Hierarchy of Surface Models and Irreducible Triangulations*

Siu-Wing Cheng[†] Tamal K. Dey[‡] Sheung-Hung Poon[†]

July 8, 2003

Abstract

Given a triangulated closed surface, the problem of constructing a hierarchy of surface models of decreasing level of detail has attracted much attention in computer graphics. A hierarchy provides view-dependent refinement and facilitates the computation of parameterization. For a triangulated closed surface of n vertices and genus g, we prove that there is a constant $c > 0$ such that if $n > c \cdot g$, a greedy strategy can identify $\Theta(n)$ topology-preserving edge contractions that do not interfere with each other. Further, each of them affects only a constant number of triangles. Repeatedly identifying and contracting such edges produces a topology-preserving hierarchy of $O(n + g^2)$ size and $O(\log n + g)$ depth. Although several implementations exist for constructing hierarchies, our work is the first to show that a greedy algorithm can efficiently compute a hierarchy of provably small size and low depth. When no contractible edge exists, the triangulation is irreducible. Nakamoto and Ota showed that any irreducible triangulation of an orientable 2-manifold has at most $\max\{342g - 72, 4\}$ vertices. Using our proof techniques we obtain a new bound of $\max\{240g, 4\}$.

Keywords: level of detail, 2-manifold, homology, edge contraction, irreducible triangulation.

1 Introduction

Surface simplification has been a popular research topic in computer graphics [2, 4, 10, 12, 13, 18, 19]. Most practical surface simplification methods apply to triangulated surface models and are based on local updates including vertex decimation and edge contraction. Garland's survey [9] gives a good review of the literature. Vertex decimation removes a vertex together with its incident edges and triangles and then retriangulates the hole left on the surface. Edge contraction collapses an edge to a single vertex (often a new vertex), removing the two incident triangles of the contracted edge and deforming the other triangles touching the contracted edge. If the topology of the surface is not explicitly preserved when applying local updates, the resulting surface might be pinched at a vertex or at an edge. That is, the surface ceases to be a 2-manifold, see Figure 1. Arbitrary topology changes could easily produce visual artifacts (for example, imagine that a rod is squeezed in the

[†]Department of Computer Science, HKUST, Clear Water Bay, Hong Kong. Email: {scheng,hung}@cs.ust.hk

[‡]Department of Computer and Information Sciences, Ohio State University, Columbus, Ohio, USA. Email: tamaldey@cis.ohio-state.edu

*The research was partially supported by Research Grant Council, Hong Kong, China (HKUST 6088/99E) and by NSF under grant CCR-9988216 when the first author visited the Ohio State University.

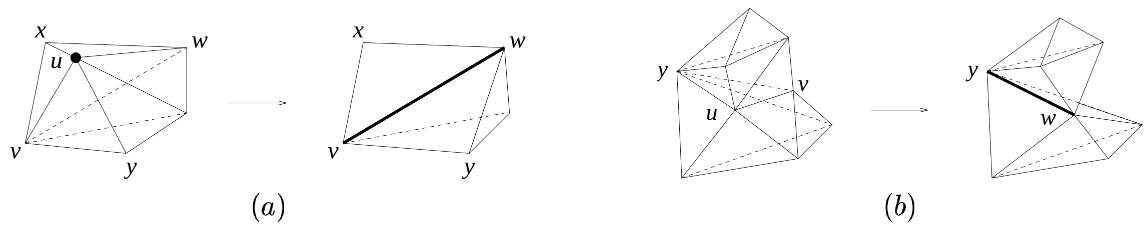

Figure 1: In figure (a), the decimation of u and a retriangulation produce a pinching at the edge vw which could be avoided if xy instead of vw is used in the retriangulation. In figure (b), the contraction of uv to w produces a pinching at the edge wy.

middle by an edge contraction). Also, some applications require that the topology be preserved. Repeated topology-preserving vertex decimation or edge contraction can produce a hierarchy of models of decreasing level of detail that is useful in many applications. For example, Lee et al. [15] compute a parameterization of the triangulated surface model using such a hierarchy, which can be used for remeshing, texture mapping and morphing. In dynamic virtual environments the hierarchy allows objects to be adaptively refined in a view-dependent manner [4, 13, 18, 19]. Basically, undoing a local update increases the local resolution and redoing a local update reduces the local resolution. These applications require the local updates to be independent, that is, they do not affect the same triangle.

A hierarchy can be conceptually viewed as a directed acyclic graph. The nodes at the topmost level are the triangles in the original surface. When applying a local update, nodes are created for the new triangles and arcs are directed from each old triangle affected to the new triangles created. A new level of detail is obtained by applying a set of independent local updates simultaneously. Each local update should affect a small number of triangles as the time complexity of undoing/redoing the local update is proportional to it [4, 19]. Further, the depth of the hierarchy should be small as it bounds the maximum time to obtain a single triangle in the original surface from the model of the lowest level of detail. Given a triangulated surface of n vertices, any hierarchy constructed by repeated applications of independent topology-preserving vertex decimations or edge contractions has depth $\Omega(\log n)$.

For planar subdivisions with straight edges and triangular finite faces, Kirkpatrick [14] and de Berg and Dobrindt [3] showed how to perform independent vertex decimations to construct a hierarchy of $O(\log n)$ depth and $O(n)$ size. Each model in the hierarchy also has straight edges and triangular finite faces. Recently, Duncan et al. [8] showed how to apply planarity-preserving edge contractions to compute a hierarchy of $O(\log n)$ depth for maximal planar graphs. This takes care of triangulated closed surfaces of genus zero as well. Our result resolves the corresponding question for triangulated closed surface of arbitrary genus g, which complements the experimental effectiveness of several existing implementations [4, 15, 19].

The problem of computing the hierarchy of surface triangulations is related to a mathematical question that has been studied before. An edge is *contractible* if its contraction does not change the surface topology. A triangulation of a 2-manifold is called *irreducible* if no edge is contractible. Is there an upper bound on the number of vertices of an irreducible triangulation in terms of the genus g? Barnette and Edelson [1] first proved that a finite upper bound exists. Later, Nakamoto and Ota [16] proved a bound of $270 - 171\chi$, where χ is the Euler's characteristic. This yields a bound of $342g - 72$ for orientable 2-manifolds. (The $270 - 171\chi$ bound works for non-orientable 2-manifolds as well.) This immediately implies that a contractible edge exists when $n > 342g - 72$. If a vertex is not incident on any contractible edge, it remains so after a topology-preserving edge contraction [17]. Thus, there are at least $\lceil (n - 342g + 72)/2 \rceil$ contractible edges. However, in order to construct a hierarchy of low depth, we require the contractible edges to be independent and we need many of them. It is tempting to adapt the analysis of the Dobkin-Kirkpatrick hierarchy [6] to argue that there are linearly many independent edges, but this argument alone is insufficient since we need to guarantee that those independent edges are contractible as well.

Our results include a new upper bound of $240g$ on the number of vertices of an irreducible triangulation. The proof techniques are different from that of Nakamoto and Ota. By using our techniques and by considering a maximal matching of contractible edges, we prove that for any constant $d \geq 444$, if $n > 9182g - 222$, a greedy strategy can identify at least $\frac{n-1310g+30}{64(d+1)}$ independent topology-preserving edge contractions. Each edge contraction affects at most $d + 2$ triangles. This produces a topology-preserving hierarchy of $O(n + g^2)$ size and $O(\log n + g)$ depth (Theorem 3). These results follow from two topological results about triangulations (Theorem 1 and Theorem 2). Since our topological results are applicable to triangulations with curved edges and curved triangles, we do not assume a piecewise linear embedding of triangulations for our topological results.

The rest of the paper is organized as follows. Section 2 provides the basic definitions. Section 3 introduces a family of crossing cycle pairs which is the main tool for obtaining our results. We prove the new upper bound on the number of vertices of an irreducible triangulation in Section 4. Section 5 presents our topological and algorithmic results on constructing a hierarchy.

