Photogrammetrie • Fernerkundung • Geoinformation 1/2009, S. 53 – 63

3D Filtering of High-Resolution Terrestrial Laser Scanner Point Clouds for Cultural Heritage Documentation

Clemens Nothegger & Peter Dorninger, Vienna

Keywords: Point Processing, 3D Reconstruction, Visualization, Surface Normals

Summary: Terrestrial Laserscanning has proved to be an important tool for documentation of cultural heritage objects. The latest generation of phase shift scanners features an extremely high scanning speed and improved accuracy, thus making it possible to capture surface detail in the millimetre range. This was previously the exclusive domain of close range triangulation scanners. Triangulation scanners, however, usually have a very limited field of view, thus requiring a large number of scans even for relatively small objects. This can be economically prohibitive. Phase shift scanners, on the other hand, produce huge amounts of data, which commonly used modelling software cannot handle properly. In this paper we present a chain of pre-processing steps which utilizes redundancy to reduce the amount of data without loosing detail at edges. We compare the results obtained by applying our method with the results obtained with commercially available software packages. Since our method is computationally intensive, it is designed to be applied in a batch process, rather than interactively. Therefore, we present a method for estimating necessary parameters. Furthermore we show that a global set of parameters is not suitable, since the point density varies significantly within a single scan and suggest a way to set these parameters adaptively.

Zusammenfassung: *3D-Filterung hoch auflösender Terrestrischer Laserscanner-Punktwolken zur Kulturgutdokumentation.* Im letzten Jahrzehnt haben sich terrestrische Laserscanner als Werkzeug zur Geometrieerfassung von Kulturgütern etabliert. Die rasante Steigerung von Messrate und Messgenauigkeit bei Phasenvergleichs-Scannern ermöglich die Erfassung von Details, die bislang ausschließlich mit Nahbereichsscanner erfasst werden konnten. Allerdings ist der Aufnahmebereich dieser Scanner stark limitiert, wodurch selbst für die Erfassung verhältnismäßig kleiner Objekte viele Einzelaufnahmen erforderlich sind. Dies macht den Einsatz derartiger Geräte für große Objekte unwirtschaftlich. Eine Herausforderung bei Phasenvergleichs-Scannern sind die enormen Datenmengen. So können kommerzielle Softwareprodukte derartige Punktmengen oft nicht angemessen bearbeiten. In diesem Beitrag wird eine Prozessierungskette zur Vorverarbeitung hoch auflösender Laserscanner-Daten vorgestellt. Diese ermöglicht durch Ausnutzung der Redundanz die Punktanzahl signifikant zu reduzieren und gleichzeitig alle Details – beispielsweise scharfe Kanten – zu erhalten. Wir vergleichen unsere Ergebnisse mit jenen, die mittels kommerzieller Softwareprodukte generiert wurden. Da die vorgestellte Methode rechenintensiv ist, ist sie nicht zur interaktiven Bearbeitung geeignet. Daher wird ein automatischer Ansatz zur Schätzung der notwendigen Parameter vorgestellt, wodurch eine automatische Batch-Prozessierung ermöglicht wird. Außerdem wird gezeigt, dass eine adaptive Anpassung der notwendigen Parameter – wie dies die vorgestellte Methode anwendet – notwendig ist, da die Punktdichte innerhalb eines Scans stark variieren kann.

DOI: 10.1127/0935-1221/2009/0006 0935-1221/09/0006 $ 2.75

1 Introduction

For the management of cultural heritage sites it is essential that their assets and all provisions aimed at their preservation are documented. Detailed three dimensional geometrical models are becoming more important, supplementing the traditional photogrammetric and tachymetric records. Terrestrial laser scanners have recently been successfully used to capture the data needed for the construction of such models.

There are three different types of terrestrial laser scanners: instruments utilizing the time-of-flight (TOF) measurement principle for ranges up to one kilometre, instruments utilizing the phase shift measurement principle for ranges of a few dozens of meters and instruments using triangulation for ranges of a few meters or less. Currently available TOF scanners are either slow or have a limited accuracy with their high range being of relatively little importance for the task of cultural heritage documentation at mm-resolution. Triangulation scanners on the other hand have a very small field of view, requiring a large number of individual scans and labour-intensive post processing, which can be economically prohibitive.

The latest generation of phase shift scanners features an extremely high scanning speed and improved accuracy, thus making it possible to capture surface detail with a resolution of up to one millimetre. This was previously the exclusive domain of close range triangulation scanners. Further advantages of phase shift scanners are the much larger field of view and greater potential for automated post-processing which is economically advantageous, thus making them a viable choice if sub-millimetre accuracy is not required.

