Photogrammetric Week '09 Dieter Fritsch (Ed.) Wichmann Verlag, Heidelberg, 2009

The 3D Berlin Project

MARTIN KADA, Stuttgart

ABSTRACT

The original 3D city model of Berlin is to date still the largest one transported to the Google Earth platform. As of March 2009, the second generation model is available to the public. It includes a total of 474,000 buildings spanning an area of 857 km². All buildings are fully textured from oblique images, which gives the whole city model a near photo-realistic look. In this paper, we give insight into the geometric 3D reconstruction of the buildings from given LIDAR data and building footprints and describe the approach implemented in the software tool that was used for this project.

1. INTRODUCTION

In March 2009, the Senator of Economics, Technology and Women's Issues and Mayor of Berlin, Harald Wolf, and the Berlin Partner GmbH presented the second generation of the 3D city model of Berlin to the world. The Institute for Photogrammetry (ifp) of the University of Stuttgart was involved in the project as a software developing partner of virtualcitySYSTEMS, the company who performed the geometric reconstruction of Berlin in 3D. As one of the pioneers in the automatic reconstruction of 3D city models (Haala, 1996), (Ameri Shahrabi, 2000), (Brenner, 2000), the ifp was able to provide the know-how necessary for such a tremendous project.

3D building reconstruction has been a topic for quite some time now. Many research papers have been published; commercial services and software are available. (Brenner, 2005), e.g., gives a good overview of reconstruction methods and points out that "research is still far from the goal of the initially envisioned fully automatic reconstruction systems". This situation has not yet changed much, although a lot of research is still devoted to this topic, as can be seen in the multitude of recent publications (e.g. (Arefi et al., 2008), (Möser et al., 2009), (Sohn et al., 2008)).

The 3D Berlin project was the first one to make use of our collaboratively developed building reconstruction tool. Its underlying approach is based on earlier work on the decomposition of building footprints into non-overlapping, disjoint cells for the purpose of simplifying the geometry of 3D building models in the context of cartographic visualizations (Kada, 2007). As proven in (Haala et al., 2006), the decomposition is also useful as the basis for the reconstruction of roof geometries from LIDAR points. The approach was therefore adapted and developed further to allow for an automatic reconstruction of buildings at level of detail 2 (LOD 2), as defined in the official OGC standard CityGML (see e.g. (Kolbe, 2009)). At this level of detail, buildings have distinctive roof structures and flat facades that are textured with roof and façade images.

A frequent requirement, especially from customers within the mainland Europe, is that the provided building outlines are to be preserved with only little tolerance and that ridge and eaves heights must be very accurate. This is especially important so that the facades and roofs can be properly mapped with roof and façade textures. In this project, the resulting building models were automatically textured from oblique aerial images with known interior and exterior orientation (see Fig. 1). This was done by 3D Geo, a company recently acquired by Autodesk. The final model of Berlin is now publically available as a KMZ geo content file for use in Google Earth at www.3d-stadtmodell-berlin.de.

Fig. 1: 3D city model of Berlin textured from oblique aerial images.

2. RECONSTRUCTION ALGORITHM

In our approach, we assume that the majority of residential houses have either one main section or multiple connected sections with additional smaller extensions and that a partition thereof can be properly derived from the outline polygon. An integral part of this work lies therefore on the method to decompose a 2D building footprint into a set of nonintersecting cells (see section 2.1). However, the difficulty to generate correct facade and roof shapes from a partition increases with the number, shape and arrangement of its elements. We therefore generate only a small set of non-overlapping, mostly quadrilateral shaped polygons that together approximate the original footprint. Although the result is only an approximation, it is still accurate enough for reconstruction purposes. The benefit is, though, that the sections are separated nicely, especially for residential houses with gabled or hipped roofs. It sometimes happens that the ground shapes are not quadrilateral. As not all roof shapes produce a valid solid in that case, these cells are then restricted to only bear certain roof shapes.

Once such a footprint partition is found, a general geometrical description of the roof can be constructed by assigning a parameterized standard shape to each section (see section 2.2). The shape of each cell is determined from the LIDAR points with regard to the neighbor cells to better fit adjacent cells. After identifying the points inside a cell, the normal vectors from the local regression planes of the points are tested against all possible roof shapes. Here, only the orientation is used to speed up the comparison process against the high number of shapes we support. The shape that best fits the points is then chosen and its parameters estimated from the 3D LIDAR point coordinates. Cells whose neighbor configurations suggest corner-, t- and cross-junctions are examined again and replaced if a junction shape can be fitted according to the neighbor shapes and parameters.

