

Restricted Mesh Simplification Using Edge Contractions

Mattias Andersson *

Joachim Gudmundsson †

Christos Levkopoulos‡

Abstract

We consider the problem of simplifying a triangle mesh using edge contractions, under the restriction that the resulting vertices must be a subset of the input set. That is, contraction of an edge must be made onto one of its adjacent vertices. In order to maintain a high number of contractible edges under this restriction, a small modification of the mesh around the edge to be contracted is allowed. Such a contraction is denoted a 2-step contraction. Given m “important” points or edges it is shown that a simplification hierarchy of size $O(n)$ and depth $O(\log n/m)$ may be constructed in $O(n)$ time. Further, for many edges not even 2-step contractions may be enough, and thus, the concept is generalized to k -step contractions.

1 Introduction

In computer graphics objects are commonly represented using triangle meshes. One important problem regarding these meshes is how to efficiently simplify them, while maintaining a good approximation of the original mesh. As an example, scanners often produce information redundant meshes containing millions of points and triangles. Further, often the simplification should be performed in several rounds, such that a level-of-detail hierarchy is constructed. One application of such a hierarchy is that an appropriate level may be chosen depending on viewing distance, as finer details tend to be unnecessary as the distance increases. Other applications include progressive transmission and efficient storing. It is common to represent the level-of-detail hierarchy as a directed, acyclic and hierarchical graph, where each level in the graph corresponds to a level in the level-of-detail hierarchy, and where each node in the graph corresponds to a triangle, in the natural way. The first, top-most, level in the graph corresponds to the input mesh. When a contraction is made two triangles disappear, and one or more triangles are affected in such a way that their appearance change. In the graph this is represented with edges between disappearing triangles at

some level i , and the affected triangles at level $i + 1$. Such a graph will simply be denoted a *hierarchical graph*, and the efficiency of a simplification algorithm is directly related to the size [3, 11] and depth of the hierarchical graph it may produce. Simplification algorithms constructing hierarchies of size $O(n)$ and depth $O(\log n)$ have been presented for several problem variants [1, 2, 4, 10]. Mesh simplification is generally regarded as a mature field (see [5, 8] for surveys), consisting of several suggested methods and problem variants. In this paper the method of iteratively contracting edges [1, 6, 7, 9] is considered, where contractions are made such that no crossing edges result from the contraction. This method is examined under the restriction that the set of output points is required to be a subset of the input points, i.e., contraction of an edge must be done onto one of its adjacent vertices. In order to maintain a high number of contractible edges a small modification of the mesh around an edge to be contracted is allowed. This method is denoted a 2-step contraction, and it is shown that a hierarchical graph of size $O(n)$ and depth $O(\log n/m)$ may be produced under the restriction that several “important” points or edges may not be contracted. Further, in order to enable contraction of edges that are not even 2-step contractible the concept is generalized to k -step contractions. We show that k can be bounded by either $\deg(v) - 4$ of a vertex v to be contracted, or by the number of concave corners on the hull of v (see Section 2 for definition).




Figure 1: (a) A planar triangulation with a rectangular outer hull is given as input. (b) Illustrating an edge contraction (merge-operation).

2 Contracting in k steps

As input we are given a planar triangulation T . We can assume that the outer hull of T is a rectangle, as illustrated in figure 1a.

*Department of Computer Science, Lund University, mattias@cs.lth.se

†National ICT Australia Ltd., IMAGEN Program, Joachim.Gudmundsson@nicta.com.au

‡Department of Computer Science, Lund University, christos@cs.lth.se




Figure 2: Illustrating a 2-step contraction of a degree 6 node v .