The use of phase shift scanners is not without challenges, however. With measurement rates of up to 500 kHz, a high resolution point cloud contains a huge number of points (up to 600 mio.), which commercial of the shelf (COTS) modelling software cannot handle directly. It is therefore necessary to apply thinning to reduce the quantity of data. While being fairly homogeneous locally, the point density at different parts of the scan can vary significantly, depending mainly on the polar range and angle of incidence. This makes it necessary to set parameters adaptively if optimal results are to be achieved in the processing of the point cloud. This, however, is usually not possible with COTS software. If this pre-processing is not done carefully loss of detail or even larger scale model deformation can occur, which can be hard to detect at later stages of the processing.

In this paper we propose a chain of pre-processing steps for thinning and smoothing, i. e., 3D filtering, of point clouds acquired with phase shift terrestrial laser scanners. The pre-processing steps include surface normal estimation, surface roughness classification, outlier detection and removal, thinning, and smoothing. Our goal is to avoid loss of detail, especially around edges, while also avoiding model deformations, to achieve millimetre resolution and accuracy of the geometric models.

Since this pre-processing is time-consuming it is designed to be applied in a batch process, rather than interactively. This means that all parameters of the computation must be known in advance. We show how the parameters needed in the computation can be estimated. We also show how to set these parameters adaptively.

To evaluate whether we achieved our goals, we compare the models generated after applying the pre-processing to the original point cloud to see whether there are any major differences. We also compare our results with the results obtained by using COTS software.

2 Related Work

A typical workflow from data (point cloud) acquisition to the final, geometrical surface model, e. g., triangulation (Amenta et al. 2001), or manual construction of primitives (Böhm et al. 2007) comprises the following steps:

- *Data acquisition:* Measurement of the point cloud
- *Calibration:* Elimination of systematic effects
- *Thinning:* Reducing the amount of data
- *Smoothing:* minimization of measurement noise

- *Registration:* Transformation of the individual point clouds (scans) into a project coordinate system (in photogrammetric context often referred to as orientation)
- *Merging:* Combining multiple point clouds into one cloud, possibly including thinning
- *Surface modelling:* Triangulation, freeform surface estimation, or primitive fitting

Note that the order of the steps can vary, especially the registration and surface modelling is sometimes applied before the thinning and smoothing steps.

Within the introduction, we gave a coarse overview on laser scanning techniques for point cloud acquisition. A survey of some TLS instruments can be found in (Kersten et al. 2008).

Calibration, i. e., the correction of systematic errors, is extremely important to achieve good accuracy. Ideally this is done by the instrument manufacturer and applied transparently during data acquisition. However, sometimes there are systematic errors which are not handled properly by the built in calibration (Nothegger & Dorninger 2007). In this case additional calibration functions need to be determined and applied (Dorninger et al. 2008, Gielsdorf et al. 2004, Lichti 2007).

While calibration tries to eliminate systematic errors, smoothing is used to reduce the effect of random measurement errors by using averaging. Smoothing can be applied either to the point cloud (Levin 2003), or to a mesh. Mesh smoothing is well studied. For a recent comparison of common methods see for example (Belyaev & Ohtake 2003).

Thinning is used to reduce the amount of data to process or store. Just like smoothing it can either be applied to the point cloud or to a polygonal mesh. Mesh simplification or reduction is well studied and has also been studied in the context of cultural heritage documentation (Lindstaedt et al. 2008). Thinning of point clouds ranges from simple resampling to sophisticated point cloud simplification methods (Pauly et al. 2002, Moenning & Dodgson 2003, Song & Feng 2007). The advantage of using the mesh based methods for thinning and smoothing is that they can take advantage of topological relations given by the mesh. This is also their drawback since it is not easy to generate a sufficiently good mesh from the raw point cloud.

The registration of point clouds can be realized by means of tie points (signalized points or points that can be identified within a scene), relatively (i. e., minimizing the distances of the points belonging to the individual point clouds), or as a combination of both. For registration based on tie points, the achievable accuracy increases with the number and geometric distribution of given tie points. In many cases, extensive use of signalized points is not feasible (too few natural points and the placement of signalized points on the object is prohibited), possibly leading to insufficient results. For relative registration, the iterative closest point (ICP) algorithm (Rusinkiewicz & Levoy 2001) is commonly used. For that, overlapping scenes and an approximation of the solution are required. A method, combining tie point observations and relative registration is described by (Akca & Gruen 2007).

The triangulation of point clouds has been actively studied in computer graphics. Numerous algorithms exist that can be used to reconstruct surfaces from noisefree point samples (Hoppe et al. 1992, Amenta & Bern 1999, or Dey & Goswami 2003). Some of these algorithms have been extended such that they also work with noisy data (Dey & Goswami 2006, Kolluri et al. 2004, Kazhdan et al. 2006). They work best when the noise level is low, which makes it necessary to reduce measurement noise in situations where noise levels are high. Another problem with the latter two algorithms is that they optimize globally, which makes them unsuitable for large datasets. Currently efforts are being made to derive localized algorithms (Schall et al. 2007).