2 Preliminaries

A surface $\mathcal{M}$ is a *2-manifold* (without boundaries) if for each point $x \in \mathcal{M}$, a sufficiently small open neighborhood of x is homeomorphic to a topological disk. A *triangulation* K of $\mathcal{M}$ is a popular representation of $\mathcal{M}$ in solid modeling and computer graphics. K consists of vertices, edges, and triangular faces. The vertices and edges of K form a graph embedded on $\mathcal{M}$. The embedding partitions $\mathcal{M}$ into regions and each region is bounded by three edges. These regions are the *triangular faces* of K. A triangular face is *oriented* if directions are assigned to its edges so that they form a directed cycle. The surface $\mathcal{M}$ is an *orientable 2-manifold* if the faces in K can be oriented such that each edge is assigned two opposite directions. Orientable 2-manifolds are a popular class of surfaces.

We use uv to denote the edge connecting the vertices u and v. For each vertex v in K, the collection of edges and triangular faces incident to v is the *star* of v denoted by $\mathrm{St}(v)$. For every face σ in $\mathrm{St}(v)$, if we collect the edge of σ not incident to v and also the endpoints of this edge, we obtain the *link* of v denoted by $\mathrm{Lk}(v)$. Given an edge uv, for every face σ in $\mathrm{St}(u) \cup \mathrm{St}(v)$, if we collect the edge of σ not incident to u or v and also the endpoints of this edge, we obtain the *neighborhood* of uv denoted by $\mathrm{N}(uv)$. Figure 2 shows examples of star, link and neighborhood.

The *contraction* of uv is a local transformation of K. A new vertex w is introduced to replace uv. $\mathrm{St}(u) \cup \mathrm{St}(v)$ is replaced by a local triangulation: for each vertex $x \in \mathrm{N}(uv)$, we get the edge wx; for every edge $xy \in \mathrm{N}(uv)$, we get the triangular face wxy. The underlying surface may cease to be a 2-manifold after an arbitrary edge contraction. For example, see Figure 1(b).

We are interested in cycles in K that are slightly more general than in the graph-theoretic

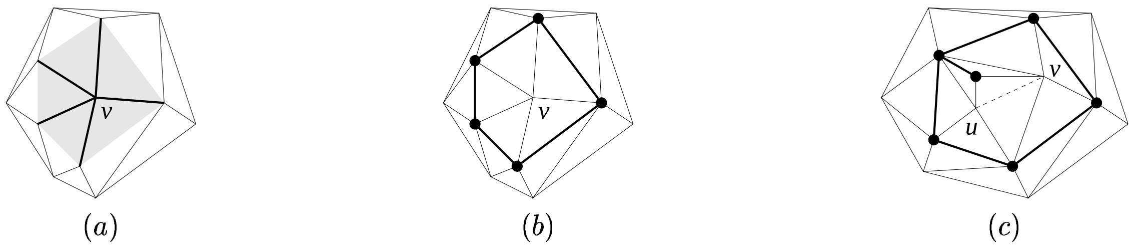

Figure 2: In (a), the bold line segments and the shaded triangles are edges and faces in $\mathrm{St}(v)$. In (b), the black dots and bold line segments are the vertices and edges in $\mathrm{Lk}(v)$. In (c), the black dots and bold line segments are vertices and edges in $\mathrm{N}(uv)$. Note that uv is non-contractible.

setting. A *cycle* in K is a subgraph of K such that the degree of every vertex in the subgraph is positive and even. A *simple cycle* is a connected subgraph of K such that the degree of every vertex in the subgraph is two. We denote a simple cycle by listing the vertices in order around the cycle. Since $\mathcal{M}$ is a 2-manifold, $\mathrm{Lk}(v)$ is a simple cycle which induces a circular ordering of edges and triangular faces in $\mathrm{St}(v)$.

We call a simple cycle in K *critical* if it consists of three edges and it does not bound a face in K. For example, the cycle uvy in Figure 1(b) is critical. If K is combinatorially equivalent to the boundary of a tetrahedron, no edge can be contracted without changing the topology type of $\mathcal{M}$. Otherwise, the contraction of an edge e is topology-preserving if and only if e does not lie on a critical cycle. Dey et al. [5] discussed a more general definition of topology-preserving edge contraction that works for non-manifolds.

3 Family of cycle pairs

We introduce a special family of crossing cycle pairs and prove several properties of these cycle pairs. They are the main tool in obtaining our results in Sections 4 and 5.

The cycles in K form a group under *addition* [11]. When cycles are added, edges that appear an odd number of times remain while edges that appear an even number of times are removed. The empty set is the identify element of this group. Given a triangular face σ, we use $\partial\sigma$ to denote the simple cycle formed by the edges of σ. Given a collection F of triangular faces in K, ∂F denotes $\sum_{\sigma \in F} \partial\sigma$, where cycles are added as described above. Figure 3 shows some examples of cycles. Two cycles A and B are *homologous* if there is a collection F of triangular faces such that $A = B + \partial F$. For example, in Figure 3, the cycles $uvwsxyz$ and $rvsy$ are homologous. It is known that there are exactly $2g$ mutually non-homologous cycles in K if $\mathcal{M}$ has genus g [11].

3.1 Crossing and crossing characteristic

Let ξ_1 and ξ_2 be two simple closed curves on $\mathcal{M}$ that intersect only at isolated points. For each intersection point $p \in \xi_1 \cap \xi_2$, ξ_1 divides a small open neighborhood of p on $\mathcal{M}$ into two open topological disks. If ξ_2 intersects both topological disks, we say that ξ_1 and ξ_2 *cross at p*. We call p a *crossing of ξ_1 and ξ_2*. The *parity of crossings* between ξ_1 and ξ_2 is the number of crossings

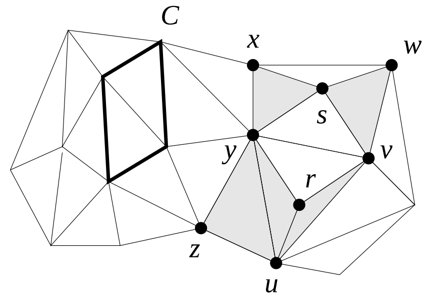

Figure 3: C shown as bold line segments is a simple cycle. Let F be the collection of shaded triangular faces. $\partial F = uvwsxyz + rvsy$.

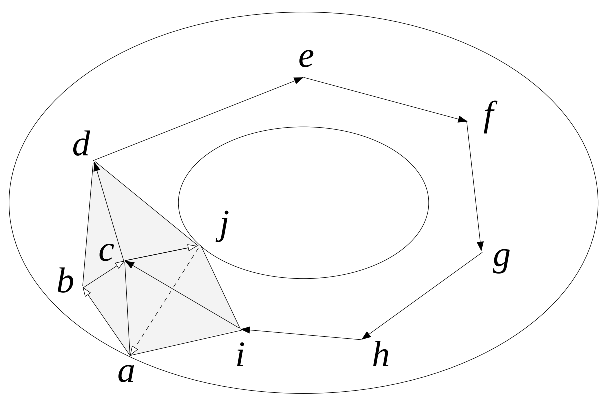

Figure 4: The shaded area is a patch of triangles on the torus $\mathcal{M}$. Two cycles are shown and they are labeled using two different arrows. The dotted arrow is an edge on $\mathcal{M}$ behind the shaded patch. The two cycles cross at the vertex c.

modulo 2. The parity of crossings is related to the concept of *intersection number* in algebraic topology [7]. Figure 4 shows an example. There is exactly one crossing at the vertex c in the figure.

