

Terrestrial Laser Scanning - A Supplementary Approach for 3D Documentation and Animation

JAN BÖHM, Stuttgart

ABSTRACT

While aerial methods undisputedly form the backbone of modern photogrammetry, we have seen an increased use of terrestrial methods in the past years to supplement aerial data especially in the field of 3D city modelling and the animation thereof. This article describes an approach to use terrestrial laser scanning in order to increase detail and realism of datasets derived from aerial methods.

When devising a new strategy for the integration of data from multiple sources two aspects come to focus. For one methods for the registration of the datasets are needed. Secondly a suitable data model has to be introduced to hold multiple aspects of the data. Consequently we present our methods for the automated georeferencing of terrestrial laser data and we introduce the concept of LASERMAPs, which are extracted from terrestrial laser data and mapped onto the 3D building model.

1. INTRODUCTION

Airborne methods, either aerial photogrammetry or airborne LIDAR, are ideally suited for mapping and they are the method of choice when large areas need to be covered. The products derived from airborne sensors such as orthophotos and DEMs are well established in the geodata community. With the emergence of new computer graphics hardware and the broadening of the user community in the past years the desire grew for new methods of presenting geodata using virtual reality technology. Especially in virtual city modeling, where walkthroughs and real-time animations are commonly encountered, the discrepancy of the birds-eye view of data collection to the first-person view used for the presentation of the data is apparent. Many of the details needed to create a realistic visual impression are missing, when only aerial sensors are employed. It is therefore inevitable to add data collected by terrestrial methods to supplement aerial data.

A major idea for the enhancement of the visual impression of objects represented by simple geometry has been presented by Catmull (1974). He proposed to use a mapping of intensity values to the surface representation in order to suggest a higher level of detail to the viewer. Texture mapping today has become one of the primary tools in virtual city modeling to increase the level of realism, as will be described in further detail below.

As well known as the method itself are its shortcomings. Since texture maps only change the intensity values of the surface, the geometric appearance is still given by the underlying coarse geometry. For virtual city modeling this means that the facades are still planar surfaces. While this fact can be ignored when the surface is rendered in an orthogonal view, it becomes more evident the lower the angle of view is to the surface. A second shortcoming of texture maps is related to the illumination of the object. Textures acquired by real images capture the objects intensity properties at the moment the image is taken. Therefore texture maps are bound to the lighting conditions encountered at the moment of the acquisition. At the rendering stage artificial light sources can be introduced at arbitrary locations. If these locations do not correspond with the original image the textures appear unrealistic. Furthermore, any illumination effects in the texture map, that were caused by the interaction of the direction of the light source and the surface orientation, will appear unrealistic when the model's orientation is changed with respect to the light source.

In order to overcome these shortcomings it is obviously desirable to add geometric detail to the models. Three-dimensional geometric detail at very fine resolution can ideally be captured using

terrestrial laser scanning (TLS). In order to be able to integrate TLS data with existing data obtained from airborne methods a crucial step preceding the data fusion is the geometric alignment of the separate dataset into a common reference frame. In the case of geodata fusion this is known as georeferencing. Our strategy for automated georeferencing of TLS data is given in section 3. Our data model for the integration of terrestrial and aerial data, which we call LASERMAPs, is introduced in section 4. To prepare for this data model and to emphasize the continuity in the processing pipeline we first quickly summarize on the current technology for 2D texturing of facades in the next section.

2. 2D FACADE TEXTURING

A large number of systems exist for the 3D modeling of urban environments. Based on either aerial imagery or laser scan data they usually create polyhedral descriptions of individual buildings (see Figure 1 on the left). Texturing these polyhedral structures greatly enhances their visual realism and suggests a much higher level of detail to the viewer as can be seen in Figure 1 on the right. When aerial imagery is used, roof structures can be textured automatically. For the texturing of the vertical facades terrestrial imagery is needed.

Facade texturing has received some attention in both the photogrammetric and the computer vision community. Approaches for manual, semi-automatic and fully automated texture extraction have been presented in the past, see for example Coorg and Teller (1999) and Wang et al. (2002). Some of the systems combine the reconstruction process with the texture extraction, see for example Bornik et al. (2001).

