Hierarchy of Surface Models and Irreducible Triangulation

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

June 2, 2002

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. In practice, the genus g is very small when compared with n for large models and the selection of edges can be enhanced by measuring the error of their contractions using some known heuristics. 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, abstract simplicial complex, 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 noticeable bad visual effects (for example, imagine that a

[†]Department of Computer Science, HKUST, Clear Water Bay, Hong Kong. Email: {scheng,hung}@cs.ust.hk. The research was partially supported by RGC CERG HKUST 6088/99E.

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

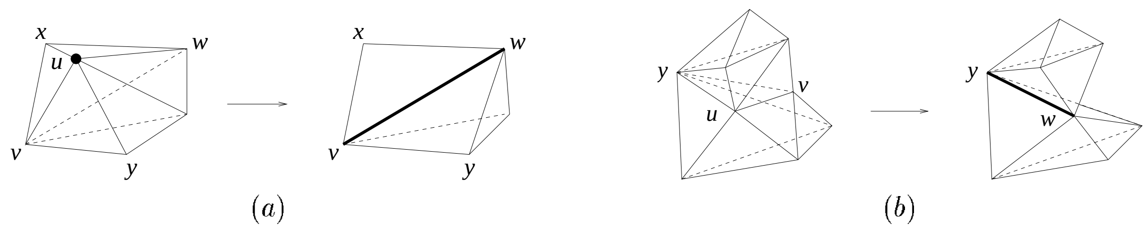

Figure 1: In figure (a), the decimation of u and a retriangulation produce a pinching at the edge vw which could be avoided if xy instead of vw is used in the retriangulation. In figure (b), the contraction of uv to w produces a pinching at the edge wy.

rod is squeezed in the 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 environment 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. In this paper, we resolve 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. 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.

In this paper, we prove a new upper bound of $240g$ on the number of vertices of an irreducible triangulation. Our 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 $d > 380$, if $n > \frac{(6008+1310d)g-888-30d}{d-380}$, 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 11). These results follow from two topological results about triangulations (Theorem 8 and Theorem 10). 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. We may sometime use a piecewise linear embedding as a tool in the proofs and we state this explicitly. In practice, when constructing a hierarchy, the surface models are linearly embedded[1] and the edge contractions are selected to keep the geometric approximation error small. There are known heuristics in the computer graphics literature for measuring the error of an edge contraction. For example, our greedy strategy can be enhanced to select edges in increasing order of quadric error [10].[2]

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

Triangulated 2-manifolds (without boundaries) are popular representations of object boundaries in solid modeling and computer graphics. The combinatorial structure of a triangulated 2-manifold can be represented using an *abstract simplicial complex* $\mathsf{K} = (V, S)$, where V is a set of vertices and S is a set of subsets of V. Each element $\sigma \in S$ has cardinality $k + 1$, $0 \leq k \leq 2$, and σ is called a k-*simplex*. S is required to satisfy the following two conditions. First, for each $v \in V$, $\{v\} \in S$. Second, For each $\sigma \in S$ and $\tau \subset \sigma$, $\tau \in S$. Each proper subset of σ is called a *face* of σ. Two simplices are *incident* if one is a face of the other. For simplicity, we write a 1-simplex $\{u, v\}$ as uv, and a 2-simplex $\{u, v, w\}$ as uvw. We also call the 1-simplices *edges*. The *star* of σ, $\mathrm{St}(\sigma)$, is the collection of simplices $\{\tau : \sigma \subset \tau\}$. If we collect the faces of τ, for all $\tau \in \mathrm{St}(\sigma)$, that are neither σ nor incident to σ, we obtain the *link* of σ denoted as $\mathrm{Lk}(\sigma)$. For each edge uv, its *neighborhood* $\mathrm{N}(uv)$ is $\{\tau \in \mathrm{Lk}(u) \cup \mathrm{Lk}(v) : u \not\subseteq \tau, v \not\subseteq \tau\}$. Figure 2 shows examples of star, link and neighborhood.

We use M_K to denote the underlying space of K, which is a *2-manifold* if the link of each vertex is a simple cycle. The circular ordering of vertices and edges in $\mathrm{Lk}(v)$ of a vertex v induces a circular ordering of edges and 2-simplices in $\mathrm{St}(v)$. A 2-simplex is *oriented* if directions are assigned to its edges so that they

[1]The surface may intersect itself during repeated edge contractions. Nevertheless, in computer graphics literature self-intersection has not been reported as a nuisance unless the complexity of the simplified surface is very tiny compared to the complexity of the original surface.

[2]There is no worst-case guarantee on the geometric approximation error if a surface is simplified using quadric error based edge contraction. Nevertheless, the experimental results are often good [10].

