1 Radiosity with Voronoi Diagrams

Wolfgang Stürzlinger [1]

ABSTRACT
Current radiosity methods rely on the calculation of geometric factors, known as form-factors which describe energy exchange between the surfaces of an environment. The most widely used method for storing the illumination across a surface is a mesh of quadrilaterals and/or triangles. For more exact computations these meshes need to be subdivided adaptively. The subdivision methods create artifacts which have visible results. A new technique for storing the surface is presented, based on Voronoi diagrams, which are well suited for the task, and can be subdivided without introducing artifacts.

1.1 Introduction

Radiosity has become a popular method for image synthesis due to its ability to generate images of high realism. It was first introduced to computer graphics by Goral et al. [8]. Further research resulted in the progressive refinement method which is able to produce good approximations of the final solution very quickly [5]. The radiosity method was extended to include specular reflection through the so called two-pass approach. For recent developements see [11], [15], [16].

Common to all these methods is a representation of the surfaces of the environment by a mesh of quadrilaterals and triangles. These "patches" are used to store the radiosity on the respective part of the surface.

The geometric form-factors were first calculated by the use of a hemicube [6]. A hemicube is placed around the center of a patch and all other patches are projected onto its surfaces. The projected area gives an estimate for the geometric form-factor between the patches. This method, although very efficient because it can be done in hardware, has two major shortcomings.

First, serious inaccuracies can occur if the size of the patch is large relative to the distance between patches. Second, a number of aliasing problems may occur. Other methods for computing the form-factors were suggested and/or implemented. Baum used a hybrid method involving both numerical and analytic methods [3] and other methods use ray-tracing to compute the form-factors [19], [15], [11], [17]. Wallace subdivides the shooting patch until an analytic solution to approximate the form-factor of the delta-areas can be used. Then the form-factors for all visible delta-areas are summed up giving a good approximation to the form-factor of the shooting patch.

To reduce the computational expense without sacrificing image quality, a widely used two-level adaptive subdivision approach has been formulated by Cohen et al. [7]. The surface is subdivided into patches and these are used to compute a coarse solution for the global illumination. Then the patches are further subdivived into elements. The elements act as receivers of the light coming from the coarse patch solution, and provide the sample values for the final rendering process.

The advantage of this method is that the element subdivision can be continued adaptively as high radiosity gradients are discovered without recomputing the patch radiosity.

[1] Institut für Computergraphik, Technical University of Vienna, Vienna, Austria, wrzl@eichow.tuwien.ac.at

Effectively, this means that the picture quality improves by sampling the environment more densely where the changes are greatest.

1.1.1 Problems with Meshing

A problem inherent to the meshing with quadrilaterals is that radiosity discontinuities which cross the polygon can only be approximated. Because the intensity is interpolated from the corners of the polygon at rendering time, visual artifacts may appear. By adaptive subdivision the discontinuity can only be approximated but it is not guaranteed that the subdivision level is adequate for the rendering process. This holds especially if we subdivide regularly (e.g. in a quadtree like manner) and disregard the direction of the discontinuity. Recently there has been a lot of research on how to represent these discontinuities in the mesh [4].

In some situations the opposite problem can occur, that is, shading discontinuities appear where they don't belong. This can happen when smooth surfaces in the environment are actually represented by a collection of independent polygons in the model. If the individual polygons are meshed independently, the elements may not line up correctly along the boundaries between neighbouring polygons and T-vertices will appear. This will also appear if the subdivison level of two neighbouring polygons is different.

These T-vertices will produce discontinuities in the shading because the shading is interpolated along the neighbouring element edges during rendering. Also it is difficult to calculate the correct area for the sampling points, especially for points lying on the common border of neighbouring polygons, when their subdivision level is different.

For even more problems with meshing see [2], [10], [18].

