

MODEL REFINEMENT BY INTEGRATED PROCESSING OF LASER SCANNING AND PHOTOGRAHMETRY

Jan Böhm, Susanne Becker, Norbert Haala

Institute for Photogrammetry, Universitaet Stuttgart, Geschwister-Scholl-Str. 24D, D-70174 Stuttgart
first.last@ifp.uni-stuttgart.de

Commission V, WG V/4

KEY WORDS: LIDAR, Point Cloud, LASERMAP, Modeling, Facade, Urban

ABSTRACT:

For tasks such as the generation of realistic visualizations from pedestrian viewpoints in virtual reality and multimedia applications, urban models extracted from airborne data are too coarse and therefore have to be refined. Within the paper, terrestrial LIDAR data as well as facade imagery is used to increase the quality and amount of detail for the respective 3D building models. These models, typically available from airborne data collection, provide a-priori information, which can be integrated efficiently both for the georeferencing of the terrestrial data and the subsequent geometric and visual enhancement. We introduce two approaches to increase the geometric detail of facades. One approach is based on a planar mapping of point clouds, called LASERMAP, the other is based on cell decomposition.

1. INTRODUCTION

For a convincing visualization of complex architecture detailed three-dimensional recording is essential. Besides the coarse overall geometry of a building the facades are of special interest. Fine geometric details of windows, balconies, stonework and ornaments give each building its individual character. Modern laser scanners are able to acquire densely sampled point clouds of a facade consisting of several million points. Using state-of-the-art meshing techniques, though not without problems, these point clouds can be converted to polygonal meshes. However, the handling of huge polygon meshes is difficult both for rendering and further editing. Furthermore a dense polygon mesh is not grouped into meaningful entities, such as different architectural structures, for example wall, roof and so on.

Deriving such a grouping leads to the task of modeling, where geometric primitives are fit to segmented portions of the data. For TLS data this is typically an interactive process, involving the manual segmentation of the point cloud and the selection of appropriate primitives. This approach is currently applied for the representation of industrial scenes, where objects can be described with a reduced set of geometric primitives and graphical realism is of little interest. For the representation of general architectural scenes this approach of modeling does not seem suitable. For one the work load of manual interaction limits the approach to single buildings or even portions of a facade. Secondly most modeling environments for TLS data are limited to simple surfaces types such as ruled surfaces and seldom higher order curved surfaces are used. This set of geometric primitives is not suitable for the representation of the highly detailed facades of historical buildings.

In this paper we explore two different ways to model the facade of a historical building. Both approaches rely on the existence of a coarse building model that describes the overall geometry of the building in a polyhedral approximation. The first approach is to convert the point cloud generated by the scanner into a LASERMAP. The term is composed from two terms

describing the source of the data, a laser scanner, and the use of the data, as a displacement or bump map. The LASERMAP is a re-interpolation of the point cloud into a regular raster centered on the planar polygon describing a single portion of the facade very similar to a digital elevation model. Each LASERMAP computed is then mapped onto the facade polygon for rendering. The LASERMAP map can either be directly used as a displacement map, when a suitable rendering engine is available or a normal map can be computed using simple normal vector computation from a height field or a bump map can be applied. This approach preserves nearly all details of the original point cloud and can represent all two-and-a-half-dimensional features of a facade. It is ideally suited for visualization, direct interaction and virtual reality applications and is especially efficient on modern graphics hardware.

The second approach attempts to derive a true CAD description of facade primitives such as windows, doors and balconies. The approach is based on a cell decomposition process. The initial portion of the point cloud covering the portion of the facade which is to be modeled is segmented into cells belonging to the main plane of the facade and out-of-plane cells which belong to windows or other structures on the facade. Geometric primitives are then fitted to the out-of-plane cells. This approach achieves an effective segmentation of the facade into meaningful entities. Geometrically it is only limited by the number of primitives implemented. Combinations of the two approaches are possible, where CAD primitives are fit to the data where suitable primitives are available and LASERMAPs are used otherwise.

2. LASERMAPS FOR FACADE MODELLING

When terrestrial laser data is to be combined with pre-existing building models, the data has to be registered first. There are several methods to provide the registration, ranging from direct georeferencing, to manual and automatic alignment. Details of the approaches to registration and georeferencing of terrestrial laser data are described by Schuhmacher and Boehm (2005).

