Reconstruction of Smoothed Polyhedral Surfaces from Multiple Range Images

G. Häusler and S. Karbacher

University of Erlangen, Physics Institute, Chair for Optics
Staudtstr. 7/B2, D-91058 Erlangen, Germany
Email: {haeusler,karbacher}@physik.uni-erlangen.de

Abstract

In order to digitize the whole surface of a three-dimensional object by means of an optical range sensor, usually multiple range images are acquired from different viewpoints. We demonstrate how the range images can be accurately merged into a single triangular mesh with curvature dependent density by the use of local topological mesh operations. A new filter, that is specially adapted to the requirements of geometrical data, has been designed. This enables smoothing of measuring errors like noise, aliasing, outliers, and registration errors with minimum interference of real object features like edges. Curvature variations are minimized and surface undulations are avoided in order to produce high quality surfaces for rendering and NC milling.

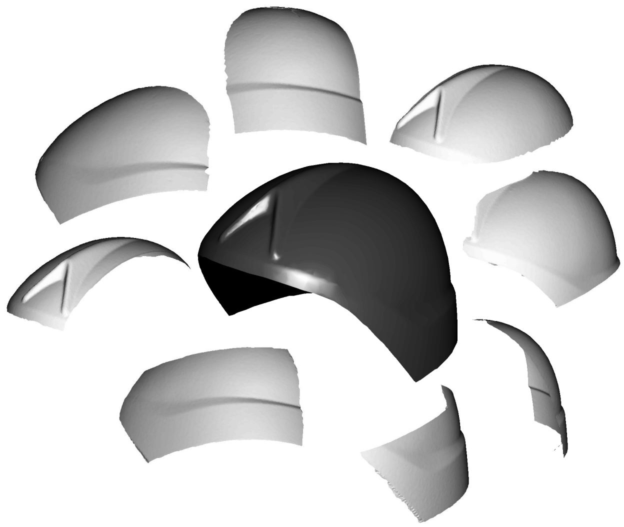

Figure 1: CAD-surface of a helmet, reconstructed from 8 range images.

1 Introduction

The last few years optical 3D-sensors have become powerful tools for *reverse engineering*. The shape of a three-dimensional object is sampled and turned into a CAD description. Figure 1 shows 8 single range images of a helmet and its reconstructed CAD surface (center). The reconstructed model can be handled like synthetic CAD data. This enables the processing of old design models on a computer. Using CAM techniques like NC milling or stereolithography, three-dimensional replicas of the digitized objects can be made. In dentistry such methods are used to scan teeth or plaster casts and to produce crowns and inlays from the data automatically.

Registered range images are not well suited to be used directly as input for CAD systems. They do not really describe surfaces, but clouds of point coordinates in 3D-space, in particular if multiple range images were taken from different views. The amount of data points may be very large (from millions to hundreds of millions). Furthermore the data are usually distorted by measuring errors like noise, aliasing, outliers, etc. Mainly two problems have to be solved for *surface reconstruction*: the reconstruction of the topology of the sampled object (*triangulation*) and the processing of the surface geometry in order to eliminate measuring errors and reduce data (*surface modeling*). At present, the most frequently used method is approximation of tensor product surfaces[1] (*reconstruction of smooth surfaces*). Unfortunately such methods require much interactive control. A simpler and more accurate way is to generate a polyhedral surface (e.g. a triangular mesh), which is sufficient for visualization (virtual reality) or CAM. Tensor product surfaces

[1] rectangular patches of spline or Bézier surfaces

are necessary only for "real" reverse engineering, where designers want to handle free form surfaces of digitized objects equal to synthetic surfaces (e.g. surface manipulation by control points).

We propose a new method for modeling range data by using meshes of curved or flat triangles with curvature dependent density. Our method allows to eliminate data distortions by measuring, calibration or registration errors with minimum interference of real object features. Curvature variations are minimized and surface undulations are avoided in order to produce high quality surfaces for rendering and NC milling or as a preprocessing step for "real" reverse engineering. A more detailed description can be found in [8].