1.2 Polygonal meshes with Voronoi Diagrams

In section 1.1 we highlighted some problems with meshes. If we use a different type of mesh we can eliminate these problems. In the following we assume that we can assign a two dimensional coordinate system to each surface. We initialize each surface with a regular grid (in world space) of sampling points and construct the Voronoi diagram for these points. The sampling points will be referred to as sites further on. Care has to be taken that some sites lie on the edge of the parametric space. This gives a mesh of either quadrilaterals or hexagons depending on the placement of the sites. But we stress that our method also works with irregular placement of sites, and is especially well suited for adaptive subdivision methods.

1.2.1 Voronoi Diagrams

Voronoi diagrams are well known geometric data structures (for an introduction see e.g. [12], [13] or [1]), which are used mostly for planar point location problems. The plane is subdivided into polygons where each polygon contains all points nearest to a given site (see Fig. 1.1).

The geometric dual of the Voronoi diagrams is the Delauny triangulation where two sites are connected by an edge iff they have a common polygon edge in the Voronoi diagram.

The Voronoi polygons are well suited for meshing and the radiosity method. They are always convex and the average number of edges in a Voronoi polygon does not exceed 6, because the whole diagram cannot have more than $3n-6$ edges, with n being the number of sites. The Delauny triangulation will consist of at most $2n-4$ triangles.

The Voronoi diagram can be represented either directly by a winged-edge data struc-

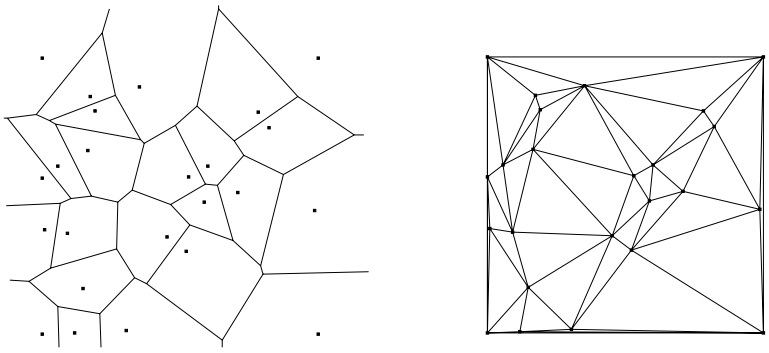

FIGURE 1.1. Voronoi diagram (left) and Delauny triangulation (right).

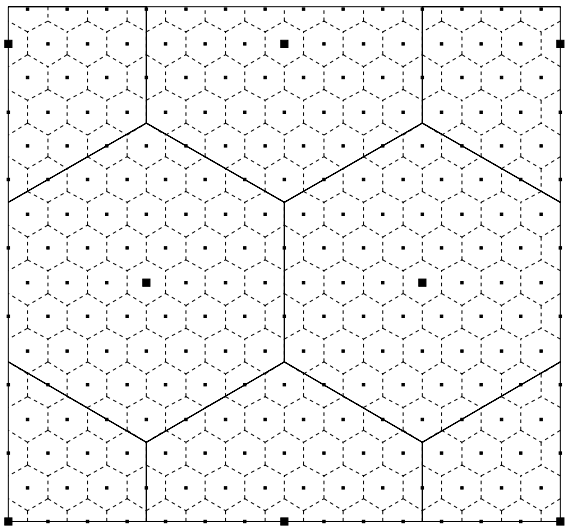

FIGURE 1.2. Patches (solid) and Elements (dashed).

ture, or the Delauny triangulation can be used to calculate the Voronoi diagram. Both representations allow the incremental insertion of sites. (see e.g. [9] or [12] Chapter VIII.7 p.260).

1.2.2 Subdivison into elements

To generate the elements we insert more sites into the Voronoi diagram. We choose these (element-) sites in an identical way to the original (patch-) sites, namely arranged in a regular grid. The insertion of these sites will destroy the original edges of the patches, therefore the area of an element will not neccesarily be enclosed by its respective patch polygon, but the sampling site of an element will always be inside the patch polygon (shown for elements arranged in a hexagonal grid in Fig. 1.2). Because the radiosity values are calculated only for a delta-area around each site, this will lead to no problems.

