Digital Anastylosis of the Octagon in Ephesos

BARBARA THUSWALDNER, SIMON FLÖRY, ROBERT KALASEK, MICHAEL HOFER
Vienna University of Technology
and
QI-XING HUANG
Stanford University
and
HILKE THÜR
Austrian Academy of Sciences

Anastylosis is the archaeological and architectural reconstruction of a ruined monument at the historic site after careful study of the remaining original elements. We present research results concerning 3D technologies that are used in the digital anastylosis of cultural heritage monuments. Based on current state-of-the-art research we present 3D data collection, digital artefact reconstruction, and digital reassembly of existing fragments, illustrated by means of the Octagon monument in Ephesos, Turkey. Our focus is on methods that belong to geometry processing.

Categories and Subject Descriptors: I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling, Geometric algorithms, languages, and systems; Curve, surface, solid, and object representations.

General Terms: Algorithms, Measurement

Additional Key Words and Phrases: Anastylosis, 3D data acquisition, global and local registration, digital surface reconstruction, digital reassembly, visual computing.

Author's address: Barbara Thuswaldner, Institute of History and Theory of Architecture and Historic Building Survey, TU Vienna, Karlsplatz 13, A-1040 Vienna; email: barbara.thuswaldner+e251@tuwien.ac.at; Simon Flöry and Michael Hofer, Geometric Modeling and Industrial Geometry, TU Vienna, Wiedner Haupstr. 8-10/104, A-1040 Vienna; email: floery@geometrie.tuwien.ac.at; email: hofer@geometrie.tuwien.ac.at; Robert Kalasek, Department of Spatial Development, Infrastructure, and Environmental Planning, TU Vienna, Operngasse 11, A-1040 Wien, Austria; email: kalasek@srf.tuwien.ac.at; Qi-Xing Huang, Computer Graphics Laboratory, Stanford University, 353 Serra Mall, Stanford, CA 94305; email: huangqx@stanford.edu; Hilke Thür, Institute for Studies of Ancient Culture, Austrian Academy of Sciences, Bäckerstr. 13, A-1010 Vienna, email: hilke.thuer@oeaw.ac.at;

The authors acknowledge support by the Austrian Science Fund (FWF) under grant P18865, by the Austrian Archaeological Institute, and the Austrian Academy of Sciences. Our special thanks go to Friedrich Krinzinger and Marina Döring for their ongoing support in this project.

The historic photos of 1905 shown in the article are courtesy of the Austrian Archaeological Institute (ÖAI).

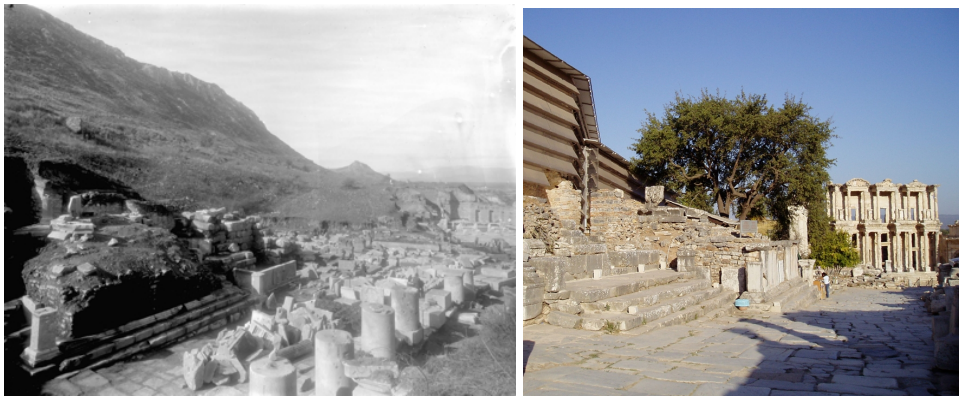

Fig. 1. Photos of the monument site in Ephesos, Turkey (left) 1905 and (right) 2005. In the left hand side of each photo we see the remainders of the Octagon, in the right hand side we see the Celsus library after excavation and the anastylosis result.

1. INTRODUCTION

The digital reconstruction of archaeological findings has attracted many researchers to work on this challenging topic. Several ambitious interdisciplinary projects have been undertaken to push the frontiers of research on hardware, software, and human computer interaction to develop tools that digitally support the traditional work of archaeologists and their co-workers in related fields. Nowadays a typical excavation site is interdisciplinary with researchers ranging from archaeologists and anthropologists to geologists and geodesists. Research in computer science and especially geometry processing provides important modern tools to support the traditional work in all of the above fields **(“smart tools for experts”)**.

In the present article we present research results concerning technologies that support the *anastylosis* of cultural heritage monuments. Anastylosis is the process of reconstructing a ruined monument at the historic site using the available original fragments and complementary “new” fragments. Besides the intrinsic research challenge a major motivation to re-erect historic buildings is tourism. The famous historic city Ephesos in Turkey—whose origins date back over 5000 years and which is being excavated for more than 100 years already—is visited by more than 2 million tourists every year. Of course, the attractiveness of a historic site raises significantly with the amount of architecture that has been reconstructed and makes the visit a more enjoyable experience for visitors from all over the world. For example, in Ephesos the Celsus library was built around 100 AD, re-erected from 1970 to 1978, and soon became an awe-inspiring landmark for every visitor (Fig. 1). “Best-viewed in-situ with your own eyes” still outperforms (and possibly always will) even the most advanced visualization technologies that use the latest hard- and software to provide a purely virtual reality experience. Thus there is a need to rebuild architecture that was ruined by influences ranging from earth quakes to simple decay over time.

As a matter of course, *anastylosis* is highly controversial. How can we faithfully re-erect a ruined historic monument if there are no building plans or drawings that

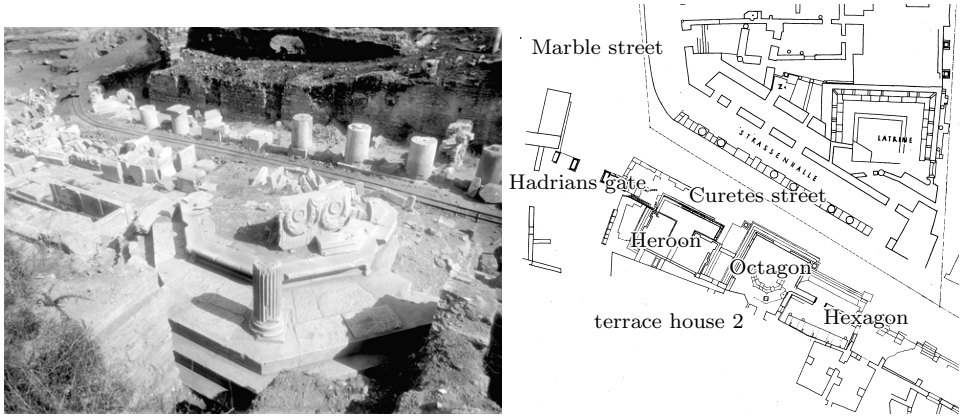

Fig. 2. (Left) Photo of the Octagon (1905). (Right) Map of the monument site in Ephesos (courtesy of [Miltner 1959]).

could guide the reconstruction? The best one can ask for is that the building under consideration has been mentioned and hopefully described in some historic text so that we know about its existence and what it roughly looked like. For a faithful *anastylosis* one has to carefully study the remaining original fragments and try to solve the posed "3D-puzzle" of how to put together the remaining fragments in the correct way. Compared to a regular jigsaw puzzle, where we usually have all the parts and just need to put them together in a correct way, the challenge posed to reassemble a monument from a small number of remaining deteriorated fragments is a good deal bigger. The needed algorithms and tools to digitally solve the posed problem in **combination with the archeologists' expertise** are the topic of this paper. Today, a fully automatic solution seems possible only if the available data is almost complete, which is hardly ever the case. We illustrate the current state-of-the-art by means of the *anastylosis* of the *Octagon* monument (Figs. 1 and 2) in Ephesos, Turkey.

Using 3D digital technologies for the whole *anastylosis* process is not only helping the process itself, but is also beneficial for virtual or augmented reality representations that can be generated as a side product. While until a few years ago virtual cultural heritage was often simply the result of skilled geometric modeling and computer graphics, it is now becoming a more faithful result of *visual computing* which combines computer vision, computer graphics, and human computer interaction technologies. Computer vision provides the basic technologies to acquire and reconstruct three-dimensional digital models, and computer graphics has the tools to visualize these enormous amounts of data. Human computer interaction is necessary in several steps of the *anastylosis* process and advanced interfaces simplify and improve the workflow considerably.

1.1 The Octagon in Ephesos

At the beginning of the 20^{th} century three monuments were excavated at the western end of the Curetes Street in Ephesos under the guidance of R. Heberdey [Heberdey 1905]. These monuments are lined up from west to east on the southern side of the

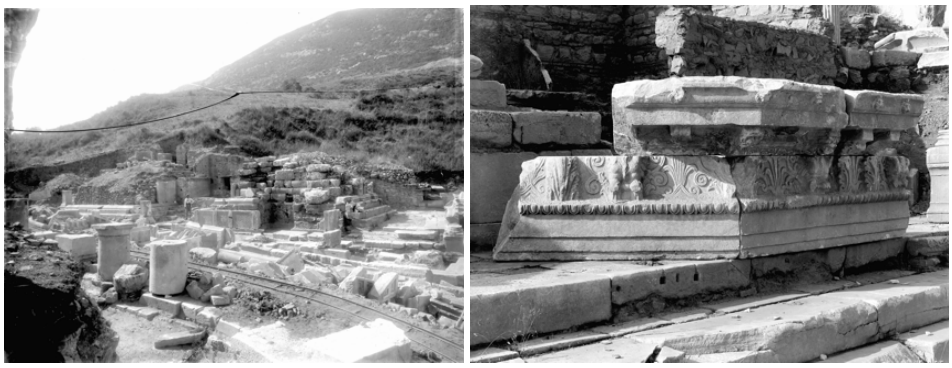

Fig. 3. Photos from the year 1905 showing components of the Octagon. The photo at right shows manually reassembled parts. (Photos are courtesy of the Austrian Archaeological Institute).

street (see the site plan in Fig. 2 (right)). In the first systematic description by W. Wilberg, the monuments already received the names they are referred to commonly today [Alzinger 1974]. The most western one is called *Hadrians gate*, the one in the middle *Heroon* and the eastern one *Octagon*. The object of our interest is the *Octagon*. An early assessment of the excavations showed that many components of the building still exist (Fig. 3); therefore a first reconstruction could be established (Fig. 4). Essentially, this reconstruction is still valid until today. Despite references to the *Octagon* in numerous scientific treatments and publications since its discovery [Heberdey 1905; Alzinger 1974; Oberleitner et al. 1978; Thür 1990; Berns 2003] a thorough reconstruction including all remaining fragments located on their original position does not seem to exist until now.

In the present paper we outline such a reconstruction strongly relying on the aid of modern 3D scanners and extensive computer calculations. Before we give the details we start with a description of the *Octagon*, which serves as a paradigm for the application of these new techniques for the reconstruction of ancient architecture.

The *Octagon*, which was a mausoleum, was built on a quadratic basal surface. In total, it is 13 meters high. Its front elevation is subdivided into three parts as follows. Beyond a substructure consisting of three steps there is a base building measuring nine times nine meters. The interior of this structure contained the sarcophagus. Beyond the base profile, the outside of the base building is ornamented with slightly curved rectangular surfaces whose upper part is finished by a moulding. Behind the marble face, the entire base building is made of *opus caementicium*, which can be regarded as the ancient predecessor of concrete. The grave chamber containing the sarcophagus is covered by a barrel vault. It can be reached via the rear part of the base building through a small corridor. The corridor as well as the grave chamber is built up by marble blocks.