The aim is to simplify T by iteratively performing edge contractions, as shown in Fig. 1b, where an edge (u, v) can be contracted such that u is moved to v , or v is moved to u . A problem that often occurs during edge contractions of triangulations is that the resulting triangulation might not be planar. An edge contraction is said to be *valid* if the resulting triangulation still is planar. In this paper we consider the problems of defining valid contractions and computing valid contractions. Below some basic operations and notations are defined:

Consider a vertex v and assume the degree of v is d , and let u_1, \dots, u_d be the vertices adjacent to v in clockwise order around v . The *hull* of v , denoted $\mathcal{H}(v)$, is the cycle described by the edges $(u_i, u_{(i+1) \bmod d})$ for $1 \leq i \leq d$. We need some additional definitions. A *split* operation splits a vertex v of degree d into two vertices v and v_1 connected by an edge such that the resulting triangulation is planar, v has degree d' and v_1 has degree $d - d' + 4$, as illustrated in Figure 2a-b. Next, a *merge* operation contracts one vertex u onto another vertex u' , both connected by an edge in T , into one vertex u' , as illustrated in Figure 2c-d. Let $d(u')$ denote the degree of u' before the contraction. After the contraction u' has degree $d(u) + d(u') - 3$. And, finally, a *split-and-merge* operation on vertices v and u first performs a split operation of v , followed by a merge operation on v_1 and u . The split-and-merge operation is said to be *valid* if the triangulation is planar at each step of the operation. Note that an edge contraction is obtained by a single merge operation.

Next we define 1-step contractible using the merge concept, and we then generalize this concept into k -step contractible. If two vertices u and v , connected by an edge in the triangulation T can be merged into a new vertex w placed at u such that the contracted triangulation still is planar, then v is said to be *1-step contractible* (at u). A vertex v is said to be *k -step contractible* if and only if one can perform at most $k - 1$ valid split-and-merge operations followed by a 1-step contraction of v .

With regards to guaranteeing a hierarchical simplification graph of small size and depth, mainly 2-step contractions will be considered. Figure 2 shows a vertex v that is 2-step contractible since a valid split-and-




Figure 3: No edge in region A can be contracted, unless using k -step contractions.




Figure 4: The two cases of Theorem 1.

merge operation is followed by a 1-step contraction of v . However, to be able to change the level of details for models where the user is allowed to perform the standard operations - rotate, zoom and so on, we would ideally perform edge contractions in small specified areas. For example, in configurations like the one shown in Figure 3, if one wants to perform k -step contractions which change or contract only edges inside the area A , then k has to be at least proportional to the number of edges inside A (divided by some constant).

Thus, the generalized concept of a k -step contraction is needed, and below (Theorem 1 and 3) upper bounds on k are shown.

Theorem 1 Any vertex v , not on the hull of T , with degree at most m is k -step contractible, where $k = \max\{1, m - 4\}$.

Proof. The theorem is proven by induction on the degree of v .

Base case: Vertices of degree at most four can easily be 1-step contracted. We thus assume that v has degree five, which immediately implies that $\mathcal{H}(v)$ contains five points. Consider the interior of $\mathcal{H}(v)$. If there exists a corner v' of $\mathcal{H}(v)$ which can see all other corners of $\mathcal{H}(v)$ then v is 1-step contractible at v' , and thus, the theorem holds. Next, since $\mathcal{H}(v)$ has at least three convex corners, $\mathcal{H}(v)$ has at most two concave corners. If $\mathcal{H}(v)$ has only one concave corner u then this corner must see all the vertices of $\mathcal{H}(v)$ and, hence, v is 1-step contractible to u . If $\mathcal{H}(v)$ has two concave corners we have two cases, as shown in




Figure 5: Example of a vertex v that is not 1-step contractible.

Figure 4. Note that edges between v and $\mathcal{H}(v)$ are not included in order to avoid cluttering of the figure. In the first case, Figure 4a, the concave corner points p and u are not incident, while in the second, Figure 4b, they are.

First case First note that p and u must lie inside the triangle defined by the three convex corners in $\mathcal{H}(v)$, and also that p and u always see each other. There exist two incident convex corner points p' and u' , such that p' is incident on p and u' is incident on u . It is straightforward to see that p sees all points of $\mathcal{H}(v)$ if the edge (p, p') does not cross the line-extension of the segment (u, u') , as p then can see u' . The same holds for u , the edge (u, u') and the line extension of (p, p') . However, both cases can not occur simultaneously as this implies that the edges (u, u') and (p, p') must cross. Thus, either p or u can see all of $\mathcal{H}(v)$.