After the cell shapes have been determined, the cells are glued together to form the final geometry of the model. The whole process is exemplarily depicted in Fig. 2. Next, the resulting building models are textured from oblique areal images. Any lack of geometric detail that is due to our rather restricting model oriented approach is then hardly noticeable in the result.

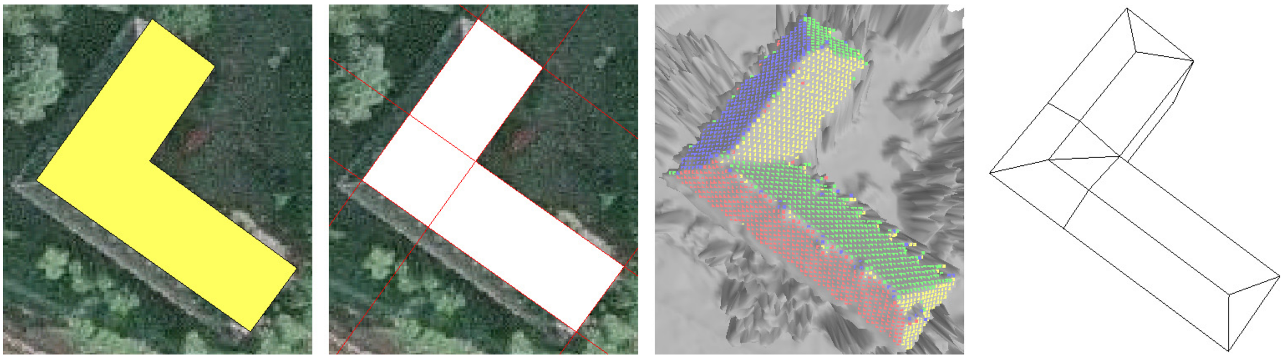

Fig. 2: The reconstruction algorithm decomposes a given building footprint (1) into mostly quadrilateral cells (2), classifies the LIDAR points according to their local regression planes (3) and gives the cells the best fitting roof shapes (4).

2.1. Cell Decomposition

As referred to in (Foley et al., 1996), a spatial partitioning representation in solid modeling, where solids are decomposed into nonintersecting, typically parameterized primitives, is called cell decomposition. Serving as the basis for the building reconstruction process, we first of all generate such a partition for each building footprint. As mentioned above, this is done solely from information found in the building's outline. The big challenge herein is to avoid decomposing the area in too many small cells, for which it becomes increasingly difficult to reconstruct a well-shaped roof, especially if the building outline is very detailed and consists of many short line sections. So instead of using all the available lines from the outline polygon and infinitely extend them to split the footprint, an adequate subset must be found that results in a set of primitives that together reflects well the characteristic shape of the building (see Fig. 3).

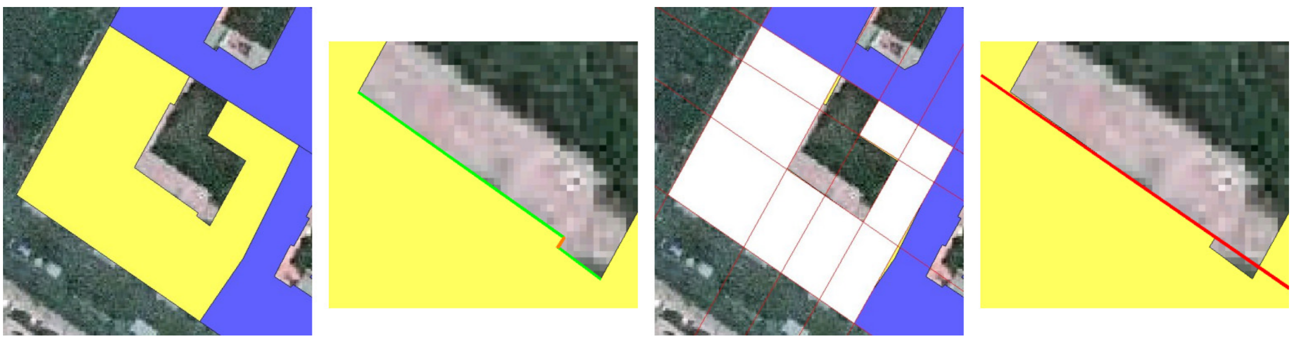

Fig. 3: Overview and close-up view of a building footprint (left) and its cell decomposition (right).