On the top of the base building we find the Octagonal main structure that is responsible for the name *Octagon*. Beyond three Octagonal steps a massive *cella*, surrounded by an eight sided *peristasis* (i.e., a ring of columns) is built; the columns have Attic styled bases and Corinthian capitals. The base of the wall of the *cella* is surrounded by a bench; its upper part contains a garland frieze consisting of representations of fruits as well as bulls and flowers (the latter being known as

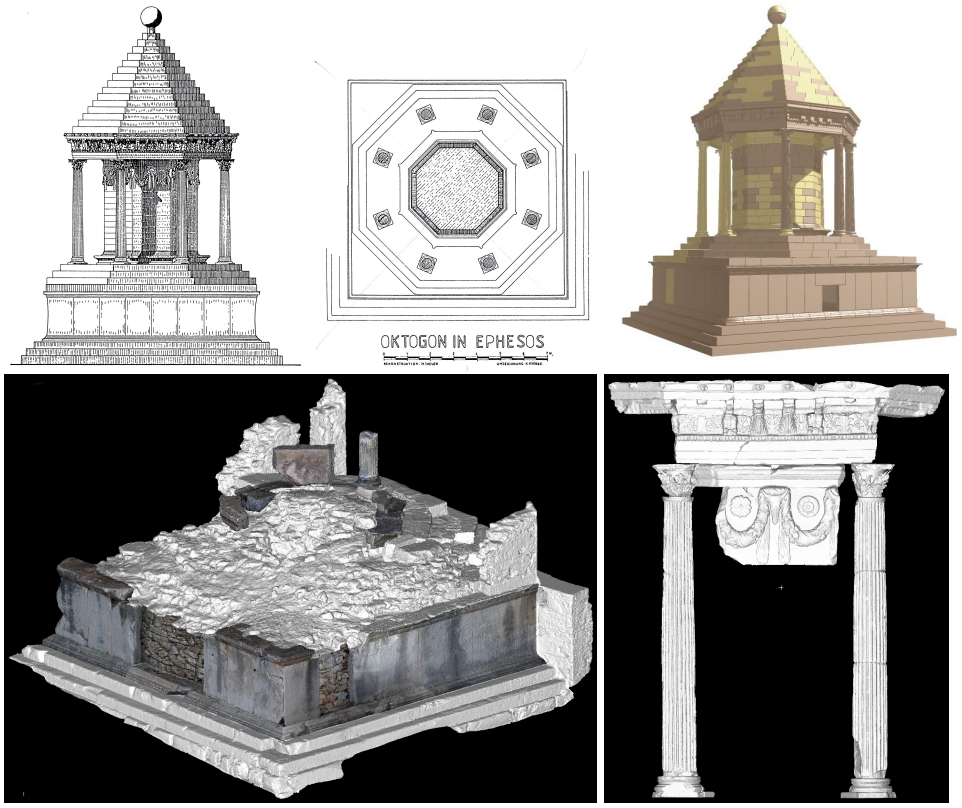

Fig. 4. (Top left) Hand drawing of the Octagon by W. Wilberg. (Top right) Modern visualization as a 3D model generated using a CAD system without 3D measurements. (Bottom left) The remaining base of the Octagon building as a partially textured 3D model generated from 3D measurement data. (Bottom right) Digital 3D model generated from 3D measurement data of that part of the Octagon which was re-erected in the Ephesos museum in Vienna. Note that the objects are scaled differently in the bottom two images.

bucrania; part of it can be seen in Fig. 5). The halls surrounded by the columns described above have a waffle-slab ceiling (Fig. 4) consisting of trapezium-shaped elements. The space beyond the capitals of the columns is subdivided into three parts, each of which contains mouldings and representations of griffins and leaves. The upper part of the *Octagon* consists of a steep Octagonal pyramidal roof. On the top of the roof a ball made of stone is located. Finally, from the point of view of constructional engineering it is important to mention that the whole building is made of hewed marble blocks that are built without using any kind of mortar (the *cella* being the only exception). This required a very exact hewing of the stone blocks in order to keep the slits between the blocks small. Since dowels and clamps have often been used in order to fix one stone onto the other, the corresponding holes in the stones are an important hint to the original position of the stone in the building.

In the following section we illustrate why we chose our documentation methods

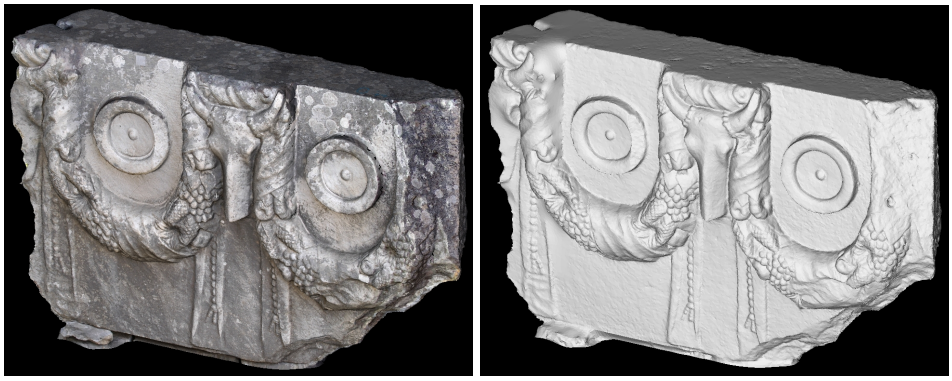

Fig. 5. For virtual reality applications a textured 3D model as shown on the left might be sufficient since geometric mistakes are nicely hidden. However, for digital reassembly purposes the flawed geometric model generated by commercial software from 3D measurement data is definitely not good enough. Our digital 3D models are of a much higher accuracy and quality.

as well as the particular building in order to give a paradigm for the reconstruction of ancient architecture. Since the modern protective roof for the Terrace House 2 with its transparent façade considerably influences the phenotype of the Curetes Street, there is a strong desire to rebuild some of the ancient monuments in Ephesos in a way that the ancient sense of the site remains established. It is our intention to rebuild ancient constructions by using as much of the original building substance as possible. This way of reconstruction is called *anastylosis*. In using *anastylosis*, as described in the Charta of Venice [Gazzola et al. 1964], the usage of modern building elements in order to achieve the stabilization of the original parts has to be kept to a minimum. Moreover, one desires to put the available building parts at their original locations. From a certain amount of building parts upwards it is quite hopeless to gain such a construction "by hand", though. However, the computational tools and methods we are going to use provide us with a new way for performing such complex constructions in an optimized way. The *Octagon* is well suited as a paradigm for the computational reconstruction process for several reasons. First of all, many components of the building still exist. Secondly, many of the blocks are in quite a good condition and, moreover, the *Octagon* is a fairly small building making it easy to survey.

The basis for the reconstruction of the Octagon at its original location—regardless of the way of reconstruction; a partial one or a complete anastylosis—is formed by the research on the development of the building up to its current state. This research is mainly done by archaeologists and architects, being supported by researchers in geometry processing in recent time. The research on the development of a building includes in particular a thorough inspection of the remaining building substance as well as a detailed categorization of all the remaining parts. All the information is then used in order to describe the development of the building starting from its original construction up to its current state [Gromann 1993].

In the summer of 2005 systematic research on the *Octagon* started again. One of the first aims was to obtain information on the suitability of this building to be

Fig. 6. Finding matching fragments and reassembling them by hand is a tedious job.

a candidate for computer aided *anastylosis*. Under the usage of high end geodetic surveying methods, the available building parts have been recorded computationally with help of a 3D scanning device. These data records form the basis for establishing computer models of each building part. At the same time, a statistical validation of the building substance has been performed. The architectural pieces were allocated to the respective parts of the building. Furthermore, originally matching pieces that could be matched by hand have been fitted together. However, some of the fragments (Fig. 4) are located in the Ephesos museum in Vienna, Austria, and only modern 3D technologies allow as part of our planned work to digitally put together fragments that are physically located in completely different places (Turkey–Austria). Performing these preliminary tasks allowed to estimate the complexity of the *anastylosis* process and lead to the decision to undertake this challenging project.

1.2 Outline of the paper

The paper follows the main steps of an anastylosis process as outlined in Fig. 7 and focuses on steps 3, 4, and 5. Besides a fairly detailed description of the digital anastylosis of the Octagon in Ephesos, the main technical contributions of the present paper are:

—new methods for fully-automatic out-of-core global registration and local registration of large scan data sets,

—a geometric optimization algorithm for penetration free registration of fragments that meets certain geometric constraints.

After a review of the existing literature, the remainder of the paper is organized as follows. Section 3 deals with the acquisition of large amounts of 3D data under outdoor conditions. In Section 4 we present a new algorithm for the automatic generation of highly accurate meshes from multi-view 3D scan data. Section 5 deals with the reassembly of the remaining fragments and presents a new algorithm that allows the matching of non-fractured parts. In Section 6 we present our results of the digital anastylosis of the Octagon in Ephesos.

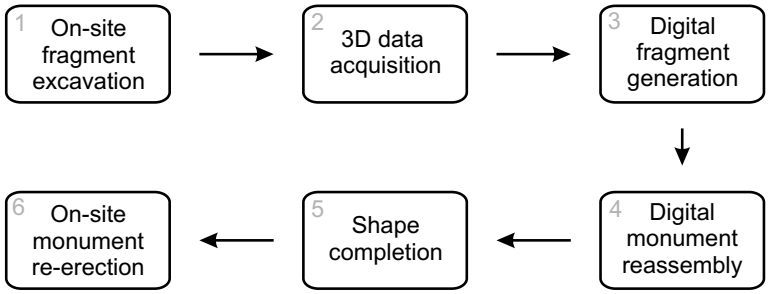

Fig. 7. High-level overview of the anastylosis steps.

2. RELATED WORK

There is a huge amount of related literature for several of the sub-steps of the outlined procedure (Fig. 7) and we cite only the most relevant ones for our work. Other steps have received far less attention in the scientific community and are interesting topics of current research in geometry processing and related fields. There are a few papers in the visual computing literature that deal with anastylosis per se. Other related work concerns the individual steps of the whole process.

Anastylosis. The brute-force destruction of the great Buddha statues in the Bamiyan valley in Afghanistan in March 2001 raised the public awareness for the preservation and restoration of important ancient monuments. According to [Petzet 2002] an anastylosis is the only appropriate solution for this unique place. The feasibility of such an undertaking was studied by [Zou and Unold 2003]. While in building research anastylosis is an "old topic" [Hueber 2002], recently researchers in computer graphics and computer vision contributed to the *virtual* anastylosis of ancient monuments [Zalesny et al. 2003; Mueller et al. 2004; Gool et al. 2004]. These contributions mainly deal with the creation of realistic textures for virtual 3D models.

Data acquisition and 3D scanning of cultural heritage. Several research groups turned their attention to 3D scanning techniques for cultural heritage conservation. There are various publications promoting the use of 3D scanners for acquisition of 3D object information, textures and colors [Böhler et al. 2001; Guarnieri et al. 2004; Salemi et al. 2005; Valzano et al. 2005]. One of the most prominent papers that deals with the 3D scanning of cultural heritage is [Levoy et al. 2000] describing Stanford's Digital Michelangelo project. A good overview of the 3D model acquisition pipeline for computer graphics applications is given in the article [Bernardini and Rushmeier 2002], and for post-processing of scanned 3D surface data we refer to [Weyrich et al. 2004]. The task of 3D data acquisition of buildings and monuments is of interest in the civil engineering community as well [Abmayr et al. 2005; Bahmutov et al. 2006]. A large amount of work can in the area of 3D-scanning cultural heritage can be found in the publications of Roberto Scopignio (see e.g. [Spinelli et al. 2006; Scopigno 2007] and the references therein) and coworkers and also in the papers by Ioannis Stamos and coworkers (see e.g. [Allen et al. 2003; Stamos et al. 2007] and the references therein).

Digital fragment generation. To process the huge amounts of acquired 3D

data takes an tremendous effort. It took Stanford University 6 years of paid and volunteer student labor to generate 3D models of the 1186 fragments of the Forma Urbis Romae [Koller and Levoy 2006]. Some of the data from the Digital Michelangelo Project [Levoy et al. 2000] still needs to be processed and **requires improved** algorithms to resolve difficult alignment and mesh completion problems. Due to technical constraints, we are usually not able to acquire 3D information of the whole fragment at once, but we produce numerous partial scans instead. It is the task of *global* registration to determine the way these single scans need to be arranged to represent the original object roughly. Most global registration work is based on features, that are matched against each other [Johnson and Hebert 1999; Li and Guskov 2005; Gal and Cohen-Or 2006; Gelfand et al. 2005; Huang et al. 2006]. Other practical approaches are the following papers (and references therein). [Pingi et al. 2005] take the scan order into consideration for scan alignment. [Chao and Stamos 2005] perform semi-automatic range image registration based on extraced features. The rough alignment achieved with global registration is further improved in a *local* registration step [Besl and McKay 1992; Chen and Medioni 1991; Rusinkiewicz and Levoy 2001] where multi-view registration avoids accumulation errors [Neugebauer 1997]. Additional constraints to the registration play an important role to achieve a proper alignment for reconstruction tasks [Huang et al. 2006].

