FINITE MANIFOLDS AND MINIMAL FINITE MODELS OF CLOSED SURFACES

DORON L GROSSMAN-NAPLES

Abstract. We describe Stong minimality and absolute minimality of finite models of topological spaces, exhibit Stong minimal models of all closed surfaces, and derive several elementary lower bounds for the size of absolutely minimal models. We define the notion of a finite manifold and characterize finite surfaces, then use this characterization to show that a finite model of a closed surface is a finite surface if and only if it is induced by a regular CW structure on the surface. Finally, we use this result to deduce a better lower bound for the size of models which are finite surfaces and construct minimal finite surface models of orientable surfaces whose genera satisfy nice number-theoretic properties.

Contents

1. Introduction

The study of finite spaces is based on a beautiful isomorphism between the category of finite posets and the category of finite T_0 spaces which allows the study of certain combinatorial properties with topological techniques and vice versa. We assume that the reader is familiar with the basic theory of these spaces, and use "poset" and "finite T_0 space" interchangably. An introduction to the subject written by J.P. May can be found in the first three chapters of [6].

When studying these strange new spaces, it is natural to attempt to relate them to more familiar spaces such as CW complexes. In particular, if we could relate the homotopical information of posets and ordinary spaces, we could both gain a better understanding of how finite T_0 spaces look and study homotopical properties of ordinary spaces using combinatorial techniques. While outright homotopy equivalence between ordinary spaces and finite spaces is easily seen to be impossible ([1], Theorem 1.3), this interest is justified by a pair of results of McCord ([7]).

Date: November 23, 2018.

Definition 1.1. Given a finite poset X, its *order complex* is the simplicial complex $\mathcal{K}(X)$ whose n-simplices are $(n+1)$-chains $\{x_0 < x_1 < \cdots < x_n\}$ of X.

Definition 1.2. Given a regular CW complex K, its *face poset* $\mathcal{X}(K)$ is the poset of cells in K in which $A \leq B$ if $A \subset \overline{B}$.

Remark 1.3. If we consider X to be a category, the simplicial set associated to $\mathcal{K}(X)$ is just its nerve, and the geometric realization $|\mathcal{K}(X)|$ is its classifying space. In fact, both of these constructions are functorial on the appropriate categories.

Now we can state the results.

Theorem 1.4. *If X is a finite poset, there is a weak homotopy equivalence from $|\mathcal{K}(X)|$ to X.*

Theorem 1.5. *If K is a simplicial complex, $\mathcal{X}(K)$ is weak homotopy equivalent to $|K|$.*

These results give the desired correspondence for simplicial complexes; in fact, the second theorem can be generalized to regular CW complexes (with $|K|$ interpreted as K) as proven in [2]. This correspondence motivates the following definition.

Definition 1.6. If X is a topological space, a *finite model* of X is a finite T_0 space which is weak homotopy equivalent to X.

Since every space is weak homotopy equivalent to a regular CW complex[1], this implies that every space has a T_0 Alexandroff model, and if the regular CW complex is finite, so is the model. One basic example which illustrates the relationship between ordinary spaces and finite models is that of the non-Hausdorff suspension, introduced by McCord in [7]. Given a poset X, this is defined to be the poset obtained by adjoining two new points greater than each point in X. As suggested by the name, this is a finite version of the ordinary suspension functor for topological spaces, and in fact the normal suspension of a poset and the non-Hausdorff suspension are naturally weak homotopy equivalent. As S^0 is a finite T_0 space, this construction can be used to construct a finite model of S^n with $2n+2$ points for each n (Figure 1).

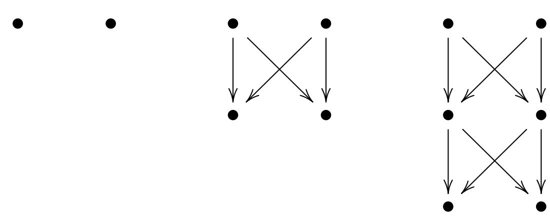

Figure 1: Hasse diagrams for finite models of S^0, S^1, and S^2.

Notice that a given space will have many finite models. For example, every finite simplicial structure gives rise to a finite model, and we can always enlarge a model

[1]Every space is weak homotopy equivalent to a CW complex, while every CW complex is homotopy equivalent to a simplicial complex of the same dimension; see, for example, Theorem 2C.5 in [5].

by adding beat points. To avoid superfluous information, reduce complexity, and gain a better understanding of these models, it is desirable to find finite models which are minimal in one of two senses.