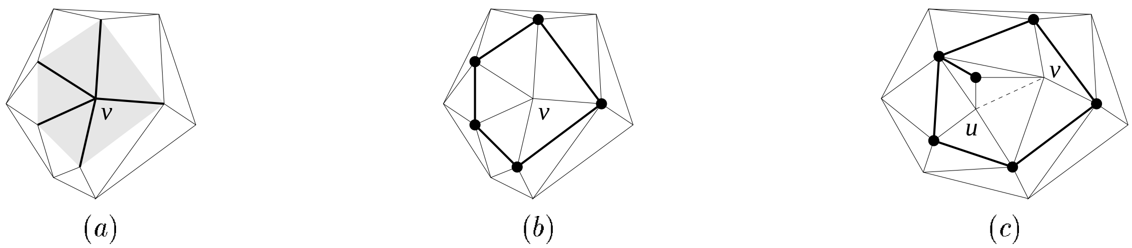

Figure 2: In (a), the bold line segments and the shaded triangles are simplices in $\mathrm{St}(v)$. In (b), the black dots and bold line segments are the vertices and edges in $\mathrm{Lk}(v)$. Note that $\mathrm{St}(v) \cap \mathrm{Lk}(v) = \emptyset$. In (c), the dashed line segment is uv and the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

form a directed cycle. M_K is an *orientable 2-manifold* if the 2-simplices of K can be oriented such that each edge is assigned two opposite directions. There are two ways to orient a 2-simplex, so there are two ways to orient K. Orientable 2-manifolds are a popular class of surfaces.

The *contraction* of an edge 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 vx; for every edge $xy \in \mathrm{N}(uv)$, we get the 2-simplex vxy. This yields a new abstract simplicial complex. An arbitrary edge contraction may result in an abstract simplicial complex whose underlying space is not a 2-manifold. For example, see Figure 1(b). We call a cycle in K *critical* if it consists of three edges and it does not bound a 2-simplex in K. For example, the cycle through u, v and y in Figure 1(b) is a critical cycle. If K is combinatorially equivalent to the boundary of a tetrahedron, no edge can be contracted without changing the topology type of M_K. 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.

3.1 Chain, cycle and crossing

We reexamine cycles using concepts from algebraic topology. For $0 \leq k \leq 2$, a *k-chain* is a formal sum of a set of k-simplices with coefficients 0 or 1. The addition is commutative. Terms involving the same k-simplex can be added together by adding their coefficients using modulo 2 arithmetic. The modulo 2 arithmetic implies that a k-simplex appears in the final sum when it appears an odd number of times. The *boundary of a k-simplex* σ is the sum of $(k-1)$-simplices that are faces of σ. The *boundary of a k-chain* is the sum of the boundaries of its k-simplices. We use ∂ to denote the boundary operator. Figure 3(a) shows some examples.

A 1-chain is a *cycle* if its boundary is empty. The boundary of a 2-chain is always a cycle. The *length* of a cycle is the number of edges in it. We call a cycle *simple* if it is simple in the graph-theoretic sense. For example, in Figure 3(a), A' is a simple cycle but $\partial\Sigma$ is not. Recall that a cycle is *critical* if it consists

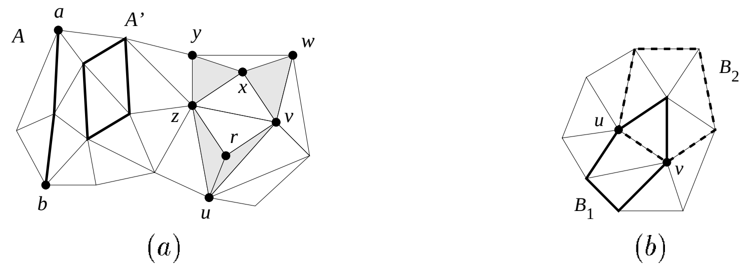

Figure 3: In (a), there are two 1-chains A and A' shown as bold line segments and there is a 2-chain Σ shown as shaded triangles. $\partial A = a + b$. $\partial A' = 0$. $\partial\Sigma = uv + vw + wx + xy + yz + uz + rv + vx + xz + rz$. In (b), there are two cycles B_1 and B_2 shown as bold solid and dashed line segments respectively. B_1 and B_2 cross at the vertices u and v.

of three edges and it does not bound a 2-simplex. So a critical cycle is always simple. Two cycles B_1 and B_2 are *homologous* if there exists a 2-chain Σ such that $B_1 = B_2 + \partial\Sigma$. For example, in Figure 3(a), $rv + vx + xz + rz$ and $uv + vw + wx + xy + yz + uz$ are homologous.