Reassembly of broken objects. Motivated by the challenge of reassembling broken archaeological artifacts, several approaches have been developed for specialized reconstruction problems. These include the matching of planar 2D fragments (e.g. fractured tiles) [Hori et al. 1999; Kong and Kimia 2001; da Gama Leitão and Stolfi 2002; Goldberg et al. 2004] and objects that are roughly surfaces of revolution (sherds of pottery), see [Cooper et al. 2001; Cooper et al. 2002; Willis and Cooper 2004] and the references therein. A solution for geometric reconstruction of 3D solids was first developed by [Papaioannou et al. 2000; 2001; 2002; Papaioannou and Karabassi 2003]. The underlying assumption of this method is that the fracture faces are nearly planar and they match each other completely. The algorithm of [Huang et al. 2006] is able to automatically reassemble fractured objects by geometric matching of the fracture surfaces. A different feature-based method employing spectral techniques is discussed in the paper [Parikh et al. 2007]. Stanfords Digital Forma Urbis Romae project [Koller et al. 2006; Koller and Levoy 2006] deals with heavily eroded fragments, whose fracture surfaces sometimes do not even touch each other. In this case, instead of using the geometry of the fracture faces, reconstruction is done by matching manually annotated incisions on the fragments' top surfaces. The latter marking has been performed manually on 2D digital photos of the fragments top surfaces showing traces of the map of ancient Rome. **Recently, [Brown et al. 2008] presented a comprehensive processing pipeline for digitzing and matching ancient Theran fresco fragments in a hybrid 2D/3D approach.** Besides the literature on automatic and semi-automatic reassembly there are also publications about the non-automatic reassembly of cultural heritage objects, and approach that is fairly common, see e.g. [Laugerotte and Warze 2004].

Shape completion. It is very likely that not all fragments of a ruined ancient

Fig. 8. (Left) The scanning tent that gave some protection from the elements. (Right) The Riegl and the PT-M 1024 scanner.

monument are available. This poses an interesting new challenge in the area of 3D shape completion, namely to automatically retrieve complete new fragments that fill gaps in the reassembly of existing parts and thus complete the shape of the overall building. Shape completion is an important topic of active research in fields ranging from image and video processing via computer vision to geometry processing. Previous work on shape completion of geometric models focused on the completion of single 3D models and can roughly be divided into three categories: *surface based*, *volumetric*, and *example based* methods. A literature survey of the first two approaches is given in the article [Davis et al. 2002] which contains references up to the year 2002. Ideas from image inpainting are used in the volumetric approach by [Verdera et al. 2003]. Filling holes in point set surfaces is also discussed in [Weyrich et al. 2004]. Recently, example based methods [Sharf et al. 2004; Pauly et al. 2005; Bendels et al. 2005; Park et al. 2006] emerged in geometry processing. A geometric method to fill holes in digital elevation data via a constrained smoothing process was presented by [Hofer et al. 2006]. In the shape completion step one could use the discrete symmetries present in the Octagon. For that purpose recent advances in similarity detection, symmetrization, and symmetric shape completion [Thrun and Wegbreit 2005; Mitra et al. 2006; Venkatesh and Cheung 2006; Podolak et al. 2006; Shilane and Funkhouser 2006; Mitra et al. 2007] would be a good starting point.

3. DATA ACQUISITION

The currently used techniques for measuring fragments (also referred to as 'stones' in this article) and buildings by archaeologists and architects at a typical excavation site comprise traditional methods including the "measurement by hand" as well as photogrammetric and laser scanning methods. The building blocks of the Octagon still standing at their original position (mainly including the pedestal part, Fig.1) as well as all the other existing parts of the building were scanned in the summer of 2005 by the company ArcTron 3D GmbH. Newly found fragments were scanned on-site in Turkey by a team of TU Vienna in November 2006 and October 2007. Depending on the location, the material, as well as the condition of the building parts under consideration, three different 3D-scanning systems were employed, (i) the time-of-flight Riegl LMS-Z 420i 3D scanner, and (ii) the structured-light triangulation scanner PT-M 1024 (Fig. 8).

The Riegl scanner allows to process objects at a distance of 2 up to 800 meters

and is well suited to scan large objects such as whole ensembles of buildings, landscape topography, or to generate reference scans of single architectural objects. In vertical direction the device is able to process objects within an angle of 80 degrees at most. Horizontally the whole surrounding (360 degrees) can be processed at once. The maximal resolution is as low as 0.002 degrees while the accuracy for the measurement of distances is approximately 7 millimeters. With these properties, the scanner is able to process fairly big scenes in a reasonable amount of time. Depending on the resolution, the scanning procedure for a 360 degree scan lies between two and 90 minutes. In our case the scanning periods hardly exceeded the duration of 10 minutes. As a particular feature, the scanner can be equipped with a high resolution digital camera which is attached to the scanner body. The camera produces a sequence of color pictures that overlap the whole scanning region. Combining the 3D information provided by the scanner with the color information provided by the calibrated digital camera, each measurement point is also assigned color information.

In order to produce more detailed scans of smaller objects the PT-M 1024 scanner was used in combination with the software package QT Sculptor. This device is able to process objects in a distance range of 0.5 to 3 meters. Its accuracy is about 0.1 millimeters. The scanner consists of an LCD projector and two high-speed cameras being mounted on an aluminium bar. Neither the cameras nor the projector contain any mechanical movable parts. This property increases the accuracy of the scans considerably and makes the system especially suited for measuring smaller parts of a building, especially parts with "fine structure" such as ornaments.

According to the various locations of the different parts of the building substance of the Octagon, the scanning team had to deal with various environmental conditions. Essentially, one can subdivide these conditions into three groups. The following paragraphs are intended to describe the particular problems and the way of solving them for these different situations:

The in situ parts of the building, i.e. the parts which are located on their original position, were scanned with the Riegl scanner. Posed difficulties include e.g. the accessibility. Today we find the remainders of the Octagon tightly wedged in between other monuments. Moreover, in the times of Late Antiquity, other constructions were built over the remainders of the Octagon. The Riegl scanner can be used in regions with poor light conditions which was especially important for the treatment of the grave chamber. On the other hand, the Riegl scanner can also be used under strong sunlight. As the region surrounding the Octagon is highly frequented by visitors of Ephesos, especially during summer time, the fact that the Riegl scanner allows to scan quite quickly was another advantage. All in all we acquired 3D data from 41 different viewpoints in order to get panoramic views of the entire situation. Since some of the in situ parts of the Octagon contain elaborated fine structures, the PT-M 1024 system was used to scan these parts in full detail.

Most of the fragments of the Octagon that are located somewhere at the excavation site in Ephesos—but not at their original position—are now stored in an open space behind the famous Library of Celsus (Fig. 8). In the summer of 2005 about 150 building blocks were scanned with the PT-M 1024 scanner to guarantee the utmost accuracy. Another set of recently found 15 additional fragments were

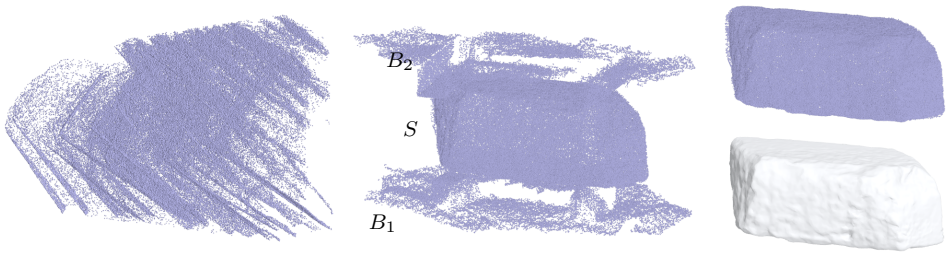

Fig. 9. The original scanning data (left) is aligned including background information (center), which is removed after registration (right, top). From this final point cloud, a mesh is generated (right, bottom). The typical scan data for a stone contains the stone S itself and two copies B_1, B_2 of the background (middle).

scanned in November 2006. In this context it is worthwhile to mention that already in antiquity parts of collapsed buildings were used in order to build new ones. This explains why some of the blocks of the Octagon were quite a distance away from its original location. For this reason, the quest for more such parts of the Octagon can hardly be regarded as completed.

Some essential building parts of the Octagon are located in the Ephesos-Museum in Vienna since 1978. These parts are collected in the museum in order to provide the visitor with a reconstruction of a small part of the original building and thus giving a feeling for the whole structure. This arrangement of the building blocks makes it difficult to reach each of them in a way that is optimal for scanning. This problem occurs especially for parts built in the rear as well as the upper part of the reconstruction. Moreover, by assembling the building parts some of the boundary faces are hidden. Thus one research challenge is the following. Develop algorithms that use mirrors attached to the walls surrounding the inaccessible parts to also scan areas that are directly inaccessible. Inaccessible parts are a particular problem as these surfaces would give valuable hints for the original neighbors of the blocks and therefore of their original position. After a general validation of the entire situation, the ensemble was processed in a similar way as the in situ parts in Ephesos. After some panoramic scans done with the Riegl scanner in order to digitalize the entire ensemble, PT-M 1024 was used in order to gain 3D data of details. In fact, contrary to the in situ ensemble in Ephesos, this ensemble was entirely covered by PT-M 1024 scans. So, the main purpose of the less accurate Riegl scan data is to provide a "map" for the detail-scans.

Compared to 3D data acquisition in a lab under controllable conditions, 3D data acquisition in the field—such as an archaeological excavation site—is a much more challenging job. Ideally, for archaeological findings one wants to get under outdoor conditions the same high quality resolution as e.g. [Levoy et al. 2000] achieved in an indoor environment. The typical scanning scenario for a stone from the Octagon monument was the following: scan everything possible from one side, then turn the stone over only once and scan the remaining parts. This can be seen after the global registration of the scan data before the background is removed. The point cloud consists of the object (foreground) and two copies of the background (Fig. 9).

4. DIGITAL FRAGMENT GENERATION

Once a fragment has been scanned, a digital model is created from the scanning data. As optical 3D scanners are not capable of capturing whole fragments at once, several partial scans (more than 7000 scans in total for the Octagon) are obtained that remain to be fitted together in order to retrieve the original object. Furthermore, the quality of the scanner data is low due to environmental influences during scanning. The digital fragment generation step works on the data to eliminate these factors and gives high quality point clouds and mesh models of the fragments in return. The problem of merging single scans of an object turns out to be a challenging task. It is known as *registration* and is a frequently encountered topic in geometry processing. For the construction of digital fragments, such a registration will be encapsulated by two steps to prepare and improve the data: a preprocessing stage and a final mesh generation stage.

4.1 Preprocessing the data

As archaeological fragments are typically scanned outdoors, the scanning data includes outliers and background points (not belonging to the object but to the ground, ...) which need to be eliminated. We use the approach introduced in [Weyrich et al. 2004] to remove salient outliers and points from the background. We keep background data points for the registration step as this additional information stabilizes and accelerates the data registration process. Moreover, we compute oriented normals of each data point according to [Mitra and Nguyen 2003] after removing outliers.

4.2 Registration of fragment scan data

It is a common approach to split the problem of merging scans into two subtasks. In a *global* registration step it is decided, which scans overlap and belong together. A *local* registration step then refines these rough initial alignments. We will return to the registration problem from a different point of view in Sec. 5 when we tackle the issue of reassembling the original building from its fragments: first we need to fit partial scans together (this section) and later on we are searching for matching building blocks. For this reason, we want to outline our registration steps in more general terms first before we consider digital fragment generation in detail.

It has been suggested (see e.g. [Pingi et al. 2005]), that one should keep track of the ordering of the scans while scanning the data. This usually simplifies the problem of finding overlapping matching scans. However, the challenge in our project was that we were given the scan data without such information and it was impossible to establish the scan order afterwards. Thus we needed to develop an algorithm that automatically performs the global and local registration. An additional advantage of computing all pairs of overlapping scans is that it improves the stability of the registration process.

