

Data compression of multiresolution surfaces

Reinhard Klein, Stefan Gumhold

WSI/GRIS University of Tübingen, 72076 Tübingen, Germany

Abstract. In this paper we introduce a new compressed representation for multiresolution models (MRM) of triangulated surfaces of 3D-objects. Associated with the representation we present compression and decompression algorithms. Our representation allows to extract the surface at variable resolution in time linear in the output size. It applies to MRMs generated by different simplification algorithms like local vertex deletion or edge and triangle collapse. The time required to transmit models over communication lines and the space needed to store the MRMs is significantly reduced.

1 Introduction and previous work

Triangle meshes are one of the most popular representations of surfaces for computer graphics applications. On the one hand, rendering of triangles is widely supported by hardware and, therefore, fast. On the other hand, there is an increasing set of data acquisition techniques which generate triangle meshes as output. However, most of these techniques generate much more triangles than necessary to represent the given object with a small approximation error. Iso-surface generation can create 1-10 millions of polygons. A digital map of Germany with a resolution of 40 meters in North-to-South and West-to-East direction results in about 500M points. These huge amounts of data lead to problems with data storage and post-processing programs. Animation and real-time rendering of such data sets is almost impossible even on high performance graphics hardware.

Various techniques were published that aimed to reduce surface complexity in order to speed up rendering time [HH92,MSS94,KT96,SZL92,CVM⁺96,RKH96],[CCMS96,HDD⁺93,Hop96,RR96,RB93,Tur92,EDD⁺95]. Aside from the mesh simplification algorithms recent research also focuses on multiresolution representations of triangle meshes. A complex mesh is replaced by several levels of detail (LODs). There are already a number of MRMs allowing view-dependent extraction of simplified meshes at variable resolution [dBD95,dFP95,CPS95],[KLR⁺95,Pup96,Hop96,KS96].

The models for terrain surfaces and parameterized free-form surfaces in our previous work [KHK96] allow to refrain from storing the connectivity¹ of the meshes as well as the hierarchy between the different LODS. This leads to a

¹ The connectivity of a triangle mesh comprehends all adjacency relationships of the mesh.

massive data reduction for the MRM. In the other, more general models, the connectivity and the hierarchy have to be stored explicitly. Since the number of triangles of the multiresolution model is about three times the number of original triangles, it is often impossible to keep large models in memory. Data reduction of the connectivity and the hierarchy is indispensable for real-time applications.

A first step into this direction were the progressive meshes proposed by Hoppe [Hop96]. Although with this approach high reduction rates are feasible only restricted selective refinement is possible. His recent paper [Hop97] concentrates on this deficiency but at the expense of storage efficiency. Note that the progressive meshes are based on a special mesh simplification technique, the edge collapse operation. The Multi-Triangulation (MT) introduced by Puppo [Pup96] is more general in the sense that it can be combined with different simplification techniques. In [KK97] we generalized the MT proposed by Puppo for surface meshes embedded in 3D, but for large models the storage needed to store the connectivity and the hierarchy is too high. Therefore, in this paper we propose a new approach, which significantly reduces the storage costs of the model.

2 The simplification algorithm and the multiresolution model

2.1 The simplification algorithm

The simplification algorithm simplifies the input triangulation by successively removing vertices [SZL92]. The minimal one-sided Hausdorff distance between the current triangulation and the simplified one determines which vertex is removed next. A priority queue is used to speed up the computation of the next vertex [KLS96]. All triangles incident to a removed vertex are eliminated from the current triangulation and the resulting hole is retriangulated. From the several possible retriangulations an application dependent optimal one is chosen. A closer look on the different retriangulation strategies reveals that the edge collapse operation can be considered as a special retriangulation technique [KK97]. The simplification algorithm stops if no further vertices can be removed from the simplified triangulation without violating an approximation criterion. The criterion used in our reduction algorithm limits the one-sided Hausdorff distance between the input triangulation and the simplified one.