Let B_1 and B_2 be two simple cycles in K. Suppose that B_1 and B_2 share a vertex v such that the edges of B_1 and B_2 incident to v are distinct. If the edges of B_1 alternate with the edges of B_2 in $\mathrm{St}(v)$ (recall that there is a circular ordering of edges and 2-simplices in $\mathrm{St}(v)$), we say that B_1 and B_2 *cross at* v. We call v a *crossing of* B_1 *and* B_2. Figure 3(b) shows an example. The above definition of crossing might not be applicable when two cycles share edges. So we will perturb cycle edges in order to proceed further. We remark that the crossing of cycles, as discussed here, is related to the concept of *intersection number* in algebraic topology [7]. However, the definition of intersection number does not cater to edge sharing.

Since perturbation is a geometric operation, we need to work with a geometrical realization of K which is a simplicial complex $\widehat{\mathsf{K}}$ (embedded without self-intersection in a space of sufficiently high dimension [11]). $\widehat{\mathsf{K}}$ is a triangulation of a piecewise linear surface: each edge appears as a line segment and each 2-simplex appears as a triangle. Since K and $\widehat{\mathsf{K}}$ have identical combinatorial structure, we do not distinguish corresponding cycles in K and $\widehat{\mathsf{K}}$. We would like to emphasize that $\widehat{\mathsf{K}}$ is only a tool. Our results are topological and independent of the geometric realization.

Let ξ_1 and ξ_2 be two simple closed curves on the underlying piecewise linear surface of $\widehat{\mathsf{K}}$. We say ξ_1 and ξ_2 *cross at a point* p if there is a small region $R(p)$ around p such that $\xi_1 \cap \xi_2 \cap R(p) = \{p\}$ and $\xi_1 \cap R(p)$ contains points on both sides of $\xi_2 \cap R(p)$ locally. We also call p a *crossing of* ξ_1 *and* ξ_2.

Let B_1 and B_2 be two simple cycles in K. We treat B_1 and B_2 as two simple closed curves on the underlying surface of $\widehat{\mathsf{K}}$. We perturb B_1 to another simple closed curve ξ_1 on $\widehat{\mathsf{K}}$ 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 $\mathrm{int}(\gamma)$ lies in the interior of a triangle of $\widehat{\mathsf{K}}$ incident to e, $\gamma \cap e$ consists of the endpoints of e, and $\mathrm{int}(\gamma)$ 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, so the definition of crossings of two simple closed curves is applicable. We use $B_1 \circ B_2$ to denote the *parity of the number of crossings of* ξ_1 *and* B_2. We can generalize the definition 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 we define $B_1 \circ B_2 = (\sum_{j=1}^{q} B_1 \circ B_{2j}) \bmod 2$. The following lemma shows that $B_1 \circ B_2$ is well defined and its proof can be found in Appendix I.[3]

[3]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.

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

Lemma 1 leads to the following lemma concerning the crossing between 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. By definition, $B_1 = B_2 + \partial\Sigma$ for some 2-chain Σ. So $A \circ B_1$ = $(A \circ B_2 + A \circ \partial\Sigma) \bmod 2$. Clearly, $A \circ \partial\tau = 0$ for any 2-simplex τ. Thus, $A \circ \partial\Sigma = 0$ 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: (1) each C_i is a critical cycle, (2) each D_i is a simple cycle of length at most ℓ, (3) 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, and (4) 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. The following lemma is the main result of this subsection.

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 Appendix II. First, we use the following lemma to select a subset $\mathcal{S}_\ell \subseteq \mathcal{F}_\ell$.

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 $\mathcal{S}_\ell = \{(C_i, D_i)$: the anchor of C_i 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. ◻

Next, we partition the C_i's in $\mathcal{S}_\ell$ into equivalence classes of mutually homologous cycles. We pick one cycle from each 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 prove that $|\mathcal{F}'_\ell| = \Omega(|\mathcal{F}_\ell|)$ by showing that each equivalence class has $O(1)$ cycles. Then the fact that K has at most $2g$ mutually non-homologous cycles yields an upper bound on $|\mathcal{F}_\ell|$. We need some definitions and an utility lemma (Lemma 5). Define a *whisk* to be a collection of mutually homologous C_i's in $\mathcal{F}_\ell$ such that they share a common edge xy and neither x nor y is the anchor of any C_i in the collection. We call xy the *axis* of the whisk. Given a whisk W, we use W^* to denote the set of vertices and edges in W, i.e., the graph formed by the union of the cycles in W. We also call W^* a whisk for convenience.

Then $B_1 \circ B_2$ can be defined to be $(\sum_{i=1}^{p} \sum_{j=1}^{q} B_{1i} \circ B_{2j}) \bmod 2$. However, this generalization is not needed for obtaining our results.