2 Related Work

The reconstruction of the object topology from a cloud of sampled data points can be solved by means of graph theory. At present this approach has only little importance, as it is difficult to handle the amount of data, which are provided by modern optical 3D-sensors. Furthermore it is possible only to interpolate the measured data points exactly, but not to smooth errors. The best results are achieved using α-shapes [5] and γ-graphs [15].

Since a few years mainly volumetric approaches are used. These are based on well established algorithms of computer tomography like *marching cubes* [10] and therefore are easy to implement. They produce approximated surfaces, so that error smoothing is carried out automatically. Hoppe et al. [7, 6] and Curless and Levoy [2] have achieved good results. The method of Hoppe et al. is able to detect and model sharp object features. It generates thinned, CAD-incompatible meshes of curved triangles, which approximate the original point cloud with high accuracy. Unfortunately the computational costs allow only a few 10 000 points to be processed, even on fast work stations. The approach of Curless and Levoy is able to handle millions of data points. In contrast with the previous methods, it cannot process unstructured point clouds. It requires topological information provided by the matrix-like structure of range images. The results are nearly the same as for Hoppe et al., but sharp features are not modeled exactly. As neither curved triangles nor mesh thinning techniques are used, dense meshes containing a huge amount of small triangles are usually produced.

A new approach uses a special kind of deformable polyhedral mesh, the simplex mesh [3]. A coarse initial mesh is "shrunk" upon the point cloud, until the desired compromise of smoothness and accuracy is achieved. This method requires user interactions, in particular when surfaces with holes or objects with non-zero genus are to be modeled.

The usage of topology information provided by the range images enables faster algorithms and more accurate results. For that reason, researchers have proposed several methods for merging multiple range images into a single triangular mesh. The *mesh zippering* method of Turk and Levoy [14] generates dense meshes of flat triangles, whereas our approach produces meshes of flat or curved (but CAD-incompatible) triangles with curvature dependent density. Merging methods usually work incrementally. Furthermore, pure topology reconstruction without any interference of the data points is possible. On the other hand, special efforts for error smoothing are necessary. Our method includes an effective smoothing filter. In contrast to other surface reconstruction methods it is able to smooth single images without significant loss of details. All the other methods require redundant information. High quality smoothing is possible only in overlapping areas of different images.

Filters for smoothing polyhedral meshes without usage of redundant information are still in the state of intense research. Lounsbery [11] uses a generalization of a multiresolution analysis based on wavelets for this purpose. Unfortunately this approach works solely on triangular meshes with subdivision connectivity.[2] A filter that works on general meshes was proposed by Taubin [13]. He has generalized the discreet Fourier transformation, in order to realize low pass filters. However, translation of concepts of linear signal theory is not the optimal choice. Surfaces of three-dimensional objects usually consist of segments with low bandwidth and transients with high frequency between them. They have no "reasonable" shape, as it is preconditioned for linear filters. "Optimal" filters like *Wiener* or *matched* filters usually minimize the RMS-error. Oscilla-

[2] All vertices (with singular exceptions) have the same number of neighbors.

tions of the signal are allowed, if they are small. For visualization or milling of surfaces curvature variations are much more disturbing than small deviation from the ideal shape. A smoothing filter for geometric data should therefore minimize curvature variations and try to reinduce an error, that is smaller than the original distortion of the data. These are the requirements we considered, when we designed our new smoothing method.

Beyond that, for "real" reverse engineering the reconstruction of tensor product surfaces like NURBS[3] is necessary. Stoddart et al. [12] and Eck and Hoppe [4] have proposed solutions that work fully automatically. Unfortunately, such methods cannot model highly structured surfaces. The semi-automatic method of Krishnamurthy and Levoy [9] uses distance maps to overcome this limitation. It requires triangular meshes as input data. Therefore polyhedral surface reconstruction is needed for data preprocessing.