2.2 The Multi-Triangulation

In the following we briefly describe a data structure for the MT as proposed in [Pup96, KK97]. Besides a list of vertices and a list of triangles the data structure contains a list of *fragments* and a coarse and therefore small starting triangulation, e.g. a tetrahedron. The fragments consist of two small sets of connected triangles, the coarse *floor* and the finer *ceiling* with the same border polygon, see Figure 1. As the ceiling can replace the floor in order to refine the triangulation, the

fragments implicitly define the hierarchy between different levels of detail. In the case of vertex removal, each fragment represents one vertex elimination. The indices of the triangles incident to the removed vertex are listed within the fragment in the list *ceiling* and the indices of the triangles in the retriangulation are stored in the list *floor*. Additionally, the global one-sided Hausdorff distance between the mesh without the removed vertex and the original triangulation is stored in each fragment. Each triangle keeps two pointers to the *upper* and *lower fragment* it belongs to. These interconnect the fragments that overlap and complete the storage of the hierarchy among the fragments.

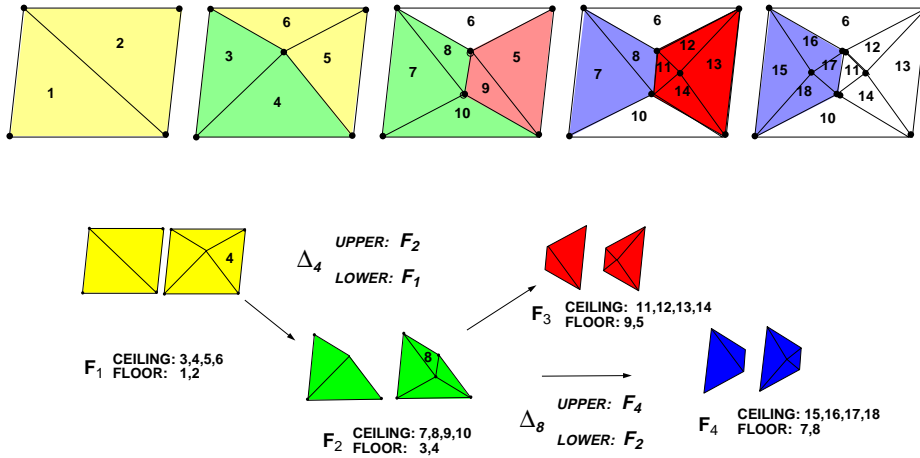


Fig. 1. A fragment contains the indices of the triangles in the floor and the ceiling. Each triangle stores pointers to its upper and lower fragment.

3 Using patterns for compression

The main idea of the compression algorithm is based on the following observations:

1. The number of triangles in the fragments is small and bounded by the maximum order of the removed vertices.
2. The fragments with the same number of vertices only differ in the triangulation of the floor.

Therefore, only a few different triangulation types appear in the fragments. For a pentagon all possible triangulations can be obtained from one of its triangulations by rotations. The storage needs for the connectivity of the MT can be significantly reduced if the connectivity information is stored exclusively for the

*equivalence classes*² of triangulations. To reconstruct a fragment in the current triangulation beside the index of its equivalence class, two indices of adjacent vertices must be known, see figure 2. From these so called *anchor vertices* together with the connectivity of the equivalence class the complete fragment in the current triangulation can be reconstructed, see figure 2. The first anchor vertex determines the position of the fragment in the current triangulation and the second fixes the orientation. In the following the different equivalence classes of triangulations are called *patterns*. In the more general case, where the con-

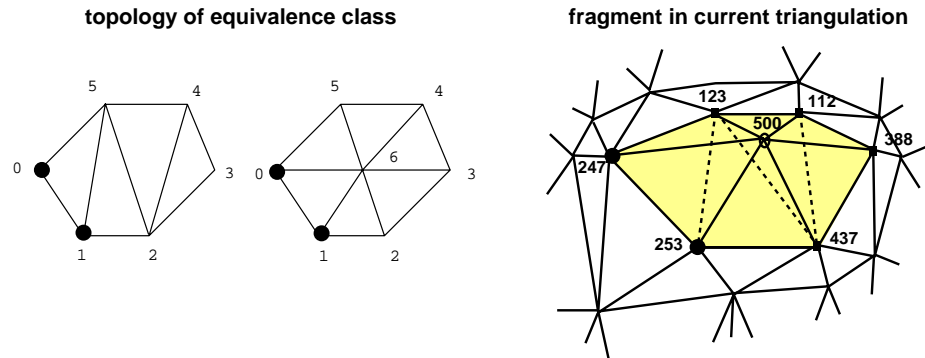


Fig. 2. A fragment is reconstructible from two anchor vertices together with the connectivity of the equivalent class. The anchor vertices 0,1 are equivalent to the vertices 247 and 253 in the current triangulation.

nectivity of the ceiling is not predetermined, the connectivity of the fragment is stored in form of two patterns composing a *rule*.