3 Pre-processing

The goal of our proposed pre-processing chain is to reduce point density and measurement noise considerably while preserving richness in detail as best as possible. The first step in the chain is the analysis of the surface and a classification according to the local curvature. The second step consists of a computation of robust surface normal vectors with outlier detection and elimination. The third step is a

thinning step, which is followed by a forth and final smoothing step.

All steps require the local neighbourhood of a point *p*. As the local neighbourhood of a point *p* we use the *k* points which are closest to *p*. These *k* points can be found efficiently by using the *kd-tree* data structure (Berg et al. 2000). The choice of *k* is crucial, however, if consistent results are to be achieved. (Mitra et al. 2004) therefore suggest using all points within a certain radius instead. This radius depends on local point density and curvature. For locally homogeneous point distribution, however, this is comparable to choosing an individual *k* for each point. The choice of *k* is described in more detail in Sections 3.4 and 3.5.

3.1 *Surface Analysis*

The surface analysis we perform is based on the covariance analysis of the local neighbourhood. (Pauly et al. 2002) have shown that the surface variation – which can be computed from the eigenvalues of the covariance matrix – can be used instead of the local surface curvature for estimating surface roughness. In some cases it is even advantageous to use surface variation instead of curvature estimation based on function fitting. The surface is then classified into smooth, rough and intermediate areas by using the maximum surface variation in the neighbourhood and applying some thresholds. Note that these thresholds heavily depend on the choice of *k*.

3.2 *Normal Vector Estimation*

Normal vectors can be estimated quite easily from the covariance analysis of the local neighbourhood (Dey et al. 2005). The problem with this approach is that it only works for continuous and smooth surfaces, i. e., surfaces not containing any sharp features. This problem can be circumvented by first segmenting the point cloud into piecewise smooth patches. (Fleishman et al. 2005) proposed a method based on region growing in conjunction with robust estimation to reconstruct these piecewise smooth surfaces.

We propose a different approach, which is also based on robust estimation, but it avoids the need to explicitly segment the point cloud. We use the Fast Minimum Covariance Determinant (FMCD) estimator (Rousseeuw & van Driessen 1999) for the highly robust estimation of surface normal vectors. The FMCD estimator robustly determines a covariance matrix C from the local neighbourhood k. It works by starting with a minimal subset of known good, i. e., belonging to the same smooth surface patch, points needed to compute a non-singular covariance matrix C_0. This set is then enlarged to contain at least $k/2+2$ points by including those points which have the highest probability of being good points as well and afterwards, a new covariance matrix C_1 is computed. This step is iterated by again selecting the $k/2+2$ points having the highest probability of being good points and computing C_{i+1} until convergence. That last covariance matrix is used to select all points from the original neighbourhood that are also good points, up to a certain error bound. The surface normal vector is then computed from this set of points which probably belongs to the same smooth surface patch by taking the eigenvector corresponding to the least eigenvector of the covariance matrix C_{good}.

As described above the algorithm needs at least four known good points to start with. The original algorithm uses random sampling to get as many starting sets as are needed to have at least one set of only good points, up to a certain probability, and iterating the most promising sets. Depending on the neighbourhood size this would require up to several hundred start sets. This is not really much, but since the iteration is computationally expensive and hundreds of millions of points need to be processed this would result in quite unacceptable computation times.

Thus we replaced the random sampling with a heuristic which utilizes the fact that the points are in fact point samples of a surface. If the surface was classified as being smooth, only as single start set is used consisting of the points having the least distance to an adjusting plane. This is sufficient since in this case we do not expect any points belonging to different surface patches. If the surface was classified

as being rough – thus possibly containing sharp edges – we subdivide the points into octants arranged around the centroid and use one starting set from each octant by choosing the points closest to an adjusting plane for that octant.

Robust estimation can be used to detect outliers, i. e., data not originating from the same underlying process as the majority of the other data, more reliably; since it is not affected by leverage points (ROUSSEEUW & LEROY 1987). Outliers in TLS may be caused, for example, by specular reflections, or if the laser beam hits multiple surfaces. For outlier determination we use the result from the normal vector estimation. A point is classified as an outlier and deleted, if it is not contained in the final set of points which contains all points that are likely to be part of the same surface.

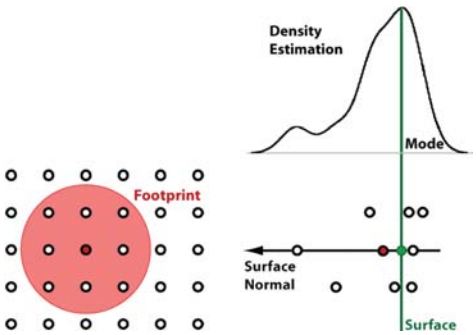

Fig. 1: Illustration of the thinning step. Left: point distribution and footprint extension. Right: distribution of the points within the footprint in a profile view, local surface normal and finally estimated surface at the mode of the locally estimated density.