3 Overview

The new method requires "calibrated" range images as input data. These consist of a matrix of coordinate triples $(x, y, z)_{ij}$. They are arranged in the same manner as the original chip matrix (i, j). The sensor may be placed in arbitrary positions, the object surface may be sampled incompletely, and The sampling density may vary, but should be as high as possible. Beyond that, the object may have arbitrary shape, and the field of view may contain several objects. The images must be registered, which means that the different views must be adjusted to each other by a matching process. The following processing steps are used to turn this data into a single mesh of curved or flat triangles:

1. **Mesh Generation:** Because of the matrix-like structure of the range images, it is trivial to turn them into triangular meshes with the data points as vertices. For each vertex the surface normals are calculated from the normals of the surrounding triangles. If necessary, the normals are smoothed by weighted averaging. They are mainly used for interpolation of curved triangles and for curvature computation.

[3]Non Uniform Rational B-Spline

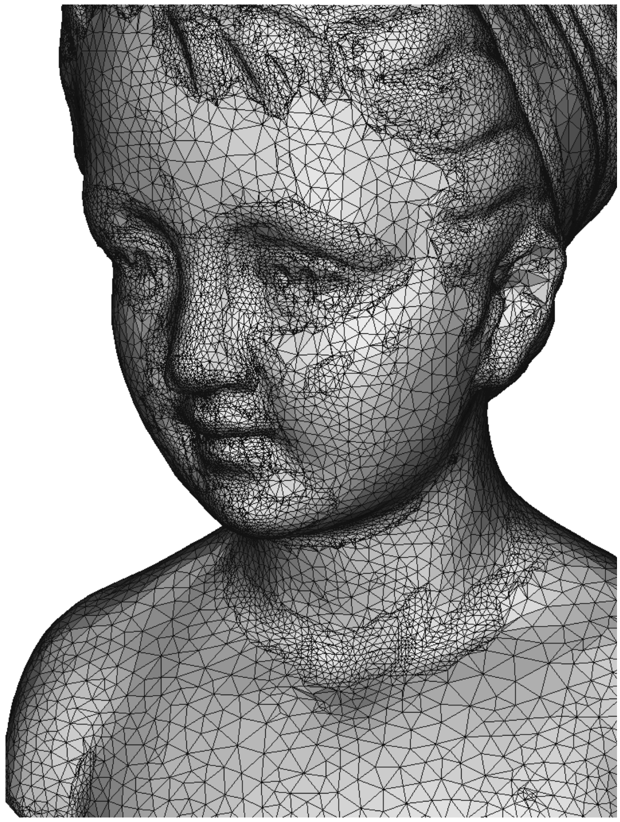

Figure 2: Thinned mesh of a bust with curvature dependent density. The permitted approximation error in highly curved areas (shoulders, forehead, cheeks, etc.) is less than in areas with low curvature.

2. **First Smoothing:** In order to utilize as much of the sampled information as possible, the smoothing of measuring errors like noise and aliasing is done before mesh thinning.
3. **First Mesh Thinning:** Ideally the complete data should be preserved for the following steps. Mesh thinning should be done until the end of the processing chain, but merging of dense meshes usually requires too much memory, so that mesh reduction must be carried out in advance. Therefore the permitted approximation error should be as small as possible. Figure 2 shows a thinned mesh with curvature dependent density.
4. **Merging:** The meshes from different views are merged by pairs using local mesh operations like *vertex insertion*, *gap bridging* and *surface growth* (see Fig. 3). Usually one starts with a master image, which includes as much object surface as possible. The other

