Building Reconstruction – Outside and In

Christopher Gold, Rebecca Tse and Hugo Ledoux.
School of Computing, University of Glamorgan
Pontypridd, Wales, CF37 1DL UK
cmgold@glam.ac.uk

Abstract

The modelling or reconstruction of buildings has two aspects – on the one hand we need a data structure and the associated geometric information, and on the other hand we need a set of tools to construct the building incrementally. This paper discusses both of these aspects, but starts from the simpler exterior model and geometry determination, and then looks at representations of the building interiors.

Our starting point is a set of raw LIDAR data, as this is becoming readily available for many areas. This is then triangulated in the x-y plane using standard Delaunay techniques to produce a TIN. The LIDAR values will then show buildings as regions of high elevation compared with the ground. Our initial objective is to extrude these buildings from the landscape in such a manner that they have well defined wall and roof planes. We may have already been provided with the building footprint from national mapping information, or we may need to extract it from the triangulation. We do this by superimposing a coarse Voronoi cell structure on the data, and identifying wall segments within each. We then examine the triangulated interior (roof) data, identify planar segments and connect them to form the final surface model of the building embedded in the terrain. This is done using Euler Operators and Quad-Edges. Building interiors are added by using an extension of these.

1. Previous work

We have previously demonstrated (Tse and Gold, 2002) that a TIN may be represented with advantage using the Quad-Edge data structure of Guibas and Stolfi (1985), and that this structure is closely related to the basic Euler Operators used in CAD systems for boundary representation (b-rep) modelling of exterior surfaces (e.g. Mantyla, 1988; Lee, 1999). Other structures may also be used (e.g. Baumgart, 1972; Weiler, 1986). Based on this equivalence the TIN model may be modified using these Euler Operators, to permit the modelling of seamless exterior surfaces of buildings or other structures, for example by the splitting or merging of faces, the creation or deletion of bridges or tunnels, etc. Fig. 1 shows the basic Quad-Edge element, and Fig. 2 shows the elementary Splice Operator. Fig. 3 shows a simple

triangulation network, the Quad-Edges and the topological loops around faces and nodes. Note that this represents both the primal and the dual structure: loops around Delaunay nodes are equivalent to Voronoi cell boundaries, and loops around Voronoi nodes are equivalent to Delaunay triangle boundaries.

Quad-Edge data structure

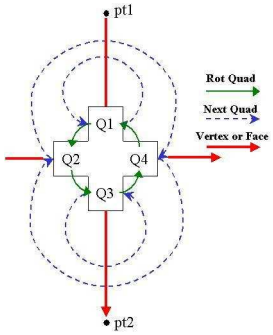

Fig. 1: Quad-Edge structure

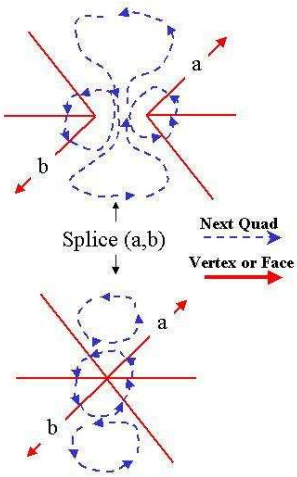

Fig. 2: Splice operation

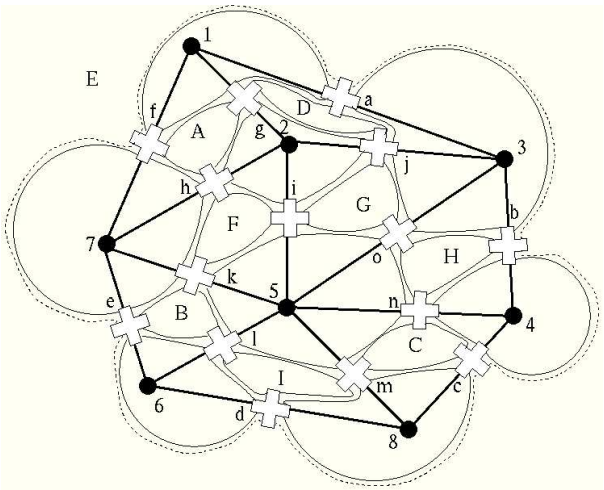

Fig. 3: Quad-Edge navigation

Euler Operators to create a tunnel or bridge

In the CAD industry the most common elementary operations on surface models (b-reps) are called Euler Operators. These have been shown (Tse and Gold, 2002) to be simply constructed from the Quad-Edge operations Make-Edge (for a new edge) and Splice (to join or split two Quad-Edges). Each Euler Operator has an inverse. Shown here are MEV (Make Edge and Vertex, Fig. 4), MEF (Make Edge and Face, Fig. 5) and how to construct a tunnel or a bridge (Fig. 6). Fig. 7 shows two simple examples.