However, the resulting outline will not be identical to the original one, but rather be a generalization thereof. So to best resemble the outline, the set of decomposition lines should approximate well the original points and line segments.

Our algorithm for generating cell decompositions from given outlines has been thoroughly described in the context of 3D building generalization (see e.g. (Kada, 2007)). But instead of generating 3D decomposition planes from the facade polygons of a 3D building model, the 2D decomposition lines are now generated from the 2D outline. In a nutshell, the line segments are grouped into subsets of "parallel" lines that are pair wise a maximum distance away from each other. This is the generalization distance, which means in this context, that the cells resulting from the footprint partitioning will not have sides that are shorter than this length. Line segments are considered parallel if the angle between their directions is below an angle threshold. This allows for a better generalization of connected line segments and therefore helps to keep the number of generated cells low. For each subset of line segments, the associated decomposition line is computed by averaging the line equations of its elements. Short line segments of arbitrary direction, but whose endpoints are both closer to the decomposition line than the parallel line segments, are associated with this subset, but will not contribute to the averaging of this or any other decomposition line.

For example, the green line segments on the left side of Fig. 3 are considered parallel under the chosen angle threshold of 15 degrees. The added perpendicular distance of any two endpoints to the red decomposition line, which is the average of the green line segments, is below the generalization distance. While the connecting orange line segment is not parallel to any green line segments, its endpoints also falls under the distance threshold and therefore does not contribute to any decomposition line.

2.2. Roof Shape Determination

Now that a cell decomposition of the footprint is available, the parameterized roof shapes of all cells need to be found. We do this by examining the normal vectors of all points inside the same cell. As point normal vectors are usually not given in surface models, they first have to be generated. If the surface model is structured as a grid, we compute the normal vector of each point from the eight triangles fanned around it and average their normal vectors. However, if the raw data is available in form of an unstructured point cloud, we estimate a point's local plane of regression from its k-nearest neighbors (five in our case) and take the resulting surface normal vector.

For the construction of the building's roof, we classify the roof shapes that we use in our approach into three types: basic, connecting and manual shapes. Whereas the shapes of the first two classes can be determined in an automatic process, the last class of roof shapes is only available for manual editing. Among the basic roof shapes are flat, shed, gabled, hipped and Berliner roof. As not all houses have only one section, there is a need to connect the roofs of the sections with specific junction shapes. Fig. 4 shows a variety of elemental roof shapes and their combination to more complex buildings.

In summary, we automatically determine a cell's roof type by comparing the points' normal vectors with the roof faces of all possible shapes and compute the percentage of points that fit the direction of the roof part they are inside. For a gabled roof, e.g., we divide the cell into two equal parts, distribute the points accordingly and count the number of points whose normal vectors are in accordance with the respective side. Each roof type defines one or more parts, whose size may or may not be dependent on the roof parameters. E.g., the ridge line length of a hipped roof is variable and therefore affects the size of the four roof parts. The longer the ridge line grows, the smaller the two side hips become. This affects how accurately the shape can be determined.