LEMMA 5 *Let W be a whisk. Let xy be the axis of W. Let $\mathcal{Z}$ be a set of whisks such that*

(i) any two cycles in $W \cup \bigcup_{V \in \mathcal{Z}} V$ are homologous,

(ii) $W \cap V = \emptyset$ for any $V \in \mathcal{Z}$,

(iii) $U^* \cap V^* \subseteq \{x, y\}$ for any two distinct whisks $U, V \in \mathcal{Z}$.

Then $|\mathcal{Z}| \leq \ell - |W|$ and $|W| \leq \ell$.

Proof. Let C_i be a cycle in W. Let D_i be the cycle that pairs up with C_i in $\mathcal{F}_\ell$. By definition, $D_i \circ C_i = 1$. Since C_i and C_j are homologous for any cycle C_j in any whisk in $\mathcal{Z}$, $D_i \circ C_j = D_i \circ C_i = 1$ by Lemma 2. It follows that D_i contains a vertex w of C_j. The vertex w cannot be x or y as D_i does not share an edge with C_i. Since $U^* \cap V^* \subseteq \{x, y\}$ for any two distinct whisks $U, V \in \mathcal{Z}$, each whisk 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 in W. Since D_i has length at most ℓ, we conclude that $|\mathcal{Z}| \leq \ell - |W|$. As $|\mathcal{Z}| \geq 0$, rearranging terms yields $|W| \leq \ell$. ◻

We are ready to bound $|\mathcal{F}'_\ell|$ from below.

LEMMA 6 *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 $\mathcal{S}_\ell$. We first bound $|\mathcal{H}|$. We pick maximal whisks $W_r \subseteq \mathcal{H}$, $1 \leq r \leq m$, in a greedy fashion such that $W_r^* \cap W_s^* = \emptyset$ for $1 \leq r \neq s \leq m$. By Lemma 5 ($W = W_r$ and $\mathcal{Z} = \{W_1, \cdots, W_m\} - \{W_r\}$), $m - 1 \leq \ell - |W_r|$ which implies that $m \leq \ell$ and

$$|W_r| \leq \ell + 1 - m. \tag{1}$$

We partition $\mathcal{H} - \bigcup_{r=1}^m W_r$ into a collection $\mathcal{Y}$ of maximal whisks. By the property of $\mathcal{S}_\ell$, no cycle in $\mathcal{S}_\ell$ contains the anchor of another cycle in $\mathcal{S}_\ell$. By greediness, for any $V \in \mathcal{Y}$, $V^* \cap W_r^* \neq \emptyset$ for some $1 \leq r \leq m$. If $V^* \cap W_r^* \neq \emptyset$, the maximality of W_r implies that $V^* \cap W_r^* = \{x\}$ for some endpoint x of the axis of W_r. Take any whisk $V \in \mathcal{Y}$. By Lemma 5 ($W = V$ and $\mathcal{Z} = \{W_1, \cdots, W_m\}$), we have $m \leq \ell - |V|$ which implies that

$$|V| \leq \ell - m, \text{for any } V \in \mathcal{Y}. \tag{2}$$

Let $x_{r1}x_{r2}$ be the axis of W_r. Let s_{rj}, $1 \leq j \leq 2$, be the number of whisks V in $\mathcal{Y}$ such that $V^* \cap W_r^* = \{x_{rj}\}$. By Lemma 5 ($W = W_r$ and $\mathcal{Z}$ = the set of whisks in $\mathcal{Y}$ that share x_{rj} with W_r^*), we have

$$s_{rj} \leq \ell - |W_r|. \tag{3}$$

Thus, $|\mathcal{H}| \overset{(2)}{\leq} \sum_{r=1}^m (|W_r| + (s_{r1} + s_{r2}) \cdot (\ell - m)) \overset{(3)}{\leq} \sum_{r=1}^m (|W_r| + 2(\ell - |W_r|)(\ell - m)) = \sum_{r=1}^m (2\ell(\ell - m) - (2\ell - 2m - 1)|W_r|)$. If $m = \ell$, then $|\mathcal{H}| \leq \sum_{r=1}^m |W_r| \leq \ell$ by (1). 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 $|\mathcal{S}_\ell| \geq |\mathcal{F}_\ell|/20$, $|\mathcal{F}'_\ell| \geq \frac{2}{\ell^3} \cdot \frac{1}{20} \cdot |\mathcal{F}_\ell| = |\mathcal{F}_\ell|/(10\ell^3)$. ◻

Proof of Lemma 3: If M_K has genus g, K contains at most $2g$ cycles that are mutually non-homologous. Thus, $|\mathcal{F}'_\ell| \leq 2g$. The result then follows from Lemma 6. The bound for $|\mathcal{F}_3|$ is provided in Appendix II. ◻

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 7 *Assume that* M_K *has positive genus. Let* A *be a critical cycle passing through vertices* v, x *and* y. *Then one of the following holds.*

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

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

Proof. Observe that $x, y \in \mathrm{Lk}(v)$. Let L be the list of vertices in $\mathrm{Lk}(v)$ in clockwise order starting at x (recall that $\mathrm{Lk}(v)$ 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 A is a critical cycle, x and y are not 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 M_K is positive, uv lies on a critical cycle B that passes through u, v and some vertex $w' \in \mathrm{Lk}(v)$. If A and B 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 8 *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 7(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 $(\mathrm{St}(u) \cup \mathrm{St}(v)) \cap (\mathrm{St}(r) \cup \mathrm{St}(s)) = \emptyset$. 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 4 shows an example.