Definition 1.7. We say a finite T_0 space is *Stong minimal* if its cardinality is minimal in its homotopy class. We say a finite T_0 space is *absolutely minimal* if its cardinality is minimal in its weak homotopy class.[2]

Note that the second notion of minimality is stronger than the first; the sphere is a rare case where the evident Stong minimal model is absolutely minimal. It is also perhaps a more natural notion of minimality when it comes to the study of finite models, since it is equivalent to being the smallest finite model of a space. However, Stong minimality is easier to check, and Stong minimal models are easier to find: two finite T_0 spaces are homotopy equivalent if and only if their cores are homeomorphic, so a space is Stong minimal if and only if it has no beat points, and a Stong minimal model can be obtained from any finite model simply by removing beat points. In contrast, at the time of the writing of this paper, there is no known algorithm for reducing an arbitrary finite T_0 space to an absolutely minimal space, or even for determining whether a space is absolutely minimal.

For this reason, results regarding absolutely minimal models have typically involved exhibiting a particular model for a space and showing that no smaller space can have the same homotopy or homology groups. Barmak and Minian take this approach in [1] in which they show that the finite models described above are the unique absolutely minimal models of S^n for each n. Having found such models for the most basic topological spaces, we turn next to another well-known countable collection of spaces with simple homology: closed surfaces. Cianci and Ottina exibit absolutely minimal models of the torus, the projective plane, and the Klein bottle in [4], but their methods for bounding model size below are not related to the genus of these surfaces, and hence do not generalize directly to surfaces of higher genus. In this paper, we begin by describing Stong minimal models for all closed surfaces[3]. We then find some lower bounds for the size of arbitrary finite models of closed surfaces using results from [4] together with some elementary combinatorial facts. Having derived some minor results for the general case, we specialize to a particularly well-behaved class of finite T_0 spaces called finite manifolds, characterize them in dimension 2, and conclude by deriving a much stronger bound for finite models of this type, which we use to find some finite models which are minimal among finite surfaces.

2. Stong minimal models of closed surfaces

In this section, we construct regular CW models for all closed surfaces and show that the associated posets have no beat points, making them Stong minimal. These models are generalizations of absolutely minimal models presented in [4].

Given an orientable surface S of genus g, the usual CW structure for S is a regular $4g$-gon with edge identifications represented by the word $a_1b_1a_1^{-1}b_1^{-1}...a_gb_ga_g^{-1}b_g^{-1}$.

---

[2]These definitions are standard, but the terminology is not: the first is typically called a "minimal finite space" or simply "minimal", and the second a "minimal finite model". This nomenclature allows for such peculiar entities as spaces which are both minimal and finite models, but are not minimal finite models. We use different terminology to avoid confusion.

[3]We assume throughout that our surfaces are connected, as all our results can be immediately generalized to disconnected closed surfaces by taking coproducts.

However, this structure is not regular. To fix this issue, we add in the perpendicular bisectors of each edge with no new identifications, splitting each external edge into two 1-cells attached by a 0-cell. We also gain one new vertex in the center at the intersection of all the bisectors. This gives a regular CW structure for S and thus a finite model (see Figure 2). It is easy to check that this model has $14g+2$ points. Note that every 0-cell is contained in at least two 1-cells, every 1-cell contains exactly two 0-cells and is contained in at least two 2-cells, and every 2-cell contains exactly two 1-cells. Consequently, no vertex in the Hasse diagram of the model has in-degree or out-degree 1, so there are no beat points. Thus, this model is Stong minimal.