For a given set of partial data, it is intuitive to first identify features on every scan. Then, for each pair of data (S, M), we obtain a rough alignment by matching these features and improve it with a local registration. By measuring the alignment quality of these intermediate results, we end up with a rating for every single pairwise match (S, M). Based on this information we derive a consistent set of

partial data sets we finally align simultaneously. It turns out, that just a single iteration as outlined here might fail to reconstruct the whole fragment (or building) but results in several merged subgroups instead. However, it is obvious to repeat these steps with those groups as input data until a final reconstruction is achieved.

4.2.1 *Global pairwise registration.* As stated above, global and local *pairwise* registration will only serve as auxiliary tool for computing consistent matches. For a final alignment, we employ a multi-view registration to achieve better global matching results and to avoid any error propagation (see Sec. 4.2.3 and 4.2.4). For a given partial scan S, we start by computing various feature descriptors in each element of the point cloud (except the boundaries). In [Huang and Pottmann 2005; Huang et al. 2006] new integral invariants prove both as reliable and powerful tools for robust matching. Considering the smooth case first, let $\mathbf{p}$ be a point on a surface Φ, which itself is the boundary of a domain $D \subseteq \mathbb{R}^3$. Furthermore, χ_D denotes the characteristic function of D ($\chi_D(\mathbf{x}) = 1$ for $\mathbf{x} \in D$ and $\chi_D(\mathbf{x}) = 0$ otherwise) and $B_r(\mathbf{p})$ a ball of radius r, centered at $\mathbf{p}$, with boundary sphere $S_r(\mathbf{p})$. Then, the *volume descriptor* $V^r(\mathbf{p})$ and *area descriptor* $A^r(\mathbf{p})$,

$$V^r(\mathbf{p}) = \frac{3}{4\pi r^3} \int_{B_r(\mathbf{p})} \chi_D d\mathbf{x}, \qquad A^r(\mathbf{p}) = \frac{1}{4\pi r^2} \int_{S_r(\mathbf{p})} \chi_D d\mathbf{x}, \tag{1}$$

describe the ratio between the volume of $B_r(\mathbf{p}) \cap D$ and $B_r(\mathbf{p})$ or between the area of $S_r(\mathbf{p}) \cap D$ and $S_r(\mathbf{p})$, respectively. Both, $V^r(\mathbf{p})$ and $A^r(\mathbf{p})$, relate to the mean curvature in $\mathbf{p}$; for a proof, further details and a discretization we refer to [Pottmann et al. 2005].

Depending on the size of the features (in relation to the scanning resolution) we compute either $V^r(\mathbf{p})$ or $A^r(\mathbf{p})$ at four different radii. For each radius value, we cluster the points of S according to their descriptor value and compute additional geometric properties such as the barycenter for each cluster (called *features* in the following as well). Considering two partial scans S and M, these clusters are used to establish a pairwise match between the two data sets. For the sake of brevity, we only sketch our method here, details of the algorithm are described in [Huang et al. 2006]. From a set of initial feature correspondences (clusters computed for the same radius with similar descriptor values), several pruning steps comprising geometric and topological tests remove false pairs of features. From the remaining correspondences, a set of consistent feature pairs is extracted by a robust statistical method called the *forward search* (see [Atkinson et al. 2004]). These final feature links are used to align the two scans roughly. Note that using patch features for matching is actually a very efficient method (for details see [Huang et al. 2006]).

4.2.2 *Local pairwise registration.* Local pairwise registration is the most important criterion in rating the quality of the obtained match. Given two scans S and M, the goal is to compute that rigid body motion α minimizing the squared distance between the *fixed* or *target* system S and the *moving* system M. For this purpose, we approximate the unknown α by its linear velocity vector field $\mathbf{v}(\mathbf{x}) = \bar{\mathbf{c}} + \mathbf{c} \times \mathbf{x}$. Moreover, we compute for each point $\mathbf{x}_i \in M$ the closest point $\mathbf{y}_i \in S$, the distance $d_i = \|\mathbf{x}_i - \mathbf{y}_i\|$ and the normed connecting vector $\mathbf{r}_i = (\mathbf{x}_i - \mathbf{y}_i)/\|\mathbf{x}_i - \mathbf{y}_i\|$. Then, we obtain α as in [Pottmann et al. 2004] as that uniquely defined helical motion,

for which its velocity vector field minimizes

$$f_{pw}(\bar{\mathbf{c}}, \mathbf{c}) = \sum_{\mathbf{x}_i \in M} (d_i + \mathbf{v}^T(\mathbf{x}_i) \cdot \mathbf{r}_i)^2. \tag{2}$$

We repeat this optimization 5 to 10 times to achieve a final alignment. Eventually, we use a combination of the average distance between S and $\alpha(M)$, $\varepsilon_{pw}(S, M) = (\sum_{\mathbf{x}_i \in \alpha(M)} \|\mathbf{x}_i - \mathbf{y}_i\|^2/|M|)^{1/2}$, and the deviation in the normals of corresponding closest points (see [Huber and Hebert 2003]) as a quality measure for rating the pair (S, M)

4.2.3 *Global multi-view registration.* Once all pairs of scans have been matched and rated, we are able to proceed with putting the whole fragment together. Assuming really bad matches have already been discarded, we have several hints which scans will match. By converting this information into a graph representation (where each node represents a scanned data set and each edge a possible match), we employ a heuristic method extracting constistent groups of matching scans. Again, we refer to [Huang et al. 2006] instead of describing this step in detail. However, we want to point out some issues arising in our specific application: the size of archaeological data sets is typically huge and thus all the data cannot be loaded into memory at the same time. For the 150 fragments of the Octagon, a single fragment can comprise up to 300 partial scans, counting typically more than hundreds of thousands of points each. Moreover, every single scan overlaps with 6 to 30 other scans, depending on the scanning process. For the feature cluster computation, only one scan needs to be in memory. As the number of features for a data set is relatively small (typically between 500 and 1000), features of all scans can be kept in memory. For the pairwise local registration, a scan is loaded and compared to all other scans, one by one. Thus, not more than two scans need be loaded into main memory at the same time.

4.2.4 *Local multi-view registration.* This doesn't hold for local multi-view registration, which is used to verify and refine the results of global mutli-view registration. As successive pairwise local registration for a total reconstruction would lead to significant accumulation errors, we generalize Equ. (2) to multiple systems $\{S_i\}$. For two systems S_j and S_k, we locally regard without loss of generality S_j as fixed and S_k as moving system. For a point $\mathbf{x}_i \in S_k$, we define $\mathbf{y}_i \in S_j$, distance d_i and residue $\mathbf{r}_i$ as before and approximate the motion of S_k towards S_j by its velocity vector field $\mathbf{v}_{jk}(\mathbf{x}) = \mathbf{v}_{j0}(\mathbf{x}) - \mathbf{v}_{k0}(\mathbf{x})$ (where $\mathbf{v}_{i0}(\mathbf{x}) = \bar{\mathbf{c}}_i + \mathbf{c}_i \times \mathbf{x}$ as before towards a globally fixed system S_0). Then, we could obtain the parameters $\mathbf{c} = (\mathbf{c}_1, \bar{\mathbf{c}}_1, \ldots, \mathbf{c}_n, \bar{\mathbf{c}}_n)$ describing the scans' motions by minimizing an approximation of the squared distance between adjacent systems S_j and S_k

$$f_{scan}(\mathbf{c}) = \sum_{\mathbf{x}_i \in S} (d_i + \mathbf{v}_{jk}^T(\mathbf{x}_i) \cdot \mathbf{r}_i)^2, \tag{3}$$

where $S = \cup_i \{S_i : S_i \text{ is moving system}\}$. However, this would require to load all scans into main memory at the same time, which is impractical in our case. Instead, we store the closest point $\mathbf{y}_i^{jk} \in S_j$, distance d_i^{jk} and residue $\mathbf{r}_i^{jk}$ for every sample $\mathbf{x}_i^{jk} \in S_k$ in the final local pairwise registration step (Sec. 4.2.2) of S_j and S_k and

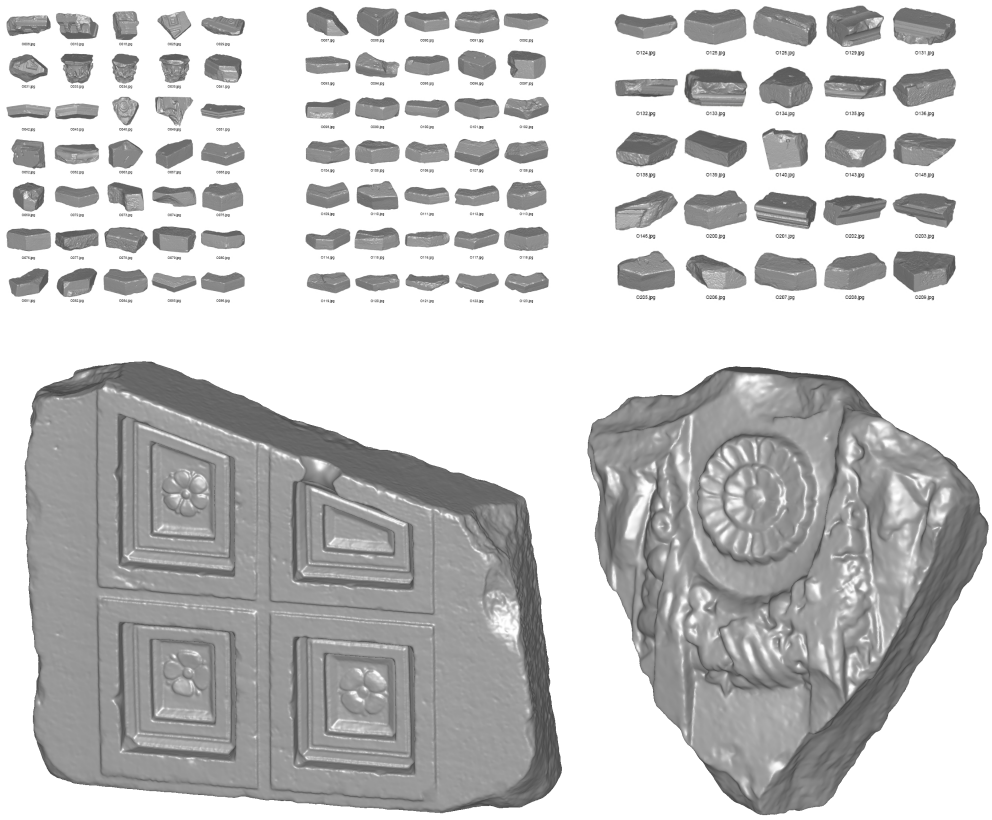

Fig. 10. Digital 3D models of selected fragments of the Octagon. (Top row) Thumbnails of 95 different blocks. (Bottom row) Close ups of two blocks.

modify Equ. (3) to

$$f_{mv}(\mathbf{c}) = \sum_{\mathbf{x}_i^{jk} \in S} (d_i^{jk} + \mathbf{v}_{jk}^T(\mathbf{x}_i^{jk}) \cdot \mathbf{r}_i^{jk})^2, \tag{4}$$

in order to avoid the closest point computation which would require to keep all the systems in memory. As the rough alignment of the pairwise global matching step are usually already very good, this simplification is justified. However, one disadvantage of optimizing Equ. (4) is that a biased pairwise match propagates its error to the final result. We use an regression approach to handle such cases: after minimizing Equ. (4), we compute the current residual $\varepsilon_{mv}(S_j, S_k) = \sum_{\mathbf{x}_i^{jk} \in S_k} (d_i^{jk} + \mathbf{v}_{jk}^T(\mathbf{x}_i^{jk}) \cdot \mathbf{r}_i^{jk})^2$ of each pairwise match and discard those with residuals larger than a user defined threshold. Similar to pairwise registration, we iterate through this process 10 to 20 times until a final alignment is obtained.

For the Octagon data sets we report the following running times of our code. Computing features and normals takes 1-2 mins per scan, surface matching needs 0.1-1s per scan pair, and multiple global and local registration takes between 1-10

Fig. 11. Stones of the Octagon that were matched manually by means of their high-level features.

mins depending on the complexity of the current stone.