Second case Consider the one convex corner point q not incident on either p or u . Since q is connected to both p' and u' , q will see all corners of $\mathcal{H}(v)$ if and only if q sees both p and u . Next, consider the line-extension l of the segment (p, u) . Since p and u are concave corners p' and u' must lie on the same side of l . Further, q must connect to p' and u' such that p' and u' form convex angles and p and u form concave angles. This means that q must lie on the opposite side of p and u , with regards to l , which immediately implies that q can see both p and u .

Induction hypothesis: Assume that the theorem holds for all vertices of degree at most $m - 1$.

Induction step: Assume that v has degree m . There exists a point u on the hull of v that can see at least four consecutive vertices of the hull including itself. Denote these vertices u_1, \dots, u_4 . Split v into v_1 and v such that v_1 is connected to u_1, \dots, u_4 and v . Now, v_1 can be 1-step contracted at u . The degree of v is now $m - 1$, thus applying the induction hypothesis on v proves the theorem. \square

Note that if v has degree 6 then it might not be 1-step contractible as shown in Figure 5.

Corollary 2 *At least two edges in T are 1-step contractible.*

Proof. Follows from Theorem 1 and the fact that any planar graph has total degree at most $6n - 12$ (Euler's theorem). Details omitted. \square

Note that if only few edges are 1-step contractible then almost all vertices in T must have degree 6. The bound stated in Corollary 2 is probably very conservative. If the number of 1-step contractible edges is small that implies that almost all vertices of T have degree 6. However, we have not been able to construct any examples where almost all vertices have degree 6 while simultaneously being not 1-step contractible. Next we present an alternative bound on k . As only concave corners restrict visibility, intuitively it should be easier to contract a vertex with few concave corners on its hull. The following theorem can be shown.

Theorem 3 *Every vertex v , not on the hull of T , with at most c concave vertices on its hull is k -step contractible, where $k = c$.*

Proof. Proof omitted. \square

3 The number of k -contractible edges

Allowing k -step contractions increases the flexibility of simplification since it allows a greater fraction of the edges to be contracted. In this section a lower bound on the number of k -contractible edges is given, where Theorem 1 and the fact that the total degree is bounded is used.

Observation 1 *At least $\binom{k-1}{k+2}n$ vertices are k -step contractible, for any $k \geq 2$.*

Proof. Let L be the set of interior vertices of T that are k -step contractible. From Theorem 1 we know that L at least consists of all vertices of degree at most $d = k + 4$. Let N be the remaining set of interior vertices. Recall that the sum of the degrees over all vertices is at most $6n - 12$. This implies that the size of L is minimized when all vertices in N has degree $(d + 1)$ and the vertices in L and the four vertices on the hull has degree 3. We have the following equation:

$$4 \cdot 3 + (d + 1)|N| + 3|L| = 6n - 12 \text{ where } |N| + |L| = n - 4.$$

As a result it holds that $|L| \geq \left(\frac{d-5}{d-2}\right)n \geq \left(\frac{k-1}{k+2}\right)n$. \square

4 The hierarchical graph

In this section we show that using 2-step contractions we can achieve a hierarchical graph, as defined in the introduction, of size $O(n)$ and depth $O(\log n/m)$, given m important points or edges that may not be contracted. In order to do this several edges must be simultaneously contracted in each round, that is, at




Figure 6: Initially x and y are contractible, but after x has been contracted y is no longer contractible.

each level of the graph. Next, note that a previously valid 1-step contractible edge might become invalid after other edges have been contracted, as shown in Figure 6. In order to avoid this problem, for the purpose of finding simultaneously contractable edges, we consider independent edges. Let S'_2 be the set of 2-step contractible vertices of degree at most six. Combining Theorem 1 and Observation 1 it is straightforward to see that $|S'_2| \geq \frac{n}{4}$. Since a vertex in S'_2 has at most six neighbors we can choose at least $\frac{n}{4 \cdot 7} = \frac{n}{28}$ vertices from S'_2 such that none of these chosen vertices has a neighbor from S'_2 . Thus, there exists a constant fraction γ of independent 2-step contractible vertices, and the following theorem can be shown.