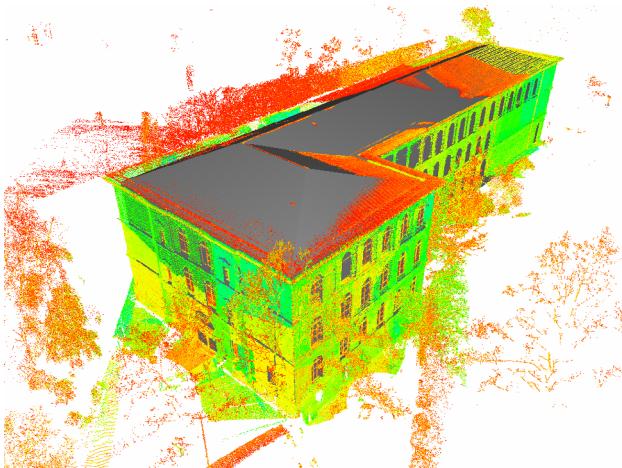


Figure 1. A coarse model of the building and the point cloud acquired by terrestrial LIDAR.

For the remainder of this paper we assume that the registration has been performed and the LIDAR data is in the same coordinate frame as the building model. We demonstrate this situation on a case study using the central historical building of the Universitaet Stuttgart as an architectural example in Figure 1. The point cloud was acquired with a Leica HDS 3000 from more than 15 stations. The data covers the facades of the building at a point density better than 20 mm. Special care was taken to cover all details of the facade and avoid shadowing of occluding objects.

To derive a combined model of the data, which is efficient for rendering it is helpful to look at the concepts for modeling developed in computer graphics. In computer graphics the duality of coarse over-all geometry and fine detail has long been noted. The separation of the two is a fundamental modeling principle. Starting with the observations of Blinn (1978), that the effect of fine surface details on the perceived intensity is “primarily due to their effect on the direction of the surface normal ... rather than their effect on the position of the surface”, modeling concepts were developed, which keep fine surface detail separate as a perturbation of the normal direction or a displacement to the underlying coarser geometry.

In the case of a building model the situation is rather simple, since the coarse geometry is described by polyhedral boundary representation and facades are typically planar polygons. This property is also exploited when intensity images are used for texture mapping. Texture coordinates map two-dimensional coordinates of the local planar coordinate system centered on the facade polygon to two-dimensional pixel coordinates of the texture image. In order to use this 2D-2D mapping for the integration of terrestrial LIDAR data, we have to derive a two-dimensional representation of the point cloud. Obviously this is not possible for the whole point cloud. We rather have to split the point cloud into groups with respect to the facades of the building using a simple buffer operation for each facade polygon.

The portion of the point cloud that belongs to a particular facade can now be interpolated into a regular raster, in a fashion very similar to digital elevation models derived from aerial LIDAR. We refer to such a re-interpolated point cloud as a LASERMAP (Böhm, 2005). The term is composed from two terms describing the source of the data, a laser scanner, and the use of the data as a source for 2D mapping. Figure 2 shows a

LASERMAP of the front facade of the aforementioned building. The gray values correspond to offsets relative to the plane of the facade. The map was computed at a resolution of 10 mm to preserve details, which gives an image of 2878 x 1778 pixels. Each pixel stores the offset in 16 bits.

2.1 Normal Map

The LASERMAP can be processed like any other image. One simple processing step is the generation of a normal map. The normal map stores the perturbations of the normal vector at each pixel due to surface irregularities. We can compute unit normal vectors from the partial derivatives of the surface functions as represented by the LASERMAP. This is easily done applying derivative filters to the LASERMAP. Figure 3 shows the rendering of a single facade polygon using a normal map derived from the above LASERMAP. While a normal map gives the impression of fine surface detail, this is only achieved by varying the shading of each output pixel, the actual geometry is still a flat polygon. This is advantageous as it does not increase the polygon count of the full, when compared to the original model. However, the flatness of the surface is revealed to the observer under very oblique viewing angles.



Figure 2. LASERMAP of a single facade derived from the point cloud.