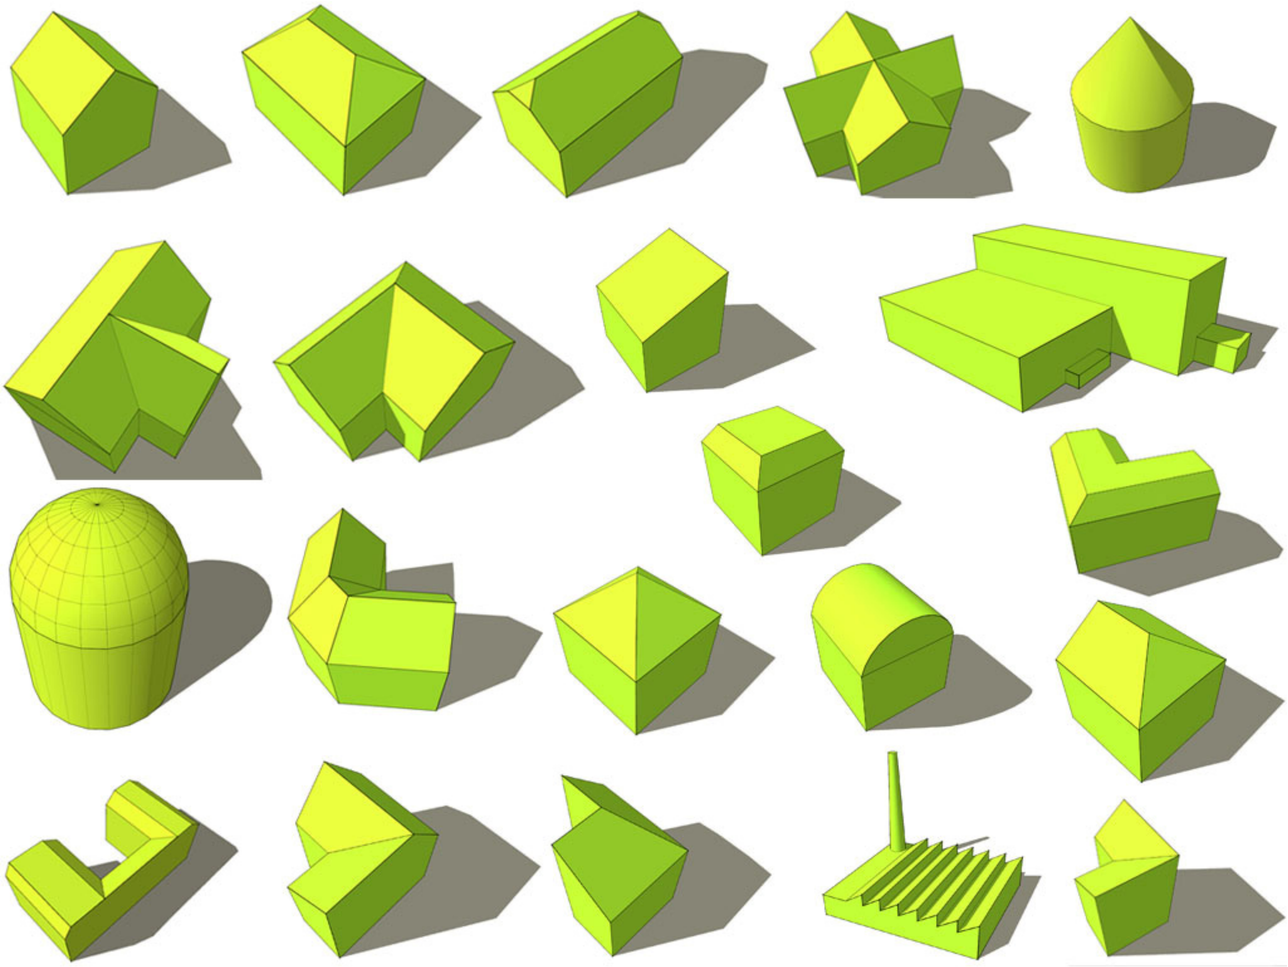

Fig. 4: A selection of elemental roof shapes and their combination to more complex building models.

3. THE 3D BERLIN PROJECT

The reconstruction of Berlin in 3D was the first project that was performed with the new building reconstruction software that implements the described cell decomposition approach. The project was an extension of the original city model, which is to date still the largest one transported to the Google Earth platform. Due to the projects time constraints and the enormous number of buildings, an image based reconstruction was quickly deemed as being too time-consuming and costly and was therefore ruled out as an option. Consequently, the choice was made to use the existing LIDAR data, as a reconstruction based on elevation information allows for a high degree of automation. In addition to the digital surface model, which was provided with a density of four points per square meter, the building footprints from the authoritative real estate cadastre were used.

The project was conducted in two phases. During the first phase, which lasted from March to July 2008, the eastern districts were completed. The area spans around 359 square kilometers and contains over 227,000 buildings. As we kept adapting and improving both workflow and software during production time, the second phase could be completed in only four months. From November 2008 to February 2009, the 498 square kilometer area of the western districts with its 247,000 buildings was reconstructed.