Theorem 4 *Given m important points $S'' \subset S$ in a triangulation T one can perform $O(\log n/m)$ rounds of 2-step contractions to obtain a triangulation T' of a point set S' with complexity $O(m)$ such that $S'' \subseteq S' \subset S$.*

Proof. Let n_i denote the number of vertices before round i and consider an arbitrary constant $\delta < \gamma$. Perform rounds until $m \geq \delta n_i$, that is until the resulting point set S' have complexity $O(m)$. This is possible, since as long as $m \leq \delta n_i$, there are at least $\gamma n_i - \delta n_i = (\gamma - \delta)n_i$ 2-step contractible vertices remaining, containing no important point. Thus, T' can be obtained using at most $O(\log_{\frac{1}{\gamma-\delta}} n - \log_{\frac{1}{\gamma-\delta}} m) = O(\log n - \log m) = O(\log n/m)$ rounds of contractions. \square

Corollary 5 *Using rounds of 2-step contractions a hierarchical graph of size $O(n)$ and depth $O(\log n/m)$, given m important points, may be produced in $O(n)$ time.*

Proof. Note that the above theorem immediately enables the construction of hierarchical graph of depth $O(\log n/m)$. Next, consider the size. Note that the number of nodes in the hierarchical graph is $O(n)$ and only 2-step contractible vertices of degree at most six are used during the rounds of contractions. This means that at most four triangles are affected by a contraction, which implies that each node in the hierarchical graph has at most four incident edges. Thus, the hierarchical graph has size $O(n)$

Next, consider the time complexity of creating the hierarchical graph. Note that the Theorem 4 was shown using only 2-step contractible edges of constant degree (at most six). Thus, in each round i the set of γn_i independent 2-step contractible edges can be found in $O(n_i)$ time. This means, since $n_i \leq n(\gamma)^{i-1}$, that the total running time is $O(n + n\gamma + n(\gamma)^2 + \dots + n(\gamma)^{O(\log n/m)}) = O(n(\gamma + \gamma^2 + \dots + \gamma^{O(\log n/m)})) = O(n)$. \square

Finally, note that the above results also holds for m important edges (or m edges and vertices, in total), since each important edge restricts possible contraction for only a constant (two) number of vertices.

References

- [1] S. Cheng, T. Dey and S. Poon. Hierarchy of Surface Models and Irreducible Triangulations. *Computational Geometry Theory and Applications (CGTA)*, 27:135-150, 2004.
- [2] M. de Berg, K. T. G. Magillo. On Levels of Detail in Terrains. In *Proc. ACM Symposium on Computational Geometry*, pp.26-27, 1995.
- [3] L. De Florian, P. Magillo and E. Puppo. Building and Traversing a Surface at Variable Resolution. In *Proc. IEEE Visualization'97*, pp.103-110, 1997.
- [4] C. A. Duncan, M. T. Goodrich and S. G. Kobourov. Planarity-Preserving Clustering and Embedding for Large Planar Graphs. In *Proc. Graph Drawing '99*, pp.186-196, 1999.
- [5] M. Garland. Multiresolution Modelling: Survey and Future Opportunities. *Eurographics '99, State of the Art Report (STAR)*. <http://graphics.cs.uiuc.edu/garland/papers.html>
- [6] S. Gumhold, P. Borodin and R. Klein. Intersection Free Simplification. In *Proc. 4th Israel-Korea Bi-National Conference on Geometric Modeling and Computer Graphics*, pp 11-16, 2003
- [7] P. Heckbert and M. Garland. Surface Simplification Using Quadric Error Metrics. In *Proc. SIGGRAPH'97*, pp 209-216, 1997
- [8] P. Heckbert and M. Garland. Survey of Polygonal Surface Simplification Algorithms. *Multiresolution Surface Modelling Course, SIGGRAPH'97* <http://graphics.cs.uiuc.edu/garland/papers.html>
- [9] H. Hoppe. Progressive Meshes. In *Proc. SIGGRAPH'96*, pp. 99-108, 1996
- [10] D. G. Kirkpatrick. Optimal Search in Planar Subdivisions. *SIAM Journal of Computing*, 12:28-35, 1983.
- [11] J. C. Xia, J. El-Sana and A. Varshney. Dynamic View-Dependent Simplification for Polygonal Models. In *Proc. IEEE Visualization'96*, pp.327-334, 1996.