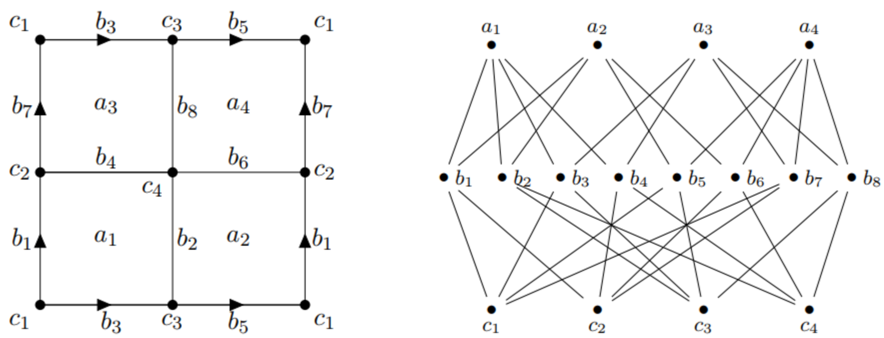

Figure 2: The regular CW structure and associated finite model for the orientable surface of genus 1. Taken from [4].

The construction of the models for nonorientable surfaces is similar. Given a nonorientable surface S of genus g, we begin with the usual CW model: the regular $2g$-gon with edge identifications given by the word $a_1^2...a_g^2$. To make this regular, we add in both the perpendicular bisectors of the edges and the line segments between opposing vertices. This yields a regular CW structure with $11g+2$ points, and the face poset is Stong minimal for the same reason as in the orientable case.

3. Elementary bounds

In this section, we use the weak homotopy invariance of Euler characteristic and several results from [4] to derive lower bounds for the size of arbitrary finite models of closed surfaces other than S^2 and $\mathbb{R}P^2$, whose absolutely minimal finite models are already known. We denote the Euler characteristic of a space X by $\chi(X)$ and the cardinality of X by $\#X$.

In [4], Cianci and Ottina define what they call a *splitting property* (S2) for finite posets. The details are not relevant, but the following result they derive is.

Proposition 3.1. *Let X be a finite T_0 space which is Eilenberg-MacLane of type $(G,1)$. If X satisfies (S2) then $H_1(X)$ is free abelian and $H_n(X)=0$ for $n>1$.*

We obtain the following corollary.

Corollary 3.2. *No closed surface other than S^2 or $\mathbb{R}P^2$ can have a model satisfying (S2).*

Proof. Let S be a closed surface that is not S^2 or $\mathbb{R}P^2$. Then S is covered by $\mathbb{R}^2$, so it is Eilenberg-MacLane of type $(\pi_1(S),1)$. If S is nonorientable, then $H_1(S)$ is not free abelian, and if S is orientable, $H_2(S)$ is nontrivial. Since homology is a weak homotopy invariant, the same is true of any finite model of S, so the result follows by the previous proposition. □

There are two relevant consequences derived in [4] of not satisfying (S2).

Proposition 3.3. *If X is a finite T_0 space not satisfying (S2) that is connected and Stong minimal, then $\#X \geq 16$.*

Proposition 3.4. *If X is a finite T_0 space with at most two maximal points or at most two minimal points, X satisfies (S2).*

These give us our lower bounds.

Theorem 3.5. *Let X be a finite model of a surface S other than S^2 or $\mathbb{R}P^2$. Then $\#X \geq \max(16, \log_2(|\chi(S)|))$.*

Proof. Since we can reduce any model to a Stong minimal model by removing beat points, we may assume without loss of generality that X is already Stong minimal. Furthermore, because path-connectedness is detected by homology, X must be path-connected and thus connected. It follows that $\#X \geq 16$.

To obtain the other bound, note that since Euler characteristic is weak homotopy invariant, $\chi(S) = \chi(\mathcal{K}(X))$, which is the alternating sum of the number of chains in X of various lengths, $\sum_{k \text{ a chain}}(-1)^{\#k+1}$. By the triangle inequality, the absolute value of the Euler characteristic must be less than or equal to the total number of chains in X, $\sum_{k \text{ a chain}} 1$. Since chains are subsets of X, this is less than or equal to the number of subsets of X, $2^{\#X}$. The result follows. □

We can improve our logarithmic bound to a square root bound in the case where X has height 3.

Proposition 3.6. *Let X be a height-3 finite model of a surface other than S^2 or $\mathbb{R}P^2$. Then $\#X \geq \sqrt{2|\chi(S) - 7|}$.*

Proof. Let $n = \#X$. The only negative contribution to $\chi(\mathcal{K}(X))$ is from the edges, of which there are at most $\binom{n}{2}$ since they are 2-chains in X. We know that there are at least 6 vertices since X does not satisfy (S2): there are at least three maximal points and three minimal points by Proposition 3.4, and no point can be both maximal and minimal since X is connected. We also know there must be at least one face because X is of height 3. Thus, there must be at least $|\chi(S) - 7|$ edges, so $n^2 \geq n^2 - n = 2\binom{n}{2} \geq 2|\chi(S) - 7|$, from which the result follows. □

It is conceivable that this method could be extended to posets of greater height. (It is trivial from the simplicial homology of $\mathcal{K}(X)$ that any finite model must have height at least 3.)