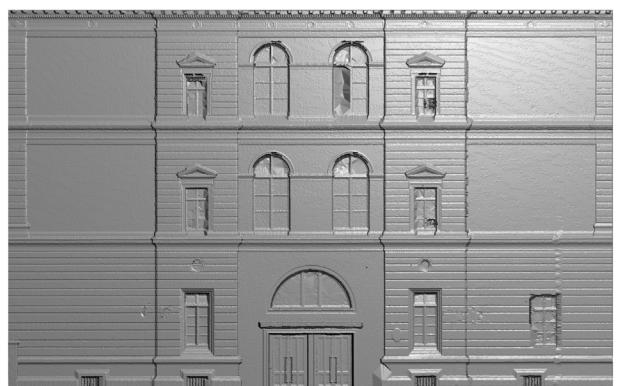


Figure 3. 3D Rendering of a single LASERMAP as a normal map.

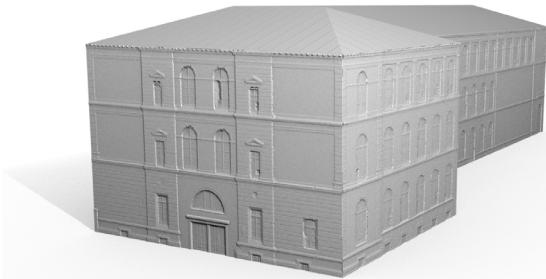


Figure 4. 3D Rendering of the complete model with full displacement maps.

2.2 Displacement Map

To generate true 3D geometry is more involved. In order to apply the offset values stored in the LASERMAP, first we have to generate vertices (triangles) which we can apply the offsets to. But instead of explicitly generating vertices and meshing them, we generate the vertices implicitly by using a subdivision surface (Catmull, 1974). A subdivision surface iteratively subdivides the faces of a given mesh, in our case the facade polygon, into smaller faces. Due to the recursive nature of subdivision surfaces, the number of generated faces is typically a power of two. After the subdivision is defined the LASERMAP can be applied as a displacement map. Figure 4 shows the model of the University building using displacement mapping. For the front facade mentioned above, we created 1024 x 512 faces using subdivision surfaces.

2.3 Automated Repair

It is a common problem in real-world terrestrial laser scanning projects, that scanning artifacts impair the quality of the point cloud. The artifacts can be created by simple occlusions, varying surface reflectivity, beam deflections and other typical problems. It can be very difficult to correct these situations directly in the point cloud, since we can only interact on individual points. Point cloud editing is typically limited to removing points. The situation improves with meshed point clouds, but still interactive editing can be very cumbersome and time consuming.

Here we show that using LASERMAPs, the removal of artifacts can be greatly simplified and can be partially automated. We capitalize on the fact that facade architecture does not consist of purely random geometry, but is composed of repetitive elements. If one such element is defect due to artifacts in the point cloud, we can simply replace it with a copy of an intact element. Since the representation of the LASERMAP is essentially the same as an image, image processing operations can be employed to aid in this task.

In Figure 5 we show an example of an automated repair process. In the rendering shown in Figure 3, one of the windows in the top row clearly has a defect, due to the window being half-opened at the time of scanning. A detailed view is shown in Figure 5 at the top. The respective detail of the LASERMAP is shown underneath. We start by interactively marking the defective area in the LASERMAP, shown as a white box in the third image. The repair process then

automatically searches for a similar area in the LASERMAP. This is done using a simple image matching technique. The best match is used as a replacement a copied onto the defect area. The result of the repair operation is shown in the third image and the respective rendering underneath.

The process is a quick and easy to use tool to remove defects caused by scanning from the facade model. While it is limited to repetitive structures, it can cope with any kind of shape and is not limited to certain geometries.