We now associate each patch with the elements whose sites it is enclosing. The generation of the elements necessitates a recalculation of the area of the patches.

1.2.3 Radiosity with Voronoi diagrams

The area of interest for each sampling site can be found from the Voronoi diagram easily by calculating the area of the Voronoi polygon. The unbounded Voronoi polygons at the

edge of the surface have to be clipped at the border before their area can be calculated correctly (e.g. see Fig. 1.3).

For the radiosity calculation we use the normal progressive radiosity algorithm with raytraced form-factors [19], but the form-factor calculation for the shooting patch should be modified analogous to Tampieri [17], as the patch polygon does not enclose the elements. The latter method takes the element areas and element radiosities of the shooting patch into account and interpolates radiosities when neccessary (e.g. when the shooting patch is close to the receiver). This algorithm adapts easily to our method, if we use an n-ary tree of triangles. Note that the area associated with each of the internal nodes of the tree should be the sum of all element areas whose sampling points lie inside the bordes of the node.

The Delauny triangulation is used for interpolation of radiosity values. A property of the Delauny triangulation is that it maximizes the minimal angle over all triangulations for a given set of sites, which makes it well suited for computer graphics algorithms. The Delauny triangulation is also very well suited for the final output of the scene.

1.2.4 Adaptive subdivison on radiosity discontinuities

During the progressive radiosity algorithm we can check if an element is crossed by a radiosity discontinuity by comparing the radiosity value of the sampling site with the radiosity values of the neighbours. If we find such an element we can subdivide it, interpolate the radiosity for the newly created element sites and recalculate the contribution of the current shooting patch for the sites.

The adaptive subdivision uses a subdivision scheme for arbitrary polygons. If an element is subdivided we have to check that the subdivision level of the neighbouring elements differs by at most one, otherwise the neighbouring elements have to be subdivided recursively, too. This leads to a mesh with a gradual change between regions with different sampling densities. Thus interpolation artifacts which would be visible in the output are avoided.

This method will not introduce any T-vertices, as we still work with a Voronoi diagram. The adaptive subdivision will change the area of an element and the patch, therefore the areas have to be recalculated. Note that the area of all neighbouring elements and patches may have to be modified, too.

See Fig. 1.3 and 1.4 for an example where the algorithm has subdivided the elements along the shadow edges of a triangle. The initial regular grid used was 9 by 9 vertices and a quadtree-like subdivision scheme was used for simplicity.

1.3 Results

If we use a regular interleaved grid of points then the resulting polygons will be hexagons which are a good approximation for a disk. Wallace's method of computing the form-factors with raytracing [19] will therefore yield better results as his method approximates the shooting patch by a disk.

Note that the Delauny triangulation of the sampling sites (Fig. 1.4) produces almost the same results as Baum's mesh elements [2].

This method was implemented using the Delauny triangulation as the internal representation. Initial experiments showed that the shading of the mesh shown in Fig. 1.4 produced undesireable artifacts along the shadow edge (Fig. 1.5, top). The mesh generation algorithm was modified to produce a mesh like in Fig. 1.6 which yielded a better result (Fig. 1.5, bottom). This was achieved by modifying the incremental insertion algo-