Our proof proceeds in two steps. First, we focus on the contractible edges of K by considering a subgraph G_K that contains *all vertices* of K and the contractible edges of K. (So G_K might be disconnected.) We prove that a maximal matching of G_K has linear size. Second, we prove that any maximal matching of G_K 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* M_K. *Assume that* $g > 0$. *Any maximal matching of* G_K *matches at least* $(n - 1310g + 30)/16$ *vertices.*

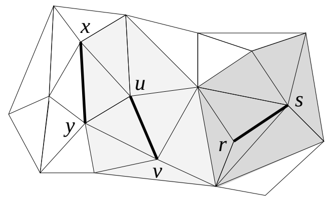

Figure 4: xy and uv are not independent, but both are independent from rs. The open regions covered by $\mathrm{St}(u) \cup \mathrm{St}(v)$ and $\mathrm{St}(r) \cup \mathrm{St}(s)$ are shaded differently. $\mathrm{N}(uv)$ and $\mathrm{N}(rs)$ share two vertices and one edge.

Proof. Our proof uses some geometric operations, so we again work with a geometric realization $\widehat{\mathsf{K}}$ of K. We use S to denote the underlying surface of $\widehat{\mathsf{K}}$, i.e., S is the set of points on $\widehat{\mathsf{K}}$ without the triangulation structure. We obtain an embedding of G_{K} on S by first drawing all vertices and edges of $\widehat{\mathsf{K}}$ on S and then erasing all the non-contractible edges. G_{K} induces a subdivision of S which we denote by $G_{\mathsf{K}}(\mathsf{S})$.

Pick a maximal matching of G_{K}. Let H_{K} be the subgraph of G_{K} (embedded on S) consisting of matched vertices and the edges of G_{K} between them. So H_{K} contains all matching edges but H_{K} may contain some non-matching edges as well. As our argument proceeds, we will create some segments on S, called *purple segments*, that connect matched vertices. The purple segments can be straight or curved. The purple segments will be used later to form a new graph with H_{K}.

We bound the number of unmatched vertices by charging them to edges in H_{K} and 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_{K} is at most 1, then Lemma 7(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_{K} is larger than 1. Since v is unmatched, all neighbors of v in G_{K} are matched. Let u and w be two consecutive neighbors of v in G_{K}. Let R be a region in $G_{\mathsf{K}}(\mathsf{S})$ such that uv and vw lie consecutively on the boundary of R. R covers some triangles in $\mathrm{St}(v)$, see Figure 5. We pick a subset $T \subseteq \mathrm{St}(v)$ of triangles such that $\bigcup_{t \in T} t \subseteq R$ and uv and vw lie in the boundary of the closure of $\bigcup_{t \in T} t$. Let R_{uvw} denote the closure of $\bigcup_{t \in T} t$. Figure 5 shows an example of R_{uvw}. Note that any incident edge of v in $\widehat{\mathsf{K}}$ that lies inside R_{uvw} is non-contractible. There are three different ways to charge for v.

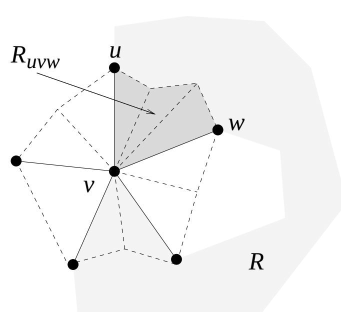

Figure 5: The solid line segments are incident edges of v in G_{K}. The dashed line segments are edges in $\mathrm{Lk}(v)$ and $\mathrm{St}(v)$ that are not in G_{K}. The shaded region is R. The darker subregion is R_{uvw}.

Case 1: v is not incident to any edge in $\widehat{\mathsf{K}}$ that lies inside R_{uvw}. It follows that $R_{uvw} = uvw$. So uw is an edge in $\widehat{\mathsf{K}}$. Since $R_{uvw} \subseteq R$, either uw lies inside R or uw lies on the boundary of R.