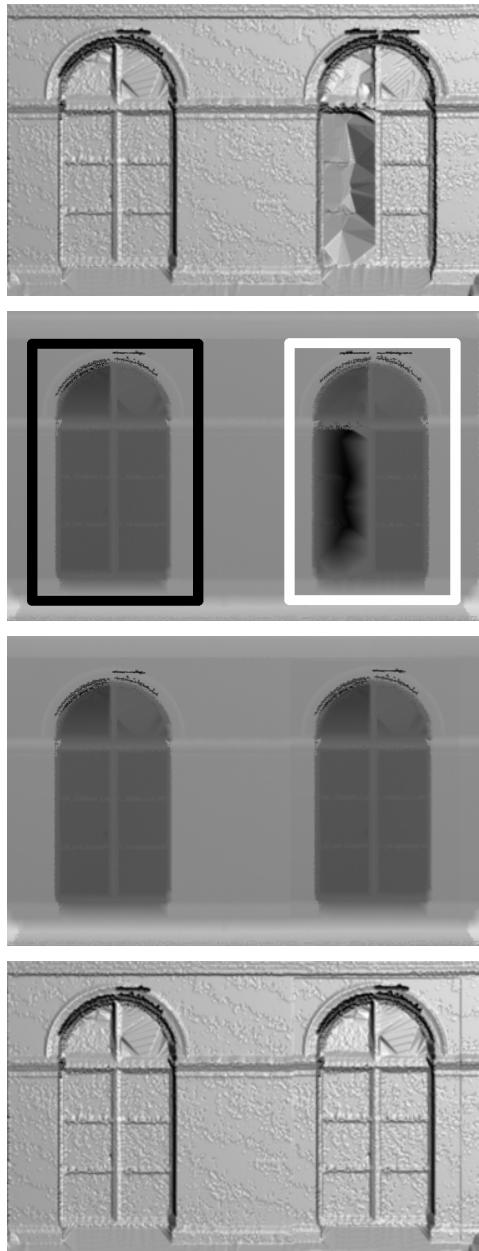


Figure 5. Semi-automatic repair of an artifact (a half-opened window). A detail of the original 3D rendering is shown on top. The interactive marking of the defect (white) and the automated match (black) of a replacement is shown second, followed by the result of the automated repair. The bottom image shows the 3D rendering using the repaired LASERMAP.

The concept of LASERMAPs is a very efficient representation of detailed facade geometry, which is ideally suited both for high-quality rendering and interactive visualization. However the resulting model is tied to the resolution of the LIDAR point cloud. When the viewers distance is very close to the surface the limited resolution becomes apparent. To overcome these limitations in resolution, explicit CAD modeling is unavoidable.

3. GENERATION OF CAD DESCRIPTION FOR REFINED 3D BUILDING MODELS

Similar to the application of LASERMAPs, the generation of a true CAD description of facade primitives such as windows, doors and balconies also benefits from a close integration of existing building models as they are provided from airborne data collection and terrestrial data. Again, this avoids data collection from the scratch and facilitates the modeling process.

Usually, CAD descriptions of 3D building models are based on representations by constructive solid geometry (CSG) or boundary representation (B-Rep). In contrast, we apply a representation of the buildings by cell decomposition (Haala et al 2006). By these means, the problems to correctly implement the generation of topological correct boundary representations for the polyhedral building models can be avoided. These problems, which are well known from solid modeling (Hoffmann 1989), especially occur, if geometric constraints, such as meeting surfaces, parallelism and rectangularity have to be guaranteed for measured and thus error-prone data in the context of building reconstruction (Brenner 2004). The formulation of such regularization conditions is also simplified, if an object representation based on CSG is used. However, while CSG is widely used in computer aided design since it allows for powerful and intuitive modeling (Mäntylä 1988), most visualization and simulation applications require the additional derivation of a boundary representation. While this is conceptually easy, its correct and efficient implementation can be difficult. Problems can arise from error-prone measurements, limited numerical precision and unstable calculation of intersections.

Within our approach based on cell decomposition, complex solids are described similar to CSG by a combination of relatively simple, basic objects in a bottom up fashion. In contrast to CSG, which combines simple primitives by means of regularized Boolean set operators, decomposition models are limited to adjoining primitives. Since the basic primitives must not intersect, they are thus ‘glued’ together to get the final model. In this sense, cell decomposition is similar to a spatial occupancy enumeration, where the object space is subdivided by non overlapping cubes of uniform size and orientation. However, cell decompositions are based on a variety of basic cells, which may be any objects that are topologically equivalent to a sphere i.e. do not contain holes. This allows for a simplified combination of the respective elements, while the disadvantages of exhaustive enumeration like large memory consumption and the restricted accuracy of the object representation can be avoided.

Our algorithm for the generation of CAD descriptions is a two-step approach based on terrestrial LIDAR and image data. It aims on the generation of explicit 3D facade geometry for an existing coarse 3D building model by the integration of window objects. At first, cell decomposition is used to model windows and doors from the LIDAR data. In a second step, the window

frames are further refined by photogrammetric analysis of additional image data.