4.3 Mesh generation

For computational as well as for visualization reasons it is advantageous to generate the final point cloud representation of the reconstructed fragment into a triangular mesh. In practice we find that recomputing and orienting normal directions prior to mesh generation improves the final quality considerably, as irregular sampling biases the original normal information. For some of the data sets the Poisson surface reconstruction technique of [Kazhdan et al. 2006] worked fine to compute watertight solid mesh models from a set of oriented input surfels. However, many of the data sets could not be meshed satisfactorily with this technique. Thus we had to develop a whole new strategy that has been published in the paper [Huang et al. 2007]. This paper describes a Bayesian surface reconstruction technique via iterative scan alignment to an optimized prototype. We used the algorithm described therein to create high-quality meshes for all stones of the Octagon. In Fig. 10 we show selected images of the digital 3D models we created.

5. DIGITAL REASSEMBLY

Digital reassembly is another aspect 3D digital technologies are able to support researchers. At this point of our processing pipeline, the single fragments have already been digitized and reconstructed successfully and we have point cloud and mesh models of all the parts. We now want to find out which fragments fit together and align those properly. As mentioned in Sec. 4, our reassembly algorithm will follow the same general ideas as the digital reconstruction of the fragments: at first, features on the fragments are extracted that are matched globally and finally registered locally.

The task of reassembling a building from its stones differs significantly from that of reassembling a broken solid. Fracture surfaces of broken solids contain rich intrinsic geometric features that are sole sufficient to reassemble the broken solid, as demonstrated in [Huang et al. 2006]. On the contrary, the entities from which one has to reconstruct an ancient building are the original building blocks. These building blocks are fragments in a much looser sense as opposing faces of neighboring stones do not contain the rich geometric features of fracture surfaces that would be

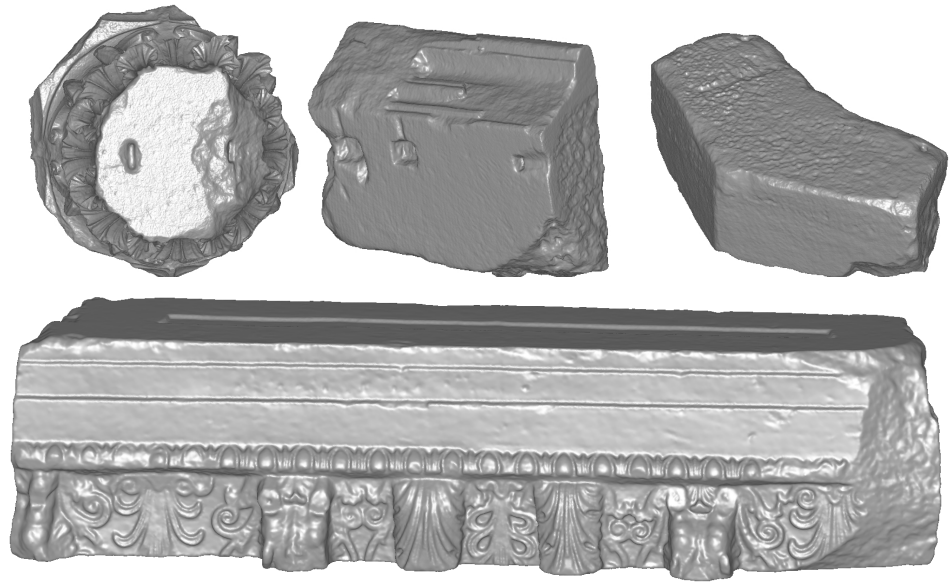

Fig. 12. Clamp holes, different surface roughness, lines and planes that can be identified and need to extend over adjacent stones.

advantageous for matching them. For this reason, we require external, high level features such as ornaments, edges or clamping holes for a successful reconstruction (Figs. 11,12). Thus the quality of the 3D digital models has to be high enough so that the necessary features can be extracted in an accurate and robust way.

5.1 Feature extraction

While a person immediately recognizes distinct features such as holes or other traces of human processing on the fragments, we need to define an abstract scheme that lets the computer identify and represent such features in a convenient way. This part of our work is steady work in progress and we want to propose a possible framework in the following and document subtasks already realized.

As stated above, matching faces carry too few information for a digital reassembly of the *Octagon*. Instead, we require features beyond these matching faces for an effective reconstruction. The choice of features described in the following is closely related to the *Octagon* anastylosis and might be different for other projects. However, we aim at a general concept that can be applied in other projects as well. The Octagon stone faces show different surface roughnesses that let us classify the faces roughly. Front faces (by means of being part of the facade) are smooth and may show ornaments. Any other faces exhibit larger surface roughness, and may carry traces of human processing as well, e.g. edges framing trimming areas in order to make the stones better fit together or holes of clips stabilizing the construction work in ancient times.

Our feature extraction algorithm closely follows these observations (see Fig. 13). At first we break down the fragments in several planar faces. Such *segmenta-*

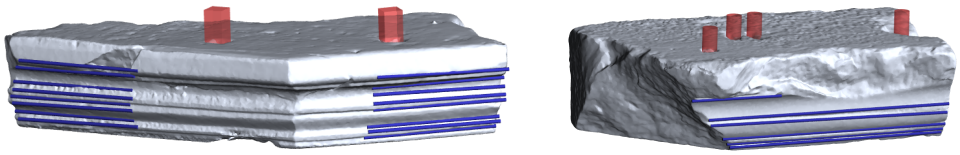

Fig. 13. Features on the surface of fragments comprise straight lines (blue) and clipping/clamping holes (red).

tion problems are very common in Computer Graphics and Image Processing (see [Várady et al. 1997; Dey et al. 2003; Gelfand and Guibas 2004; Pottmann et al. 2004]). As we are extracting planar regions only, we employ a simple RANSAC strategy ([Fischler and Bolles 1981]). The set of planar regions $\{F_i\}$ of a fragment $\mathcal{F}$ forms the basic structure of our feature representation. With simple distance computations we determine adjacent faces and derive a directed graph $G = (V, E)$ with the faces F_i as vertices, connected by two edges e_{ij} and e_{ji} if F_i and F_j are adjacent (Fig. 14). We will see later why we prefer two directed edges instead of a single undirected edge between two nodes.

For each face F_i we compute a handful of attributes. We store the dimension of the planar region (by means of width w_i and height h_i) as adjacent fragments should be of similar height as well as they may have to meet some total length constraint. In addition, we characterize the roughness of a surface similar to [Huang et al. 2006] by computing for each point $\mathbf{p}$ a local bending energy, based on the k-nearest neighbors $\mathbf{q}_i$ of $\mathbf{p}$ (with k typically set to 50) and their normals $\mathbf{n_p}$ and $\mathbf{n}_{\mathbf{q}_i}$,

$$e_k(\mathbf{p}) = \frac{1}{k} \sum_{i=1}^{k} \frac{\|\mathbf{n_p} - \mathbf{n}_{\mathbf{q}_i}\|^2}{\|\mathbf{p} - \mathbf{q}_i\|^2}.$$

We obtain the surface roughness $\rho(\mathbf{p})$ in a point $\mathbf{p}$ of F_i by summing up e_k over the neighborhood $N_r(\mathbf{p}) = \{\mathbf{q} : \|\mathbf{p} - \mathbf{q}\| < r\}$ of $\mathbf{p}$

$$\rho(\mathbf{p}) = \frac{1}{|N_r(\mathbf{p})|} \sum_{\mathbf{q} \in N_r(\mathbf{p})} e_k(\mathbf{q})$$

and define the surface roughness ρ of F_i as the mean surface roughness of all $\mathbf{p} \in F_i$. **Here, r denotes the radius of the neighborhood ball centered at p, for its choice see e.g. [Huang et al. 2006].**

As for the segmentation of the original fragments, *feature extraction* is a well surveyed topic in geometric processing ([Pauly et al. 2003; Li and Guskov 2005; Gal and Cohen-Or 2006; Huang et al. 2006]). Ornaments and trimming lines are extracted by examining and summarizing local curvature extrema. As we are only interested in features spanning over more than one fragment, we fit lines to those high curvature points close to the edges of a fragment and store one line element per feature. For the clamping holes, we determine clusters of high residue to a face's fitting plane and derive approximate position and radius for these (Fig. 13).

We store this information about feature lines and clipping holes in the edges e_{ij} of G instead of the nodes as it will be used to verify matching faces by checking the geometric properties of *adjacent* faces. Assuming two co-planar faces F_i and

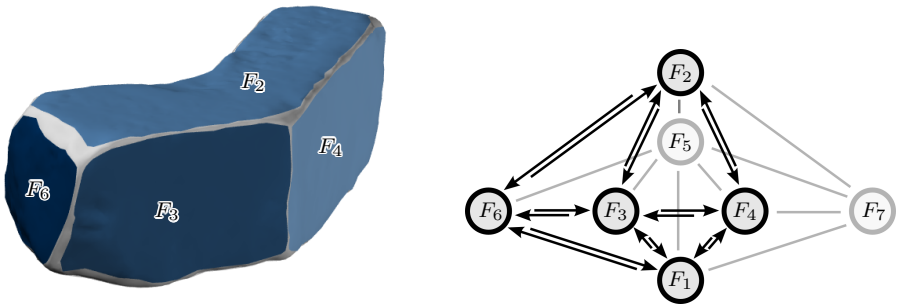

Fig. 14. The segmentation of a fragment into its planar regions (left) is used to build an adjacency graph G (right) storing feature descriptors.

F'_j of two adjacent fragments, any feature line will pass continuously over the gap between the parts just as any clamping holes will be located symmetrically with respect to the edge.

In summary, each fragment is represented as a directed graph $G(V, E)$. The nodes of G are the planar faces of the fragment $F_i = \{w_i, h_i, \rho_i\}$. Vertices of adjacent faces are connected by two directed edges e_{ij} and e_{ji}, where e_{ij} stores feature line and hole information necessary to verify whether F_i matches another face.

5.2 Global alignment of the digital stones

Once we have an abstract representation of the fragments' features we tackle the issue which fragments were originally neighbors. For this purpose, we compare the features of all parts and try to find meaningful pairs. As this step only gives a general assignment, it is often referred to as *global registration* or as *global alignment* (see [Gelfand et al. 2005; Li and Guskov 2005; Huang et al. 2006]). It will be the task of *local registration* to finally align the fragments (see Sec. 5.3). Note the subtle difference between global alignment of digital stones to the global alignment of digital scans as described in Sec. 4.

Our primary goal is to find matching faces (pairs of faces belonging to different fragments that were located face to face originally). We iterate through all possible pairs of faces (F, F') and rate the quality of each pair by comparing the associated feature descriptors. An important aspect of our global matching algorithm is that a correct match needs not to fulfill *all* criteria strictly, as this is nearly impossible because of eroded and missing features. Instead, each passed consistency check increases the score of the match under consideration while failed tests do not discard a match immediately. We check features on the faces themselves for consistency before we examine the faces' neighborhoods. Let $N(F)$ be the set of adjacent faces of F. Then, for a good match, $|N(F)| = |N(F')|$ holds as well as the areas of the faces wh and $w'h'$ are of approximately the same size. Likewise, we require the surface roughness of both to be similar. For adjacent faces $N(F)$ and $N(F')$ of F and F', we can't tell a priori which faces of $N(F)$ and $N(F')$ will be adjacent. Thus, we check all possible combinations for similar surface roughness, matching feature lines and symmetric hole positions. We choose the best pair, and compare

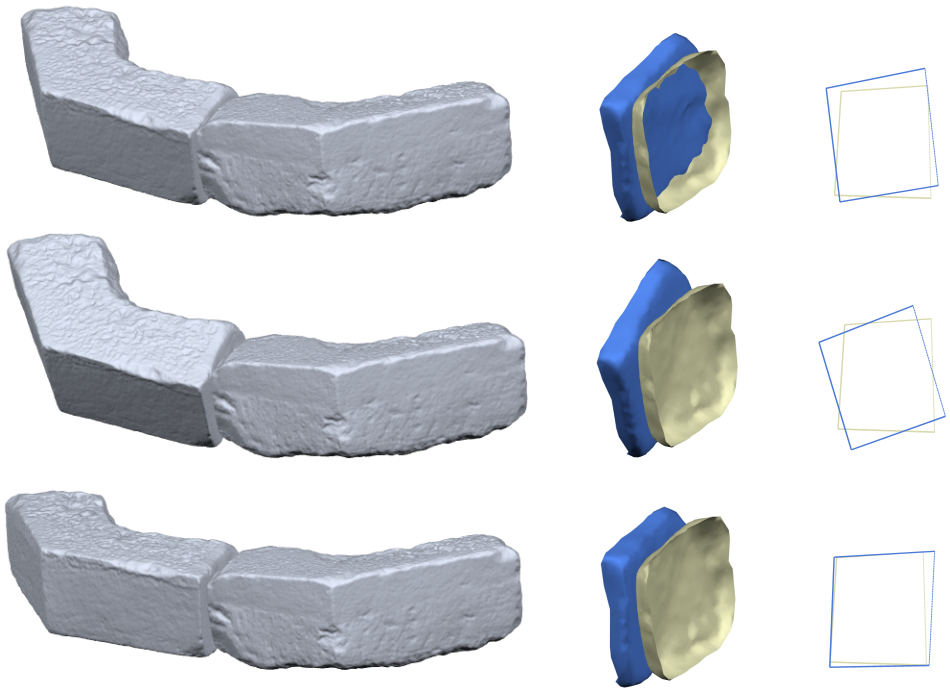

Fig. 15. Registration without any constraints (top row), without mutual penetration (middle row), without mutual penetration and the additional enforcement of "co-planar faces" (bottom row). For each scenario, final alignment (left column), penetration/non-penetration of matching faces (center column) and a cut through the adjacent faces' fitting planes (right column) are shown. In the right column figures, solid lines visualize those planes with enforced co-planarity.

the ratings of the remaining pairs of faces in order.