Case 1.1: uw is an edge in G_{K}. It follows that $R = uvw = R_{uvw}$. Since u and w are matched vertices, uw belongs to H_{K} too. We charge for v by putting a red pebble at uw. Since uw bounds at most two regions in $G_{\mathsf{K}}(\mathsf{S})$, case 1.1 can be applied at most twice to uw producing at most two red pebbles on uw. Thus, each edge of H_{K} receives at most two red pebbles.

Case 1.2: uw is not an edge in G_{K}. So uw is non-contractible. If we have created a purple segment γ_{uw} connecting u and w before, we put a green pebble at γ_{uw} to charge for v. Otherwise, we create the straight purple segment $\gamma_{uw} = uw$ and put a green pebble at γ_{uw} to charge for v.

We claim that the purple segment γ_{uw} receives at most two green pebbles overall. If γ_{uw} receives a second green pebble, the boundary of R is a closed polygonal line connecting four vertices. The four vertices are u, v, w, and an unmatched vertex v' such that v', u and w satisfy the conditions of case 1 and case 1.2 (with v replaced by v'). That is, $\gamma_{uw} = uw$ and two green pebbles are put at γ_{uw} to charge for v and v'. Since R does not have any unmatched vertex on its boundary other than v and v', γ_{uw} cannot receive a third green pebble. In all, each purple segment receives at most two green pebbles.

Case 2: Some edge vx in $\widehat{\mathsf{K}}$ lies inside R_{uvw}. Recall that any incident edge of v that lies inside R_{uvw} is non-contractible. So vx lies on a critical cycle A. Let vy and xy be the other two edges of A. If vy lies inside R_{uvw} or on the boundary of R_{uvw}, then Lemma 7(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 $\widehat{\mathsf{K}}$, we set $\gamma_{uw} = uw$. Otherwise, we draw γ_{uw} as a segment, curved if necessary, inside R_{uvw}. Clearly, γ_{uw} does not cross any edge of H_{K}. Moreover, by our drawing strategy, γ_{uw} does not cross any other purple segments created before.

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, $1 \le i \le 3$, other than v and each v_i is adjacent to both u and w. We pick v_k such that $v_k \ne x$ and $v_k \ne y$ (recall that x and y are vertices of the critical cycle A passing through v). Let B be the cycle consisting of the edges uv, vw, wv_k, and $v_k u$. Since vx lies inside R_{uvw} and vy lies outside R_{uvw}, A and B cross at v. We add the cycle pair (A, B) to $\mathcal{F}_4$ to charge for v.

By Lemma 3, $|\mathcal{F}_4| \le 1280g$. It remains to bound the total number of pebbles on the edges of H_{K} and the purple segments. Recall that there is no crossing among the edges of H_{K} and the purple segments. We add the purple segments as edges to H_{K} and we add more edges, if necessary, to obtain a connected graph H^* that is embedded on S 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, $E \le 3N - 6 + 6g$. Since each edge of H_{K} carries at most two red pebbles and 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 $5E \le 15N - 30 + 30g$.

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

THEOREM 10 *Let n be the number of vertices of K and let g be the genus of M_{K}. Assume that $g > 0$. For any constant $d > 380$, if $n \ge \frac{(6008+1310d)g-888-30d}{d-380}$, 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 $d > 380$. 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 $(d+4)|M|/2 \geq (d+4)(n - 1310g + 30)/64 \geq 6n - 12 + 12g$ by our choices of d and n, contradicting the Euler's relation. We pick the independent contractible edges as follows. First, we mark all matching edges in M. We pick a marked matching edge e whose neighborhood size is at most d, unmark e as well as all other matching edges that are not independent from e. We repeat the above until no more matching edge can be picked. Since at most $d + 1$ matching edges can be unmarked in each iteration, at least $|M|/(2(d+1))$ matching edges must be picked. ◻

Although the proof of Theorem 10 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 11 *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, the edge contractions should be selected to keep the geometric approximation error small. 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.[4] 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.

[4] Garland and Heckbert studied surface simplification and did not consider the computation of a hierarchy in their paper [10].

[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*, 10 (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 I

First, we prove that $B_1 \circ B_2$ is well defined when both B_1 and B_2 are simple cycles.

LEMMA 12 *Given two simple cycles B_1 and B_2 in K, $B_1 \circ B_2$ is independent of the perturbation of B_1.*

Proof. Let ξ_1 and ξ_1' be two simple closed curves on $\widehat{\mathsf{K}}$ obtained by different perturbations of B_1. Let e be an edge of B_1. Let γ and γ' be the two perturbed versions of e in ξ_1 and ξ' respectively. We modify ξ_1 by replacing γ with γ' and examine the parity of crossings between the new closed curve and B_2. If B_2 does not contain e, it is clear that the crossing status at the endpoints of e remain unchanged after the replacement (i.e., an endpoint is a crossing after the replacement iff it was a crossing). Consider the case where B_2 contains e. If γ and γ' lie inside the same triangle of $\widehat{\mathsf{K}}$ incident to e, the crossing status at each endpoint of e remains unchanged after the replacement. If γ and γ' lie inside different triangles of $\widehat{\mathsf{K}}$ incident to e, the crossing status at each endpoint of e is switched after the replacement (i.e., an endpoint is a crossing after the replacement iff it was not a crossing). Thus, the parity remains unchanged. ◻

We are ready to prove that $B_1 \circ B_2$ is well defined when B_2 is a sum of simple cycles.