4. Characterization of finite manifolds

We now describe a particularly well-behaved class of finite spaces and characterize them in dimension 2.

Definition 4.1. A finite T_0 space X is a *finite n-manifold* if $|\mathcal{K}(X)|$ is a topological n-manifold.

Remark 4.2. We can extend much of the usual language for topological manifolds to finite manifolds. For example, we call a finite 2-manifold a finite surface, and we can define an analogous notion of a finite manifold with boundary.

Remark 4.3. It is immediate from invariance of dimension that a finite n-manifold must be of height $n+1$.

The first of the following definitions is standard, while the second extends the idea of the first to finite T_0 spaces.

Definition 4.4. An n-dimensional simplicial complex K is called *pure* if every simplex in K is contained in an n-simplex.

Definition 4.5. A finite T_0 space X of height n is called *pure* if every maximal chain in X is of height n.

As suggested by the terminology, these notions are equivalent.

Proposition 4.6. *A finite T_0 space X is pure if and only if $\mathcal{K}(X)$ is pure, and a finite simplicial complex K is pure if and only if $\mathcal{X}(K)$ is pure.*

Proof. Suppose X is a finite T_0 space of height n, so $\mathcal{K}(X)$ is a simplicial complex of dimension $n-1$. A k-simplex in $\mathcal{K}(X)$ is a chain of length $k+1$, so every simplex in $\mathcal{K}(X)$ is contained in an $(n-1)$-simplex if and only if every chain in X is contained in a chain of length n.

Now suppose K is a finite simplicial complex of dimension n, so $\mathcal{X}(K)$ is a poset of height $n+1$. The height of a maximal chain in $\mathcal{X}(K)$ is one greater than the dimension of the largest simplex it contains, so every maximal chain is of height $n+1$ if and only if every simplex in K is contained in an n-simplex. □

The reason for introducing pureness is that it plays an important role in the characterization of finite surfaces.

Theorem 4.7. *A finite T_0 space X is a finite surface if and only if it satisfies the following conditions:*

(i) X is pure of height 3;
(ii) For each height-2 point x, there are exactly two points greater than x and two points less than x; and
(iii) For each maximal point x_m and each minimal point x_n, the set $(x_m, x_n) = \{x \in X \mid x_n < x < x_m\}$ contains either zero or two points.
(iv) For each extremal point x, the set of points other than x which are comparable to x is connected.

The bulk of the proof of this theorem is based on the corresponding result for simplicial complexes. Stating it requires the following standard definition.

Definition 4.8. If v is a vertex in a simplicial complex K, the *link* of v, $\mathrm{Lk}(v, K)$, is the undirected graph whose vertices are the 1-simplices of X with v as a face, and where there is an edge between two vertices if they are faces of a common 2-simplex.

Lemma 4.9. *The geometric realization of a finite simplicial complex K is a surface if and only if K satisfies the following conditions:*

(i) K is pure and 2-dimensional;
(ii) Each 1-simplex of K is a face of exactly two 2-simplices; and
(iii) For each vertex v of K, $|\mathrm{Lk}(v, K)|$ is homeomorphic to S^1.

Proof. If (i) fails, $|K|$ is not a surface by invariance of dimension. If (ii) fails, removing a line from any sufficiently small connected neighborhood of a point in the edge yields three components, so it is not locally Euclidean. If (iii) fails, removing v from a sufficiently small connected neighborhood yields two components, so it is not locally Euclidean.

Suppose now that all three conditions hold. Then (i) guarantees that we only need to check the interior of 0-, 1-, and 2-simplices. The last is trivial. Since gluing together two polygons at an edge yields a Euclidean neighborhood for points on the edge, 1-simplices follow by (ii). Finally, 0-simplices follow by (iii), since it implies that at a 0-simplex v, the realization is locally homeomorphic to the disk obtained by gluing together triangles along their edges circularly. □

Now we can prove the theorem.