3.3 *Thinning and Smoothing*

In the thinning step (cf. Fig. 1), the estimated normal vectors are used to determine points having the highest probability of being closest to the real surface. We start by selecting a random point. From this point we take all points within a cylinder defined by a radius r around that point and having normal vectors differing by no more than a certain angle. The cylinder axis is defined by the normal vector. For optimal results r should be chosen such that the points within this radius can be considered to be redundant measurements of more or less the same area. This can be assumed if the point spacing is smaller than the laser footprint. The points within the cylinder are projected onto the cylinder axis. Along the cylinder's axis a univariate density is estimated and the point being closest to the mode of this distribution is chosen as the representative point. The other points are removed from the dataset. The mode of the distribution is found by applying the mean-shift algorithm (CHENG 1995).

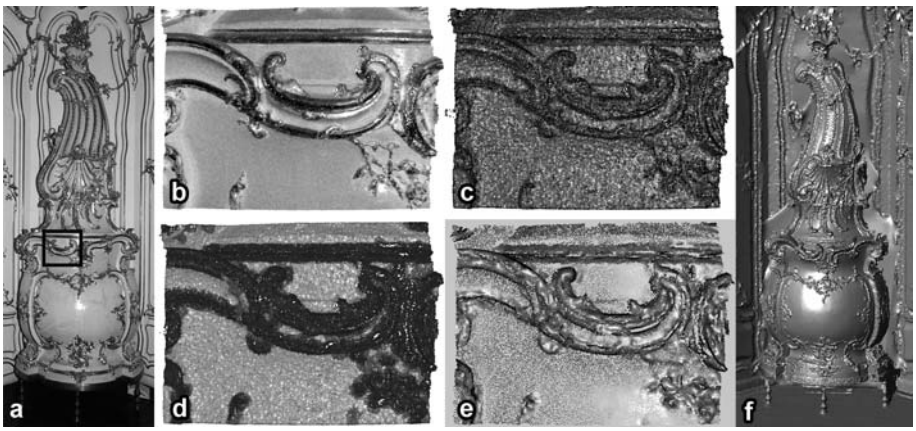

Fig. 2: Photo of a Rocaille stove showing the test area (black rectangle) covering about 45 by 35 cm (a: whole object, b: test area). c shows the original point cloud, d the classification result and e the filtered point cloud. A triangulation of the whole object is shown in f.

This is repeated until the desired point density is reached. The result is a resampled point cloud in which random measurement noise is reduced and which has a more homogeneous spacing of points while preserving detail. This is because in areas with higher point densities – either because of being closer to the scanner, because of the incidence angle, or because of overlapping scans – more points are deleted than in areas of low point densities. The radius r can be chosen adaptively depending on the roughness classification (cf. Section 3.1) and thus allowing for curvature based thinning, i. e., more points are retained in rough areas than in smooth areas.

Fig. 2 shows the described processing steps applied to a test dataset. Images of the whole object and the testing area are shown in (a) and (b). A rendering of the original point cloud is shown in (c). The result of the classification step is shown in (d) and the result of the 3D filtering, i. e., smoothing and thinning, in (e). Fig. 2f shows a triangulation of the whole object.

3.4 *Parameter Estimation*

As explained above, the choice of k is crucial for consistent and reliable results for the entire scan. We want to choose k such that for all flat surfaces the surface variation is constant, i. e., the ratio of the smallest to the largest eigenvalue of the covariance matrix of the neighbourhood must be constant:

$$\lambda_3 / \lambda_1 = \text{const (with } \lambda_1 > \lambda_2 > \lambda_3) \quad (1)$$

Note that for larger k, λ_3 converges to the random measurement error σ^2. σ^2 is also assumed to be constant, meaning that λ_1 must also be constant. λ_1 depends on the neighbourhood size k, the local point density, and the local curvature. But since we are only considering flat areas, we can disregard the curvature. In other words k is proportional to λ_1 up to a factor depending on the local density.

For a given start point, we can compute the ratio λ_3 / λ_1 for increasing k. At a certain point this ratio will drop below a threshold τ. k is then chosen as the point where the ratio falls below the threshold τ. The value of τ can be chosen empirically. It depends on the error bounds for the normal vectors.