LEMMA 13 *Given a simple cycle B_1 and a sum B_2 of simple cycles in K, $B_1 \circ B_2$ is independent of the sum expression of B_2 and the perturbation of B_1.*

Proof. Let v be a shared vertex between B_1 and B_2. Let ξ_1 be the simple closed curve on $\widehat{\mathsf{K}}$ obtained by a perturbation of B_1. ξ_1 divides a small region around v into two topological disks N_1 and N_2. Since B_2 is a sum of simple cycles, the number of edges of B_2 incident to v is even. It follows that the numbers of edges of B_2 in N_1 and N_2 have the same parity. The parity of the crossings of ξ_1 and B_2 at v is completely determined by whether N_i contains an odd or even number of edges of B_2. If the number is odd, the parity of crossings of ξ_1 and B_2 at v is odd. Otherwise, the parity is even. So the sum expression of B_2 is unimportant. We can also argue, as in the proof of Lemma 12, that the choice of ξ_1 is unimportant. ◻

Appendix II

We first bound $|\mathcal{F}_3'|$ from below.

LEMMA 14 *There is a subset $\mathcal{F}_3' \subseteq \mathcal{F}_3$ of cardinality at least $|\mathcal{F}_3|/120$ such that for 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 $\mathcal{S}_3$. Our goal is to bound $|\mathcal{H}|$. We pick maximal whisks $W_r \subseteq \mathcal{H}$, $1 \leq r \leq m$, in a greedy fashion such that $W_r^* \cap W_s^* = \emptyset$ for $1 \leq r \neq s \leq m$. By greediness, $\{W_1, \cdots, W_m\}$ is maximal. We partition $\mathcal{H} - \bigcup_{r=1}^{m} W_r$ into a collection $\mathcal{Y}$ of maximal whisks. For any whisk $V \in \mathcal{Y}$, observe that:

- Since $\{W_1, \cdots, W_m\}$ is maximal, $V^* \cap W_r^* \neq \emptyset$ for some $1 \leq r \leq m$.
- If $V^* \cap W_r^* \neq \emptyset$, then $V^* \cap W_r^* = \{x\}$ for some endpoint x of the axis of W_r and x is not the anchor of any cycle in V by the property of $\mathcal{S}_3$.
- For any two distinct $U, V \in \mathcal{Y}$, U^* does not contain the anchor of any cycle in V by the property of $\mathcal{S}_3$.

By Lemma 5 ($W = W_r$ and $\mathcal{Z} = \{W_1, \cdots, W_m\} - \{W_r\}$) , we have $m - 1 \leq 3 - |W_r|$ and $|W_r| \leq 3$. Since $|W_r| \geq 1$, we have $m \leq 3$. We conduct a case analysis to show that $|\mathcal{H}| \leq 6$.

Case 1: $m = 1$. Let $x_1 x_2$ be the axis of W_1. We partition $\mathcal{Y}$ into $\mathcal{Y}_1 \cup \mathcal{Y}_2$, where

$$\mathcal{Y}_j = \{V \in \mathcal{Y} : V^* \cap W_1^* = \{x_j\}\}.$$

If $|W_1| = 3$, then by Lemma 5 ($W = W_1$ and $\mathcal{Z} = \mathcal{Y}_j$), $|\mathcal{Y}_j| \leq 3 - |W_1| = 0$. So $|\mathcal{H}| = |W_1| = 3$. Consider the case where $1 \leq |W_1| \leq 2$. For any whisk $V \in \mathcal{Y}_j$, by Lemma 5 ($W = V$ and $\mathcal{Z} = \{W_1\} \cup (\mathcal{Y}_j - \{V\})$), we have $(|\mathcal{Y}_j| - 1) + 1 \leq 3 - |V|$. It follows that $|\mathcal{Y}_j| \leq 3 - |V|$. Since $|V| \geq 1$, $|\mathcal{Y}_j| \leq 2$. This implies that there are at most four cycles in $\mathcal{H} - W_1$. So $|\mathcal{H}| \leq |W_1| + 4 \leq 6$.