We use the parity of crossings between simple closed curves (that intersect only at isolated points) to define the *crossing characteristic* $B_1 \circ B_2$ of two simple cycles B_1 and B_2. Note that B_1 and B_2 may share edges, so they do not necessarily intersect at isolated vertices. Figure 5(a) and (b) show two such cases. We perturb B_1 to another simple closed curve ξ_1 on $\mathcal{M}$ as follows. Fix the vertices of B_1. For each edge e of B_1, perturb e to a closed curved segment γ such that $\text{int}(\gamma)$ lies in the interior of any triangular face of K incident to e, $\gamma \cap e$ consists of the endpoints of e, and $\text{int}(\gamma)$ does not intersect any curved segment obtained by perturbing other edges of B_1. Consequently, ξ_1 and B_2 intersect only at the vertices of B_1. $B_1 \circ B_2$ is defined as the parity of crossings between ξ_1 and B_2. We can generalize to the case where B_2 is a sum of simple cycles. Let $B_2 = \sum_{j=1}^{q} B_{2j}$, where B_{2j} are simple cycles. Then $B_1 \circ B_2 = (\sum_{j=1}^{q} B_1 \circ B_{2j}) \bmod 2$.[1] Lemma 1 shows that $B_1 \circ B_2$ is well-defined, i.e., independent of the perturbation of B_1 and the sum expression of B_2. Later, we will make use of the simple property that if $B_1 \circ B_2 = 1$, B_1 and B_2 must share a vertex.

[1] The generalization can be taken further. Let $B_1 = \sum_{i=1}^{p} B_{1i}$ and let $B_2 = \sum_{j=1}^{q} B_{2j}$, where B_{1i} and B_{2j} are simple cycles. Then $B_1 \circ B_2 = (\sum_{i=1}^{p} \sum_{j=1}^{q} B_{1i} \circ B_{2j}) \bmod 2$. However, this generalization is not needed for obtaining our results.

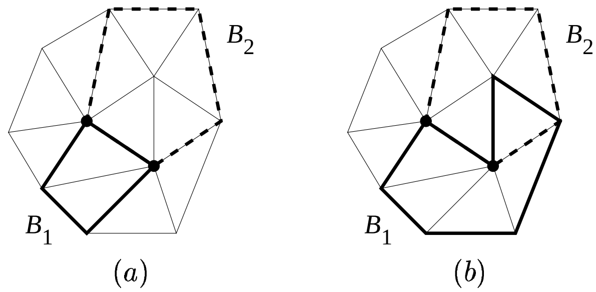

Figure 5: There are two cycles B_1 and B_2 shown in bold solid and dashed line segments respectively. The crossings between B_1 and B_2 are not defined at shared edges.

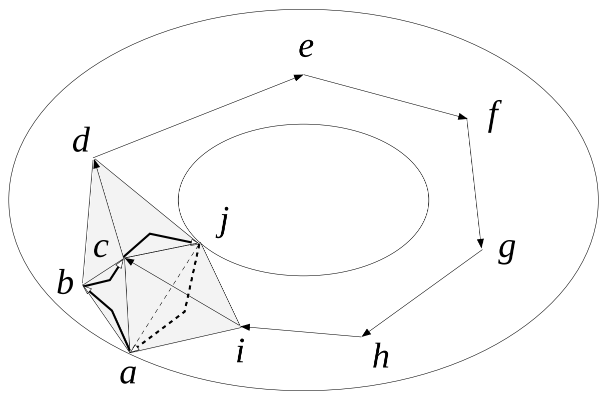

Figure 6: The bold line segments form the perturbed $abcj$.

As an example, the crossing characteristics in Figure 5 (a) and (b) are zero.

When B_1 and B_2 are simple cycles that do not share any edge, the perturbation of B_1 is redundant and has no effect. In this case, $B_1 \circ B_2$ coincides with the parity of crossings between B_1 and B_2. Note that even if B_2 is a simple cycle that does not share any edge with B_1, it can be rewritten as a sum of simple cycles and these simple cycles may share edges with B_1. Our definition is consistent and gives the same answer in both cases. For example, the simple cycle $cdefghi$ in Figure 4 can be written as $abcdefghi + abci$. We perturb $abcj$ as shown in Figure 6. The perturbed $abcj$ crosses $abcdefghi$ at a, so $abcj \circ abcdefghi = 1$. The perturbed $abcj$ crosses $abci$ at a and c, so $abcj \circ abci = 0$. Thus, $abcj \circ (abcdefghi + abci) = 1$ which agrees with $abcj \circ cdefghi$. It remains a legitimate question whether $B_1 \circ B_2$ has any relation with the parity of crossings between B_1 and B_2 when B_1 and B_2 share edges. To answer this question, one needs to define crossings along shared edges appropriately. Such a study is irrelevant to our application and outside the scope of this paper. In fact, we bypass defining crossings in the degenerate cases by using crossing characteristic.

Lemma 1 *Given a simple cycle B_1 and a sum B_2 of simple cycles, $B_1 \circ B_2$ is independent of the sum expression of B_2 and the perturbation of B_1.*

Proof. We first argue that the sum expression of B_2 is unimportant. Let v be any vertex of B_1.

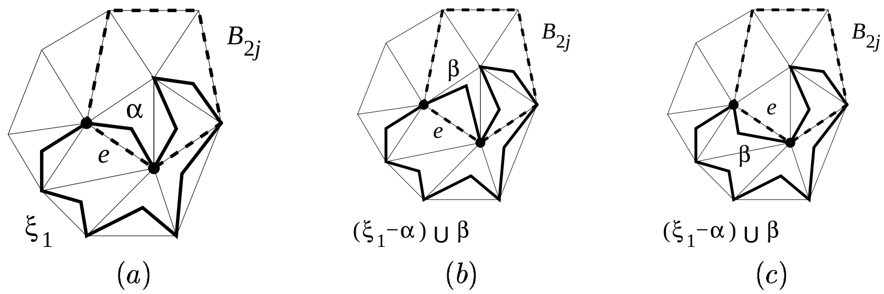

Figure 7: Figure (a) shows a perturbation ξ_1 of B_1 and B_{2j}. Figure (b) shows the substitution of α by β when they lie in the same face. Figure (c) shows the substitution when α and β lie in different faces. In both cases, the parity of crossings between ξ_1 and B_2 is the same as that between $(\xi_1 - \alpha) \cup \beta$ and B_2.

Let ξ_1 be the simple closed curve obtained by a perturbation of B_1. ξ_1 divides the edges in $\mathrm{St}(v)$ into two subsets. Let n_1 and n_2 denote the numbers of edges of B_2 incident to v in the two subsets. Since B_2 has an even number of edges incident to v, n_1 and n_2 have the same parity. (Although v may not be a vertex in B_2, v could appear in some simple cycles in the sum expression of B_2 when there is cancellation of edges. In this case, $n_1 = n_2 = 0$.) If n_1 is even (resp. odd), the number of B_{2j}'s crossing ξ_1 at v is even (resp. odd) and so the parity of crossings at v between ξ_1 and B_2 is 0 (resp. 1). Hence, the parity of crossings between ξ_1 and B_2 is independent of the sum expression of B_2.

Next, we argue that the perturbation of B_1 is unimportant. Let ξ_1 and η_1 be two simple closed curves obtained by different perturbations of B_1. Let e be an edge of B_1. Let α and β be the two perturbed versions of e in ξ_1 and η_1 respectively. We examine the parity of crossings between $(\xi_1 - \alpha) \cup \beta$ and B_{2j} for some j. If B_{2j} does not contain e, then $(\xi_1 - \alpha) \cup \beta$ and B_{2j} cross at an endpoint v of e iff ξ_1 and B_{2j} cross at v. Consider the case where B_{2j} contains e. Figure 7(a) shows ξ_1 and B_{2j}. If α and β lie inside the same triangular face, then as before, $(\xi_1 - \alpha) \cup \beta$ and B_{2j} cross at an endpoint v of e iff ξ_1 and B_{2j} cross at v. See Figure 7(b). If α and β lie inside different triangular faces, then $(\xi_1 - \alpha) \cup \beta$ and B_{2j} cross at an endpoint v of e iff ξ_1 and B_{2j} do not cross at v. See Figure 7(c). So we can incrementally transform ξ_1 to η_1 while the parity of crossings between any intermediate curve and B_{2j} remains the same. Thus, $B_1 \circ B_{2j}$ is independent of the perturbation of B_1, and so is $B_1 \circ B_2$. ◻

Lemma 1 leads to the following lemma concerning the crossing characteristics of a simple cycle and two homologous simple cycles.