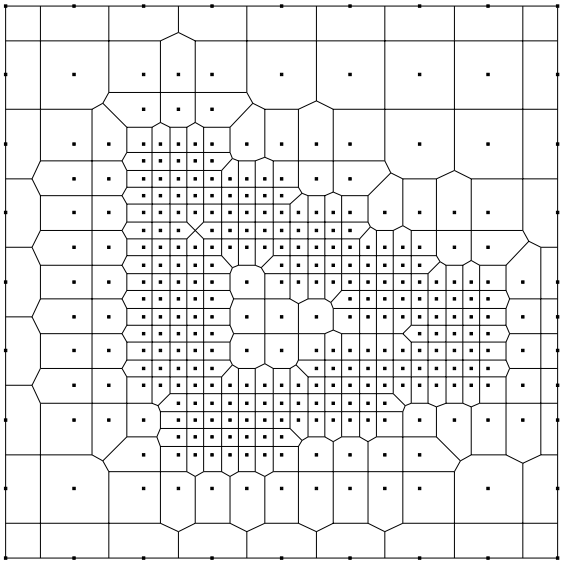

FIGURE 1.3. Adaptive subdivision with quadrilaterals (Voronoi Diagram).

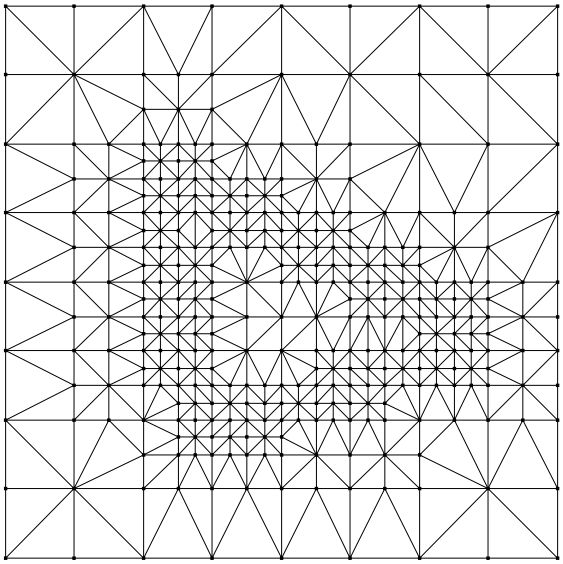

FIGURE 1.4. Adaptive subdivision with quadrilaterals (Delauny Triangulation).