3.5 *Parameter Adaptation*

As described above, the optimal k depends on the relation between measurement noise and point density. For TLS the measurement noise is fairly constant, increasing just slightly with

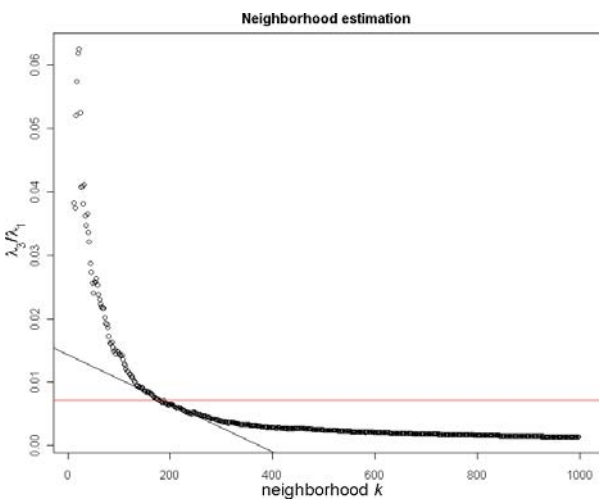

Fig. 3: Neighborhood estimation based on λ_3/λ_1 for increasing k. Threshold τ (red line) and the tangent fitted to the curve for the intersection point determination (black line) are shown.

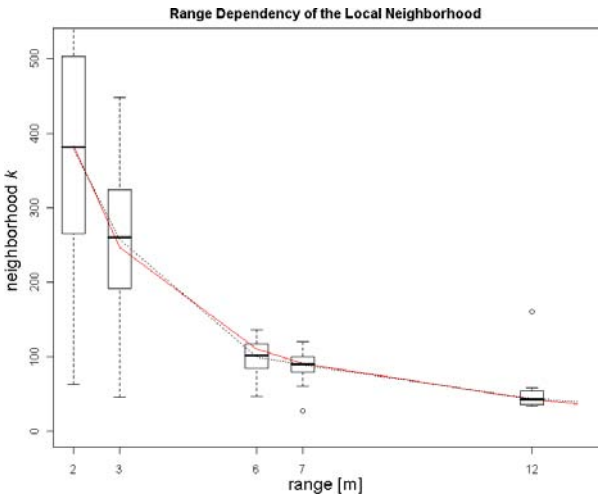

Fig. 4: Range dependency of local neighbourhood. Shown are neighbourhood estimates (black dotted line), fitted function (red line), and boxplots of the samples for each interval.

increasing range. The point density, however, decreases rapidly, proportional to the square of the range. Another influence is the angle of incidence. As the angle of incidence increases, the point density decreases. Some scanners do not correct the convergence of scan lines at the zenith and nadir, thus they exhibit an increased point density towards these points.

We propose the following empirical formula to compute k:

$$k = a \cdot \cos(\alpha) / (r \cdot \sin(\theta)) + b \qquad (2)$$

where α is the angle of incidence, r is the range, θ is the zenith angle and a and b are coefficients. The term $sin(\theta)$ must only be used if the scanner does not compensate the convergence of scan lines. Fig. 4 shows the estimations for k at different ranges and the fitted function. The estimation was performed as described above using 250 random samples for which $cos(\alpha)$ and $sin(\theta)$ were approximately 1.0. The samples were divided into 1 m range intervals and the median in each interval was taken as the estimate for that range. The large variation in the estimation for short ranges is to be expected because for these ranges the aggregation interval is fairly large.

4 Application and Discussion

All point clouds presented in this paper were acquired in *Schönbrunn Palace*, Vienna, Austria using *Faro* scanners (www.faro.com). The detail of a pillar was acquired using a *Faro Photon* and the point cloud representing the historic stove (see also Section 3.3) was acquired using a *Faro LS 880HE*. Both scanners apply the phase shift measurement principle. They emit a laser beam with a wavelength of 785 nm, their maximum sampling rate is 120000 Hz, and the distance measurement range is between 1 and 70 m. The beam divergences and diameters are slightly different (0.16 mrad and 3.3 mm at exit for the *Photon* resp. 0.25 mrad and 3 mm for the *880HE*). The achievable precision of the *Photon* (rms@25 m: 2.2 mm) is better than that of the *880HE* (rms@10 m: 3 mm).

We processed the data on a dual Xeon X5355 workstation (2.66 GHz, 8 cores). On this machine our algorithm processes 0.5 million points per minute on average. The estimation of the parameters takes about an hour, however, this has to be done only once for each type of application.

To evaluate the quality of a triangulated model, we compare the model to the original point cloud using difference models. All triangulations were determined by *Geomagic Stu-*

dio 10 (www.geomagic.com). For a best fitting model we expect randomly distributed differences with a magnitude equal to that of the measurement noise level, no systematic deformations, but possibly large differences at erroneously measured points.

The first test dataset is an 80 by 40 cm part of a pillar of an arcade of *Schönbrunn Palace.* The point cloud consists of 1.2 million points acquired from four scan positions at a scanning distance of 1.5 m. The mean point spacing was 2 mm. This object mainly consists of planar faces which are bounded by sharp edges. We chose this object to test our method's behaviour at these edges. Fig. 5, left, shows the result of a processing with *Geomagic*, whereas the right image shows the result obtained by applying our pre-processing prior to triangulation. The noise level of the scanner is approximately 2 mm, thus we would expect that a mesh which faithfully models the original point cloud show only green, yellow or cyan colours in the difference model, whereas orange, red or blue indicate significant deviations. It can clearly be seen that we mostly achieve our goal, whereas the model derived with *Geomagic* only shows deviations of approximately 3 mm around the sharp edges.