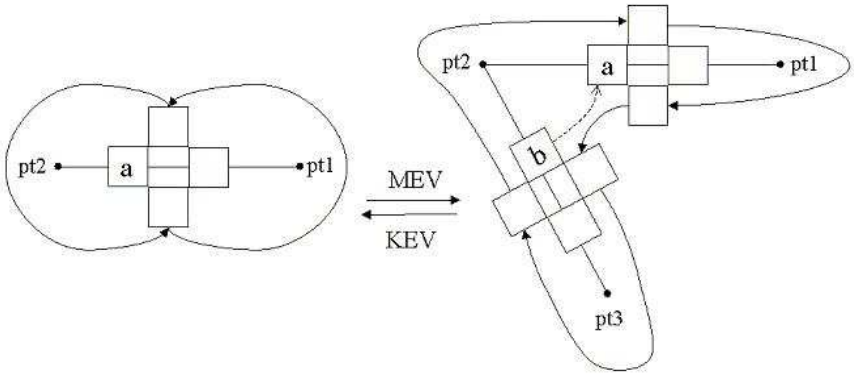

Fig. 4: Connecting two Quad-Edges with MEV

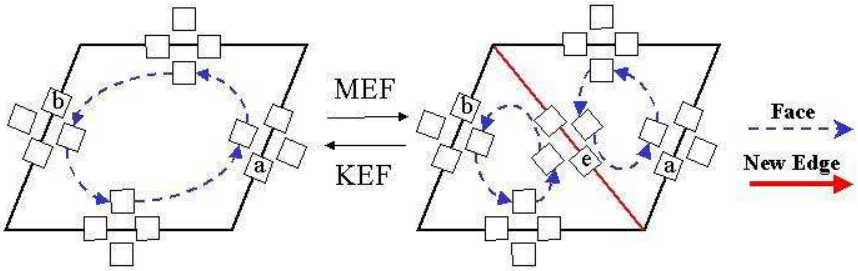

Fig. 5: Splitting a polygon with MEF

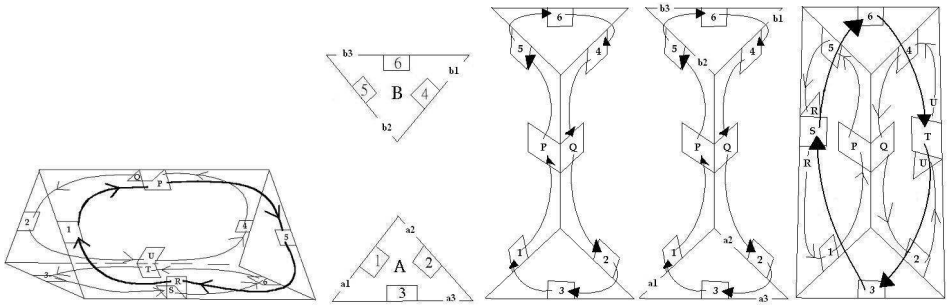

Fig. 6: Tunnel Construction

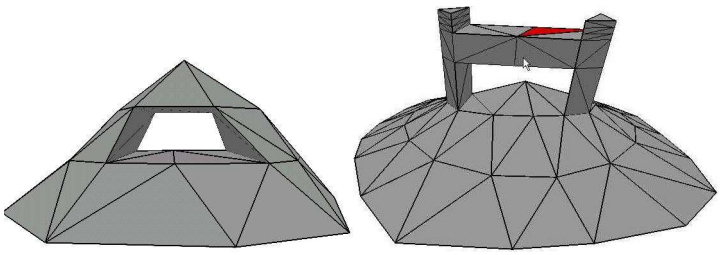

Fig. 7: Bridges and tunnels

Building Extraction with Provided Mapping Information

If we are given the building footprint we may insert points into the terrain model along these boundary lines, and use the Euler Operators to extrude the building vertically (Tse and Gold, 2001). This will give a flat roof at the average height of the LIDAR data within the boundary, with the building modelled by Quad-Edges. Fig. 8 shows the terrain with the building footprints, and Fig. 9 shows the extruded buildings.