Proof. The first condition for the poset is equivalent to the first condition for the simplicial complex. Given pureness, the second and third poset conditions together are equivalent to the second simplicial complex condition, because (together with the pureness) they are equivalent to the statement that for any two comparable points p and q, there are exactly two ways of extending the 2-chain $\{p, q\}$ to a 3-chain. Finally, the second, third and fourth poset conditions together are equivalent to the third simplicial complex condition, since a graph is a circle if and only if it is connected and each vertex has degree 2. □

There is an alternate characterization of finite surfaces which is also useful. While it is ultimately just a more compact rephrasing of Theorem 4.7, we will see that it is convenient for a number of purposes. The proof is given by point-counting together with the above criterion for a graph to be a circle, and comparing to the conditions of our original classification.

Definition 4.10. Let X be a finite poset and $x \in X$. Then the *link* of x, $\mathrm{Lk}(x)$, is the set of points other than x which are comparable to x.

Corollary 4.11. *A finite T_0 space X is a finite surface if and only if for each $x \in X$, $|\mathrm{Lk}(x)|$ is homeomorphic to S^1.*

One of the reasons this statement of the theorem is advantageous is that it can more easily describe the higher-dimensional version of the theorem. Although we have written it out specifically for finite surfaces, the proof of this theorem generalizes directly to higher dimensions[4] , so we obtain the following.

Corollary 4.12. *A finite T_0 space X is a finite n-manifold if and only if for each $x \in X$, $|\mathcal{K}(\mathrm{Lk}(x)|)$ is homeomorphic to S^{n-1}.*

Another benefit of this form of the theorem is its relationship to the following result of A. Björner in [3].

Theorem 4.13. *Let P be a finite poset, and for each $x \in P$, denote the set of points less than x by $\hat{U}_x$. Then P is the face poset of a regular CW complex if and only if for each $x \in P$, $|\mathcal{K}(\hat{U}_x)|$ is homeomorphic to a sphere.*[5]

[4]In two dimensions, the conditions guarantee precisely that we have triangles glued in a circular fashion, which yields a Euclidean neighborhood of every point. The higher-dimensional equivalent is for n-simplices to be glued so as to form a ball in the neighborhood of a vertex, which is expressed via the condition that the indicated poset has order complex homeomorphic to S^{n-1}.

[5]We take the empty space to be the sphere of dimension -1.

This gives us a final characterization of finite surfaces which will be crucial in obtaining our bound in the next section.

Theorem 4.14. *A finite T_0 space X is a finite surface if and only if it is the face poset of a regular CW structure on some closed surface.*

Proof. Firstly, suppose $X = \mathcal{X}(Y)$, where Y is a regular CW structure on some closed surface. Then $|\mathcal{K}(X)|$ is nothing more than the cellular subdivision of Y, so the two are homeomorphic.

Suppose conversely that X is a finite surface, and let x be some point in X. If x is minimal, then by definition $|\mathcal{K}(\hat{U}_x)| \cong S^{-1}$. If x is on the second level, there are exactly two points below it by Theorem 4.7, so $|\mathcal{K}(\hat{U}_x)| \cong S^0$. Finally, if x is maximal, then $|\mathcal{K}(\hat{U}_x)| \cong S^1$ by Theorem 4.13. □

Before moving on, it is worth taking a moment to consider the common element between Corollary 4.11 (more generally Corollary 4.12) and Theorem 4.13 which allowed us to prove this relationship: subposets whose order complexes are simplicial spheres. It is generally a nontrivial problem to determine whether a finite poset has this property, although the case in dimension 1 is simple: a height-2 poset has geometric realization S^1 if and only if it is connected and every vertex has degree 2. Given this fact, the following theorem suggests the idea of taking an inductive approach to the problem.

Theorem 4.15. *If X is a finite n-manifold, then $|\mathcal{K}(X)|$ is homeomorphic to S^n if and only if X is a finite model of S^n.*

Proof. One direction is obvious: if $|\mathcal{K}(X)|$ is homeomorphic to S^n, then X is a finite model of S^n by Theorem 1.4.