Case 2: $m = 2$. For $1 \leq r \leq 2$, let $x_{r1} x_{r2}$ be the axis of W_r. We partition $\mathcal{Y}$ into $\mathcal{Y}_{11} \cup \mathcal{Y}_{12} \cup \mathcal{Y}_{21} \cup \mathcal{Y}_{22} \cup \mathcal{X}$, where

$$\begin{aligned} \mathcal{Y}_{rj} &= \{V \in \mathcal{Y} : V^* \cap W_r^* = \{x_{rj}\}, V^* \cap W_{3-r}^* = \emptyset\} \\ \mathcal{X} &= \{V \in \mathcal{Y} : V^* \cap W_1^* \neq \emptyset, V^* \cap W_2^* \neq \emptyset\}. \end{aligned}$$

By Lemma 5 ($W = W_r$ and $\mathcal{Z} = \{W_{3-r}\} \cup \mathcal{Y}_{rj}$), we have $|\mathcal{Y}_{rj}| + 1 \leq 3 - |W_r|$ which implies that

$$|\mathcal{Y}_{rj}| \leq 2 - |W_r|. \tag{4}$$

Also, observe that $0 \leq |\mathcal{X}| \leq 4$.

If $|\mathcal{X}| = 0$, then $|\mathcal{H}| = \sum_{r=1}^{2}(|W_r| + |\mathcal{Y}_{r1}| + |\mathcal{Y}_{r2}|)$. By (4), we have $|\mathcal{H}| \leq 8 - |W_1| - |W_2| \leq 6$.

Suppose that $1 \leq |\mathcal{X}| \leq 2$. Let V be any whisk in $\mathcal{X}$. For $1 \leq r \leq 2$, by Lemma 5 ($W = V$ and $\mathcal{Z} = \{W_1, W_2\} \cup \mathcal{Y}_{rj}$ for some choice of j such that $V^* \cap W_r^* = \{x_{rj}\}$), we have $|\mathcal{Y}_{r1}| + 2 \leq 3 - |V|$ or $|\mathcal{Y}_{r2}| + 2 \leq 3 - |V|$. So $|\mathcal{Y}_{r1}| = 0$ or $|\mathcal{Y}_{r2}| = 0$, say $|\mathcal{Y}_{r1}| = 0$. Thus, $|\mathcal{H}| = |\mathcal{X}| + \sum_{r=1}^{2}(|W_r| + |\mathcal{Y}_{r2}|)$. By (4), $|\mathcal{H}| \leq |\mathcal{X}| + 4 \leq 6$.

Suppose that $3 \leq |\mathcal{X}| \leq 4$. For $1 \leq r \leq 2$, there exists an endpoint x of the axis of W_r where there are at least two distinct whisks $U_r, V_r \in \mathcal{X}$ such that $U_r^* \cap W_r^* = V_r^* \cap W_r^* = \{x\}$. For $1 \leq j \leq 2$, by Lemma 5 ($W = W_r$ and $\mathcal{Z} = \{U_r, V_r\} \cup \mathcal{Y}_{rj}$), we obtain $|\mathcal{Y}_{rj}| + 2 \leq 3 - |W_r|$ which implies that $|\mathcal{Y}_{rj}| = 0$. By Lemma 5 ($W = W_r$ and $\mathcal{Z} = \{U_r, V_r\}$), we have $2 \leq 3 - |W_r|$ which implies that $|W_r| = 1$. Thus, $|\mathcal{H}| = |\mathcal{X}| + |W_1| + |W_2| \leq 6$.

Case 3: $m = 3$. By Lemma 5 ($W = W_r$ and $\mathcal{Z} = \{W_1, W_2, W_3\} - \{W_r\}$), $2 \leq 3 - |W_r|$ which implies that $|W_r| = 1$. We claim that $\mathcal{H} - \bigcup_{r=1}^{m} W_r$ is empty. Suppose not. Let C_j be a cycle in $\mathcal{H} - \bigcup_{r=1}^{m} W_r$ and let D_j be the cycle that pairs up with C_j in $\mathcal{F}_3$. By Lemma 2, D_j shares a vertex with W_r^* for $1 \leq r \leq 3$. For $1 \leq r \leq 3$, W_r^* does not contain the anchor $C_j \cap D_j$ by property of $\mathcal{S}_3$. So D_j must contain at least four vertices, a contradiction. Thus, $|\mathcal{H}| = \sum_{r=1}^{3} |W_r| = 3$.

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 such $\mathcal{H}$. 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 15 $|\mathcal{F}_3| \leq 240g$.

Proof. If $\mathsf{M_K}$ has genus g, then K contains at most $2g$ cycles that are mutually non-homologous. Thus, $|\mathcal{F}_3'| \leq 2g$. Then Lemma 14 implies that $|\mathcal{F}_3| \leq 240g$. ☐