The second test dataset is a 45 by 35 cm part of a historic stove (cf. Fig. 2b). The point cloud consists of 348921 points with a mean point spacing of 0.7 mm at the scanning distance of 5 m. This dataset contains surfaces of varying curvatures, some sharp edges but no flat surfaces. First, we determined triangulation models using the instrument vendor's software *FaroScene*, and *Geomagic*. We tested numerous, different workflows to determine triangulation models using these two software packages. Apparently the best results were achieved, combining functionality of both products. Model 1, shown in Fig. 6a, was generated by applying only a little smoothing in *FaroScene* and successive triangulation with noise reduction in *Geomagic*. It represents the surface detail almost properly (cf. Figure 6b: low differences at sharp structures and almost no global deformations), but the influence of the noise was not suppressed sufficiently. Hence, we increased the smoothing in *FaroScene* to generate model 2 (cf. Fig. 6c). Apparently, the noise was eliminated, but systematic differences occur at detailed structures (red and blue coloured regions) and additionally, sharp edges are smoothed (cf. Fig. 6d). Model 3 was determined from a point cloud after applying our pre-processing method resulting in a filtered point cloud of 75000 points.

Due to the slightly higher noise level (approx. 3 mm) of the *880HE* scanner used for

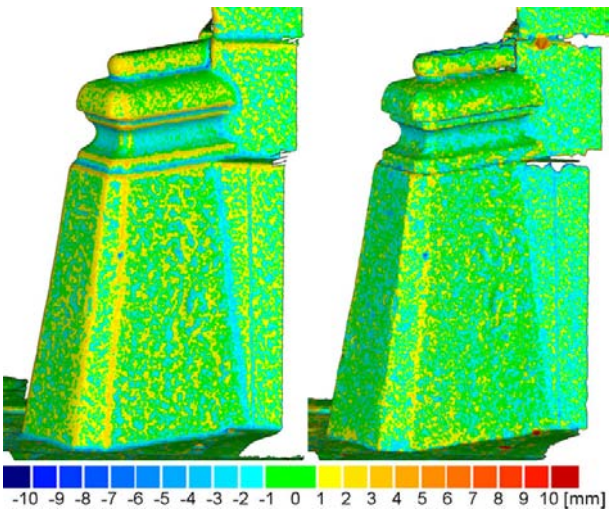

Fig. 5: Detail of a pillar. Left: Model generated using Geomagic only. Right: Model generated by combining our method with a triangulation by Geomagic. The color coding represents the differences of the models to the original point cloud and ranges from –10 (blue) to +10 mm (red) with 1 mm intervals. Differences of ±1 mm are shown in green.

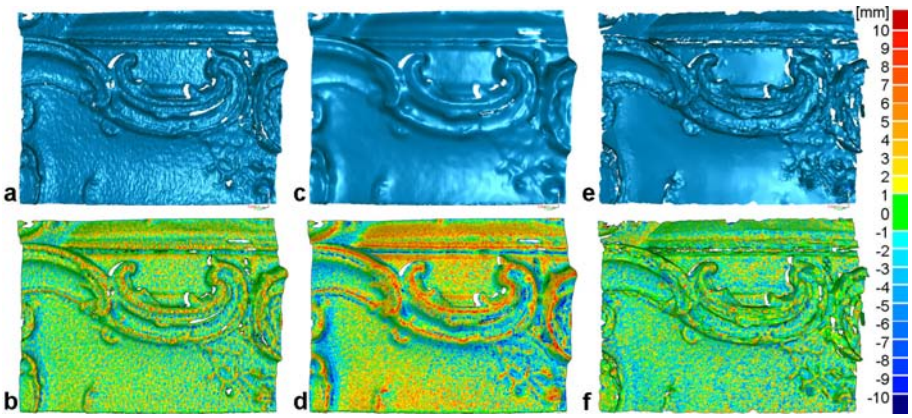

Fig. 6: Rendered triangulations of model 1 (a), 2 (c), and 3 (e). Differences of these models to the original point cloud are shown below ((b), (d), and (f)). Color coding as in Fig. 5.

data acquisition, also orange and light blue colours are acceptable. Here also we succeeded in achieving our goals as our model features the advantages of the other two models: The noise is suppressed sufficiently within the whole testing area (cf. Fig. 6e) and the mean systematic differences at detailed surface structures, i. e., stuccoes, are within the range of the noise (cf. Fig. 6f), hence preserving richness in detail. Moreover, no global deformations – as occurring in model 2 – are introduced by this approach and sharp structures (edges) are well preserved (cf. Fig. 6e, upper, left region).