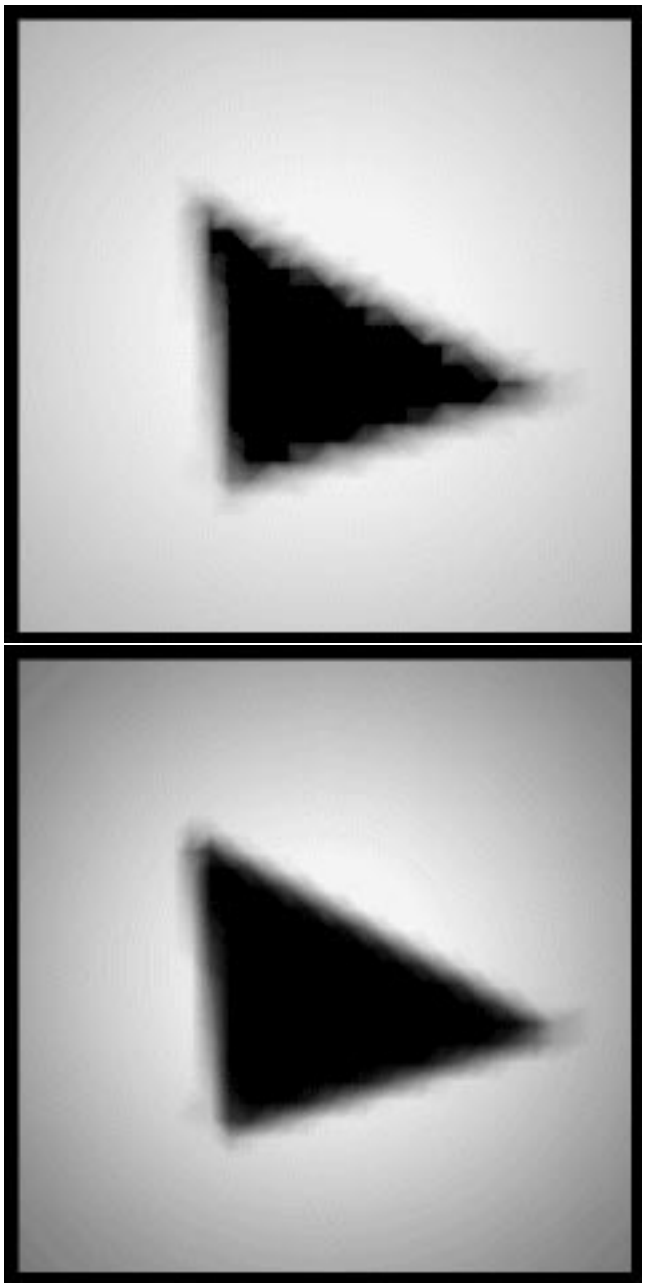

FIGURE 1.5. Shaded versions of Delauny Triangulation (top) and Modified Delauny Triangulation (bottom).

rithm to handle ambigous situations (e.g. four sampling sites lying on a circle) in such a way that aliasing artifacts are reduced, i.e. if a radiosity discontinuity crosses a quadrilateral the edge seperating the two triangles is chosen to approximate the direction of the discontinuity.

The geometry used to generate these images was similar to [19], but we used only 341 samples to generate these images.

1.4 Conclusion

A new method to store the illumination across a surface for radiosity has been presented. This method allows the placements of sampling sites anywhere on a surface and neverthe-

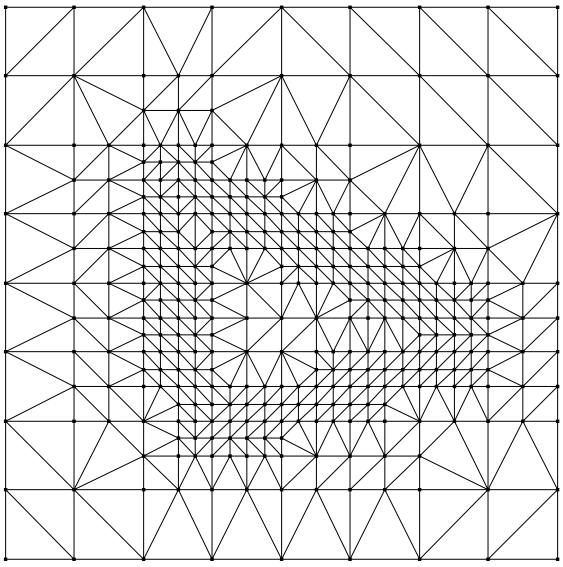

FIGURE 1.6. Modified Delauny Triangulation for adaptive subdivision with quadrilaterals.

less we are still able to calculate the area of interest for each site correctly. The adaptive subdivision algorithm presented will only produce well-formed elements, T-vertices and other artifacts never appear, so solution and interpolation artifacts are avoided. Also we are not limited to a mesh of quadrilaterals and triangles but can use arbitrary polygons.

This method is well suited for form-factor calculations using raytracing because the polygons can be chosen to approximate a disk more closely (i.e. hexagons), and therefore the computed form-factors will be more exact, although this needs to be investigated.

Implicit surface intersection are still treated like shadow edges, therefore the method is prone to artifacts like shadow- and light leaks. We are currently investigating methods that make possible the direct representation of polygons in the Delauny triangulation.

1.5 Further Extensions

- The two-dimensional coordinate system mapped onto the surface of e.g. a cylinder may produce artifacts at the 'seam'. It is possible to extend the definition of Voronoi diagrams to handle a sphere or a cylinder correctly.
- The ability to place sampling sites anywhere on the surface allows the placement of sites on or near shadow edges once these have been identified. This could be used to produce even better meshes.
- If we use a constrained Voronoi diagram [14] it is possible to insert edges directly into the Voronoi diagram. This could be used to represent shadow edges on the surfaces directly [4].