The simplest approach is to use a single image and warp this image onto the planar facade using the perspective transformation. Four points in both the planar coordinate system of the facade and the corresponding points in the image have to be determined. Assuming that all facades are bound by a rectangular curve, the approach can be further simplified and only the four corner points of the facade have to be identified in the image and no control information is required. The advantage of such an approach is that it does not require the determination of the exterior orientation and neither does it require the calibration of the camera. Furthermore no control points in object space have to be determined. The idea to use a two-dimensional rasterized map to store the values, that are to be mapped onto the surface, has shown to be quite successful and several hundred buildings have been



Figure 1: Facade textures can compensate for lack of geometric detail. The improvement in visual appearance can be seen from the above two views on a virtual city model, where the right view is based on the exact same geometry and only texture was added.

textured using several thousand terrestrial images (Kada et al., 2003). Many tools have been developed both in the research community and in the commercial domain to facilitate the work. Most modern GIS systems support the concept of storing such a map for every single face of a polyhedral description of an object. When trying to use three-dimensional data to add detail to the polyhedrons of a virtual city model this concept should be re-used in order to take full advantage of the existing processing pipeline. However before three-dimensional point clouds can be mapped onto facades the problem of alignment has to be solved, as is explained in the next section.

3. GEOREFERENCING OF TERRESTRIAL LASERSCANNING DATA

Terrestrial laserscanning has become popular for the acquisition of architectural scenes, due to the fact that it provides direct, reliable and dense surface measurement in a versatile fashion at independently varying stand-off distance and resolution. In order to integrate data recorded at multiple stations with an existing dataset, such as a virtual city model, georeferencing of TLS data is essential. The common practice for georeferencing of TLS data, as suggested by the manufacturers, requires three-dimensional control points, preferably provided by specially designed targets. However in surveying position and height are recorded separately due to historical reasons and therefore three-dimensional control points in a geodetic system are not naturally available. Georeferencing of TLS data typically is achieved by introducing a supporting measurement system, for example a total station, to transfer the geodetic control information onto the special targets. This process is regarded as the most precise way to give a geodetic datum to TLS measurements. The actual steps of the procedure are beyond the scope of this paper, it shall be noted however, that this process does not only cost extra time and effort, it also requires the user to provide a second measurement system and to be able to operate it. It is therefore of great interest to study alternative methods to establish a georeference in TLS data. We present a two-stage process for georeferencing based on integrated sensor measurement and iterative refinement.

The sensor-driven approach to georeferencing has established itself for quite some time in aerial photogrammetry, where it is known as direct georeferencing. Sensors for the direct measurement of position, typically GPS sensors, and sensors for the measurement of attitude angles - inertial measurement units - are integrated with the image measurements. This integration is either considered the final result, or it is seen as an initial value for further iterative refinement including additional control information. Considering the high dynamic characteristics of in-flight measurement, typically high-end components are used in aerial photogrammetry. For ground-based applications the price of these high-end components will often exceed the total budget for the sensor. Therefore low-cost components receive attention in many terrestrial applications. The direct orientation of TLS systems has been investigated mainly for mobile mapping applications. An example is the work by Talaya et al. (2004). In contrast the applications we are investigating in this work are concerned with stationary measurements, and therefore the system design is closer related to the approach described by Fritsch et al. (2000). We describe the integration of a GPS receiver and a simple digital compass for the direct orientation of a TLS system.

3.1. Sensor-driven Georeferencing

The advantage of sensor-driven methods for georeferencing laserscanner data is the possibility to save time and cost due to direct measurement and automation in determining the global position and orientation of every station. Compared to the conventional method of georeferencing with the use of a total station, the sensor-driven approaches require less additional equipment and are not in need of existing control points around the scene.




Figure 2: Left: Integrated orientation device, consisting of a low-cost GPS and a compass and pan / tilt sensor. The sensors have a wireless connection to a laptop via Bluetooth. Right: The system mounted on top of a HDS 3000.

A simple and inexpensive possibility for directly determining the global position of the laserscanner is to use a low-cost GPS. The GPS receiver is centered close to the laserscanner's rotation axis.

During the scanning process, the GPS position is simultaneously recorded. The results are available in real time. The expected accuracy of this absolute positioning is about 10 m and better, when a correction signal is applied.

The GPS receiver only measures the position of the stations. The orientation has to be determined by an additional procedure. One possibility is to use a digital compass and tilt sensor. The compass is co-located with the low-cost GPS. The whole construction is illustrated in Figure 2.

