, H.T. & Gibbins, D., 2008. ‘Multi-scale Feature Extraction for 3D Models Using Local Surface Curvature’, *DICTA '08: Proceedings of the 2008 Digital Image Computing: Techniques and Applications*, pp.16-23.

[16] Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F. & Ranzuglia, G.,2008. ‘MeshLab: an Open-Source Mesh Processing Tool’, Sixth Eurographics Italian *Chapter Conference*, pp.129-136.