3.1. Pre-processing

The CityGML definition for level of detail 2 (LOD 2) city models state, that only buildings with a footprint greater or equal 4 × 4 meters will be represented in the model. Smaller buildings, whose arbitrarily oriented minimum bounding rectangles of the footprints are below this threshold, were therefore excluded from the reconstruction. This is also motivated from the fact that the subsequent texturing from oblique aerial images cannot be guaranteed to accurately fit such small buildings. Structures that are located underground or were reconstructed to lie at least partially under the existing digital terrain model were also eliminated. This affected approximately 3,000 buildings, which were mostly underground parking lots and carports.

In addition, buildings with predominantly flat roofs were treated separately. By using the semantic information found in the cadastral data, kiosks, gas stations, warehouses, multi-story car parks, gas plants, sewage plants, sports halls, indoor swimming pools, etc. were collected. They were then reconstructed by extruding the ground plans to a height level where most of the associated LIDAR points could be found in their vicinity. The ground plan extrusion has the advantage, that the given ground plan could be retained unchanged (see e.g. Fig. 5). This is especially useful for buildings that have only one height level, but very complex footprints. For more complicated buildings that have flat roofs with several height levels, the ground plan had to be manually edited to get a real-life representation.

Fig. 5: The reconstruction of the “Ernst-Reuter-Platz” with its many flat roofed buildings.

3.2. Automatic Reconstruction of the Building Geometry

The remaining buildings were reconstructed with the cell decomposition approach. The Berliner roof – a particularly unusual roof type typically found on many buildings in Berlin – presented a challenge as well as numerous inner courtyards presented problems during extraction (see Fig. 6). Therefore, the reconstruction approach had to be adapted to automatically detect this unique roof structure.

Fig. 6: Reconstruction of a district with many "Berliner roofs".

The Berliner roof is an asymmetric roof shape, which is basically a shed roof disinclined slightly to the back side. By having a steep slant at the front and sometimes also at the back side, the roof appears to be gabled from a pedestrians point of view. This shape is very common for Berlin apartment houses build during the period of promoterism in the 19^{th} century.

To identify the front side of a cell with a possible Berliner roof, we seek the side closest to the building's oriented bounding rectangle. If the cell is a corner cell, or if all cells are side by side, then two or more sides of the cell should be within closest distance to the bounding rectangle. Here, the side with the highest number of normal vectors pointing towards to is determined. This is in most cases the back side. Both methods are necessary, as the second one generally fails more often, but is the only one that works for the latter case.

Then, the distances from the front and back side to the two fake ridge lines are determined using a plane sweep approach. At the front ridge line, the 2D components of the points' normal vectors show in opposite directions. As for the back ridge line, we say that all points' normal vectors with an angle below 30 degree compared to the upward vector belong to the shed part of the roof. Using these two criteria, we can accurately determine the two ridge lines that separate the three roof regions. Their height is computed from the plane equations estimated from the points of the two steep slant sections.

In addition to the Berliner roof, a total of 17 individual roof types have been integrated into the software in order to enable greater accuracy during reconstruction and to reduce the amount of manual editing needed.

3.3. Post-processing

As the roof structures of buildings cannot always be inferred only from their footprints, the automatic reconstruction could not deliver the required accuracy in all cases. Therefore, the data set had to be manually corrected in a semi-automatic process. During the project time, up to seven operators were involved in the post-editing stage and another four to five were working on quality assurance and texturing.

In our editing tool, the decomposition lines can be copied, added, deleted, translated and rotated. The result is immediately visible to the operator, as the cells' roof shapes are automatically reconstructed in a split second after every change. Once the cell decomposition fits the roof's shape, the cell parameters can be manually corrected, adjusted to neighboring cells or even copied from other buildings. If the decomposition produced too many small cells, then their number can be decreased by a merging operation.

Even though the editing of the building models using decomposition lines is at first a little uncommon, we noticed that operators got used to it very quickly and can efficiently produce even landmarks with complex geometry. The manual mode also allows for more complex roof shapes like mansard, cupola, barrel and even some detail elements like dormers, which would be impossible to model in an automatic process from the given input data.