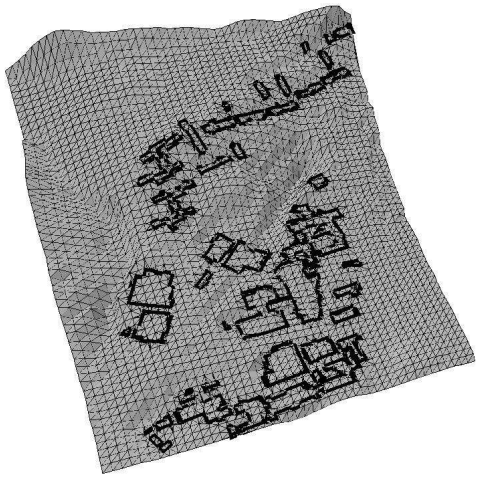

Fig. 8: Building footprints on terrain

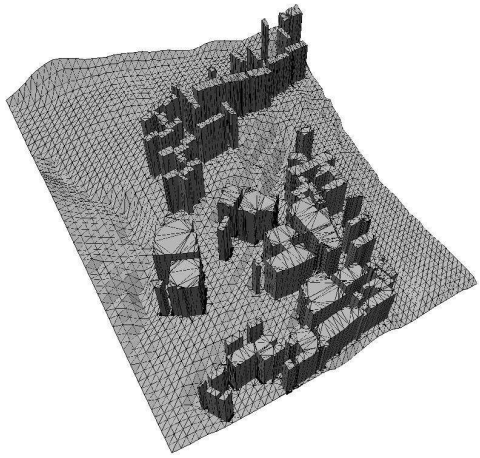

Fig. 9: Extruded buildings

2. Building Construction from LIDAR data alone

In principle it should be possible to extract good approximations of buildings from a sufficiently dense set of elevation data. In practice this is difficult. There are two steps: firstly to extract the vertical walls, and then to model the roof. There are two basic approaches: to attempt to fit a pre-defined template to the data (e.g. Vosselman, 2003; Rottensteiner and Briese, 2003); or to attempt to construct a building-like shape by extracting features from elevation data (e.g. vertical walls or roof planes). The first approach is limited by the models included in the system, while the second will only approximate the building form, and will often need subsequent rectification. We take the second approach, and only assume properties of "buildings" when absolutely necessary.

Automatic Wall Extraction

While it is not difficult to identify the near-vertical triangles in the TIN it is not a simple task to form a complete building from these segments. The remote sensing literature has many examples of attempts to first detect line segments and then glue them together.

Our approach is always to preserve a tessellation model, with connectivity, rather than attempting to connect line segments. We apply a coarse Voronoi diagram over the original data, with perhaps 50-100 LIDAR points in each cell. We then attempt to modify these cells so that the building boundaries (defined as a partition between "high" points and "low" points") are a subset of the Voronoi cell edges. We can split cells along the high/low edge to achieve this.

Proposition 1: Buildings are collections of contiguous elevations that are higher than the surrounding terrain. Their boundaries are "walls".

The approach is based on calculating the eigenvalues and eigenvectors of the 3 x 3

variance-covariance matrix of the coordinates of the points within a cell. The first eigenvector (with the largest eigenvalue) “explains” as much of the overall variance as possible, the second (perpendicular) eigenvector explains as much as possible of what is left, and the third (perpendicular to the other two) contains the residue. (For example, a wrinkled piece of paper might have the first eigenvector oriented along the length of the paper, the second along its width, and the third “looking” along the wrinkles.) Thus the eigenvector of the smallest eigenvalue indicates the orientation of a wall segment, if present, and looks along it.

The next step is to locate the line parallel to the smallest eigenvector that best separates “high” elevations from “low” ones within each Voronoi cell. This is achieved iteratively, by testing various positions of this line in order to find the greatest difference between the means of the elevation values in the Voronoi cell that are on each side of the line. (In order to minimize the effect of sloping roofs or terrain, only those elevations close to the line are used.) If this maximum difference is not sufficiently large then no wall segment was detected.

Proposition 2: Walls have a specified minimum height, and this height difference is achieved within a very few “pixels”.

The Voronoi cells are then split along these lines, by adding a generator on each side of this line, at the mid-point. This gives a set of “high” Voronoi cells surrounded by “low” ones. Building boundaries are then determined by walking around the cells forming the high region, using the topological consistency of the Voronoi tessellation. This must form a closed region, or else the high region is not considered to be a building. The building outline is then estimated from the Voronoi boundary segments – but only those that were created with the eigenvector technique, not those Voronoi cell boundaries that only connect them. Fig. 10 shows the Voronoi structure and the “high” LIDAR points, as well as the wall segments detected. Fig. 11 shows the “high” Voronoi cells before they are split, and Fig. 12 shows the approximation of the walls based on the split cells.

Proposition 3: A building consists of a high region entirely surrounded by walls.