To prove the other direction, suppose X is a finite n-manifold which is a finite model of S^n. Then $|\mathcal{K}(X)|$ is a CW space which is weak homotopy equivalent to S^n, and hence homotopy equivalent to it by the Whitehead theorem. Since $|\mathcal{K}(X)|$ is a closed n-manifold, the result follows by the Poincaré conjecture. □

5. Bounds for finite surfaces

Throughout this section, we will denote the number of height 1, 2, and 3 points by ℓ, m, and n respectively.

The problem of finding absolutely minimal finite models amounts to minimizing the sum of the number of points at each level. As the following result shows, by restricting to finite surfaces, we need consider only one number rather than three or more.

Proposition 5.1. *Let X be a finite surface which is a model of a closed surface S of genus g. If S is orientable, then $\#X = 2m + 2 - 2g$. If S is nonorientable, then $\#X = 2m + 2 - g$.*

Proof. Because X is a finite surface, $\#X = \ell + m + n$. But we also know by Theorem 4.14 that X is the face poset of a regular CW complex structure on S, so $n - m + \ell = \chi(S)$. Thus, $\#X = \ell + m + n = 2m + \chi(S)$. The result follows by the standard formula for the Euler characteristic of a closed surface. □

Using the fact that any finite model of a closed surface other than $\mathbb{R}P^2$ or S^2 must satisfy the (S2) splitting property and thus have at least three maximal and three minimal points, we can immediately derive from this the linear lower bounds $2g+10$ and $g+10$ for the size of finite surface models of orientable and nonorientable closed surfaces respectively. However, we can do slightly better than this.

Theorem 5.2. *Let X be a finite surface modelling the closed surface S of genus g. If S is orientable, then $\#X \geq 2\lceil 4\sqrt{g}\rceil + 2g + 6$. If S is nonorientable, then $\#X \geq 2\lceil 2\sqrt{2g}\rceil + g + 6$.*

Proof. Let c_i denote the degree of the i^{th} maximal point in the Hasse diagram of X. Then since each point in the middle level has up-degree 2, $\sum_i c_i = 2m$, and so for at least one i, $c_i \geq 2m/n$. Call this point x_i. Because $\hat{U}_{x_i}$ is a finite model of S^1, the number of minimal points less than x_i must be equal to the number of level 2 points less than x_i, which is just c_i. Thus, we get $c_i \leq l$, so $\lceil 2m/n\rceil \leq l$. The same argument for bottom points shows that $\lceil 2m/\ell\rceil \leq n$.

Adding these inequalities (and ignoring the ceilings), we get $2m(1/n + 1/\ell) \leq n + \ell = m + \chi(S)$, since $n - m + \ell = \chi(S)$. The smallest possible value of the left side of the inequality is achieved when $n = \ell = (m + \chi(S))/2$, and we get $8m \leq (m + \chi(S))^2$. Solving, we get $m \geq 4\sqrt{g} + 2g + 2$ in the orientable case and $m \geq 2\sqrt{2g} + g + 2$ in the nonorientable case. The result follows from Proposition 5.1. □

It is not clear that these inequalities are sharp, especially because we dropped the ceilings to derive them. However, there are some cases in which we can be certain they are achieved. To show this, we perform the following construction, illustrated in Figure 3.

Proposition 5.3. *Let n and ℓ be positive even integers and set $2m = n\ell$. Then there is a finite orientable surface with n, m, and ℓ points in its third, second, and first levels respectively.*

Proof. To construct this surface, take n ℓ-gons and identify them in the following way. Glue every other edge of the first ℓ-gon to every other edge of the second with coherent orientation, then glue the remaining edges of the second ℓ-gon to every other edge of the third (again with coherent orientation), and continue until the final ℓ-gon is glued back to the first. Because we have an even number of polygons, the final gluing will also have coherent orientation. Explicitly, we may embed the polygons in $\mathbb{R}^3$ centered at equal intervals along a circle and with parallel top edges, and glue them together via homotopies of $\mathbb{R}^3$. Then each step of gluing switches the sides which are glued between containing and not containing the top edge, so having an even number of polygons guarantees that the first and last polygons will glue properly, so the space we have constructed admits an embedding in $\mathbb{R}^3$. This construction also guarantees that the link of every vertex will be a circle (since it is connected and every vertex in the graph has degree two) and every edge will be adjacent to exactly two faces, so this will produce a closed orientable surface with a regular CW structure consisting of n faces, m edges, and ℓ vertices. We finish the construction by taking its face poset. □