3.1 Facade Structure from TLS

Within our reconstruction algorithm the flat front face of a 3D building is separated into 3D cells. Each 3D cell represents either a homogeneous part of the facade or a window area. Therefore, they have to be differentiated based on the availability of measured LIDAR points. After this classification step, window cells are eliminated while the remaining facade cells are glued together to generate the refined 3D building model.

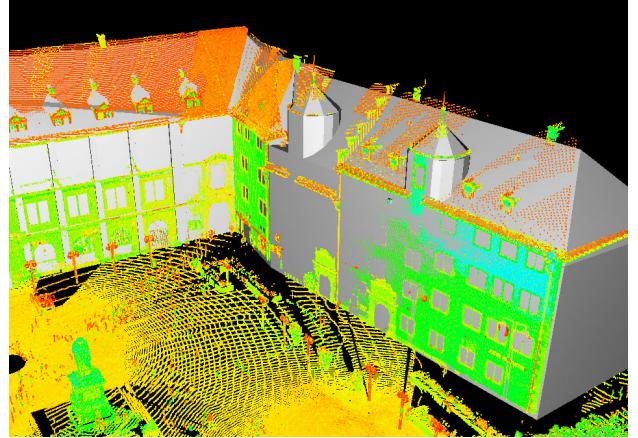


Figure 6. 3D point cloud from laser scanning aligned with a virtual city model.

Figure 6 shows an overlay of the measured TLS data and the coarse CAD model from an existing city model. As it is visible, fewer 3D points are measured on the facade at window areas. This is due to specular reflections of the LIDAR pulses on the glass or points that refer to the inner part of the building.

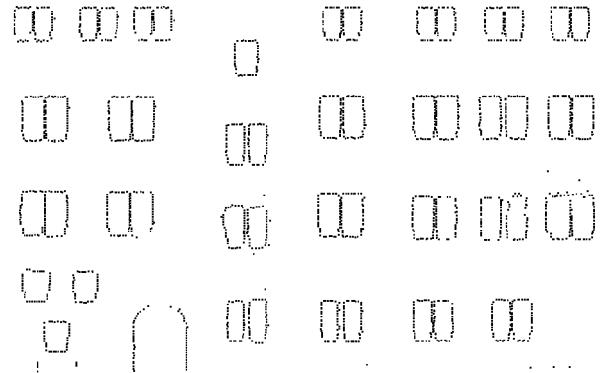


Figure 7. Detected edge points for a planar facade polygon at horizontal and vertical window structures.

Thus, as it is demonstrated in Figure 7 our point cloud segmentation algorithm detects window edges by searching for areas with no point measurements. In principle, such holes can also result from occlusions. However, this is avoided by using point clouds from different viewpoints. In that case, occluding objects only reduce the number of LIDAR points since a number of measurements are still available from the other viewpoints. Our segmentation process differentiates four types of window borders: horizontal structures at the top and the bottom of the window, and two vertical structures that define the left and the right side. As an example, the edge points of a left window border are detected if no neighbor measurements to

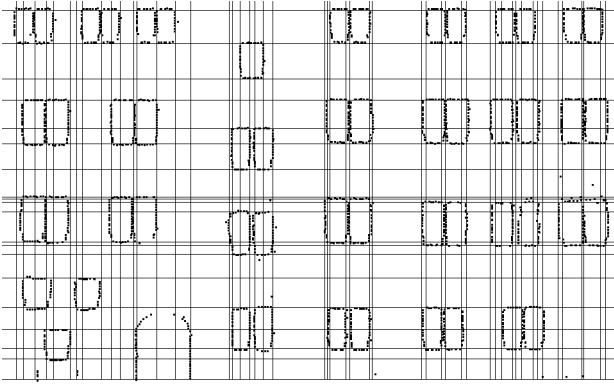


Figure 8. Detected horizontal and vertical window lines.

their right side can be found in a pre-defined search radius, depending on the expected point density at the facade plane. Based on the extracted edge points, which are again depicted in Figure 8, the window borders are then determined in the following step by estimating horizontal and vertical lines from non-isolated edge points.

As it is also demonstrated in Figure 8, a decomposition of the building facade in suitable cells can be generated from these boundary lines. Each of these cells either defines a homogeneous part of the facade or a window area. The discrimination between window and facade cells can be realized based on the availability of measured LIDAR points. After this classification window cells, i.e. areas with no point measurements, can be eliminated from the facade. Within this step, neighborhood relationships and constraints concerning the simplicity of the resulting window objects can be realized. As an example, uncertain cells are classified depending on their neighbors in order to align and adapt proximate windows in horizontal and vertical direction and to guarantee convex window objects.