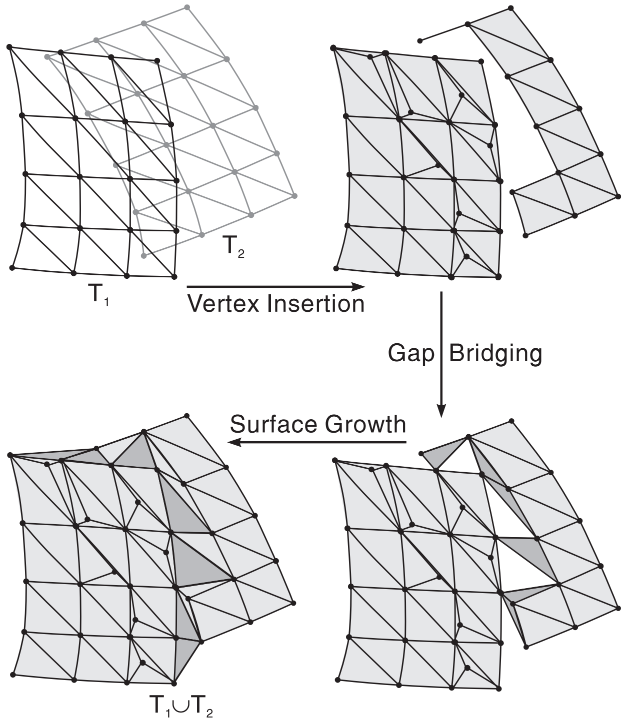

Figure 3: Merging of two meshes using *vertex insertion*, *gap bridging* and *surface growth* operations.

images are merged into the master successively. Only those new vertices are inserted, whose absence would cause an approximation error bigger than a given threshold.

5. **Final Mesh Thinning:** The mesh thinning is continued, until the given approximation error is reached.
6. **Geometrical Mesh Optimization:** The mesh thinning usually causes awkward distributions of the remaining vertices, so that elongated triangles occur. Geometrical mesh optimization moves the vertices along the curved surface, in order to produce a better balanced triangulation.
7. **Topological Mesh Optimization:** At last the surface triangulation is reorganized using *edge swap* operations (see Fig. 4), in order to optimize certain criteria. Usually, the interpolation error is minimized. If the results are to be translated into a mesh of Bézier triangles, a criterion is used, that avoids elongated triangles. Otherwise the computation of the Bézier points may become unstable.

The result of this process is a mesh of curved triangles. Our new modeling method is able to interpolate curved surfaces solely from the vertex

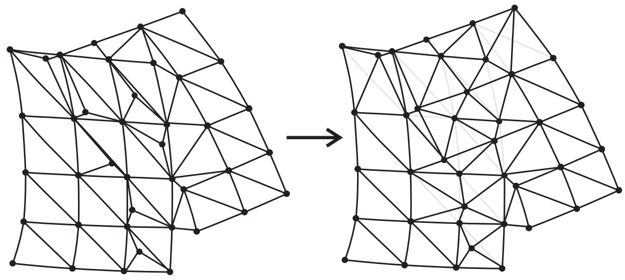

Figure 4: Topological optimization of the mesh from Fig. 3 using *edge swap* operations. Triangles, that are as equilateral as possible were aspired (*2D Delaunay triangulation*).

coordinates and the assigned normal coordinates. This enables a compact description of the mesh, as modern data exchange formats like Wavefront OBJ and VRML[4] support this data structure. Unfortunately all known external software utilizes vertex normals only for visual smoothing via Gouraud or Phong shading, not for geometrical interpolation.

4 Modeling of Scattered Data

Most of the errors that are caused by the measuring process (noise, aliasing, outliers, etc.) can be filtered at the level of raw sensor data. The theory of digital signal processing shows how to treat such grid data, therefore it is possible to derive filters which are good, or even optimal for certain purposes. A special class of errors (calibration and matching errors) first appear after merging of the different views. As the data are no longer represented on a grid, conventional filters for digital signal processing do not work. Unfortunately no basic theory exists for handling scattered data, hence smoothing is usually done by surface approximation. This approach requires time consuming user interaction, otherwise it would produce poor results because of smoothed object features like edges. We introduce a new approach based on the basic local characteristic of the sampled surface which is viewpoint independent: the curvature.

In zero order approximation it is assumed, that the sampling density is high enough to neglect the variations of surface curvature between adjacent sample points. If this is true, the underlying

[4] Virtual Reality Modeling Language

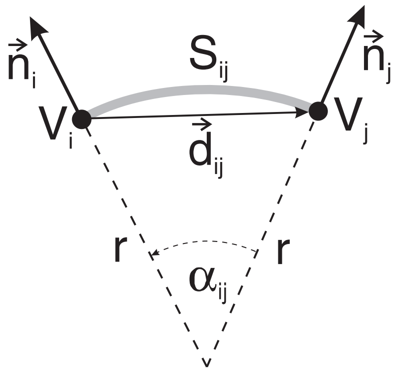

Figure 5: Cross section $\mathbf{S}_{ij}$ through a constantly curved surface.

surface can be approximated by a mesh of circular arcs. This simplified model provides a basis for all computations that our reverse engineering method requires: e.g. normal and curvature estimation, interpolation of curved surfaces, or smoothing of polyhedral surfaces.