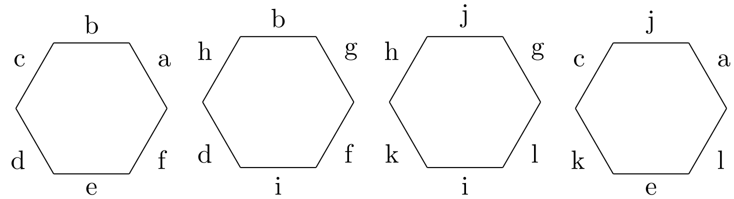

Figure 3: The polygons and identifications obtained by performing the construction with $\ell = 6$ and $n = 4$. All edges are oriented clockwise. Performing the gluing will yield the orientable surface of genus 2.

If we take $n = 4$, $\ell = 6$, this produces a model of the orientable surface of genus 2 with $n = 12$ (Figures 3,4). Geometrically, this is obtained by gluing together four hexagons in pairs to obtain two pairs of pants, then gluing together the pairs of pants to obtain the surface. By our bound above, this is minimal among finite orientable surfaces of genus 2.

It is unfeasible to explicitly construct every model individually to check if it achieves our bound. However, as the following theorem shows, for g with particularly nice number-theoretic properties, we don't need to.

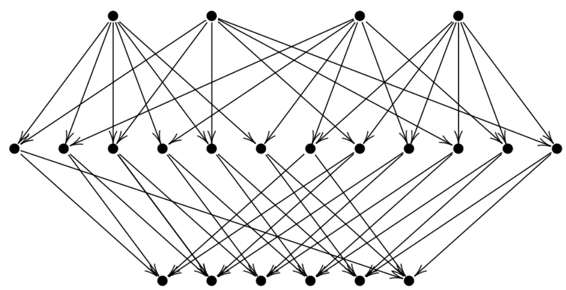

Figure 4: A minimal finite orientable surface of genus 2.

Theorem 5.4. *If g is a perfect square, then performing this construction with $n = \ell = 2\sqrt{g} + 2$ yields a minimal finite orientable surface of genus g.*

Proof. The resulting space has $m = 2g + 2 + 4\sqrt{g}$, so its Euler characteristic is $n - m + \ell = 2 - 2g$, which shows that it is indeed of genus g. Its cardinality is $\ell + m + n = 8\sqrt{g} + 2g + 6$, and since $\sqrt{g}$ is an integer, this is precisely the lower bound derived above. □

The simplest case is when g is a perfect square. However, the lower bound is more generally achieved by this construction when g is a product of two integers which are sufficiently close. For example, if g is of the form $(k-1)(k-2)$, then as long as k is at least 3, we get $4k - 7 < 4\sqrt{g} \leq 4k - 6$, so $\lceil 4\sqrt{g} \rceil = 4k - 6 = k - 1 + 2 + 4\sqrt{g}$, and setting $n = 2k$, $\ell = 2(k-1)$ yields a surface of the desired genus which achieves the bound. To further generalize this result is a problem of number theory.

Acknowledgments

It is a pleasure to thank my mentor Weinan Lin for his guidance throughout the writing of this paper. I would also like to thank Professor Peter May, both for his suggestions when I was first starting this paper and for the work he has put into running the UChicago REU.

References

[1] J.A. Barmak and E.G. Minian. Minimal Finite Models. Journal of Homotopy and Related Structures, vol. 2(1) (2007), pp. 127-140.

[2] J.A. Barmak and E.G. Minian. One-point reductions of finite spaces, h-regular CW-complexes and collapsibility. Algebraic Geometry and Topology, vol. 8 (2008) pp. 1763-1780.

[3] A. Björner. Posets, Regular CW Complexes and Bruhat Order European Journal of Combinatorics, vol. 5(1) (1984), pp. 7-16.

[4] N. Cianci and M. Ottina. Poset splitting and minimality of finite models. Preprint, 2015, arXiv:1512.06088v1.

[5] A. Hatcher. Algebraic Topology. Cambridge University Press (2002).

[6] J. P. May. Finite Spaces and Larger Contexts. http://math.uchicago.edu/ may/REU2018/ FINITEBOOK.pdf

[7] M. McCord. Singular homology groups and homotopy groups of finite topological spaces. Duke Mathematics Journal 33 (1966), pp. 465-474.