Lemma 2 *Let A, B_1 and B_2 be three simple cycles in K. If B_1 and B_2 are homologous, then $A \circ B_1 = A \circ B_2$.*

Proof. Since B_1 and B_2 are homologous, $B_1 = B_2 + \partial F$ for some collection F of triangular faces. So $A \circ B_1 = (A \circ B_2 + A \circ \partial F) \bmod 2$. Clearly, $A \circ \partial\sigma = 0$ for any triangular face σ. Thus, $A \circ \partial F = 0$ which implies that $A \circ B_1 = A \circ B_2$. ◻

3.2 Crossing cycle pairs

Let $\ell \geq 3$ be a parameter. Let $\mathcal{F}_\ell$ denote a family of cycle pairs $\{(C_i, D_i) : 1 \leq i \leq |\mathcal{F}_\ell|\}$ that satisfy four conditions:

- Each C_i is a critical cycle.
- Each D_i is a simple cycle of length at most ℓ.
- For any i, C_i and D_i cross at a vertex called the *anchor* of C_i and C_i does not share any other vertex with D_i.
- For $i \neq j$, the anchors of C_i and C_j are different. Note that for $i \neq j$, C_i or D_i may share vertices and edges with C_j and D_j.

Figure 8 shows an example of a family of three crossing cycle pairs on a torus. It is clear that if we construct a family that restricts every cycle pair to share a common D as in Figure 8, then there are at most ℓ such pairs as the anchors of the C_i's are distinct vertices. Lemma 3, which is the main result of this subsection, shows an upper bound on the family size in terms of g and ℓ in general.

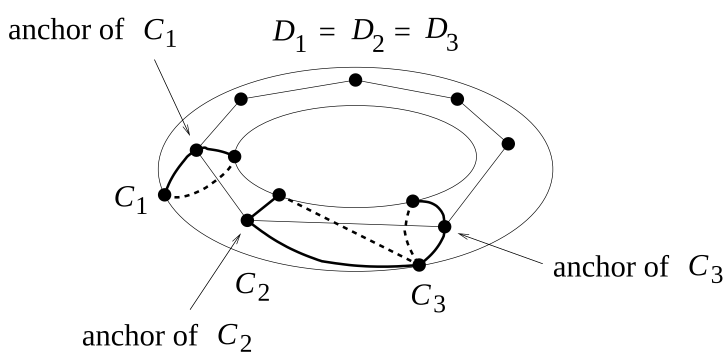

Figure 8: The bold cycles are C_1, C_2, and C_3. $D_1 = D_2 = D_3$. Note that C_2 and C_3 share a vertex.

Lemma 3 *$|\mathcal{F}_3| \leq 240g$ and for $\ell \geq 3$, $|\mathcal{F}_\ell| \leq 20\ell^3 g$.*

We will show that $|\mathcal{F}_3|$ is an upper bound on the number of vertices of an irreducible triangulation and we will use $|\mathcal{F}_4|$ to prove our results on constructing a hierarchy. We provide the proofs for the bound $20\ell^3 g$ below. The sharper bound of $240g$ for $|\mathcal{F}_3|$ can be found in the appendix.

We first sketch the main ideas in the proof. We use Lemma 4 to select a subset $\mathcal{S}_\ell \subseteq \mathcal{F}_\ell$. Then we partition the C_i's in $\mathcal{S}_\ell$ into equivalence classes of mutually homologous cycles. We pick one cycle from each equivalence class and set $\mathcal{F}'_\ell = \{(C_i, D_i) : C_i \text{ picked}\}$. So any two distinct C_i and C_j in $\mathcal{F}'_\ell$ are non-homologous. We will prove that $|\mathcal{F}'_\ell| = \Omega(|\mathcal{F}_\ell|)$ by showing that each equivalence class has $O(1)$ cycles. Since K has at most $2g$ mutually non-homologous cycles, $|\mathcal{F}'_\ell| \leq 2g$ and so we obtain an upper bound on $|\mathcal{F}_\ell|$. We provide the details in the rest of the subsection.

Lemma 4 *There is a subset $\mathcal{S}_\ell \subseteq \mathcal{F}_\ell$ of cardinality at least $|\mathcal{F}_\ell|/20$ such that for any two distinct C_i and C_j in $\mathcal{S}_\ell$, C_i does not contain the anchor of C_j.*

Proof. Let G be the graph formed by the union of C_i's in $\mathcal{F}_\ell$. Each C_i has three edges, so the degree sum of vertices in G is at most $6|\mathcal{F}_\ell|$. We claim that there are at least $|\mathcal{F}_\ell|/2$ anchors in G of degree nine or less. Otherwise, the degree sum of anchors in G is at least $10x + 2(|\mathcal{F}_\ell| - x) = 8x + 2|\mathcal{F}_\ell|$, where $x > |\mathcal{F}_\ell|/2$ is the number of anchors in G of degree ten or more. So the degree sum is greater than $6|\mathcal{F}_\ell|$ which is a contradiction. We pick a maximal independent subset of anchors in G whose degrees are at most nine. Then we set $\mathcal{S}_\ell = \{(C_i, D_i) : \text{the anchor of } C_i \text{ is picked}\}$. Clearly, $|\mathcal{S}_\ell| \geq |\mathcal{F}_\ell|/20$ and for any $C_i \neq C_j$ in $\mathcal{S}_\ell$, C_i does not contain the anchor of C_j. ◻

Let $\mathcal{H}$ be an equivalence class of mutually homologous C_i's in $\mathcal{S}_\ell$. Select all the cycles C_i's in $\mathcal{H}$ that share a common edge xy. Note that neither x nor y is the anchor of any C_i selected by the property of $\mathcal{S}_\ell$. We call the graph formed by the union of these C_i's a *bundle* of $\mathcal{H}$ and we call xy the *axis* of this bundle. Given a bundle W, we use $\text{size}(W)$ to denote the number of cycles forming W. Two bundles are *disjoint* if they do not share any vertex, otherwise they are *adjacent*. By the property of $\mathcal{S}_\ell$, two cycles in $\mathcal{H}$ can share non-anchor vertices only. If two cycles in $\mathcal{H}$ share two non-anchor vertices, they are part of the same bundle. Therefore, two adjacent bundles share exactly one common axis endpoint and we say that they are *adjacent at this vertex.* It also follows that there is a unique partition of $\mathcal{H}$ into bundles. Given a subset $\mathcal{Z}$ of bundles in $\mathcal{H}$, $|\mathcal{Z}|$ denotes the number of bundles in $\mathcal{Z}$.