As it is depicted in Figure 9, the facade geometry is modeled by eliminating the classified window cells from the existing coarse building model. For this purpose, a plane parallel to the facade at window depth is determined from LIDAR points measured at the window crossbars. The window cells are then carved out from the building model at this window depth.

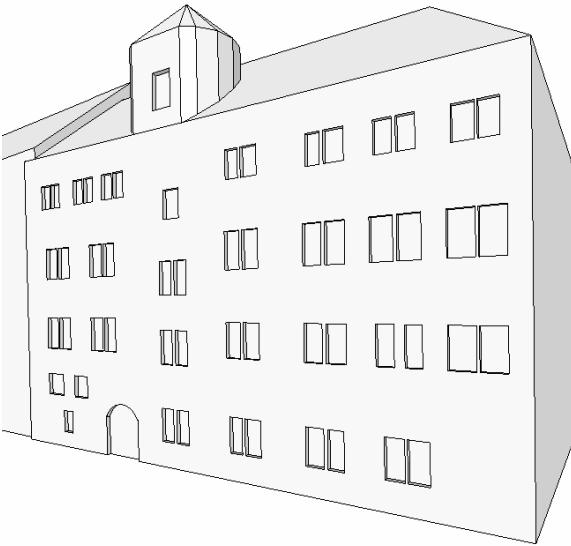


Figure 9. Refined facade of the given building model.

While the windows are represented by polyhedral cells, also curved primitives can be integrated in the reconstruction process as demonstrated by the round-headed door of the building. Furthermore, our approach is not limited to the modeling of indentations like windows or doors and can be used to add details as protrusions to the facade.

3.2 Refined Facade Structure from Stereo Image Analysis

Usually, the limited point sampling distance of TLS restricts the amount of geometric detail which can be provided by the reconstruction process. Thus, additional 3D information is optionally captured by our system from stereo images. As an example window crossbars can be reconstructed from linearly shaped facade detail. These objects are provided from edge points, which are extracted and matched by a suitable process. First, a standard Sobel filter is used for edge extraction, which are thinned and split into straight segments. Afterwards, the resulting 2D edges of the stereo image pair are matched to generate the required 3D information.

However, this process is hindered by the frequent occurrence of such facade structures, which results in a large number of mismatches. Therefore, only these 2D edges are further processed, which are within windows regions already detected within the LIDAR data. Figure 10 depicts selected 2D edges for an exemplary window in both images, which are used to generate 3D primitives by the matching step. Possible mismatches will result in 3D edges outside the available window area. Therefore, these wrong edges can be easily identified and removed. In addition, only horizontal and vertical 3D edges are considered for the further reconstruction process.

For each window, hypotheses about the configuration of the crossbars are generated and tested against the 3D edges derived from the images. Possible shapes are dynamically generated as templates by recursively dividing the window area in two or three parts. Recursion stops when the produced glass panes are too small for a realistic generation of windows. The minimum width and height of the glass panes are restricted by the same threshold value. After each recursion step, the fitting of the template with the 3D edges is evaluated. The partition is accepted if 3D edges are available within a buffer area around the dividing line. In a final step, the crossbars and the window frame are modeled. For this purpose, new 3D cells with a pre-defined thickness are generated at the accepted horizontal and vertical division lines as well as at the window borders. The result is exemplarily shown for two windows in Figure 11.

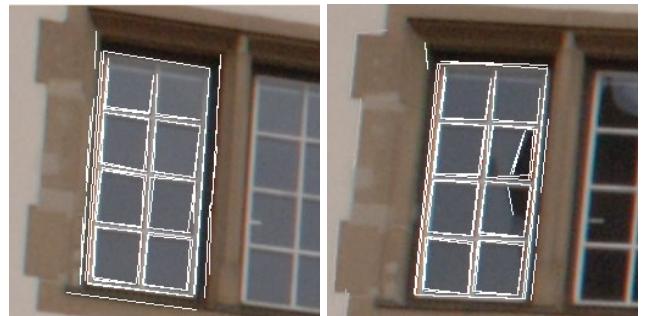


Figure 10. Selected 2D edges for an exemplary window in both images.