These findings are further illustrated by Tab. 1. The first column lists the number of points from which the model was created. In our pre-processing we reduced the number of points to about 22%. For the construction of the models 1 and 2 only smoothing and outlier elimination was used such that no loss of detail resulting from resampling was incurred. Nonetheless it can be seen that Model 3 has the lowest average deviation from the original point cloud, despite the other models being derived from more points. The deviations of Model 2 are almost 50% higher than those of Model 3. To assess the smoothness of the models we computed the angles between neighbouring triangles. We then computed the average of the 25% and 10% smallest angles respectively. Model 1 clearly has the roughest surface of all models. Considering the 10% smallest angles the Models 2 and 3 are equal, however, considering the 25% smallest angles shows that Model 2 is smoother in this respect. This is because the 10% of smallest angles are exclusively in the flat areas while the 25% also contain triangles which are within the feature rich areas. Fig. 6 also shows that in feature rich areas Model 2 is clearly smoother, if not overly smooth.

Tab. 1: Comparison of the models. The table lists the number of points from which the model was derived, the standard deviation of the distance to the original points, the average of the positive and negative distances and the average angle between neighboring triangles considering small angles only.

Model	Points used	Triangles	Standard Deviation	Av. Pos. Distances	Av. Neg. Distances	Av. Angle 25%	Av. Angle 10%
1	343010	199999	3.1 mm	2.4 mm	−2.2 mm	4.66°	0.97°
2	346295	199999	4.0 mm	3.1 mm	−2.9 mm	1.05°	0.21°
3	76377	156841	2.7 mm	2.0 mm	−2.1 mm	1.94°	0.22°

5 Conclusions

Phase shift laser scanning is an efficient data acquisition technique that is suitable for many purposes. Compared to triangulation scanners, data acquisition can be performed faster and the effort for registration of individual scans can be reduced significantly. Nevertheless the processing of the point cloud poses a number of challenges that must be overcome if the full potential of the instruments is to be tapped. These challenges are the need to verify the calibration of the instrument, the huge size of the point cloud and the relatively high noise level.

We presented a method for pre-processing point clouds from phase shift laser scanners, which can be applied fully automatically, since it only requires one parameter to be set – the size of the local neighbourhood – and this parameter can be chosen adaptively according to well known criteria. We demonstrated that it does neither introduce smoothing around edges nor any global deformations. The result of this pre-processing is a point cloud which is reduced in size and has a more homogeneous point density.

The point cloud simplification methods cited in Section 2 are much more aggressive in their reduction of the size of the point cloud. They are, however, aiming at deriving a final point set for storage and display, where efficient use of storage space is important. We consider the thinned point cloud only as the input for later processing steps, i. e., triangulation, where a homogeneous, dense sampling of the surface is required.

Our algorithm works by fitting planes to the surface. This works well in most circumstances, but for surfaces with high curvature it leads to some deformation. It remains to be examined if this algorithm can be extended to robustly fit second order surfaces if a smooth, but highly curved surface is detected. Another issue that should be investigated is whether the initial surface analysis can be made more reliable by using a hierarchical approach.

To determine 3D models by triangulating the pre-processed points, we used the commercial software package *Geomagic*. Unfortunately, the triangulation algorithm suffers from certain deficiencies. For example, the resulting model is (in general) not watertight requiring interactive post processing, very large datasets cannot be processed properly, control parameters are used globally, etc. Hence, advances in triangulation methods will improve the automation and the achievable quality of 3D models determined from terrestrial laser scanner data.

Concluding, it can be stated that the presented method allows pre-processing of points acquired by phase shift scanners for subsequent model generation at millimetre scale. This increases the economical attractiveness for the application of laser scanner technology, especially in the field of cultural heritage documentation.

Acknowledgements

We would like to thank the private management of *Schönbrunn Palace* and *Steinmetzbetriebe Bamberger* as members of the *Christian Doppler-Laboratory* for supporting our investigations.

References

Akca, D. & Gruen, A., 2007: Generalized least squares multiple 3D surface matching. – International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences **36** (3/W52): 1–7.

Amenta, N. & Bern, M., 1999: Surface reconstruction by voronoi filtering. – Discrete and Computational Geometry **22** (4): 481–504.

Amenta, N., Choi, S. & Kolluri, R.K., 2001: The power crust, unions of balls, and the medial axis transform. – Computational Geometry **19 (**2–3): 127–153.

Belyaev, A. & Ohtake, Y., 2003: A Comparison of Mesh Smoothing Methods. – Israel-Korea Bi-National Conference on Geometric Modeling and Computer Graphics **2003:** 83–87.

Berg, M.D., Krefeld, M.V., Overmars, M. & Schwarzkopf, O., 2000: Computational Geometry Algorithms and Applications. Springer.

Böhm, J., Haala N. & Becker, S., 2007: Facade modelling for historical architecture. – International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences **36** (5/C53): 147–152.