Lemma 5 *Let L be a bundle in $\mathcal{H}$ with axis xy. Let $\mathcal{Z}$ be a subset of other bundles in $\mathcal{H}$ such that any two bundles in $\mathcal{Z}$ are either disjoint or adjacent at x or y. Then $|\mathcal{Z}| \leq \ell - \text{size}(L)$.*

Proof. Let C_i be one of the cycles forming the bundle L. Let D_i be the cycle that pairs up with C_i in $\mathcal{S}_\ell$. By definition, $D_i \circ C_i = 1$. Assume that C_j is one of the cycles forming a bundle in $\mathcal{Z}$. Since C_i and C_j are homologous, $D_i \circ C_j = D_i \circ C_i = 1$ by Lemma 2. It follows that D_i contains a vertex w of C_j. Since x and y are non-anchor vertices of C_i, D_i does not contain them. So the vertex w cannot be x or y. Since any two bundles in $\mathcal{Z}$ are either disjoint or adjacent at x or y, each bundle in $\mathcal{Z}$ contributes at least one distinct vertex in D_i. By the same reasoning, D_i must contain the anchors of all cycles forming the bundle L. Moreover, by the property of $\mathcal{S}_\ell$, no such anchor belongs to any bundle in $\mathcal{Z}$. Since D_i has length at most ℓ, we conclude that $|\mathcal{Z}| + \text{size}(L) \leq \ell$. ◻

Lemma 6 *Let $\mathcal{W} = \{W_r : 1 \leq r \leq m\}$ be a maximal collection of disjoint bundles in $\mathcal{H}$. Then*

(i) $m \leq \ell$ and $\text{size}(W_r) \leq \ell - m + 1$ for $1 \leq r \leq m$.

(ii) There are at most $\ell - \text{size}(W_r)$ bundles in $\mathcal{H}$ that are adjacent to W_r at the same axis endpoint of W_r.

(iii) For any bundle V in $\mathcal{H}$ adjacent to W_r, $\text{size}(V) \leq \ell - m$.

Proof. Consider (i). We invoke Lemma 5 with $L = W_r$ and $\mathcal{Z} = \mathcal{W} - \{W_r\}$. It implies that $|\mathcal{Z}| = m - 1 \leq \ell - \text{size}(W_r) \leq \ell - 1$. It follows that $m \leq \ell$ and $\text{size}(W_r) \leq \ell - m + 1$.

Consider (ii). We invoke Lemma 5 with $L = W_r$ and $\mathcal{Z}$ equal to the set of bundles in $\mathcal{H}$ adjacent to W_r at the same axis endpoint of W_r. It implies that $|\mathcal{Z}| \leq \ell - \text{size}(W_r)$.

Consider (iii). We invoke Lemma 5 with $L = V$ and $\mathcal{Z} = \mathcal{W}$. It implies that $|\mathcal{Z}| = m \leq \ell - \text{size}(V)$. It follows that $\text{size}(V) \leq \ell - m$. ◻

We are ready to show that $|\mathcal{F}'_\ell| = \Omega(|\mathcal{F}_\ell|)$.

Lemma 7 *There is a subset $\mathcal{F}'_\ell \subseteq \mathcal{F}_\ell$ of cardinality at least $|\mathcal{F}_\ell|/(10\ell^3)$ such that for any two distinct C_i and C_j in $\mathcal{F}'_\ell$, C_i and C_j are non-homologous.*

Proof. Let $\mathcal{S}_\ell \subseteq \mathcal{F}_\ell$ be a subset satisfying Lemma 4. Let $\mathcal{H}$ be an equivalence class of mutually homologous C_i's in $\mathcal{S}_\ell$. Let $\mathcal{W} = \{W_r : 1 \leq r \leq m\}$ be a maximal collection of disjoint bundles in $\mathcal{H}$. By Lemma 6(i), we have

$$\begin{aligned} m &\leq \ell \\ \text{size}(W_r) &\leq \ell - m + 1. \end{aligned}$$

For any bundle V in $\mathcal{H}$ adjacent to W_r, $\text{size}(V) \leq \ell - m$ by Lemma 6(iii). Lemma 6(ii) implies that there are at most $2(\ell - \text{size}(W_r))$ bundles adjacent to W_r. Therefore,

$$\begin{aligned} |\mathcal{H}| &\leq \sum_{r=1}^{m} (\text{size}(W_r) + 2(\ell - \text{size}(W_r))(\ell - m)) \\ &= \sum_{r=1}^{m} (2\ell(\ell - m) - (2\ell - 2m - 1) \cdot \text{size}(W_r)). \end{aligned}$$

If $m = \ell$, then $|\mathcal{H}| \leq \sum_{r=1}^{\ell} \text{size}(W_r) \leq \ell$. If $m < \ell$, then $|\mathcal{H}| < \sum_{r=1}^{m} 2\ell(\ell - m) = 2m\ell(\ell - m)$. This bound is maximized when $m = \ell/2$. So $|\mathcal{H}| < \ell^3/2$.

We pick one C_i from each equivalence class $\mathcal{H}$ of mutually homologous C_i's in $\mathcal{S}_\ell$. Let $\mathcal{F}'_\ell = \{(C_i, D_i) : C_i \text{ picked}\}$. Since $|\mathcal{S}_\ell| \geq |\mathcal{F}_\ell|/20$, we get $|\mathcal{F}'_\ell| \geq \frac{2}{\ell^3} \cdot \frac{1}{20} \cdot |\mathcal{F}_\ell| = |\mathcal{F}_\ell|/(10\ell^3)$. ☐

Proof of Lemma 3: If $\mathcal{M}$ has genus g, K contains at most $2g$ cycles that are mutually non-homologous. Thus, $|\mathcal{F}'_\ell| \leq 2g$. It follows from Lemma 7 that $|\mathcal{F}_\ell| \leq 20\ell^3 g$. The bound for $|\mathcal{F}_3|$ is provided in the appendix. ☐

4 Irreducible triangulation

In this section, we prove that any irreducible triangulation of an orientable 2-manifold of positive genus g has at most $240g$ vertices. We need the following lemma about a vertex.

Lemma 8 *Assume that $\mathcal{M}$ has positive genus. For any critical cycle vxy, one of the following holds.*

(i) There are two contractible edges uv and vw that alternate with vx and vy in $\text{St}(v)$.

(ii) A pair of critical cycles cross at v.

Proof. Observe that $x, y \in \text{Lk}(v)$. Let L be the list of vertices in $\text{Lk}(v)$ in clockwise order starting at x (recall that $\text{Lk}(v)$ is circularly ordered). If there is a vertex u before y and a vertex w after y in L such that uv and vw are contractible, then (i) is true. Assume that all edges uv, where u precedes y in L, are non-contractible. (We can symmetrically handle the case that all edges vw, where w follows y in L, are non-contractible.) Since vxy is a critical cycle, x and y are not

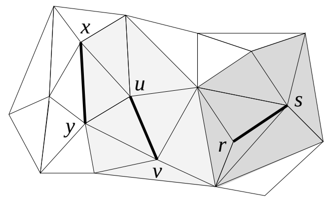

Figure 9: xy and uv are not independent, but both are independent from rs. The regions covered by $\mathrm{St}(u) \cup \mathrm{St}(v)$ and $\mathrm{St}(r) \cup \mathrm{St}(s)$ are shaded differently. $\mathrm{N}(uv)$ and $\mathrm{N}(rs)$ share two vertices and one edge.

adjacent in $\mathrm{Lk}(v)$, so we can pick an edge uv such that $u \neq x$ and u precedes y in L. Since uv is non-contractible and the genus of $\mathcal{M}$ is positive, uv lies on a critical cycle uvw' for some vertex $w' \in \mathrm{Lk}(v)$. If vxy and uvw' cross at v, then (ii) is true. Otherwise, either $w' = y$ or w' precedes y in L. We repeat the above argument with x and y replaced by u and w'. We must eventually obtain a pair of critical cycles that cross at v. ◻

Theorem 1 *Any irreducible triangulation of an orientable 2-manifold of genus g has at most $\max\{240g, 4\}$ vertices.*

Proof. The theorem is clearly true when $g = 0$. Let K be an irreducible triangulation. Assume that $g > 0$. We construct a family $\mathcal{F}_3$ of crossing cycle pairs as follows. Each vertex v in K is incident on a non-contractible edge, so v lies on a critical cycle. Since no edge of K is contractible, Lemma 8(ii) holds and a pair of critical cycles cross at v. We add this cycle pair to $\mathcal{F}_3$. The number of vertices of K is $|\mathcal{F}_3|$ which is at most $240g$ by Lemma 3. ◻

5 Hierarchy of surfaces

In this section, we prove that there are linearly many independent topology-preserving edge contractions. Moreover, a simple greedy strategy can be used to find them. Let uv and rs be two edges of K. We say that uv and rs are *independent* if $\mathrm{St}(u) \cup \mathrm{St}(v)$ and $\mathrm{St}(r) \cup \mathrm{St}(s)$ do not share any triangle. Although $\mathrm{N}(uv)$ and $\mathrm{N}(rs)$ might share vertices and edges, the contractions of uv and rs do not affect the same triangle. Figure 9 shows an example.