3.2. Iterative refinement

Since the alignment obtained from direct georeferencing is not sufficient for many applications this solution needs further refinement. This can be accomplished using the iterative closest point (ICP) algorithm introduced by Besl & McKay (1992). The result of the direct georeferencing is used as an initial value for the iterative registration of the laser scans. Once the registration of the TLS data has converged, it is kept fixed. Then the complete dataset is registered with the city model or airborne LIDAR using the same algorithm. Since the initial approximation of the direct georeferencing is within the convergence radius of the ICP algorithm this approach allows for an automated georeferencing of TLS data.

Figure 3 depicts the super-imposition of the terrestrial laserscanner data, which was directly georeferenced by the low-cost components to the available 3D building model. Based on this information, the full alignment could be computed. More details and variations of our approach on automated georeferencing of TLS data are published in Schuhmacher & Böhm (2005).




Figure 3: An example for automated georeferencing of TLS data. Top left: The unregistered point clouds collected at five stations. Top right: Results obtained from the integrated orientation device. Bottom left: Result after iterative improvement using ICP. Bottom right: Super-imposition of georeferenced TLS data and a single building from a virtual city model.

4. 3D FACADE TEXTURING

A modern laserscanner is able to acquire densely sampled point clouds 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, immediately suited for graphic rendering. While dense polygon meshes might be a suitable representation for applications in computer graphics, in the world of geomatics this representation lacks desirable features. Let alone the size of the data makes it unrealistic to derive such a representation for even a part of a city. 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 workload of manual interaction limits the approach to very small scenes or even single buildings.

Secondly most modeling environments for TLS data are limited to simple surfaces types such as ruled surfaces, seldom higher order curved surfaces are used. This set of geometric primitives is not suitable for the representation of highly detailed ornaments of historical buildings, which are of particular interest in any city model. This discrepancy in the ability to capture fine detail and the representation of surfaces in a model has long been noted in computer graphics. A solution, which is often chosen, is to separate the representation of coarse geometry from the representation of fine surface detail as shown in Figure 4 (a) to (b).

An important idea in the field of computer graphics addressing this issue was introduced by Blinn (1978). He observed 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”. This idea is a first step in separating the coarse and fine geometry of a surface. While the fine details are not modeled as explicit geometry, their perturbation on the direction of the surface normal is recorded as shown in Figure 4 (c) and introduced to the computation of surface intensity. This idea of generating normal maps to enhance the realism of computer-generated surfaces is known as bump mapping and has become a standard in computer graphics. On nearly all up-to-date computer graphics card this feature is implemented in hardware on so called graphics processor units (see for example Kilgard, 2000), and therefore ideally suited for high performance rendering and real-time animation.

In order to capture the true geometry of the surface on top of the coarse geometry, displacement mapping introduced by Cook (1984) can be used. For this modeling method the difference of the true geometry to the simplified representation is stored in a separate displacement map (see Figure 4




Figure 4. The true geometry of a surface (a) is separated into its coarse geometry (b), the normal map (c) and the displacement map (d) which contains the difference of (b) to (a).

(d)). In the form of a height field the displacement map can be handled like a color texture image. This approach has been particularly successful in applications where a representation based on NURBs is desired (Krishnamurthy and Levoy, 1996) and in applications where automated mesh simplification is performed but geometric detail shall be preserved (Cignoni et al, 1999).

In the case of adding detail to a virtual city model, the coarse geometry is already given by the model extracted from aerial LIDAR data or photogrammetry. The coarse geometry is represented by the planar polygons of the boundary representation of a building. The task is to use the terrestrial laser data to add the detail. Since terrestrial laser scanners capture mainly the vertical structures, we concentrate on the facades of buildings. Following the approach described above, we aim to extract the displacement map of the CAD model surfaces to the terrestrial point cloud. In order to be able to bring the terrestrial laser data into relationship to the city model, the laser data has to be georeferenced as described in the previous sections. Once the data is brought into a common geometric reference frame, the difference of the point cloud to a planar facade can be computed easily, as is shown in Figure 5.