As the software was constantly improved during the duration of the project, the amount of manual editing needed for the reconstruction was reduced from 30 percent in denser areas to 20 percent; manual editing for the outer lying areas also experienced a sharp improvement: from 20 percent to 15 percent.

3.4. Texturing

After the geometry had been reconstructed, the models were textured from oblique aerial images. A total of 50,000 images were used. Although the interior and exterior orientations were known, all images had to be manually adjusted to guarantee a perfect fit. The texture images were then automatically extracted, filtered and stored with the geometry.

4. CONCLUSION

The paper presents an automatic approach for the reconstruction of 3D building models from LIDAR data and existing ground plans. It is based on an algorithm to decompose given footprints into sets of nonintersecting cells, for which roof shapes are then determined from the normal directions of the LIDAR points. For more in-depth information on the reconstruction algorithm, see (Kada and McKinley, 2009).

The validity of this approach has been proven effective, as can be judged by the 3D city models of Berlin. Even though the sheer size of the project posed a tremendous challenge, it was successfully completed and the resulting model is publicly available for everyone for personal judgment.

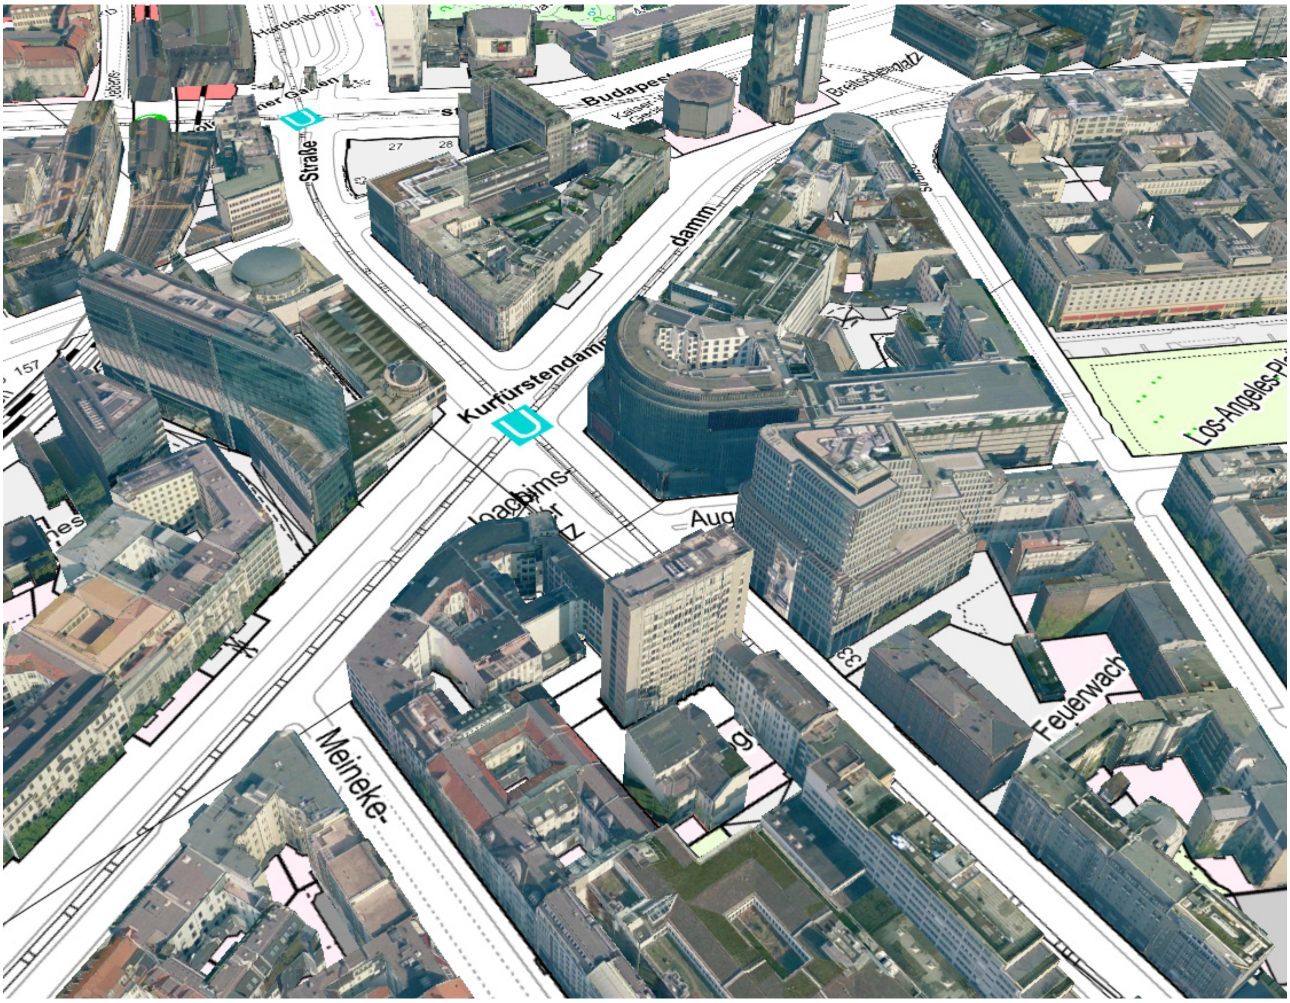

Fig. 7. The reconstruction of the prominent Kurfürstendamm textured from oblique aerial images.

In a later stage, around eighty landmarks and places of interest were modeled by hand exactly down to the last detail by graphic artists. Five buildings can even be virtually entered and visited like the Olympic Stadium, Sony Center, Reichstag building, DZ Bank building, Berlin Central Station.

5. AKNOWLEDGEMENTS

The author wants to thank virtualcitySYSTEMS for providing images, feedback and additional information about the 3D Berlin project.

6. REFERENCES

Ameri Shahrabi, B. (2000): Automatic Recognition and 3D Reconstruction of Buildings through Computer Vision and Digital Photogrammetry (Ph.D. thesis). Deutsche Geodätische Kommission, Reihe C, Nr. 526, München.