Our proof proceeds in two steps. First, we focus on the contractible edges of K by considering a subgraph G that contains *all vertices* of K and the contractible edges of K. (So G might be disconnected.) We prove that all maximal matchings of G have linear size. Second, we prove that any maximal matching of G contains an independent subset of edges of linear size. Moreover, they can be found using a greedy strategy.

Lemma 9 *Let n be the number of vertices in K and let g be the genus of $\mathcal{M}$. Assume that $g > 0$. Any maximal matching of G matches at least $(n - 1310g + 30)/16$ vertices.*

Proof. We obtain an embedding of G on $\mathcal{M}$ by erasing the non-contractible edges in K. G induces a subdivision of $\mathcal{M}$ which we denote by $G(\mathcal{M})$. Pick a maximal matching of G. Let H be the subgraph of G consisting of the matched vertices and the edges of G between them. So H contains all matching edges but H may contain some non-matching edges as well. As our argument proceeds, we will create some segments on $\mathcal{M}$, called *purple segments*, that connect matched vertices. The purple segments can be straight or curved and some purple segments may be edges of H. There will be no crossing among purple segments and edges of H. The purple segments will be used later to form a new graph with H.

We bound the number of unmatched vertices by charging them to the purple segments as well as by forming a family $\mathcal{F}_4$ of crossing cycle pairs. We charge for the unmatched vertices one by one in an arbitrary order. Let v be an unmatched vertex. If the degree of v in G is at most 1, then Lemma 8(ii) applies and a pair of critical cycles cross at v. We charge for v by adding this cycle pair to $\mathcal{F}_4$.

Suppose that the degree of v in G is larger than 1. Since v is unmatched, all neighbors of v in G are matched. The edges in G incident to v divide the triangles in $\mathrm{St}(v)$ into several intervals. Let u and w be two consecutive neighbors of v in G. Let R_{uvw} denote the region covered by the interval of triangles delimited by uv and vw. Figure 10 shows an example. There are two different ways to charge for v.

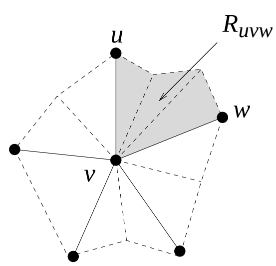

Figure 10: The figure shows $\mathrm{St}(v)$. The solid line segments are incident edges of v in G. The shaded region is R_{uvw}.

Case 1: No edge in $\mathrm{St}(v)$ lies inside R_{uvw}. So uvw bounds a triangular face in K and uw is an edge in K. If we have not created a purple segment γ_{uw} connecting u and w before, we set $\gamma_{uw} = uw$. (uw may be an edge of H.) Afterwards, we put a green pebble at γ_{uw} to charge for v. We claim that γ_{uw} receives at most two green pebbles overall. Assume that γ_{uw} receives a second green pebble. Then γ_{uw} lies inside a quadrilateral Q with vertices u, v, w, and an unmatched vertex v' such that v', u and w satisfy the conditions of case 1 (with v replaced by v'). If uw is an edge of G, uvw and $uv'w$ bound two adjacent regions in $G(\mathcal{M})$; otherwise, $uvwv'$ bound one region in $G(\mathcal{M})$. Two green pebbles have been put at γ_{uw} to charge for v and v'. Since Q does not have any other unmatched vertex, γ_{uw} cannot receive a third green pebble.

Case 2: Some edge vx in $\mathrm{St}(v)$ lies inside R_{uvw}. So vx is non-contractible and it lies on some critical cycle vxy. Since vy is non-contractible, vy cannot be uv or vw as they are are contractible. So vy lies either inside R_{uvw} or outside R_{uvw}. If vy lies inside R_{uvw} or on the boundary of R_{uvw}, then Lemma 8(ii) applies, so a pair of critical cycles cross at v. We charge for v by adding this cycle pair to $\mathcal{F}_4$.

Suppose that vy lies outside R_{uvw}. If we have not created a purple segment γ_{uw} connecting u and w before, we create γ_{uw} as follows. If uw is an edge in K, we set $\gamma_{uw} = uw$. Otherwise, we draw γ_{uw} as a curved segment inside R_{uvw}. Clearly, γ_{uw} does not cross any edge of H. (uw may be an edge of H.) Moreover, our drawing strategy can prevent γ_{uw} from crossing any existing purple segments.

After creating γ_{uw} if necessary, we check the number of blue pebbles at γ_{uw}. If γ_{uw} contains less than three blue pebbles, we add a blue pebble to γ_{uw} to charge for v. If γ_{uw} already contains three blue pebbles, these blue pebbles were introduced to charge for three unmatched vertices v_i, $1 \leq i \leq 3$, other than v and each v_i is adjacent to both u and w. We pick v_k such that $v_k \neq x$ and $v_k \neq y$. Since vx lies inside R_{uvw} and vy lies outside R_{uvw}, vxy cross the cycle $uvwv_k$ at v. We add the cycle pair $(vxy, uvwv_k)$ to $\mathcal{F}_4$ to charge for v.

By Lemma 3, $|\mathcal{F}_4| \leq 1280g$. It remains to bound the total number of pebbles on the purple segments. We add the purple segments as edges to H (only those that do not belong to H) and we add more edges, if necessary, to obtain a connected graph H^* that is embedded on $\mathcal{M}$ without any edge crossing. Let N and E be the number of vertices and edges in H^*. (So N is the number of matched vertices.) By Euler's relation, $E \leq 3N - 6 + 6g$. Since each purple segment carries at most two green pebbles and at most three blue pebbles, the total number of pebbles in H^* is at most $5E \leq 15N - 30 + 30g$.

It follows that the number of unmatched vertices is bounded by $1280g + 5E \leq 15N - 30 + 1310g$. Hence, $n \leq N + 15N - 30 + 1310g$ which implies that $N \geq (n - 1310g + 30)/16$. □

Theorem 2 *Let n be the number of vertices of* K *and let g be the genus of $\mathcal{M}$. Assume that $g > 0$. For any constant $d \geq 444$, if $n \geq 9182g - 222$, there are at least $\frac{n-1310g+30}{64(d+1)}$ independent contractible edges and for each such edge uv,* $\mathrm{N}(uv)$ *has at most d vertices.*

Proof. Let M be some maximal matching of contractible edges. We use $|M|$ to denote the number of matching edges in M. By Lemma 9, $|M| \geq (n - 1310g + 30)/32$. Given a matching edge uv, we call the number of vertices in $\mathrm{N}(uv)$ the *neighborhood size* of uv which is equal to $\mathrm{degree}(u) + \mathrm{degree}(v) - 4$. Take any constant $d \geq 444$. We claim that there are at least $|M|/2$ matching edges such that each has neighborhood size at most d. Suppose not. Then the sum of the neighborhood sizes of the matching edges is greater than $d \cdot |M|/2$. This implies that the sum of the degrees of the endpoints of the matching edges is greater than $(d+4)|M|/2 \geq (d+4)(n - 1310g + 30)/64 \geq 7(n - 1310g + 30)$ as $d \geq 444$. Since $n \geq 9182g - 222$, $7(n - 1310g + 30) \geq 6n - 12 + 12g$ which contradicts the Euler's relation. We pick a maximal independent set of contractible edges whose neighborhood sizes are at most d. This yields at least $|M|/(2(d+1))$ matching edges. □

Although the proof of Theorem 2 uses a maximal matching M, it is not necessary to compute M first. We initialize an empty output set of edges EDGE_SET. Then we examine the edges of K in an arbitrary order and grow EDGE_SET. For each edge e, we determine whether e is contractible, $\mathrm{N}(e)$ has at most d vertices, and e and the edges in EDGE_SET are independent. If these three conditions are satisfied, we add e to EDGE_SET. In all, we have the following theorem.