3.1 Progressive transmission

The model described so far allows progressive transmission. For each fragment the index of the corresponding rule, the first anchor vertex and the orientation of the fragment is needed. To store the orientation of the fragment we use the fact that the order in which the vertices are transmitted is fixed. Therefore, among the neighbor vertices of the first anchor vertex exists a vertex with smallest index. Starting from this vertex the neighbors of the first anchor vertex are indexed counterclockwise. If the order of the first anchor vertex is limited to 16, four bits are enough to uniquely define the second anchor vertex and thus the orientation of the fragment. For progressive transmission the storage of each fragment decomposes into one byte for the index of the pattern, four bytes for the index of the first anchor vertex and 4 bits to encode the orientation. This sums up to 44 bits per fragment. In the case of highly reduced models the number

² Each type of triangulation forms an equivalence class.

of fragments is the same as the number of vertices. This yields nearly the same compression rate like progressive meshes – also for the more general case of vertex removal. The only overhead compared to the progressive meshes is the storage needed for the rules, which is negligible as each floor triangulation can be encoded in one triangle strip and the ceilings are predetermined.

3.2 Compressing the hierarchy

As the hierarchy in this stage is not explicitly stored the model does not allow to extract a triangulation at variable resolution in time linear in the output size. In the original MT all fragments contain pointers to their triangles and vice versa all triangles contain pointers to their upper and lower fragments to store the hierarchy. Without further compression techniques the amount of storage required is 160 bytes per vertex ³, [Gum97].

Since in our pattern based approach the triangles are implicitly stored in the equivalence classes, the hierarchy is stored using links between the fragments. To each fragment f a closed linked list called *loop* is assigned that begins and ends at f .

The loop contains all fragments in the hierarchy whose ceilings join triangles with the floor of the loop fragment, see figure 3. These fragments are directly covered by the loop fragment. Each element of the closed linked list is stored in the fragment pointed to by the predecessor of the element. This allows bi-directional navigation among the directly overlapping fragments. The covering fragments can be found by following all the loops to the end, that visit the current fragment. As the number of triangles in a fragment is limited, also the length of the loops are limited and therefore the covering fragment is found in constant time.

To analyze the storage needs of our approach we have to count the number of links in all loops. Each link corresponds to a direct overlap of two fragments. Thus we introduce the *number of overlaps* u . The relation between u and the number of fragments f where measured in several large MRMs and yielded in good approximation $u \approx \frac{5}{2}f$ [Gum97]. In addition to the u links for each fragment the following quantities have to be stored to give full access to the connectivity and the hierarchy:

1. the indices of two anchor vertices
2. the index of the rule representing the fragment's connectivity
3. the number of loops the fragment is contained in
4. the downward link.

As the extracted triangulation is only restricted to the hierarchy, the neighbors of a vertex are not fixed. This is the reason why the second anchor vertex has to be explicitly stored.

In [KK97] it is proposed to limit the order of the removed vertices to accelerate the simplification algorithm. A maximum order of ten limits the maximal

³ We assume that vertex indices and pointers are encoded with 32 bits.