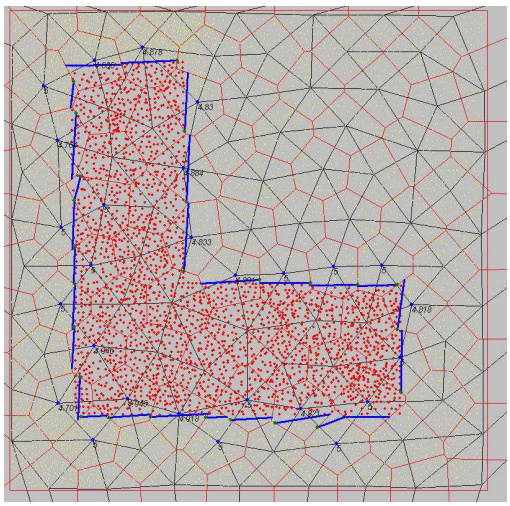

Fig. 10: Voronoi cells and wall segments

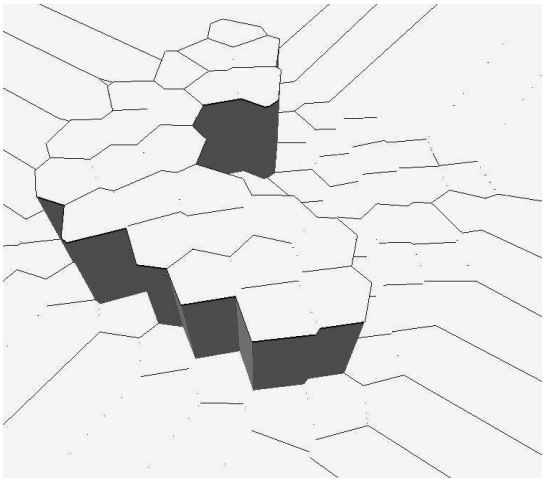

Fig. 11: Initial “high” Voronoi cells

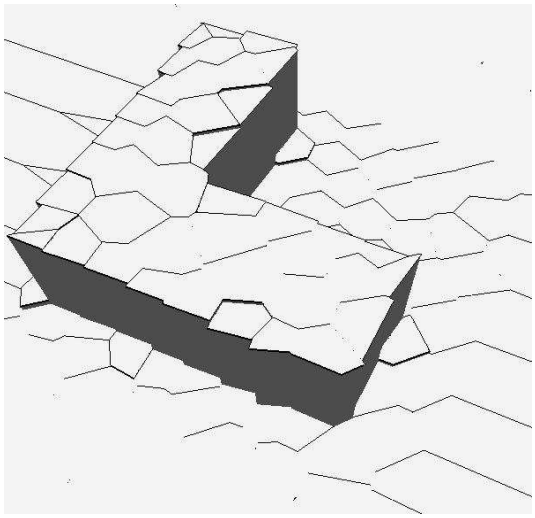

Fig. 12: “High” Voronoi cells after splitting

3. Roof Modelling

Simple Roof

Once the building boundary is determined, the interior points may be used to model the roof structure. Unlike other techniques, no assumptions are made about the form of the roof, except that it is made up of planar segments. (The method may be extended to detect other basic shapes if required.) Each of the interior triangles has an associated vector normal. The “smallest” eigenvector is used as described previously to estimate the main axis of the roof, all vector normals are plotted on a unit semicircle perpendicular to this, and clusters are found where all normals are very close to the same orientation. The mean of each cluster indicates the orientation of one or more planar roof segments with the same orientation. This works well even if the data is fairly noisy, and the scatter of the vector normals is fairly large. (If there are two or more parallel planes in the roof structure, these may be separated at this stage by constructing the Delaunay triangulation in x-y space for the data points of the cluster,

extracting the Minimum Spanning Tree, and separating the two or more parallel roof portions. The general technique is described later.) Fig. 13 shows a simple roof, Fig. 14 shows the vector normals and Fig. 15 shows the clusters of unit vectors (small circles) on the semicircle. Each roof plane is described by its vector normal plus a "visible" point on the roof plane that is within the bounds of the cluster used to estimate the vector normal. This is usually just the mean of the x-y-z values of the relevant data points. It is not necessary that all triangles within a roof segment have vector normals in the cluster – just enough to detect the plane. Fig. 16 shows the resulting roof planes for clean data, and Fig. 17 for noisy data.

Intersections of wall and roof planes are then calculated, and the building extruded or bevelled using Euler Operators to give the final b-rep building form.

Proposition 4: Roofs are made up of planar segments, most of whose constituent triangles have similar vector normals.