As in the digital fragment generation step, we translate this matching information into an abstract graph (where the nodes represent the fragments and the edges indicate possible matches between fragments) and employ the method of [Huang et al. 2006] to get consistent groups of matching fragments.

5.3 Local constrained registration without penetration

After a successful global matching phase we obtain a list of matching fragments that need to be aligned properly. The technical term for this task is *local registration*. At first, we roughly align the fragments by a registration with known correspondences ([Horn 1987]). Considering face F_i of a pair of matching faces (F_i, F'_j) we determine that face $F_k \in N(F_i)$ that will most likely be co-planar to a face $F'_l \in N(F'_j)$ (by exploiting the feature descriptors on surface roughness, feature lines and holes as before). Let $\mathbf{n}_i^F$ and $\mathbf{n}_k^F$ denote the unit normal vectors of fitting planes of F_i and F_k. Let $\langle \mathbf{n}_i^F, \mathbf{n}_k^F \rangle$ be the scalar product of two vectors in $\mathbb{R}^3$. Then $\mathbf{n}_{ik} = \mathbf{n}_k^F - \langle \mathbf{n}_i^F, \mathbf{n}_k^F \rangle \cdot \mathbf{n}_i^F$ is a vector pointing roughly from F_i in direction of F_k. Then, we obtain six points per face for a registration with Horn's quaternion approach by

computing the barycenter $\mathbf{b}_i = 1/|F_i| \sum_{\mathbf{p}\in F_i} \mathbf{p}$ and setting

$$\mathbf{p}_{0,1} = \mathbf{b}_i \pm 0.5 \cdot \mathbf{n}_i^F, \quad \mathbf{p}_{2,3} = \mathbf{b}_i \pm 0.5 \cdot \frac{\mathbf{n}_{ik}}{\|\mathbf{n}_{ik}\|}, \quad \mathbf{p}_{4,5} = \mathbf{b}_i \pm 0.5 \cdot \frac{(\mathbf{n}_i^F \times \mathbf{n}_{ik})}{\|(\mathbf{n}_i^F \times \mathbf{n}_{ik})\|}.$$

After this initial positioning, we carry out a local registration without known correspondences to fine-tune the alignment. This registration will happen simultaneously for all fragments to avoid an accumulation of errors and will be constrained to return both a penetration free and properly aligned final position. In [Huang et al. 2006], it is shown how the optimization of Equ. (3) can be constrained to avoid any mutual penetration of the matching faces $F_i \subseteq \mathcal{F}_j$ (*moving* face) and $F_l \subseteq \mathcal{F}_k$ (*target* face). By using the same notation as above, we obtain the parameters $\mathbf{c} = (\mathbf{c}_1, \bar{\mathbf{c}}_1, \ldots, \mathbf{c}_n, \bar{\mathbf{c}}_n)$ describing the fragments' motion by minimizing

$$f_{face}(\mathbf{c}) = \sum_{\mathbf{x}_i \in \mathcal{F}} (d_i + \mathbf{v}_{jk}^T(\mathbf{x}_i) \cdot \mathbf{r}_i)^2,$$

where $\mathcal{F} = \cup_i \{F_i : F_i \text{ is moving face}\}$. The constraints

$$\mathbf{n}_i^T \cdot (\mathbf{x}_i - \mathbf{y}_i + \mathbf{v}_{jk}(\mathbf{x}_i)) \geq 0, \qquad \forall \mathbf{x}_i \in \mathcal{F}$$

ensure a penetration free alignment, where $\mathbf{n}_i$ denotes the outward oriented normal vector in $\mathbf{x}_i$.

However, these constraints are insufficient for our registration problem, as the matching faces carry too few geometric features for a stable and correct local registration (see Fig. 15). Instead, the observation that certain neighboring faces will be co-planar in the final alignment lets us modify the objective function of the optimization problem. For a pair of matching faces, we know the set of corresponding neighboring faces. If surface roughness, feature line and hole information are sufficiently similar for a pair of such adjacent faces (F_r, F_s'), the registration should make these faces co-planar. Let $\varepsilon_r : e_r + \mathbf{x}^T \mathbf{n}_r^F = 0$, with $\|\mathbf{n}_r^F\| = 1$, denote the least squares fitting plane of F_r in Hessian normal form. Then, the distance to that plane is given by $d(\varepsilon_r, \mathbf{x}) = e_r + \mathbf{x} \cdot \mathbf{n}_r^F$ and we modify Equ. (5.3) by adding (with a slight abuse of notation)

$$F_{cp}(\mathbf{c}) = \sum_{\mathbf{x}\in\mathcal{F}'} (e_r + (\mathbf{x} + \mathbf{v}_{jk}^T(\mathbf{x})) \cdot \mathbf{n}_r^F)^2$$

to make adjacent faces co-planar by minimizing the distances between the corresponding points and fitting planes ($\mathcal{F}' = \cup_r \{F_r : F_r \in N(F_i) \subseteq \mathcal{F}_j\}$).

By adapting this idea, we are able to enforce two features to be co-linear as well. We represent such features as line elements $l_i = (\mathbf{p}_i, \mathbf{w}_i)$, where $\mathbf{p}_i$ denotes a point on the feature and $\mathbf{w}_i$ a directional vector. For a pair of matching feature lines l (locally fixed system j) and l' (locally moving system k), we sample l' in at least 2 points s_i'. Then,

$$F_{cl}(\mathbf{c}) = \sum_{(l,l')\in L} \sum_i (\mathbf{w}_j \times (\mathbf{s}_i' + \mathbf{v}_{jk}(\mathbf{s}_i') - \mathbf{p}_j))^2$$

describes the squared distance from the samples to the line elements, where L comprises all pairs of corresponding feature lines.

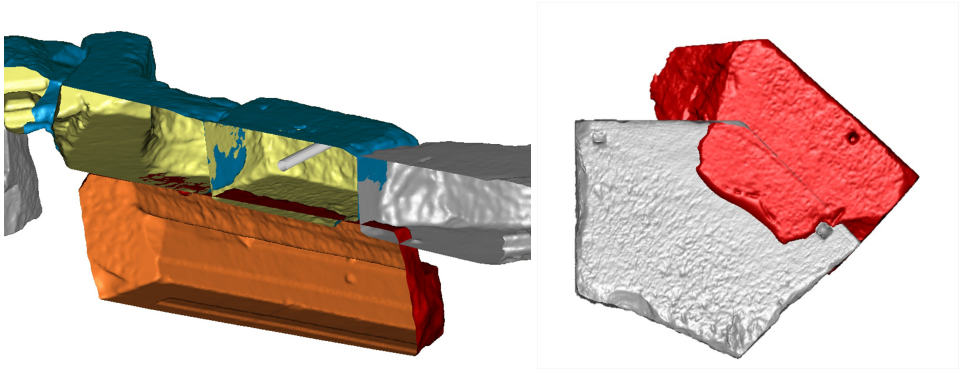

Fig. 16. Examples showing the results of human expert interactive registration of adjacent stones. Please note the unwanted mutual penetration and alignment inaccuracies.

In total, the objective function of our registration problem avoiding penetration and enforcing co-planarity and co-linearity is given by

$$\min F(\mathbf{c}) = F_{reg}(\mathbf{c}) + \mu_{cp} F_{cp}(\mathbf{c}) + \mu_{cl} F_{cl}(\mathbf{c}).$$

subject to

$$\mathbf{n}_i^T \cdot (\mathbf{x}_i - \mathbf{y}_i + \mathbf{v}_{jk}(\mathbf{x}_i)) \geq 0, \qquad \forall \mathbf{x}_i \in \mathcal{F}.$$

The weights μ_{cp} and μ_{cl} control the influence of the co-planarity and co-linearity constraints on the objective function. In general terms, we choose these factors in such a way that F_{reg} on the one hand and both F_{cp} and F_{cl} on the other hand enter the optimization equally. The solution of this linear constrained quadratic optimization problem yields affine (but not Euclidean) motions of the form $\mathbf{x} \rightarrow \mathbf{x} + \mathbf{v}_{i0}(\mathbf{x})$. However, each $(\bar{\mathbf{c}}_i, \mathbf{c}_i)$ determines a unique helical motion we use to refine the fragments position. This registration is repeated iteratively until a final alignment is achieved after 15 to 20 iterations.

6. RESULTS

Although the automatic alignment of the two fragments shown in Fig. 15 is of a high quality, a human expert immediately recognizes that these two stones were not adjacent in the original building (look at the difference in the surface finish).

Otherwise, for a human expert it is impossible to interactively position adjacent fragments in an optimized penetration free way. This is due to the fact that already for two stones you have to interactively tune the six parameters of a rigid body motion (three for the position and three more for the orientation) for one of the stones. If you have an assembly of several stones such an interactive manipulation becomes cumbersome and inaccurate (Fig. 16).

We propose the following strategy for digital anastylosis. The identification of matching blocks is based on the know-how of archaeologists or architects while reverse applicable software provides tools for manipulating the entities in a way which is very similar to the one the crane does in reality. It is an interactive and iterative process of assembling a model where the expert's knowledge about features of the model such as ornaments, edges, clamp holes or dowels is the essential input

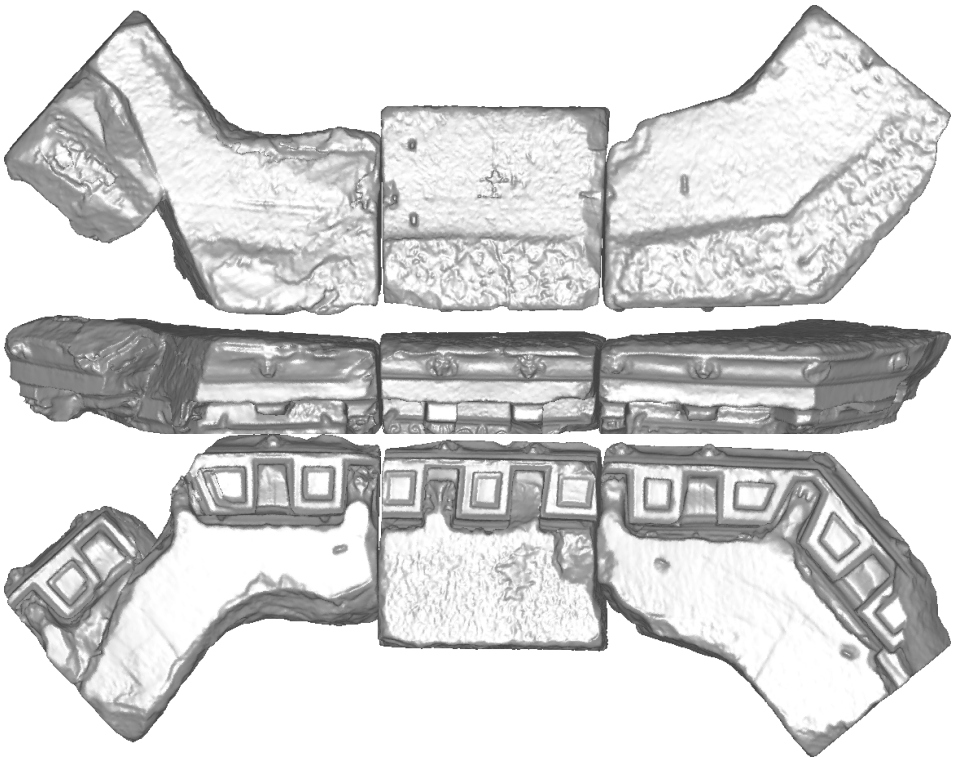

Fig. 17. The optimized mutual penetration free alignment of three adjacent blocks of the Octagon: top, front, and bottom view.

for movement, rotation and grouping of elements. Thus the way of identifying matching blocks in virtual anastylosis is not very different from the one in physical world except that there are computational tools which ease the handling of the objects.