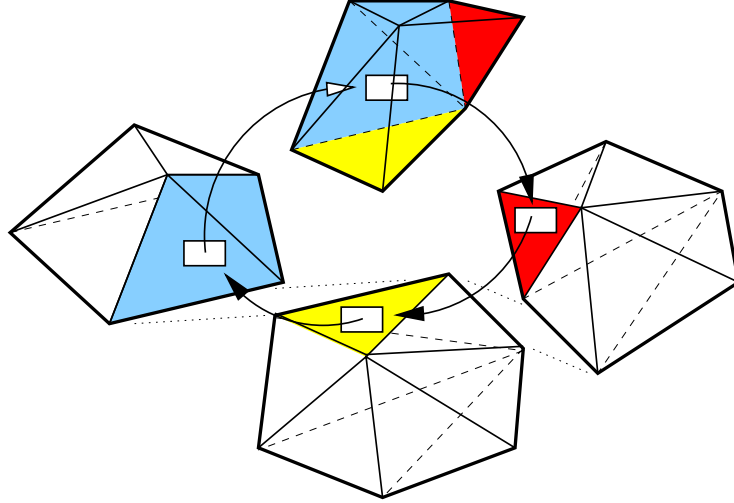


Fig. 3. The floor of the upper fragment covers three other fragments. The four fragments are linked together starting and ending at the covering fragment. Note that the elements of the single linked list are stored in the fragments.

number of rules to 7147 [Gum97]. As the number of triangles in the ceiling of the fragments is the same as the order of the removed vertex, also the maximal number of loops a fragment is contained in is limited by ten. Thus the rule index and the number of loops can together be encoded in sixteen bits.

As each loop points back to the beginning, it contains one element more as the number of overlaps it describes. Therefore, the downward links in the fragments must be counted separately. Each link consists of a fragment index and a loop index. If we assume 32 bits per vertex index, also a fragment index can be encoded in 32 bits and thus a link requires 36 bits.

Altogether, the storage needed for the connectivity and the hierarchy in the pattern based multiresolution model sums up to:

$$\begin{aligned}
 S_{pbMRM} &\approx f \times (64 + 16 + 36)\text{bits} + u \times 36\text{bits} \\
 &\approx f \times \left(116 + \frac{5}{2} \cdot 36 \right) \text{bits} = 206f\text{bits}
 \end{aligned}$$

Recalling that the number of fragments is about the number of vertices, 26 bytes per vertex are required for the pattern based multiresolution model. Compared to the MT data structure this corresponds to a reduction factor of six. Again the overhead needed to store the rules is negligible for large models.

3.3 Building a compressed multiresolution model

There are two possibilities to build a compressed multiresolution model. Either the model is extracted from a MT or it is directly build from the simplification

algorithm. The first method is slightly simpler as the hierarchy is explicitly stored in the MT. The simplification algorithm successively provides the triangles in the ceiling and the floor of the fragment corresponding to the removed vertex. If for each triangle in the current triangulation temporarily its upper fragment is stored, the covering fragments can be determined from the triangles in the ceiling. The newly produced fragment must be inserted into the loops of the covering fragments.

In both methods the difficult part is to extract the rules from the fragments. Therefore, a comparison algorithm for patterns was developed, which runs in $O(n \times o_{max}^2)$ time, where n is the number of vertices in the pattern and o_{max} the maximal order of its vertices. If the comparison of two patterns is successful, the algorithm produces a map between the pattern vertices. This allows to determine the anchor vertices of fragments whose rule was extracted earlier. As o_{max} is limited to eight ⁴ and n to ten, the comparison algorithm runs in constant time. Nevertheless to improve performance we defined a Hash-Code on patterns. Major increments are the number of vertices and the vertex orders. To achieve a distinct Hash-Code for patterns and their mirror images the order of the border vertices is exploited. Our Hash-Code is unique for all floor patterns with less than nine vertices [Gum97].

3.4 Incremental selective refinement algorithm

Within a specific application, a criterion must be available to define the required resolution of the multiresolution model. The standard approach is to define a boolean condition c on the triangles of the multiresolution model. If a triangle is invalid, the triangles in the ceiling of its upper fragment are inserted to refine the triangle. In the pattern based approach the triangles are not stored explicitly. Here the condition c is defined on each overlap of two fragments, i.e. on each upward link. This is completely sufficient and if the conditions of the triangles in an overlap can be combined to one condition, the storage space for c is reduced. For example the Hausdorff distance between a triangle and the original mesh can easily be combined by choosing the maximal distance among the triangles in an overlap. Often it is sufficient to define the condition c only on the fragments.