Theorem 3 *Given a triangulated closed surface of n vertices and positive genus g, a topology-preserving hierarchy can be constructed by repeated contractions of independent contractible edges. Each edge contraction affects $O(1)$ triangles. The hierarchy has $O(\log n + g)$ depth and $O(n + g^2)$ size.*

The algorithm as described above takes $O(n + g^2)$ time. In practice, our greedy strategy resembles existing methods employed by some computer graphics researchers to construct hierarchies [4, 15, 19]. They develop heuristic functions to measure the geometric error of local updates (vertex decimations or edge contractions). The local updates are sorted in increasing order of geometric error using such a heuristic function. Then the sorted list is scanned to pick an independent subset. There is no worst-case guarantee on the geometric approximation error of the simplified surface. However, experimental results are often good. We suggest using the quadric error proposed by Garland and Heckbert [10] for edge contractions. Evaluating the quadric error of the contraction of an edge e is done in $O(1)$ time by solving a system of three linear equations involving three variables. The solution also tells the location of the new vertex that e should be contracted to. After sorting the edges, we scan the sorted list using our greedy strategy to select independent contractible edges. Due to sorting, the time complexity of the algorithm increases to $O(n \log n + g^2 \log g)$.

Acknowledgments

The first author thanks Beifang Chen, Herbert Edelsbrunner and Min Yan for helpful discussions.

References

[1] D.W. Barnette and A.L. Edelson, All 2-manifolds have finitely many minimal triangulations. *Israel J. Math*, 67 (1989), 123–128.

[2] J. Cohen, A. Varshney, D. Manocha, G. Turk, H. Weber, P. Agarwal, F. Brooks, and W. Wright. Simplification Envelopes. *Proceedings of SIGGRAPH 96*, 1996, 119-128.

[3] M. de Berg and K.T.G. Dobrindt. On levels of detail in terrains. *Proc. ACM Symposium on Computational Geometry*, 26–27, 1995.

[4] L. De Floriani, P. Magillo and E. Puppo. Building and traversing a surface at variable resolution. *Proc. IEEE Visualization '97*, 103–110, 1997.

[5] T.K. Dey, H. Edelsbrunner, S. Guha and D.V. Nekhayev, Topology preserving edge contraction, *Publ. Inst. Math. (Beograd) (N.S.)*, 66 (1999), 23–45.

[6] D.P. Dobkin and D.G. Kirkpatrick. Determining the separation of preprocessed polyhedra. *Proc. 17th Internat. Colloq. Automata Lang. Program.*, 400-413, 1990.

[7] A. Dold. *Lectures on Algebraic Topology*, Springer-Verlag, 1972.

[8] C.A. Duncan, M.T. Goodrich, and S.G. Kobourov, Planarity-preserving clustering and embedding for large planar graphs, *Proc. Graph Drawing '99*, 186–196.

[9] M. Garland, Multiresolution modeling: survey & future opportunities. *Eurographics '99*, State of the Art Report (STAR). `http://www.cs.uiuc.edu/~garland/paper.html`

[10] M. Garland and P.S. Heckbert. Surface simplification using quadric error metrics. *Proc. SIGGRAPH '97*, 209–216.

[11] P.J. Gilbin. *Graphs, Surfaces and Homology.* Chapman and Hall, 1981.

[12] H. Hoppe, Progressive meshes. *Proc. SIGGRAPH '96*, 99–108.

[13] H. Hoppe, View-dependent refinement of progressive meshes. *Proc. SIGGRAPH '97*, 189–198.

[14] D.G. Kirkpatrick. Optimal search in planar subdivisions. *SIAM Journal on Computing*, 12 (1983), 28–35.

[15] A. Lee, W. Sweldens, P. Schröder, L. Cowsar and D. Dobkin, MAPS: multiresolution adaptive parameterization of surfaces. *Proc. SIGGRAPH '98*, 95–104.

[16] A. Nakamoto and K. Ota, Note on irreducible triangulations of surfaces. *Journal of Graph Theory*, 20 (1995), 227–233.

[17] H. Schipper, Generating triangulations of 2-manifolds. *Computational Geometry – Methods, Algorithms and Applications*, Lecture Notes in Computer Science 553, Springer Verlag, 1991, 237–248.

[18] J.C. Xia, J. El-Sana and A. Varshney. Adaptive real-time level-of-detail-based rendering for polygonal models. *IEEE Transactions on Visualization and Computer Graphics*, 3 (1997), 171–181.

[19] J.C. Xia and A. Varshney, Dynamic view-dependent simplification for polygonal models. *Proc. Visualization '96*, 327–334.

Appendix

To show that $|\mathcal{F}_3| \le 240g$, we need one more utility lemma that is similar to Lemma 6. Recall that $\mathcal{H}$ is an equivalence class of homologous C_i's in $\mathcal{S}_3$.

Lemma 10 *Let $\mathcal{W} = \{W_r : 1 \le r \le m\}$ be a maximal collection of disjoint bundles in $\mathcal{H}$. Let $\mathcal{Q}$ be a set of bundles in $\mathcal{H}$ that are adjacent to W_r at the same axis endpoint of W_r and disjoint from W_s for any $s \neq r$. Then*

(i) $|\mathcal{Q}| \le \ell - m - \text{size}(W_r) + 1$.

(ii) for any bundle $V \in \mathcal{Q}$, $\text{size}(V) \le \ell - m - |\mathcal{Q}| + 1$.

Proof. We invoke Lemma 5 with $L = W_r$ and $\mathcal{Z} = \mathcal{Q} \cup (\mathcal{W} - \{W_r\})$. It implies that $|\mathcal{Z}| = |\mathcal{Q}| + m - 1 \le \ell - \text{size}(W_r)$. It follows that $|\mathcal{Q}| \le \ell - m - \text{size}(W_r) + 1$. For any bundle $V \in \mathcal{Q}$, we invoke Lemma 5 with $L = V$ and $\mathcal{Z} = (\mathcal{Q} - \{V\}) \cup \mathcal{W}$. It implies that $|\mathcal{Z}| = |\mathcal{Q}| + m - 1 \le \ell - \text{size}(V)$. It follows that $\text{size}(V) \le \ell - m - |\mathcal{Q}| + 1$. □

We first show that $|\mathcal{F}'_3| \ge |\mathcal{F}_3|/120$.

Lemma 11 *There is a subset $\mathcal{F}'_3 \subseteq \mathcal{F}_3$ of cardinality at least $|\mathcal{F}_3|/120$ such that for any two distinct C_i and C_j in $\mathcal{F}'_3$, C_i and C_j are non-homologous.*

Proof. Let $\mathcal{S}_3 \subseteq \mathcal{F}_3$ be the set satisfying Lemma 4. Let $\mathcal{H}$ be an equivalence class of mutually homologous C_i's in $\mathcal{S}_3$. We are to show that $|\mathcal{H}| \le 6$. Let $\mathcal{W} = \{W_r : 1 \le r \le m\}$ be a maximal collection of disjoint bundles in $\mathcal{H}$. By Lemma 6(i), $m \le 3$ and $\text{size}(W_r) \le 3$ for any r. We conduct a case analysis to show that $|\mathcal{H}| \le 6$.