Cheng, Y., 1995: Mean Shift, Mode Seeking, and Clustering. – IEEE Transactions in Pattern

Analysis and Machine Intelligence **17** (8): 790–799.

Dey, T.K. & Goswami, S., 2003: Tight cocone: A water-tight surface reconstructor. – Journal of Computing and Information Science in Engineering **3:** 302.

Dey, T.K., Li., G. & Sun, J., 2005: Normal estimation for point clouds: A comparison study for a Voronoi based method. – Symposium on Point-Based Graphics **2005:** 39–46.

Dey, T.K. & Goswami, S., 2006: Provable surface reconstruction from noisy samples. – Computational Geometry: Theory and Applications **35** (1–2): 124–141.

Dorninger, P., Nothegger, C., Pfeifer, N. & Molnar, G., 2008: On-the-job detection and correction of systematic cyclic distance measurement errors of terrestrial laser scanners. – Journal of Applied Geodesy, to appear.

Fleishman, S., Cohen-Or, D. & Silva, C.T., 2005: Robust moving least-squares fitting with sharp features. – ACM Transactions on Graphics **24** (3): 544–552.

Gielsdorf, F., Rietdorf, A. & Gruendig, L., 2004: A Concept for the calibration of terrestrial laser scanners. – FIG Working Week 2004, Athens, Greece, May 22–27, on CD-ROM.

Hoppe, H., DeRose, T., Duchamp, T., McDonald, J. & Stuetzle, W., 1992: Surface reconstruction from unorganized points. – 19th annual conference on Computer graphics and interactive techniques: 71–78.

Kazhdan, m., Bolitho, m. & Hoppe, H., 2006: Poisson surface reconstruction. – Fourth eurographics symposium on geometry processing **2004:** 61–70.

Kersten, Th., Mechelke, K., Lindstaedt, m., Sternberg, H., 2008: Geometric Accuracy Investigations of the Latest Terrestrial Laser Scanning Systems. – FIG Working Week 2008, Integrating Generations, Stockholm, Sweden, June 14–19.

Kolluri, R., Shewchuk, J.R. & O'Brien, J.F., 2004: Spectral surface reconstruction from noisy point clouds. – 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing **71:** 11–21.

Levin, D., 2003: Mesh-independent surface interpolation. – Geometric Modeling for Scientific Visualization, Springer: 37–49.

Lichti, D.D., 2007: Error modeling, calibration and analysis of an AM-CW terrestrial laser scanner system. – ISPRS Journal of Photogrammetry and Remote Sensing **61** (5): 307–324.

Lindstaedt, m., Kersten, Th., Mechelke, K., Götting, m. & Heiden, R., 2008: Virtuelles 3D-Modell der antiken Tempelanlage in Sirwah/Jemen zur archäologischen Objektdokumentation durch terrestrisches Laserscanning und Photogrammetrie. – Publikationen der DGPF **17:** 59–68.

Mitra, N.J., Nguyen, A. & Guibas, L., 2004: Estimating surface normals in noisy point cloud data. – International Journal of Computational Geometry and Applications **14** (4/5): 261–276.

Moenning, C. & Dodgson, N.A., 2003: A new point cloud simplification algorithm. – 3rd IASTED International Conference on Visualization, Imaging, and Image Processing, Benalmadena, Spain: 1027–1033.

Nothegger, C. & Dorninger, P., 2007: Automated Modeling of Surface Detail from Point Clouds of Historical Objects. – International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences **36** (5/C53): 538–543.

Pauly, m., Gross, m. & Kobbelt, L.P., 2002: Efficient simplification of point-sampled surfaces. – Conference on Visualization, IEEE Computer Society: 163–170.

Rousseeuw, P.J. & van Driessen, K., 1999: A fast algorithm for the minimum covariance determinant estimator. – Technometrics **41** (3): 212–223.

Rousseeuw, P.J. & Leroy, A.M., 2003: Robust Regression and Outlier Detection. – Wiley Series in Probability and Statistics.

Rusinkiewicz, S. & Levoy, m., 2001: Efficient variants of the ICP algorithm. – 3rd International Conference on 3D Digital Imaging and Modelling, Quebec, Canada: 145–152.

Schall, O., Belyaev, A. & Seidel, H.-P., 2007: Error-guided adaptive fourier-based surface reconstruction. – Computer-Aided Design **39** (5): 421–426.

Song, H. & Feng, H., 2007: Point-cloud simplification with bounded geometric deviations. – International Journal of Computer Applications in Technology **30** (4): 236–244.

Address of the Authors:

Clemens Nothegger & Peter Dorninger, Technische Universität Wien, Christian Doppler Laboratory for "Spatial Data from Laserscanning and Remote Sensing", A-1040 Vienna, Tel.: +43-1-58801-12233, Fax: +43-1-58801-12299, e-mail: {cn, pdo}@ipf.tuwien.ac.at

Manuskript eingereicht: Juni 2008
Angenommen: November 2008