The extracted triangulation additionally must obey to the criterion of "minimal refinement", i.e. it must be the coarsest triangulation that is extractable from the multiresolution model satisfying c . Our incremental algorithm is based on the storage of the current triangulation, which satisfies c and the "minimal refinement" criterion. If the condition c is changed by the application, the current triangulation is adapted to the new condition c in two steps. First the triangulation is coarsened: top fragments ⁵ containing exclusively valid upward links are

⁴ Only the floor patterns must be compared. These only contain border vertices, the number of which is restricted by the removed vertex.

⁵ All triangles in the ceiling of a top fragment are contained in the current triangulation and thus the fragment lays on top of the hierarchy if the hierarchy is restricted to the fragments with ceiling triangles in the current triangulation.

removed from the triangulation, until no such fragments are left over. To remove a fragment from the current triangulation, the vertices of the rule are mapped to the vertices of the current triangulation by using the two anchor vertices and a modified comparison algorithm. Then the triangles in the ceiling of the fragment are replaced by the triangles in the floor.

In the second step the current triangulation is refined until all overlaps contained in the current triangulation satisfy the condition c . Refinement is achieved by inserting a fragment: in the current triangulation the floor triangles of a fragment are replaced with its ceiling triangles, introducing the previously removed vertex. Two problems have to be solved:

1. How to find the fragments to the overlaps contained in the current triangulation ?
2. How to ensure that the floor of a to be inserted fragment is part of the current triangulation ?

To solve the first problem, we enumerate the vertices and their corresponding fragments in the order the vertices are removed by the reduction algorithm. This gives us a one to one correspondence between the removed vertices and the fragments. By the way, because of this trick we don't have to store the indices of the removed vertices in the fragments, as they are given by the fragment index. Each triangle in the current triangulation belongs to an overlap that has to be checked for refinement. From the fragments corresponding to the corner vertices of the triangles inserted at last – thus the one with the smallest index – must contain the triangle in its ceiling, because all triangles around an inserted vertex are newly introduced to the current triangulation.

The second problem can be solved with the recursive algorithm *layFoundation*:

```

Algorithm layFoundation( $f$ )
for all fragments  $f_i$  in loop below of  $f$  do
    if  $f_i$  is not inserted then
        layFoundation( $f_i$ )
        insertFragment( $f_i$ )
    endif
endfor

```

Before insertion a fragment must be *founded*, i.e. all triangles in the floor of the fragment are forced into the current triangulation. This is achieved by following the downward loop to all covered fragments and insert them in the same way. As each fragment is visited exactly once by *layFoundation*, the refinement algorithm is linear in the output size.

4 Conclusion

All simplification algorithms based on vertex removal, edge or triangle collapse are well suited to build up the pattern based multiresolution model proposed in this paper. This multiresolution model supports selective refinement

as necessary for view-dependent visualization. The high reduction rates of the storage costs achieved by the use of patterns allow to process much larger models than with the other multiresolution models described in the literature ⁶.