Case 1: $m = 1$. If $\text{size}(W_1) = 3$, then Lemma 6(ii) implies that no bundle in $\mathcal{H}$ is adjacent to W_1. Thus, $|\mathcal{H}| = \text{size}(W_1) = 3$. Consider the case where $1 \le \text{size}(W_1) \le 2$. Let x be an axis endpoint of W_1. Let $\mathcal{Q}$ be the set of bundles in $\mathcal{H}$ that are adjacent to W_1 at x. Lemma 10(i) implies that $|\mathcal{Q}| \le 2$. By Lemma 10(ii), for any bundle $V \in \mathcal{Q}$, $\text{size}(V) = 1$ if $|\mathcal{Q}| = 2$; and $\text{size}(V) \le 2$ if $|\mathcal{Q}| = 1$. Thus, $\sum_{V \in \mathcal{Q}} \text{size}(V) \le 2$. Hence, the total sizes of bundles in $\mathcal{H}$ adjacent to W_1 is at most 4. So $|\mathcal{H}| \le \text{size}(W_1) + 4 \le 2 + 4 = 6$.

Case 2: $m = 2$. By Lemma 6(i), $\text{size}(W_r) \le 2$ for any r. Let $\mathcal{X}$ be the set of bundles in $\mathcal{H}$ that are adjacent to both W_1 and W_2. Observe that $0 \le |\mathcal{X}| \le 4$.

Case 2.1: $|\mathcal{X}| = 0$. Let $\mathcal{Q}$ be the set of bundles in $\mathcal{H}$ that is adjacent to W_r. If $\text{size}(W_r) = 2$, Lemma 10(i) implies that $|\mathcal{Q}| = 0$. Assume that $\text{size}(W_r) = 1$. Since $|\mathcal{X}| = 0$, Lemma 10(i) implies that there is at most one bundle in $\mathcal{H}$ adjacent to W_r at each axis endpoint of W_r. So $|\mathcal{Q}| \le 2$. For each bundle $V \in \mathcal{Q}$, $\text{size}(V) = 1$ by Lemma 6(iii). Thus, whether $\text{size}(W_r) = 1$ or 2, we have $\text{size}(W_r) + \sum_{V \in \mathcal{Q}} \text{size}(V) \le 3$. Hence, $|\mathcal{H}| \le 6$.

Case 2.2: $1 \le |\mathcal{X}| \le 2$. Let ab and cd be the axes of W_1 and W_2 respectively. Without loss of generality, we assume that there is a bundle $V \in \mathcal{X}$ that is adjacent to W_1 and W_2 at a and c, respectively. Note that $\text{size}(V) = 1$ by Lemma 6(iii).

Let $\mathcal{Q}$ be the set of bundles in $\mathcal{H}$ that are adjacent to W_1 at a and disjoint from W_2. We invoke Lemma 5 with $L = V$ and $\mathcal{Z} = \mathcal{Q} \cup \{W_1, W_2\}$. It implies that $|\mathcal{Z}| = |\mathcal{Q}| + 2 \le$

$3 - \text{size}(V) = 2$. So $|\mathcal{Q}| = 0$. That is, there is no bundle in $\mathcal{H}$ that is adjacent to W_1 at a and disjoint from W_2. Similarly, we can show that there is no bundle in $\mathcal{H}$ that is adjacent to W_2 at c and disjoint from W_1.

Let $\mathcal{Q}_b$ be the set of bundles in $\mathcal{H}$ that are adjacent to W_1 at b and disjoint from W_2. If $\text{size}(W_1) = 2$, $|\mathcal{Q}_b| = 0$ by Lemma 10(i). Assume that $\text{size}(W_1) = 1$. Lemma 10(i) implies that $|\mathcal{Q}_b| \leq 1$. For each bundle $U \in \mathcal{Q}_b$, $\text{size}(U) = 1$ by Lemma 6(iii). Thus, whether $\text{size}(W_1) = 1$ or 2, we have $\text{size}(W_1) + \sum_{U \in \mathcal{Q}_b} \text{size}(U) \leq 2$. Similarly, denoting by $\mathcal{Q}_d$ the set of bundles in $\mathcal{H}$ that are adjacent to W_2 at d and disjoint from W_1, we have $\text{size}(W_2) + \sum_{U \in \mathcal{Q}_d} \text{size}(U) \leq 2$.

It follows that

$$\begin{aligned} |\mathcal{H}| &= \text{size}(W_1) + \sum_{U \in \mathcal{Q}_b} \text{size}(U) + \text{size}(W_2) + \sum_{U \in \mathcal{Q}_d} \text{size}(U) + \sum_{V \in \mathcal{X}} \text{size}(V) \\ &\leq 4 + \sum_{V \in \mathcal{X}} \text{size}(V). \end{aligned}$$

Since $|\mathcal{X}| \leq 2$ and for each bundle $V \in \mathcal{X}$, $\text{size}(V) = 1$ by Lemma 6(iii), we conclude that $|\mathcal{H}| \leq 6$.

Case 2.3: $3 \leq |\mathcal{X}| \leq 4$. Let ab and cd be the axes of W_1 and W_2 respectively. Since $|\mathcal{X}| \geq 3$, we can assume without loss of generality that there are two bundles V and V' in $\mathcal{X}$ that are adjacent to W_1 at a.

Let $\mathcal{Q}_a$ (resp. $\mathcal{Q}_b$) be the set of bundles in $\mathcal{H}$ that are adjacent to W_1 at a (resp. b) and disjoint from W_2. We invoke Lemma 5 with $L = W_1$ and $\mathcal{Z} = \mathcal{Q}_a \cup \{V, V'\}$. It implies that $|\mathcal{Z}| = |\mathcal{Q}_a| + 2 \leq 3 - \text{size}(W_1) \leq 2$. So $|\mathcal{Q}_a| = 0$. Similarly, we invoke Lemma 5 again with $L = W_1$ and $\mathcal{Z} = \mathcal{Q}_b \cup \{V, V'\}$ to obtain $|\mathcal{Q}_b| = 0$. Thus, there is no bundle in $\mathcal{H}$ that is adjacent to W_1 and disjoint from W_2. Moreover, since there are two bundles V and V' adjacent to W_1 at a, Lemma 6(ii) implies that $\text{size}(W_1) = 1$.

Similarly, we can show that $\text{size}(W_2) = 1$ and there is no bundle in $\mathcal{H}$ that is adjacent to W_2 and disjoint from W_1.

It follows that $|\mathcal{H}| = \text{size}(W_1) + \text{size}(W_2) + \sum_{V \in \mathcal{X}} \text{size}(V) = 2 + \sum_{V \in \mathcal{X}} \text{size}(V)$. Since $|\mathcal{X}| \leq 4$ and for each bundle $V \in \mathcal{X}$, $\text{size}(V) = 1$ by Lemma 6(iii), we conclude that $|\mathcal{H}| \leq 6$.

Case 3: $m = 3$. For $1 \leq r \leq 3$, $\text{size}(W_r) = 1$ by Lemma 6(i) and Lemma 6(iii) implies that no bundle in $\mathcal{H}$ is adjacent to W_r. Thus, $|\mathcal{H}| = \sum_{r=1}^{3} \text{size}(W_r) = 3$.

This completes the proof that $|\mathcal{H}| \leq 6$ for any equivalence class $\mathcal{H}$ of mutually homologous C_i's in $\mathcal{S}_3$. We pick one C_i from each equivalence class. Let $\mathcal{F}_3' = \{(C_i, D_i) : C_i \text{ picked}\}$. Since $|\mathcal{S}_3| \geq |\mathcal{F}_3|/20$ and $|\mathcal{H}| \leq 6$, $|\mathcal{F}_3'| \geq |\mathcal{F}_3|/120$. ◻

Corollary 1 *$|\mathcal{F}_3| \leq 240g$.*

Proof. If $\mathcal{M}$ has genus g, then K contains at most $2g$ cycles that are mutually non-homologous. Thus, $|\mathcal{F}_3'| \leq 2g$. Then Lemma 11 implies that $|\mathcal{F}_3| \leq 240g$. ◻