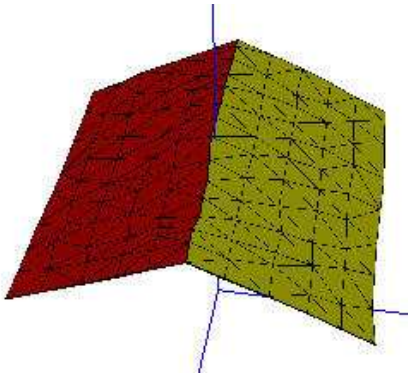

Fig. 13: Simple roof

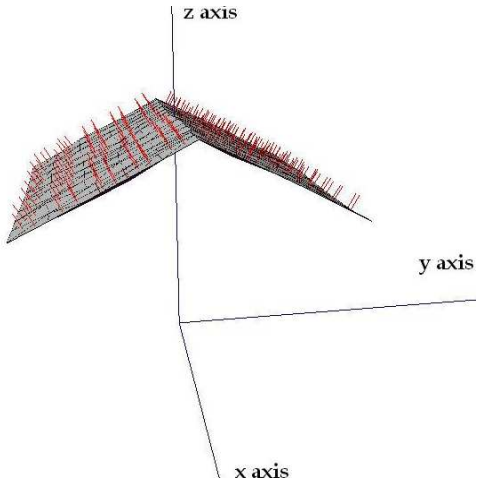

Fig. 14: Roof vector normals

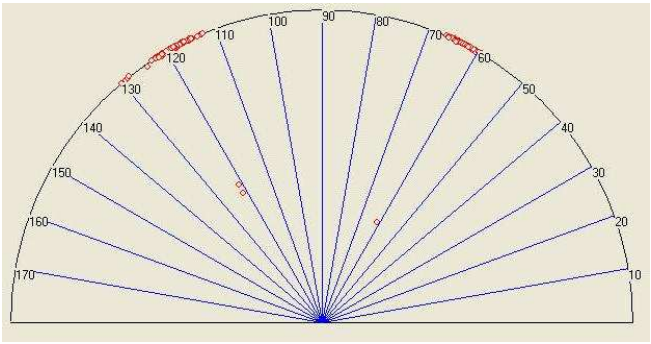

Fig. 15: Vector clusters on the unit semicircle

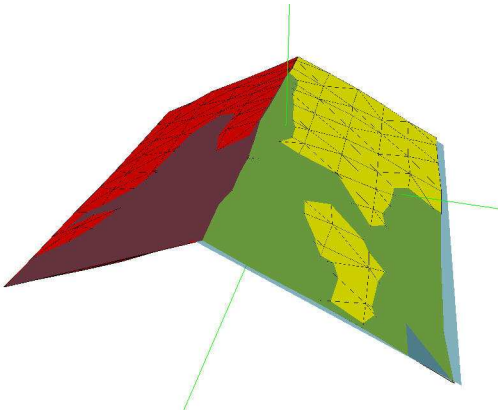

Fig. 16: Roof planes from clean data

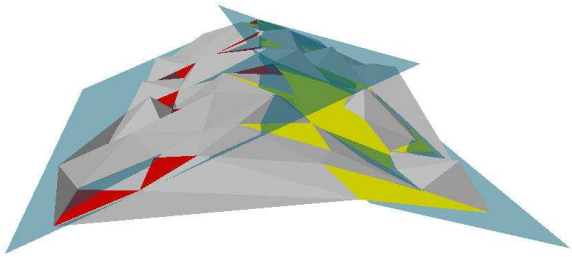

Fig. 17: Roof planes from noisy data

Complex Roof

The above method only works for roofs with a simple axis. Where planar segments have many orientations a different technique is required, using the projection of the vector normals onto the unit hemisphere. Others have used different projections (e.g. Hofmann et al., 2003). Where several triangles have similar orientations – usually from the same roof segment – there will be a cluster of points that can be determined by first constructing the Delaunay triangulation on the hemisphere and then extracting the Minimum Spanning Tree (MST, which is a subset of the Delaunay triangulation). These clusters are then used to identify planar segments as before. Fig. 18 shows a complex roof, while Fig. 19 shows a 2D projection of the triangulation of the vector normals on the hemisphere and Fig. 20 shows this in a 3D view. Fig. 21 shows the MST in 2D, and Fig 22 shows it in 3D. Four clusters can be seen, corresponding to the four roof planes. Fig. 23 shows the classified roof triangles, and Fig. 24 shows the fitted planes.

Roof planes can be considered to meet at triple junctions (cases of four planes meeting can be treated as two triple junctions with a very short edge between them). Vector algebra gives a concise calculation of these plane intersections, and thus the general form of a roof may be represented as a triangulation of the "visible points" on each roof plane. The building's form may then be extruded using Euler Operators as previously described.

Proposition 5: The relationships between roof planes may be represented as a dual triangulation.