Once a human expert has decided on which stones belong next to each other and which features should be matched, we put that information into our software. This means, in practice our global alignment of stones is aided interactively since the step of identifying possibly matching blocks of the Octagon is rather easy for a human expert. However, the optimal local penetration free alignment of matching blocks is impossible if one tries to do that by hand in a virtual environment. Here the power of our new algorithm unleashes and solves this task for us. In Fig. 17 we show the final penetration free alignment of three adjacent blocks that was computed with the algorithm described in Sec. 5. In Fig. 18 we illustrate the results for a different set of four locally adjacent Octagon blocks, again computed with our algorithm of Sec. 5.

This semi-automatic process of virtual anastylosis results in a model where each of the blocks is accurately located at its original position (Fig.19). With regard to the expected costs of the reassembly of existing and replication of missing elements we now focus on various scenarios, two of them are illustrated in Fig. 20. Based

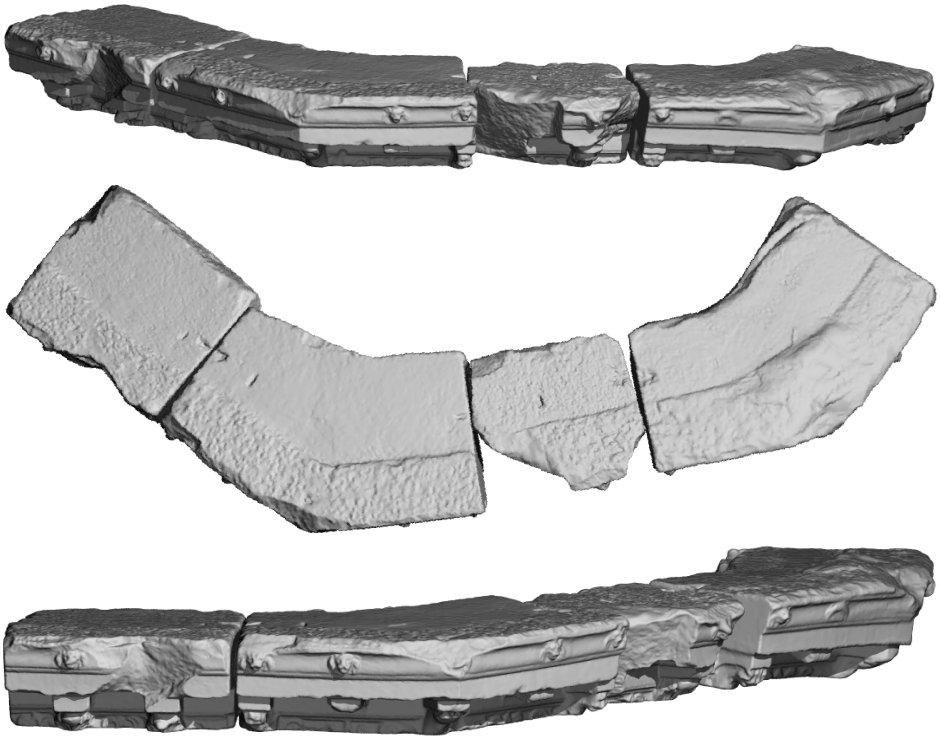

Fig. 18. Three different views of the optimized mutual penetration free alignment of four adjacent blocks of the Octagon.

on these models we can perform visual impact evaluation of the real anastylosis at this prominent spot between the modern front of the terrace house and the reconstruction of the Celsus library (Fig. 21). While the digital anastylosis is by now completed, the physical anastylosis is planed for the near future.

7. CONCLUSION AND FUTURE RESEARCH CHALLENGES

From our point of view progress in the future research on digital anastylosis of ancient monuments will allow to actually "rebuild" more ancient monuments. Currently, the available algorithms that digitally support the work of archaeologists and architects still need a lot of improvements to be of a better use with difficulties that come with real (non-lab) data. While the mathematical foundations are better and better understood also the algorithmic implementation needs to be pushed by research in computer science. The current anastylosis of the Octagon shows the benefits that result from combining automatic techniques that employ strong algorithms from geometry processing with human computer interaction by experts such as architects and archaeologists. As a future vision we see a set of digital technologies that support human experts in their daily work. In such a way computers would be at best use for our cultural heritage.

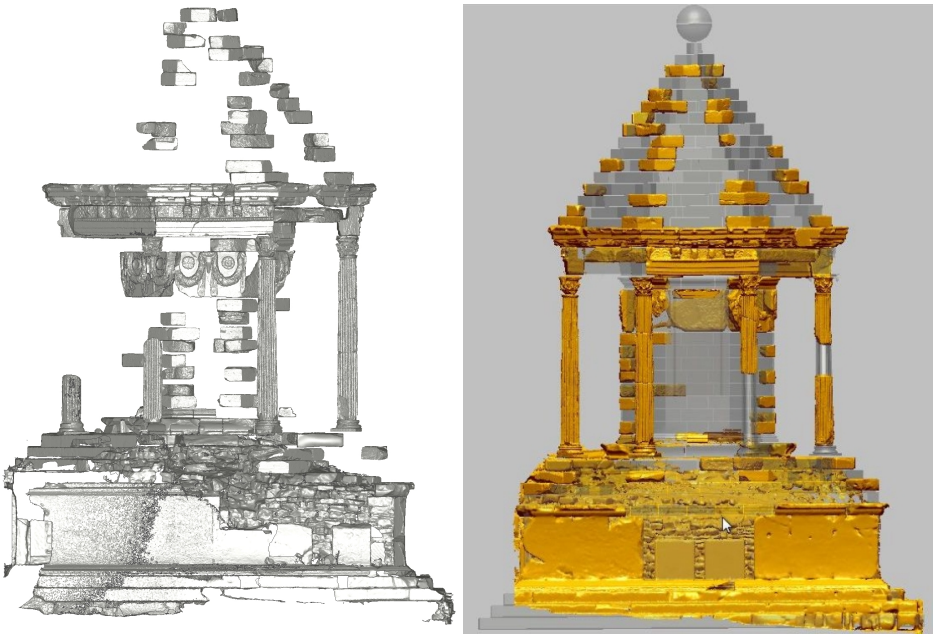

Fig. 19. The virtual anastylosis results (left) showing only the remaining blocks of the Octagon and (right) showing also replications of missing elements.

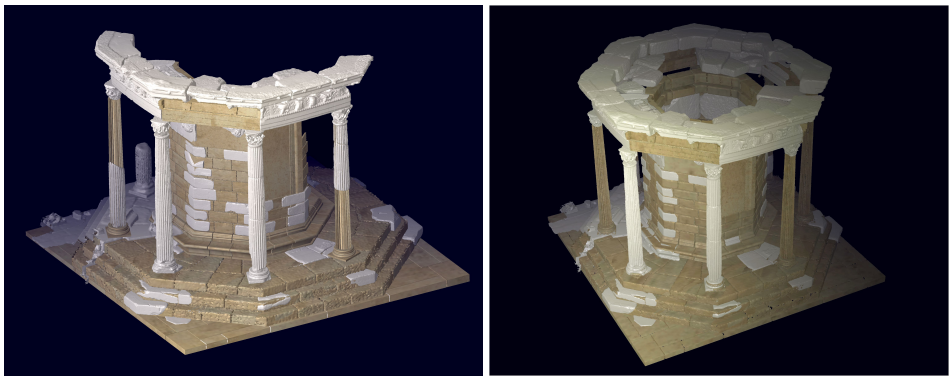

Fig. 20. Two different scenarios for a possible physical anastylosis of the Octagon.

REFERENCES

Abmayr, T., Härtl, F., Reinköster, M., and Fröhlich, C. 2005. Terrestrial laser scanning - applications in cultural heritage conservation and civil engineering. In *Proc. ISPRS 3D-ARCH 2005*.

Allen, P. K., Stamos, I., Troccoli, A., B. Smith, M. L., and Murray, S. 2003. New methods for digital modeling of historic sites using range and image data. *IEEE Journal of Computer Graphics and Applications 23,* 6, 32–41.

Alzinger, W. 1974. Augusteische Architektur in Ephesos. *SoSchrift ÖAI*, 40–42.

Atkinson, A. C., Riani, M., and Cerioli, A. 2004. *Exploring Multivariate Data With the*

Fig. 21. A study of the visual impact of the Octagon anastylosis at its prominent spot in the Curetes street in Ephesos.

Forward Search. Springer.

Bahmutov, G., Popescu, V., and Mudure, M. 2006. Efficient large scale acquisition of building interiors. *Computer Graphics Forum 25,* 3, 655–662.

Bendels, G. H., Schnabel, R., and Klein, R. 2005. Fragment-based surface inpainting. In *SGP 2005 – Poster Proceedings*, M. Desbrun and H. Pottmann, Eds. Eurographics Assoc.

Bernardini, F. and Rushmeier, H. 2002. The 3d model acquisition pipeline. *Computer Graphics Forum 21,* 2, 149–172.

Berns, C. 2003. *Untersuchungen zu den Grabbauten der frühen Kaiserzeit in Kleinasien*. Asia Minor Studien, vol. 51.

Besl, P. J. and McKay, N. D. 1992. A method for registration of 3-d shapes. *IEEE PAMI 14,* 2, 239–256.

Böhler, W., Heinz, G., and Marbs, A. 2001. The potential of non-contact close range laser scanners for cultural heritage recording. In *Proc. CIPA 2001*.

Brown, B. J., Toler-Franklin, C., Nehab, D., Burns, M., Dobkin, D., Vlachopoulos, A., Doumas, C., Rusinkiewicz, S., and Weyrich, T. 2008. A system for high-volume acquisition and matching of fresco fragments: Reassembling Theran wall paintings. *ACM Trans. Graph. 27,* 3, 84:1–84:9. (Proc. SIGGRAPH '08).

Chao, C. and Stamos, I. 2005. Semi-automatic range to range registration: a feature-based method. In *Proceedings Fifth International Conference on 3-D Digital Imaging and Modeling (3DIM 2005)*. 254–261.

Chen, Y. and Medioni, G. 1991. Object modeling by registration of multiple range images. In *Proc. IEEE Conf. on Robotics and Automation*.

Cooper, D., Willis, A., Andrews, S., Baker, J., Cao, Y., Han, D., Kang, K., Kong, W., Leymarie, F., Orriols, X., Velipasalar, S., Vote, E., Joukowsky, M., Kimia, B., Laidlaw, D., and Mumford, D. 2002. Bayesian virtual pot-assembly from fragments as problems in perceptual-grouping and geometric-learning. In *Proc. of ICPR*. Quebec, Canada.

Cooper, D. B. et al. 2001. Assembling virtual pots from 3D measurements of their fragments. In *Proc. of the VAST2001 Conference*. Greece.

da Gama Leitão, H. C. and Stolfi, J. 2002. A multiscale method for the reassembly of two-dimensional fragmented objects. *IEEE Trans. PAMI 24,* 9, 1239–1251.

Davis, J., Marschner, S. R., Garr, M., and Levoy, M. 2002. Filling holes in complex surfaces using volumetric diffusion. In *1st Int. Symp. 3D Data Processing, Visualization, and Transmission*. IEEE CS Press.