Acknowledgements:

The author thanks Ben Trumbore, Filippo Tampieri and Martin Feda for helpful comments on the paper. Special thanks to Heinz Herbeck and Michael Gervautz for their support.

1.6 References

[1] Franz Aurenhammer. A Survey of a Fundamental Geometric Data Structure. *ACM Computing Surveys, Vol 23, Number 3*, August 1991.

[2] Daniel R. Baum, Stephen Mann, Kevin P. Smith, and James M. Winget. Making Radiosity Usable: Automatic Preprocessing and Meshing Techniques for the Generation of Accurate Radiosity Solutions. In *Computer Graphics (SIGGRAPH '91 Proceedings)*, July 1991.

[3] Daniel R. Baum, Holly E. Rushmeier, and James M. Winget. Improving Radiosity Solutions through the Use of Analytically Determined Form-Factors. In *Computer Graphics (SIGGRAPH '89 Proceedings)*, July 1989.

[4] A.T. Campbell and Donald S. Fussell. Adaptive Mesh Generation for Global Diffuse Illumination. In *Computer Graphics (SIGGRAPH '90 Proceedings)*, August 1990.

[5] Michael Cohen, Shenchang Eric Chen, John R. Wallace, and Donald P. Greenberg. A Progressive Refinement Approach to Fast Radiosity Image Generation. In *Computer Graphics (SIGGRAPH '88 Proceedings)*, August 1988.

[6] Michael Cohen and Donald P. Greenberg. The Hemi-Cube: A Radiosity Solution for Complex Environments. In *Computer Graphics (SIGGRAPH '85 Proceedings)*, August 1985.

[7] Michael Cohen, Donald P. Greenberg, Dave S. Immel, and Phillip J. Brock. An Efficient Radiosity Approach for Realistic Image Synthesis. *IEEE Computer Graphics and Applications*, March 1986.

[8] Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. Modelling the Interaction of Light Between Diffuse Surfaces. In *Computer Graphics (SIGGRAPH '84 Proceedings)*, July 1984.

[9] P. J. Green and R. Sibson. Computing Dirichlet tesselations in the plane. *Computing Journal 21*, 1977.

[10] Eric A. Haines. Ronchamp: A Case Study for Radiosity. In *SIGGRAPH 91 Course Notes on Frontiers in Rendering*, 1991.

[11] Thomas J. V. Malley. A Shading Method for Computer Generated Images. Master's thesis, University of Utah, June 1988.

[12] Kurt Mehlhorn. *Data Structures and Algorithms 3: Multi-dimensional Searching and Computational Geometry.* Springer Verlag, 1984.

[13] Franco P. Preparata and Michael I. Shamos. *Computational Geometry.* Springer Verlag, 1985.

[14] Raimund Seidel. Constrained Delauny triangulations and Voronoi diagrams with obstacles. Technical report, TU Graz, Austria, 1988. Rep. 260 IIG.

[15] Francois Sillion and Claude Puech. A General Two-Pass Method Integrating Specular and Diffuse Reflection. In *Computer Graphics (SIGGRAPH '89 Proceedings)*, July 1989.

[16] Francois X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg. A Global Illumination Solution for General Reflectance Distributions. In *Computer Graphics (SIGGRAPH '91 Proceedings)*, July 1991.

[17] F. Tampieri and D. Lischinski. The Constant Radiosity Assumption Syndrome. In *Proceedings of the Second Eurographics Workshop on Rendering*, May 1991.

[18] John R. Wallace. Problems with Model Data. In *SIGGRAPH 91 Course Notes on Radiosity*, 1991.

[19] John R. Wallace, Kells A. Elmquist, and Eric A. Haines. A Ray Tracing Algorithm for Progressive Radiosity. In *Computer Graphics (SIGGRAPH '89 Proceedings)*, July 1989.