As an example we show how easy curvature estimation can be, when using this model. Figure 5 shows a cross section $\mathbf{S}_{ij}$ through a constantly curved object surface between two adjacent vertices $\mathbf{V}_i$ and $\mathbf{V}_j$. The curvature c_{ij} of the curve $\mathbf{S}_{ij}$ is ($c_{ij} > 0$ for concave and $c_{ij} < 0$ for convex surfaces)

$$c_{ij} = \pm\frac{1}{r} \approx \pm\frac{\alpha_{ij}}{d_{ij}} \approx \pm\arccos(\vec{\mathbf{n}}_i \cdot \vec{\mathbf{n}}_j), \qquad (1)$$

which can be easily computed if the surface normals $\vec{\mathbf{n}}_i$ and $\vec{\mathbf{n}}_j$ are known. The principal curvatures $\kappa_1(i)$ and $\kappa_2(i)$ of $\mathbf{V}_i$ are the extreme values of c_{ij} with regard to all its neighbors $\mathbf{V}_j$:

$$\kappa_1(i) \approx \min_j(c_{ij}) \quad \text{and} \quad \kappa_2(i) \approx \max_j(c_{ij}). \qquad (2)$$

The surface normals are computed separately, hence it is possible to eliminate noise by smoothing the normals without any interference of the data points. Therefore this method is much less sensitive to noise then the usual method for curvature estimation from sampled data which is based on differential geometry [1]. It can be shown that approximation of a mesh of circular arcs requires a sampling density which is at least four times higher than the smallest object details to be modeled. This means that the minimum sampling rate must be twice as high as the theoretical minimum given by the Nyquist frequency. Therefore further investigations are necessary to extend the new modeling method to higher orders of curvature variations, in order to get closer to the theoretical limit.

5 Smoothing Polyhedral Surfaces

We now demonstrate how this new approach can be used for smoothing of measuring errors with minimum interference of real object features like edges. If curvature variations of the sampled surface are actually negligible, but the measured data vary from the approximation of circular arcs, this must be caused by measuring errors. Therefore it is possible to smooth these errors by minimizing the variations.

For this purpose a measure δ is defined to quantify the variation of a vertex from the approximation model. Figure 6 shows a constellation similar to Fig. 5. Now the vertex $\mathbf{V}_i$ is falsely measured at a wrong position. The correct position would be $\mathbf{V'}_{ij}$ if $\mathbf{V}_j$ and the surface normals $\vec{\mathbf{n}}_i$ and $\vec{\mathbf{n}}_j$ match the simplified model perfectly (There exist different ideal positions $\mathbf{V'}_{ij}$ for every neighbor $\mathbf{V}_j$). The deviation of $\mathbf{V}_i$ with regard to $\mathbf{V}_j$, given by

$$\delta_{ij} \approx d_{ij}\frac{\cos(\beta_{ij} - \frac{\alpha_{ij}}{2})}{\cos(\frac{\alpha_{ij}}{2})} \qquad (3)$$

can be eliminated by translating $\mathbf{V}_i$ into $\mathbf{V'}_{ij}$. The sum over all δ_{ij} defines a cost function for global minimization of the variations from the approx-

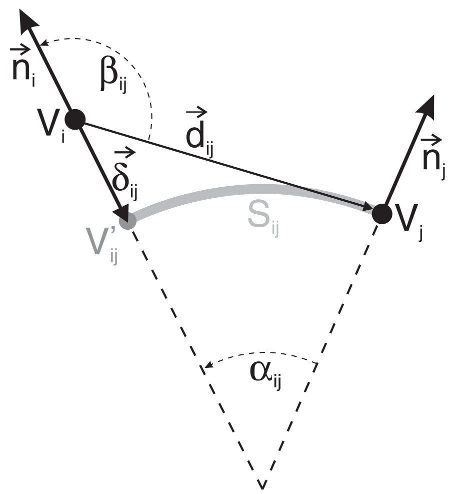

Figure 6: Cross section $\mathbf{S}_{ij}$ through a constantly curved surface. The position of vertex $\mathbf{V}_i$ is not measured correctly.