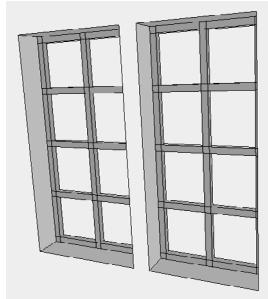


Figure 11. Reconstructed crossbars for two windows.

In order to stabilize the modeling process of crossbars, neighborhood relationships are taken into account. The crossbar configuration is assumed to be equal for all windows of similar size which are located in the same row or column. Based on this assumption, similar windows can be simultaneously processed. Thus, the crossbar reconstruction leads to robust results even for windows that are partially occluded or feature strong perspective distortions in the respective image areas.

The final result of the building facade reconstruction from terrestrial LIDAR and photogrammetric modeling can be seen in Figure 12. This example demonstrates the successful detection of crossbars for windows of medium size. In addition the dynamic generation of templates even allows for the modeling of large window areas as they often occur at facades of big office buildings.

4. CONCLUSIONS

The integration of terrestrial LIDAR data as supplementary data with existing 3D building models from airborne data collection is especially suitable for the creation of highly detailed city models. We have demonstrated two different approaches to the integration. LASERMAPS are an easy and simple modeling approach. They are especially suitable for advanced rendering engines and modern graphics hardware. They can capture any two-and-a-half-dimensional shape. We have introduced a simple repair mechanism to remove scanning artifacts. Cell decomposition derived CAD models provide geometric detail at all zoom levels. Our approach includes photogrammetric imagery into the modeling process to generate true CAD models for further handling.

ACKNOWLEDGEMENTS

We acknowledge the work of Angelos Tsakiroglou, Thanasis Georgantas and Christos Stamatopoulos of the Aristotle University of Thessaloniki, who have acquired the laser scan data of the University building during their stay in Stuttgart.

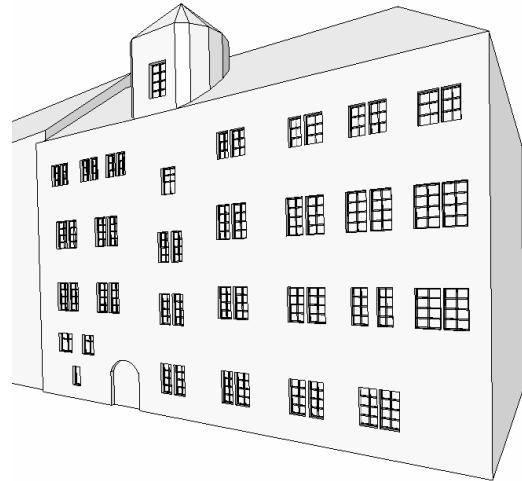


Figure 12. Refined facade with detailed window structures.

REFERENCES

Catmull, Ed. A., 1974. Subdivision Algorithm for Computer Display of Curved Surfaces. PhD thesis, Computer Science Department, University of Utah, Salt Lake City, UT, Report UTEC-CSc-74-133.

Blinn, J. F., 1978. Simulation of wrinkled surfaces. *ACM SIGGRAPH Computer Graphics* 12(3), pp. 286-292.

Böhm, J., 2005. Terrestrial Laser Scanning - A Supplementary Approach for 3D Documentation and Animation, Photogrammetric Week '05, Fritsch (ed.), Wichmann, 263-271.

Brenner, C., 2004. Modelling 3D Objects Using Weak CSG Primitives. IAPRS Vol. 35.

Haala, N., Becker, S. & Kada, M., 2006. Cell Decomposition for the Generation of Building Models at Multiple Scales. IAPRS Vol. XXXVI Part III, Symposium Photogrammetric Computer Vision, pp. 19-24.

Hoffmann, C.M., 1989. *Geometric & Solid Modelling*. Morgan Kaufmann Publishers, Inc., San Mateo, CA.

Mäntylä, M., 1988. *An Introduction to Solid Modeling*. Computer Science Press, Maryland, U.S.A.

Schuhmacher, S. & Boehm, J., 2005. Georeferencing of terrestrial laserscanner data for applications in architectural modeling 3D-ARCH 2005: Virtual Reconstruction and Visualization of Complex Architectures, XXXVI, PART 5/W17.