Since we assume a multi station configuration for the acquisition of terrestrial data, we cannot restrict ourselves to ordered point clouds. Rather we have to work with an unordered point cloud where no neighborhood relation is available. We therefore re-interpolate the portion of the point cloud within a facade's buffer region to a regular raster centered on the plane of that facade. An example of such a re-interpolated displacement map is given in Figure 5. The advantage of the re-interpolation mechanism is that we do not have to establish the neighborhood relation within the point cloud for example by triangulating the point cloud, which can be a complicated process. The disadvantage is that we are restricted to $2\frac{1}{2}$ D structures. However it is common to assume, that facades can be sufficiently described by a relief.

We call such a re-interpolated laser point cloud 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 map can either be directly used as a displacement map, when a suitable rendering engine is available or a bump map can be computed using simple normal vector computation from a height field. The application of the LASERMAP as a bump map for the CAD model is shown in Figure 6.




Figure 5. Left: Multiple georeferenced point clouds from TLS aligned with a virtual city model. Right: The LASERMAP of a single facade extracted from the data.




Figure 6. The LASERMAP applied as a bump map to the CAD model of the building. Note the change in the shading of the surface due to the different direction of the light also visible on the sphere added in the upper right corner.

5. CONCLUSIONS

Frequently, terrestrial lasers scanning is used to collect building facades at a large amount of geometric detail. The integration of TLS data as supplementary data with existing 3D building models from airborne data collection is especially suitable for the creation of highly detailed city models, also referred to as LoD3 models. As this paper demonstrates, the large effort that would be required for explicit modeling can be avoided if so-called LASERMAPs are extracted from TLS and mapped against the coarse 3D building model. This mapping can efficiently be computed when the alignment of the TLS data to the data from airborne collection is known. The automated strategy for the georeferencing of TLS data presented within this paper can solve this task.

6. REFERENCES

- Besl, P. and McKay, N. (1992): A method for registration of 3d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), pp. 239-256.
- Blinn, J.F. (1978): Simulation of wrinkled surfaces. *ACM SIGGRAPH Computer Graphics* **12**(3), pp. 286-292.
- Bornik, A., Karner, K. F., Bauer, J., Leberl, F. and Mayer, H. (2001): High-quality texture reconstruction from multiple views. *Journal of Visualization and Computer Animation* 12(5), pp. 263–276.

- 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.
- Cignoni, P., Montani, C., Rocchini, C., Scopigno, R. and Tarini, M. (1999): Preserving attribute values on simplified meshes by resampling detail textures. *The Visual Computer* **15**(10), pp. 519–539.
- Cook, Robert L. (1984): Shade trees. In *Computer Graphics (Proceedings of SIGGRAPH 84)*, volume 18, Minneapolis, Minnesota, July, pp. 223-231.
- Coorg, S. and Teller, S. (1999): Extracting textured vertical facades from controlled close-range imagery. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 625–632.
- Fritsch, D., Klinec, D. and Volz, S. (2000): NEXUS - Positioning and Data Management Concepts for Location Aware Applications. In: *Proc. of the 2nd International Symposium on Telegeoprocessing*, Nice-Sophia-Antipolis, France, pp. 171-184.
- Kada, M., Roettger, S., Weiss, K., Ertl, T. and Fritsch, D. (2003): Real-time visualization of urban landscapes using open-source software. In: *Proceedings of ACRS 2003 ISRS*, Busan, Korea.
- Kilgard, M. (2000): A Practical and Robust Bump-mapping Technique for Today's GPUs. GDC 2000: Advanced OpenGL Game Development.
- Krishnamurthy, V. and Levoy, M. (1996): Fitting smooth surfaces to dense polygon meshes. *SIGGRAPH 1996*, pp. 313-324.
- Schuhmacher, S. and Böhm, J. (2005): Georeferencing of terrestrial laserscanner data for applications in architectural modeling. to appear in *3D Arch*, Venice Italy.
- Talaya, J., Alamús, R., Bosch, E., Serra, A., Kornus, W. and Baron, A. (2004): Integration of a Terrestrial Laser Scanner with GPS/IMU Orientation Sensors. In: *ISPRS Congress 2004*, Istanbul, Turkey, Vol. 35, Part B.
- Wang, X., Totaro, S., Taillardier, F., Hanson, A. and Teller, S. (2002): Recovering facade texture and microstructure from real world images. In: *Proc. 2nd International Workshop on Texture Analysis and Synthesis*, pp. 145–149.