imation model over the whole surface. Minimizing that cost function leads to a mesh with minimum curvature variations for fixed vertex normals. This procedure can be used for surface smoothing if the surface normals describe the sampled surfaces more accurately than the data points. In case of calibration and matching errors the previous assumption is realistic. This class of errors usually causes local displacements of the overlapping parts of the surfaces from different views, while any torsions are locally negligible. Oscillating distortion of the merged surface with nearly parallel normal vectors at the sample points are the result. Figures 7 and 8 demonstrate that such errors can be smoothed without seriously affecting any object details.

In case of noise or aliasing errors (Moiré) the surface normals are also distorted, but can simply be smoothed by weighted averaging. So filtering is done by first smoothing the normals and then using the described surface filter to adapt the positions of the data points to these defaults (see Fig. 9).

6 Interpolation of Curved Surfaces

Interpolation of curved surfaces (e.g. curved triangles) can simply be done by interpolation between circular arcs. For that purpose, a new surface normal for the new vertex is computed by linear interpolation between all surrounding normals. The angles between the new normal and the surrounding ones define the radii of the arcs as it is shown in Fig. 5. Our method uses this simple interpolation scheme mainly for geometrical mesh optimization.

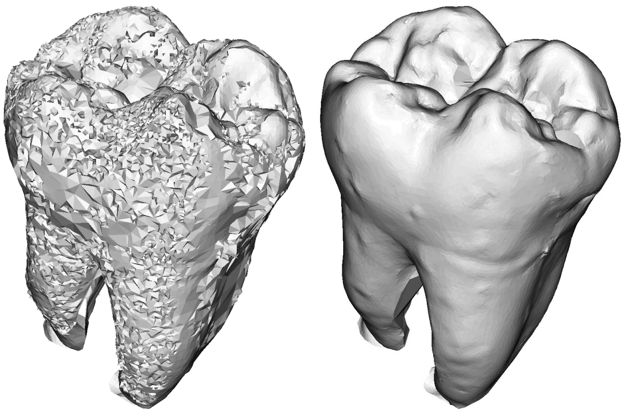

Figure 7: Distorted mesh of a human tooth, reconstructed from 7 badly matched range images (left), and the result of smoothing (right).

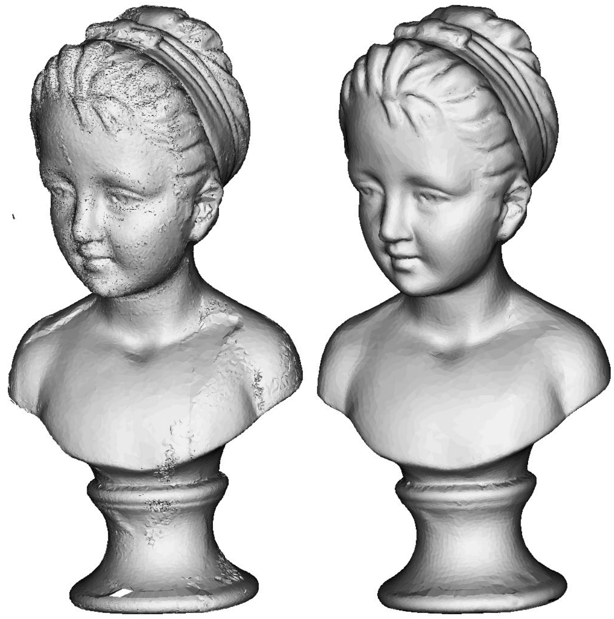

Figure 8: Smoothing of a mesh containing calibration errors (left: mesh after merging of 12 range images; right: result of smoothing).