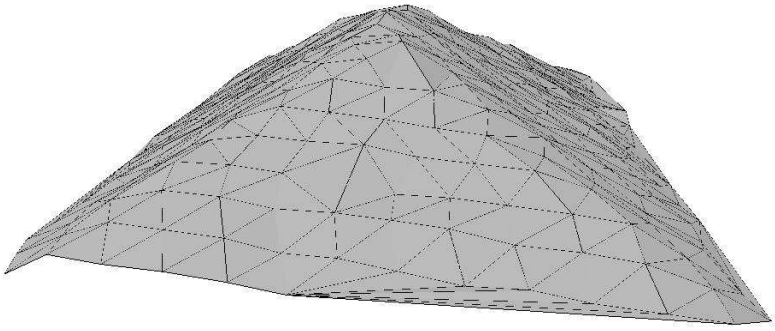

Fig. 18: Complex roof surface

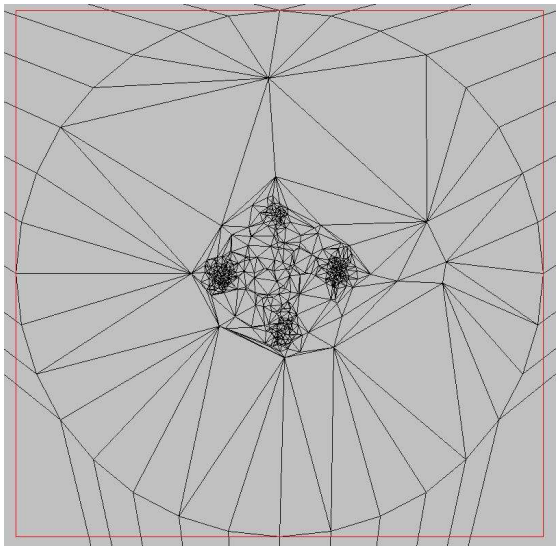

Fig. 19: Triangulation of unit hemisphere

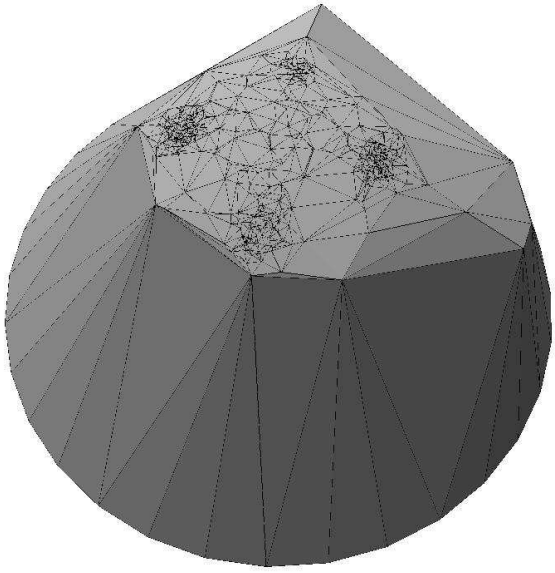

Fig. 20: Unit hemisphere – 3D view

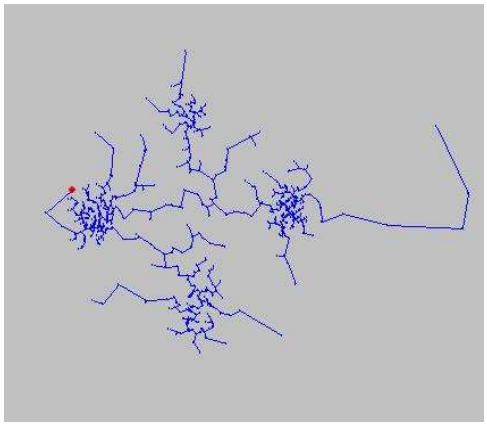

Fig 21: MST of unit hemisphere

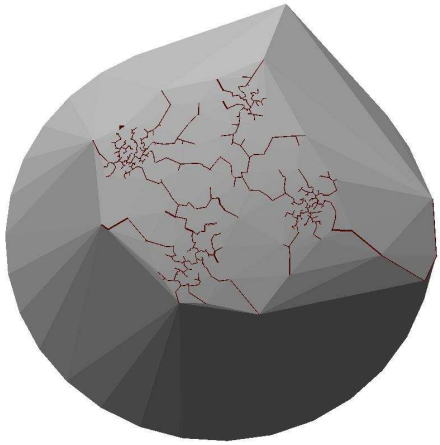

Fig. 22: MST – 3D view

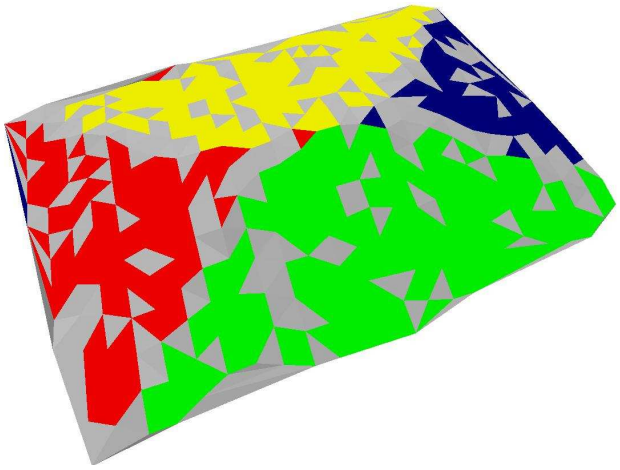

Fig. 23: Roof triangle clusters from normals

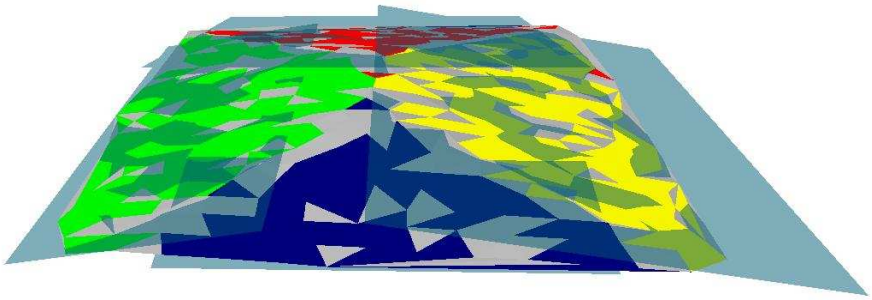

Fig. 24: Roof planes from clusters

Building Interiors

This gives the exterior building shape, embedded in the topography, defined with the Quad-Edge structure. However, where interior information is to be added this structure needs to be augmented.

The Quad-Edge structure described previously has, in addition to its topological pointers to adjacent edges, two pointers to the vertices of the primal edge and two pointers to the left and right faces (Fig. 1). However, in 3D the dual of a face is an edge (Fig. 25), so these pointers may be assigned to Quad-Edges of dual cells. The Augmented Quad-Edge structure has been developed to represent space-filling cells individually with Quad-Edge structures, linked through the equivalent set of dual cell edges. Thus if the original building (and adjacent ground) determined from LIDAR data was constructed as a single “Polyhedral Earth” model, using the standard Quad-Edge structure, the dual edges referred to by the original face pointers would each connect to “Earth” at one end and “Air” at the other.

Proposition 6: Building exteriors, together with the adjacent terrain, form a portion of the global “Polyhedral Earth”.

The Augmented Quad-Edge is designed to link face pointers in one space with Quad-Edges in the dual space – a good example of this is the 3D Voronoi diagram and dual Delaunay

tetrahedralization, where Delaunay edges represent Voronoi faces, and vice versa. This structure is ideally placed to represent cellular complexes and their adjacencies (Ledoux and Gold, In Press). Thus the addition of a hollow building interior is achieved by adding an interior cell where most of the edges correspond to those forming the previously-constructed building exterior. Additional rooms may be added by partitioning this interior cell, while the simultaneously-constructed dual cell edges form a structure for navigating the interior (Fig. 26 – two tetrahedra are linked by the dual Voronoi edges that penetrate the common face). Fig. 27 shows the basic relationships between adjacent rooms using this structure, consisting of the standard Quad-Edge operators and the "through" and "adjacent" operators from the Augmented Quad-Edge. Clearly this approach is equally valid for subterranean constructions as well as above-ground buildings.

Proposition 7: Building interiors may be constructed as individual polyhedra, linked together and to the exterior by edges of the dual graph.