References

- [CCMS96] A. Ciampalini, P. Cignoni, C. Montani, and R. Scopigno. Multiresolution decimation based on global error. Technical Report CNUCE: C96-021, Istituto per l'Elaborazione dell'Informazione - Consiglio Nazionale delle Ricerche, Pisa, ITALY, July 1996.
- [CPS95] P. Cignoni, E. Puppo, and R. Scopigno. Representation and visualization of terrain surfaces at variable resolution. In R. Scatenied, editor, *Scientific Visualization 95 (Int. Symp. Proc.)*, pages 50–68. World Scientific, 1995.
- [CVM⁺96] Jonathan Cohen, Amitabh Varshney, Dinesh Manocha, Greg Turk, Hans Weber, Pankaj Agarwal, Frederick P. Brooks, Jr., and William Wright. Simplification envelopes. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 119–128. ACM SIGGRAPH, Addison Wesley, August 1996. held in New Orleans, Louisiana, 04-09 August 1996.
- [dBD95] M. de Berg and K. T.G. Dobrindt. On levels of detail in terrains. Technical Report UU-CS-1995-12, Department of Computer Science, Utrecht University, April 1995.
- [dFP95] Leila de Floriani and Enrico Puppo. Hierarchical triangulation for multiresolution surface description. *ACM Transactions on Graphics*, 14(4):363–411, October 1995.
- [EDD⁺95] M. Eck, T. DeRose, T. D., H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 173–182. ACM SIGGRAPH, Addison Wesley, August 1995. held in Los Angeles, California, 06-11 August 1995.
- [Gum97] S. Gumhold. Compression of discrete multiresolution models. Master's thesis, Wilhelm-Schickard-Institut für Informatik, Universität Tübingen, June 1997.
- [HDD⁺93] Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Mesh optimization. In James T. Kajiya, editor, *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 19–26, August 1993.
- [HH92] Charles Hansen and Paul Hinker. Isosurface extraction SIMD architectures. In *Visualization'92*, pages 1–21, oct 1992.
- [Hop96] H. Hoppe. Progressive meshes. In *Computer Graphics Proceedings, Annual Conference Series, 1996 (ACM SIGGRAPH '96 Proceedings)*, pages 99–108, 1996.
- [Hop97] H. Hoppe. View-dependent refinement of progressive meshes, 1997. accepted for SIGGRAPH '97.
- [KHK96] R. Klein, , T. Hüttner, and J. Krämer. Viewing parameter dependent approximation of nurbs-models for fast visualization and animation using a discrete multiresolution representation. In B. Girod, editor, *Herbsttagung '96 3D Bildanalyse und -synthese*, 1996.

⁶ The reduction factor compared to Puppoo's Multi-Triangulation [Pup96] and Hoppe's Progressive Meshes [Hop97] is about 6 and 5, respectively.

- [KK97] R. Klein and J. Krämer. Multiresolution representations for surface meshes. In *Proceedings of the SCCG*, pages 57–66, 1997.
- [KLR⁺95] David Koller, Peter Lindstrom, William Ribarsky, Larry F. Hodges, Nick Faust, and Gregory Turner. Virtual gis: A real-time 3d geographic information system. Technical Report 95-14, Graphics, Visualization and Usability Center, Georgia Institute of Technology, USA, 1995.
- [KLS96] R. Klein, G. Liebich, and W. Straßer. Mesh reduction with error control. In R. Yagel, editor, *Visualization 96*. ACM, November 1996.
- [KS96] R. Klein and W. Straßer. Generation of multiresolution models from cad-data for real time rendering. In W. Straßer, R. Klein, and R. Rau, editors, *Theory and Practice of Geometric Modeling*. Springer-Verlag, 1996.
- [KT96] Alan D. Kalvin and Russel H. Taylor. Superfaces: polygonal mesh simplification with bounded error. *IEEE Computer Graphics and Appl.*, 16(3), May 1996.
- [MSS94] C. Montani, R. Scateni, and R. Scopigno. Discretized marching cubes. In R. D. Bergeron and A. E. Kaufman, editors, *Visualization '94 Proceedings*, pages 281–287. IEEE Computer Society, IEEE Computer Society Press, 1994.
- [Pup96] E. Puppo. Variable resolution of terrain surfaces. In *Proceedings Eight Canadian Conference on Computational Geometry*, August 1996.
- [RB93] J. Rossignac and P. Borrel. Multi-resolution 3d approximation for rendering complex scenes. In B. Falcidieno and T. L. Kunii, editors, *Modeling in Computer Graphics: Methods and Applications*, pages 455–465. Springer Verlag, 1993.
- [RKH96] Daniel Cohen-Or Reinhard Klein and Tobias Hüttner. Incremental view-dependent multiresolution triangulation of terrain, 1996. submitted to Pacific Graphics.
- [RR96] R. Ronfard and J. Rossignac. Full-range approximation of triangulated polyhedra. *Computer Graphics Forum*, 15(3):C67–C76, C462, September 1996.
- [SZL92] William J. Schroeder, Jonathan A. Zarge, and William E. Lorensen. Decimation of triangle meshes. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 65–70, July 1992.
- [Tur92] Greg Turk. Re-tiling polygonal surfaces. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 55–64, July 1992.