7 Experiments

In our experiments it turned out that the practicability of any surface reconstruction method depends strongly on the efficiency of its smoothing algorithms. Our method works best in case of registration and calibration errors. Figure 7 and 8 demonstrate that such errors in fact can be smoothed without seriously affecting any object details. The mesh on the left side of Fig. 7 was reconstructed from 7 badly matched range images. The mean registration error is 0.14 mm, the maximum is 1.5 mm (19 times the sampling distance of 0.08 mm!). The mean displacement of a single vertex by smoothing was 0.06 mm, the maximum was 0.8 mm. The displacement of the barycenter was 0.002 mm. This indicates, that the smoothed surface is placed perfectly in the center of the difference volume between all range images.

In Fig. 8 the meshes from the front and backside of a ceramic bust do not fit because of calibration errors (left). The mean deviation is 0.5 mm, the maximum is 4.2 mm (the size of the

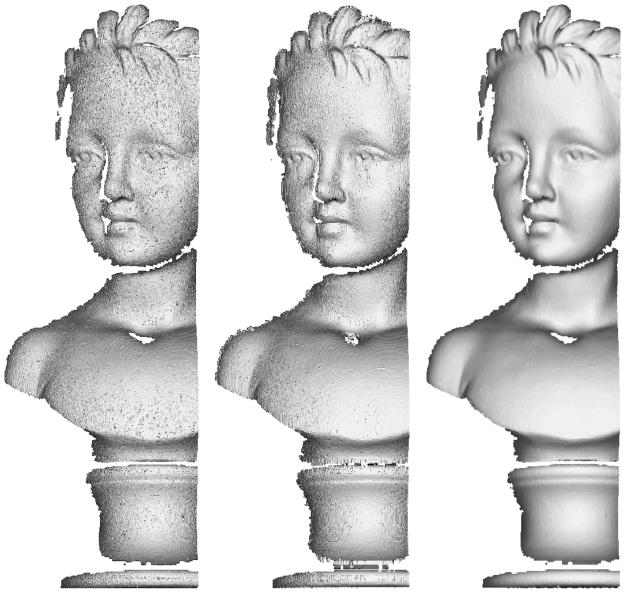

Figure 9: Noisy range image of the ceramic bust (left), smoothed by a 7×7 median filter (center), and by the new filter (right).

bounding box is $11 \times 22 \times 8$ cm^3). The mean displacement by smoothing was 0.05 mm, the maximum was 1.3 mm (right). In this example the different meshes were not really merged, but solely connected at the borders, so that a displacement, that was obviously smaller than the half of the distance between the meshes was sufficient.

Figure 9 demonstrates smoothing of measuring errors of a single range image in comparison to a conventional median filter (The simplest and most popular type of edge preserving filters). Although the errors (variations of the smoothed surface from the original data) of the median filter are slightly larger in this example, the new filter shows much more noise reduction. Beyond that, the median filter produces new distortions at the borders of the surface. The new filter reduces the noise by a factor of 0.07, whereas the median filter actually increases the noise because of the produced artifacts. The only disadvantage of our filter is a nearly invisible softening of small details.

It turned out that the errors introduced by the new filter are always less than the errors of the original data. In particular no global shrinkage, expansion or displacement takes place, a fact that is not self-evident when using real 3D-filters.

We have tested our method by digitization of many different objects, technical and natural as

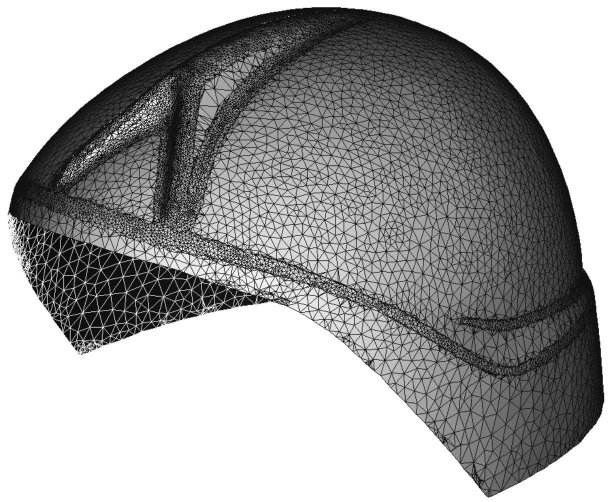

Figure 10: The triangular mesh of the helmet of Fig. 1.