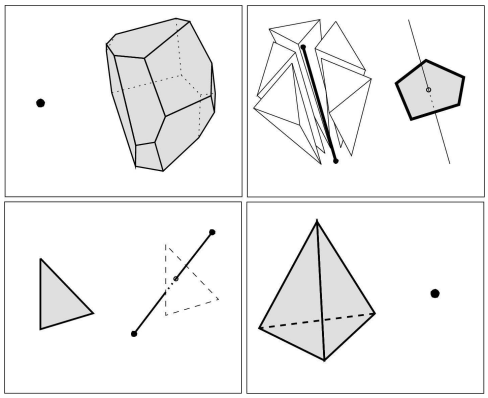

Fig. 25: 3D dual relationships

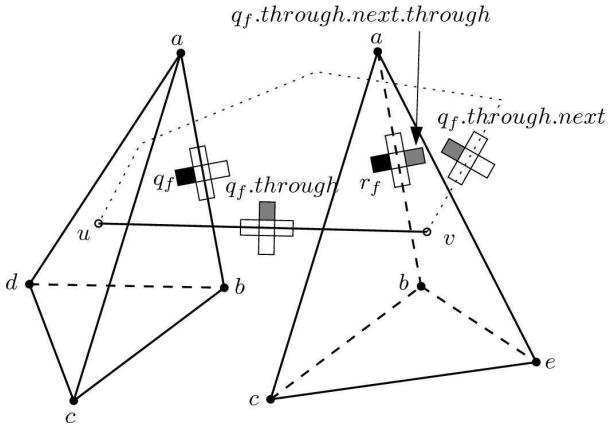

Fig. 26: 3D navigation via the dual

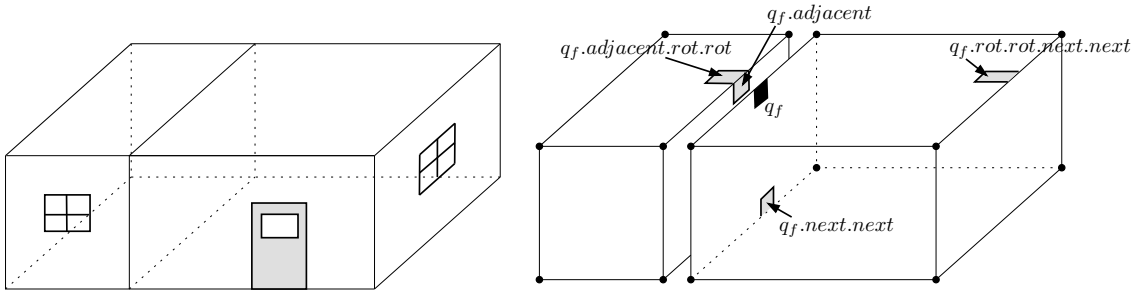

Fig. 27: Relationships between rooms

Conclusions

We have outlined a procedure for the direct extraction of building exteriors from LIDAR data without any prior knowledge of the data. This is based on five propositions that are necessary in order to recognize the basic elements of our buildings. Two additional propositions are

given in order to provide a topological context for building interiors. While the resulting building forms have only limited precision they can be rectified with additional conditions as desired in any particular urban context.

References

Baumgart, B.G. (1972): Winged edge polyhedron representation. Stanford University Computer Science Department, Stanford Artificial Intelligence Report No. CS-320.

Guibas, L., and J. Stolfi (1985): Primitives for the Manipulation of General Subdivisions and the Computation of Voronoi Diagrams. ACM Transactions on Graphics, Vol. 4, No. 2, pp. 74-123.

Hofmann, A.D., Mass, H-G. and Steilein, A. (2003). Derivation of roof types by cluster analysis in parameter spaces of airborne laserscanner point clouds. In H.-G. Maas, G. Vosselman, and A. Streilein, editors, Proceedings of the ISPRS working group III/3 workshop '3-D reconstruction from airborne laserscanner and InSAR data', volume 34, Part 3/W13, Dresden, Germany, 2003. Institute of Photogrammetry and Remote Sensing Dresden University of Technology.

Ledoux, H. and Gold, C. M. (In Press). Simultaneous storage of primal and dual three dimensional subdivisions. Computers, Environment and Urban Systems.

Lee, K. (1999): Principles of CAD/ CAM/ CAE Systems. Seoul National University. Korea.

Mantyla, M., (1988): An introduction to solid modelling, Computer Science Press, Rockville, MD.

Rottensteiner, F. and Briese, C., (2003). Automatic generation of building models from LIDAR data and the integration of aerial images. In H.-G.Maas, G. Vosselman, and A. Streilein, eds., Proceedings of the ISPRS working group III/3 workshop '3-D reconstruction from airborne laserscanner and InSAR data', volume 34 Session IV. Institute of Photogrammetry and Remote Sensing Dresden University of Technology, Dresden, Germany.

Tse, R. O. C. and Gold, C. M., (2001). Terrain, dinosaurs and cadastres - options for three-dimension modelling. In C. Lemmen and P. van Oosterom, eds., Proceedings: International Workshop on "3D Cadastres", pp. 243–257. Delft, The Netherlands.

Tse, O.C., and C. Gold (2002) "TIN Meets CAD – Extending the TIN Concept in GIS", Proceedings: ICCS 2002 Lecture Notes of Computer Science, Dongarra & A.G. Hoekstra (Eds) vol. 2331, March, Springer, Amsterdam, the Netherlands, pp. 135-143.

Weiler, K. J. (1986): Topological structures for geometric modeling. PhD thesis, Rensselaer Polytechnic Institute, University Microfilms International, New York, U.S.A.

Vosselman, G., (2003). 3d reconstruction of roads and trees for city modelling. In H.-G. Maas, G. Vosselman, and A. Streilein, eds., Proceedings of the ISPRS working group III/3 workshop '3- D reconstruction from airborne laserscanner and InSAR data', volume 34, Part 3/W13. Institute of Photogrammetry and Remote Sensing Dresden University of Technology, Dresden, Germany.