Arefi, H., Engels, J., Hahn, M., Mayer, H. (2008): Levels of Detail in 3D Building Reconstruction from LIDAR Data. In: International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVII-B3b.

Brenner, C. (2000): Dreidimensionale Gebäuderekonstruktion aus digitalen Oberflächenmodellen und Grundrissen (Ph.D. thesis). Deutsche Geodätische Kommission, Reihe C, Nr. 530, München.

Brenner, C. (2005): Building Reconstruction from Images and Laser Scanning. In: International Journal of Applied Earth Observation and Geoinformation (Theme Issue 'Data Quality in Earth Observation Techniques), Vol. 6 (3-4), pp. 187-198.

Foley, J., van Dam, A., Feiner, S., Hughes, J. (1990): Computer Graphics: Principles and Practice (2nd Edition), Addison-Wesley.

Haala, N. (1996): Gebäuderekonstruktion durch Kombination von Bild- und Höhendaten (Ph.D. thesis). Deutsche Geodätische Kommission, Reihe C, Nr. 460, München.

Haala, N., Becker, S., Kada, M. (2006): Cell Decomposition for the Generation of Building Models at Multiple Scales. In: International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVI, Part 3, pp. 19-24.

Kada, M. (2007): Scale-Dependent Simplification of 3D Building Models Based on Cell Decomposition and Primitive Instancing. In: Winter, S., Duckham, M., Kulik, L., Kuipers, B. (Eds.) Spatial Information Theory: 8th International Conference, COSIT 2007, LNCS 4736, Springer-Verlag Berlin Heidelberg, pp. 222-237.

Kada, M., McKinley, L. (2009): 3D Building Reconstruction from LIDAR Based on a Cell Decomposition Approach. In: Proceedings of the CMRT 2009 ISPRS Workshop, Paris, France.

Kolbe, T. H. (2009): Representing and Exchanging 3D City Models with CityGML. In: Lee, Zlatanova (Eds.): 3D Geo-Information Sciences, Springer-Verlag Berlin Heidelberg, pp. 15-32.

Möser, S., Wahl, R., Klein, R. (2009): Out-Of-Core Topologically Constrained Simplification for City Modeling from Digital Surface Models. In: International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII-5/W1.

Sohn, G., Huang, X., Tao, V., (2008): Using a Binary Space Partitioning Tree for Reconstructing Polyhedral Building Models from Airborne Lidar Data. In: Photogrammetric Engineering & Remote Sensing, Vol. 74, No. 11, pp. 1425-1438.