Figure 11: The helmet that was produced from the CAD-data.

well. As an example we show the results for a design model of a firefighters helmet (Fig. 1, 10 and 11). The reconstructed CAD-data were used to produce the helmet. For that purpose, the triangular mesh was translated into a mesh of Bézier triangles,[5] so that small irregularities on the border could be cleaned manually. Eight range images containing 900 000 data points (11.6 MByte) were used for surface reconstruction. The standard deviation of the sensor noise is 0.03 mm (10% of the sampling distance), the mean registration error is 0.2 mm. On a machine with a Intel P90 processor and 64 MByte RAM, the

[5]This was done by software from the Computer Graphics Group.

surface reconstruction took 49 minutes. The resulting surface consists of 33 000 triangles (800 kByte) and has a mean deviation of 0.07 mm from the original (unfiltered) range data.

8 Conclusions

In our experiments the errors, that were reinduced by the modeling process, were smaller than the errors in the original data (measuring, calibration and registration errors). The new smoothing method is specifically adapted to the requirements of geometric data, as it minimizes curvature variations. Undesirable surface undulations are avoided. Surfaces of high quality for visualization, NC milling and "real" reverse engineering are reconstructed automatically. The method is well suited for metrology purposes, where high accuracy is desired. The complexity is limited by the merging algorithm, which needs at most $\mathcal{O}(n\sqrt{n})$ operations. For virtual reality faster methods may be more suitable.

The main disadvantage of our method is its requirement for matrix-like range images as input data. It does not work with unstructured point clouds, which are produced by some 3D-sensors (e.g. point sensors). Beyond that, watertight surfaces, which are required for stereolithography, are not produced automatically, as missing data is usually not reconstructed (if desired, holes may be closed by big triangles).

References

[1] P. J. Besl and R. C. Jain. Invariant surface characteristics for 3d object recognition in range images. *Computer Vision, Graphics, and Image Processing*, 33:33–80, 1986.

[2] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *SIGGRAPH 96 Conference Proceedings*, pages 303–312. Addison Wesley, August 1996.

[3] H. Delingette. General object reconstruction based on simplex meshes. Technical Report 3111, Institut National de Recherche en Informatique et en Automatique, France, February 1997.

[4] M. Eck and H. Hoppe. Automatic reconstruction of B-spline surfaces of arbitrary topological type. In *SIGGRAPH 96 Conference Proceedings*, pages 325–334. Addison Wesley, August 1996.

[5] H. Edelsbrunner and E. P. Mücke. Three-Dimensional alpha shapes. *ACM Transactions on Graphics*, 13(1):43–72, January 1994.

[6] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction. In *Proceedings of SIGGRAPH '94*, pages 295–302. ACM Press, July 1994.

[7] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Surface reconstruction from unorganized points. In *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 71–78, July 1992.

[8] S. Karbacher. *Rekonstruktion und Modellierung von Flächen aus Tiefenbildern.* Dissertation, Friedrich-Alexander-Universität Erlangen. Shaker Verlag Aachen, 1997.

[9] V. Krishnamurthy and M. Levoy. Fitting smooth surfaces to dense polygon meshes. In *SIGGRAPH 96 Conference Proceedings*, pages 313–324. Addison Wesley, August 1996.

[10] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 163–169, July 1987.

[11] M. Lounsbery. *Multiresolution Analysis for Surfaces of Arbitrary Topological Type.* PhD thesis, University of Washington, September 1994.

[12] A. J. Stoddart, A. Hilton, and J. Illingworth. Slime: A new deformable surface. In *British Machine Vision Conference*, pages 285–294, York, 1994.

[13] G. Taubin. A signal processing approach to fair surface design. In *SIGGRAPH 95 Conference Proceedings*, pages 351–358. Addison Wesley, August 1995.

[14] G. Turk and M. Levoy. Zippered polygon meshes from range images. In *Proceedings of SIGGRAPH '94*, pages 311–318. ACM Press, July 1994.

[15] R. C. Veltkamp. *Closed Object Boundaries from Scattered Points.* Gegevens Koninklijke Bibliotheek, Den Haag, 1992.