Dey, T. K., Giesen, J., and Goswami, S. 2003. Shape segmentation and matching with flow discretization. In *Proc. Workshop Algorithms Data Strucutres (WADS 03)*, F. Dehne, J.-R. Sack, and M. Smid, Eds. LNCS 2748. Berlin, Germany, 25–36.

Fischler, M. A. and Bolles, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM 24,* 6, 381–395.

Gal, R. and Cohen-Or, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph. 25,* 1, 130–150.

Gazzola, P. et al. 1964. International charter for the conservation and restoration of monuments and sites (the venice charter 1964). online available at www. international. icomos. org/ charters/ venice_e. pdf.

Gelfand, N. and Guibas, L. J. 2004. Shape segmentation using local slippage analysis. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*. ACM Press, New York, NY, USA, 214–223.

Gelfand, N., Mitra, N. J., Guibas, L. J., and Pottmann, H. 2005. Robust global registration. In *Symposium on Geometry Processing*. 197–206.

Goldberg, D., Malon, C., and Bern, M. 2004. A global approach to automatic solution of jigsaw puzzles. *Computational Geometry 28,* 2-3, 165–174.

Gool, L. V., Waelkens, M., P, P. M., Vereenooghe, T., and Vergauwen, M. 2004. Total recall: A plea for realism in models of the past. In *Proc. XXth ISPRS Congress, Istanbul, Turkey*. Vol. XXXV, part B5. 332–343.

Gromann, G. U. 1993. *Einführung in die historische Bauforschung*.

Guarnieri, A., Vettore, A., El-Hakim, S., and Gonzo, L. 2004. Digital photogrammetry and laser scanning in cultural heritage survey. In *Proc. XXth ISPRS Congress 2004*.

Heberdey, R. 1905. Beiblatt. *ÖJh 8*, 70–72.

Hofer, M., Sapiro, G., and Wallner, J. 2006. Fair polyline networks for constrained smoothing of digital terrain elevation data. *IEEE Trans. Geosc. Remote Sensing 44,* 10/2, 2983–2990.

Hori, K., Imai, M., and Ogasawara, T. 1999. Joint detection for potsherds of broken earthenware. In *Proc. CVPR*. Vol. 2. 440–445.

Horn, B. K. P. 1987. Closed form solution of absolute orientation using unit quaternions. *J. Optical Society A 4*, 629–642.

Huang, Q.-X., Adams, B., and Wand, M. 2007. Bayesian surface reconstruction via iterative scan alignment to an optimized prototype. In *Proc. Eurographis Symp. on Geometry Processing (SGP) 2007*.

Huang, Q.-X., Flöry, S., Gelfand, N., Hofer, M., and Pottmann, H. 2006. Reassembling fractured objects by geometric matching. *ACM Trans. Graphics 25,* 3, 569–578. Proc. SIGGRAPH 2006.

Huang, Q.-X. and Pottmann, H. 2005. Automatic and robust multi-view registration. In *Geometry Preprint Series, Vienna Univ. of Technology*. Technical Report 152.

Huber, D. and Hebert, M. 2003. Fully automatic registration of multiple 3d data sets. *Image and Vision Computing 21,* 7 (July), 637–650.

Hueber, F. 2002. *Preparatory Architectural Investigation in the Restoration of Historical Buildings*. Leuven University Press, Chapter Building Research and Anastylosis.

Johnson, A. E. and Hebert, M. 1999. Using spin images for efficient object recognition in cluttered 3D scenes. *IEEE Trans. PAMI 21,* 5, 433–449.

Kazhdan, M., Bolitho, M., and Hoppe, H. 2006. Poisson surface reconstruction. In *Symposium on Geometry Processing*. 61–70.

Koller, D. and Levoy, M. 2006. Computer-aided reconstruction and new matches in the Forma Urbis Romae. *Bullettino Della Commissione Archeologica Comunale di Roma*. to appear.

Koller, D., Trimble, J., Najbjerg, T., Gelfand, N., and Levoy, M. 2006. Fragments of the city: Stanford's digital Forma Urbis Romae project. *Journal of Roman Archaeology Suppl. 61*, 237–252. Proceedings of the Third Williams Symposium on Classical Architecture.

Kong, W. and Kimia, B. B. 2001. On solving 2D and 3D puzzles using curve matching. In *Proc. CVPR*. Vol. 2. 583–590.

Laugerotte, C. and Warze, N. 2004. An environment for the analysis and reconstruction of archaeological objects. In *Proc. of the 5th International Symposium on Virtual Reality, Archaeology and Cultural Heritage VAST*, Y. Chrysanthou, K. Cain, N. Silberman, and F. Niccolucci, Eds.

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 *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 131–144.

Li, X. and Guskov, I. 2005. Multiscale features for approximate alignment of point-based surfaces. In *SGP'05*. 217–226.

Miltner, F. 1959. Vorläufiger Bericht über die Ausgrabungen in Ephesos. *ÖJh 44*.

Mitra, N. J., Guibas, L., and Pauly, M. 2007. Symmetrization. In *ACM Transactions on Graphics*. Vol. 26. #63, 1–8.

Mitra, N. J., Guibas, L. J., and Pauly, M. 2006. Partial and approximate symmetry detection for 3D geometry. *ACM Trans. Graph. 25,* 3, 560–568. Proc. SIGGRAPH.

Mitra, N. J. and Nguyen, A. 2003. Estimating surface normals in noisy point cloud data. In *Symposium on Computational Geometry*. 322–328.

Mueller, P., Vereenooghe, T., Vergauwen, M., Gool, L. V., and Waelkens, M. 2004. Photo-realistic and detailed 3d modeling: The Antonine nymphaeum at Sagalassos (Turkey). In *Computer Applications and Quantitative Methods in Archaeology (CAA2004): Beyond the artifact - Digital interpretation of the past*.

Neugebauer, P. J. 1997. Reconstruction of real-world objects via simultaneous registration and robust combination of multiple range images. *International Journal of Shape Modeling 3,* 1, 71–90.

Oberleitner, W. et al. 1978. Funde aus Ephesos und Samothrake, Katalog der Antikensammlung II.

Papaioannou, G. and Karabassi, E.-A. 2003. On the automatic assemblage of arbitrary broken solid artefacts. *Image and Vision Computing 21*, 401–412.

Papaioannou, G., Karabassi, E.-A., and Theoharis, T. 2000. Automatic reconstruction of archaeological finds — a graphics approach. In *Proc. Int'l Conf. Computer Graphics and Artificial Intelligence*. 117–125.

Papaioannou, G., Karabassi, E.-A., and Theoharis, T. 2001. Virtual archaeologist: Assembling the past. *IEEE Computer Graphics and Applications 21,* 2, 53–59.

Papaioannou, G., Karabassi, E.-A., and Theoharis, T. 2002. Reconstruction of three-dimensional objects through matching of their parts. *IEEE Trans. PAMI 24,* 1, 114–124.

Parikh, D., Sukthankar, R., Chen, T., and Chen, M. 2007. Feature-based part retrieval for interactive 3d reassembly. In *WACV '07: Proceedings of the Eighth IEEE Workshop on Applications of Computer Vision*. IEEE Computer Society, Washington, DC, USA, 14.

Park, S., Guo, X., Shin, H., and Qin, H. 2006. Surface completion for shape and appearance. *The Visual Computer 22,* 3, 168–180.

Pauly, M., Keiser, R., and Gross, M. 2003. Multi-scale feature extraction on point-sampled models. *Computer Graphics Forum 22*, 281–289.

Pauly, M., Mitra, N. J., Giesen, J., Gross, M., and Guibas, L. 2005. Example-based 3D scan completion. In *Proc. Symp. Geom. Processing*, M. Desbrun and H. Pottmann, Eds. Eurographics Assoc., 23–32.

Petzet, M. 2002. Anastylosis or reconstruction – the conservation concept for the remains of the Buddhas of Bamiyan. In *Proc. ICOMOS 13th General Assembly, Madrid*. 189–192.

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.

Podolak, J., Shilane, P., Golovinskiy, A., Rusinkiewicz, S., and Funkhouser, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Transactions on Graphics (Proc. SIGGRAPH) 25,* 3 (July).

Pottmann, H., Hofer, M., Odehnal, B., and Wallner, J. 2004. Line geometry for 3D shape understanding and reconstruction. In *Computer Vision — ECCV 2004, Part I*, T. Pajdla and J. Matas, Eds. Lecture Notes in Computer Science, vol. 3021. Springer, 297–309.

Pottmann, H., Huang, Q.-X., Kölpl, S., and Yang, Y. 2005. Integral invariants for robust geometry processing. Tech. Rep. 146, Geometry Preprint Series, Vienna Univ. of Techn.

Pottmann, H., Leopoldseder, S., and Hofer, M. 2004. Registration without ICP. *Computer Vision and Image Understanding 95,* 1, 54–71.

Rusinkiewicz, S. and Levoy, M. 2001. Efficient variants of the ICP algorithm. In *3DIM '01*. IEEE CS, 145–152.

Salemi, G., Achilli, V., Bragagnolo, D., Menin, A., and Ruzzon, F. 2005. Data acquisition for cultural heritage navigation: Integration of panoramic imaging, terrestrial laser scanning and anaglyphs. In *Proc. CIPA 2005*.

Scopigno, R. 2007. Processing methodologies and interactive visualization of 3d scanned architectures. In *From survey to the project: Heritage and Historical Town Centres*, S. P. S. Bertocci, Ed. EdiFir Edizioni Firenze, 172–177.

Sharf, A., Alexa, M., and Cohen-Or, D. 2004. Context-based surface completion. *ACM Trans. Graph. 23,* 3, 878–887.

Shilane, P. and Funkhouser, T. 2006. Selecting distinctive 3D shape descriptors for similarity retrieval. In *Shape Modeling International*.

Spinelli, A., Ganovelli, F., Montani, C., and Scopigno, R. 2006. Recovering 3d architectural information from dense digital models of buildings. In *Fourth Eurographics Italian Chapter 2006*, D. Fellner, Ed. Graz University of Technology, Eurographics Association, Graz, Austria, 177–182. Catania, 22-24 February 2006.

Stamos, I., Liu, L., Chen, C., Wolberg, G., Yu, G., and Zokai, S. 2007. Integrating automated range registration with multiview geometry for the photorealistic modeling of large-scale scenes. *International Journal of Computer Vision*. to appear.

Thrun, S. and Wegbreit, B. 2005. Shape from symmetry. In *Proc. ICCV*. Vol. 2. 1824– 1831.

Thür, H. 1990. Arsinoe IV, eine Schwester Kleopatras VII, Grabinhaberin des Oktogons in Ephesos? Ein Vorschlag. *ÖJH 60*, 43–56.

Valzano, V., Bandiera, A., and Beraldin, J.-A. 2005. Realistic representations of cultural heritage sites and objects through laser scanner information. In *Proc. Cultural Heritage and New Technologies*.

Várady, T., Martin, R. R., and Cox, J. 1997. Reverse engineering of geometric models - an introduction. *Computer-Aided Design 29,* 4, 255–268.

Venkatesh, M. V. and Cheung, S. 2006. Symmetric shape completion under severe occlusions. *Proc. IEEE Conf. Image Processing (ICIP'06)*. to appear.

Verdera, J., Caselles, V., Bertalmio, M., and Sapiro, G. 2003. Inpainting surface holes. In *Int. Conference on Image Processing (ICIP 2003)*. Vol. 2. IEEE, 903–906.

Weyrich, T., Pauly, M., Keiser, R., Heinzle, S., Scandella, S., and Gross, M. 2004. Post-processing of scanned 3d surface data. In *Symposium on Point-Based Graphics*. 85–94.

Willis, A. and Cooper, D. B. 2004. Alignment of multiple non-overlapping axially symmetric 3D data sets. In *Proc. ICPR*. Vol. IV. 96–99.

Zalesny, A., der Maur, D. A., Paget, R., Vergauwen, M., and Gool, L. V. 2003. Realistic textures for virtual anastylosis. *cvprw 01*, 14.

Zou, Y. and Unold, F. 2003. Die großen Buddha-Statuen von Bamiyan – Zustand, Stabilität und Möglichkeiten der Rekonstruktion. *Bautechnik 80,* 7, 417–422.