Classification of Minimal Separating Sets of Low Genus Surfaces*

Chris N. Aagaard, J.J.P. Veerman

August 11, 2023

1 Introduction

Given a metric space (X, d) and two distinct points $p, q \in X$, the *mediatrix of* p *and* q, which we'll denote L_{pq}, is

$$L_{pq} := \{x \in X | d(p, x) = d(q, x)\}.$$

In other work it is also known as an equidistant set, bisector, ambiguous locus, midset, medial axis, or central set ([3],[13]). In Brillouin spaces, a large class of metric spaces which includes all Riemannian manifolds, it was shown in [19] that all mediatrices have the following minimal separating property:

Definition 1.1. *A subset $M \subset X$ of a topological space X is called a separating set if $X \setminus M$ is disconnected. M is minimal separating in X if M is a separating set for X and for any $M' \subsetneq M$, $X \setminus M'$ is connected.*

Additionally, it is shown in [1] that when X is a compact Riemann surface then $L_{p,q}$ is a finite topological graph. Then, in [3], this result is generalized to show that for any compact 2-dimensional Alexandrov space X and any $\{p, q\} \subset X$, $L_{p,q}$ is a finite topological graph. That leads us to investigate which topological graphs can be realized as minimal separating sets in surface of genus g. (From here on unless otherwise specified, surface is taken to mean a connected compact oriented surface without boundary).

In this paper, we will provide a classification of precisely when a graph embedded in a surface is a minimal separating set and then use that result to classify the topological graphs

*This work was made possible in part thanks to Research Computing at Portland State University and its HPC resources acquired through NSF grant numbers 2019216 and 1624776.

which can be realized as minimal separating sets in a surface of genus g. We then use this classification to write a computer program which determines the number of homeomorphism classes of minimal separating sets in surfaces of genus at most 4.

In section 2 we define useful terminology, and establish a simple but useful lemma. Then in section 3, we prove a lemma relating minimal separating embeddings with more commonly studied cellular embeddings and then use this to establish a relationship between general (potentially disconnected) minimal separating sets in a surface and connected minimal separating sets, which are more easily compute. Section 4 introduces the language of topological and combinatorial maps to describe embeddings of connected graphs, and then more general objects called hypermaps, which we can use to speed up our computation. Section 5 describes an algorithm to find all homeomorphism classes of graphs which embed as minimal separating sets in a surface and has computational results from using that algorithm to find all graphs that embed as minimal separating sets in surfaces of genus at most 4.

Reframing the problem in terms of combinatorial maps presents an interesting connection to existing work in map enumeration, particularly the enumeration of unrooted maps discussed in [11], [14], and [15] which we discuss in the concluding remarks. We require an actual *list* of maps in order to determine numbers of homeomorphism classes of minimal separating sets for a given genus, which makes the purely enumerative methods in those papers insufficient here, but the particular set of combinatorial maps we consider presents some difficulties in applying their methods and provides a set to consider for map enumeration problems.

2 Background and Definitions

As observed in [1], a graph G embeds as a minimal separating set in a surface S of genus g, then it also embeds as a minimal separating set in any surface of genus greater than g. Such an embedding can be constructed by taking a minimal separating embedding η of G in S and adding handles within a component of $S \setminus \eta(G)$ until the resulting surface has the desired genus.

Given this observation, we define the following three sets:

Definition 2.1.
i) $\mathcal{M}_g$ *is the set of all graphs which can be realized as minimal separating sets in a surface of genus* g.
ii) $\mathcal{L}_g$ *is the set of all graphs which can be realized as minimal separating sets in a surface of genus* g*, but not in a surface of genus* $g-1$*. We say they have least separating genus* g*.*

iii) $\mathcal{C}_g$ *is the set of all connected graphs with least separating genus* g.[1]

We restate the observation from [1] as $\mathcal{L}_g = \mathcal{M}_g \setminus \mathcal{M}_{g-1}$. The relationship between $\mathcal{L}_g$ and $\mathcal{C}_g$ is more involved and we return to it later. For now we merely remark that it is often easier to deal with connected graphs so it will be convenient to focus on the connected case.

Much of what follows will consider graphs as combinatorial objects so we should mentiona a key distinction between topological graphs and combinatorial graphs which is the distinction between homeomorphism and graph isomorphism.

Any topological graph G with a vertex v of degree 2 is homeomorphic to another topological graph $\hat{G}$ where v and its 2 incident edges are replaced by a single edge. So, given a connected topological graph G (with no component homeomorphic to a circle), we can find a unique topological graph $\hat{G}$, homeomorphic to G, and with no degree 2 vertices, which we can construct by replacing all degree 2 vertices by single edges. Also, given any G we can construct an infinite family of graphs which are homeomorphic to G by repeatedly replacing single edges with 2 edges subdivided by a vertex. This means that every mediatrix is homeomorphic to infinitely many topological graphs, so, when we talk about topological graphs, we always implicitly talk about equivalence classes of graphs up to homeomorphism, and identify each equivalence class with its unique representative with no degree 2 vertices (except in the case of circles, which we identify with the graph composed of a single vertex with a loop).

For the circle, we can see immediately that this embeds as a minimal separating set any surface, and in [1] it's shown that any minimal separating set in a surface of genus 0 is homeomorphic to such a graph. Thus $\mathcal{C}_0 = \mathcal{M}_0$ consists of the graph with one vertex and one loop, and for all $g \geq 1$, we identify elements of $\mathcal{C}_g$ with their unique representative with no degree 2 vertices.

Whether a graph can be realized as a minimal separating set for a surface X is a question about *embeddings* of a graph, so we introduce some terminology to discuss graph embeddings.

Definition 2.2. *An embedding of a (topological) graph* G *into a surface* S *is a continuous function* $\eta : G \to S$ *such that* η *is a homeomorphism onto its image. We say an embedding is a cellular embedding or a topological map*[2] *if* $X \setminus \eta(G)$ *is homeomorphic to a collection of open discs.*

Given a surface X, a topological graph G may have numerous embeddings, some of which are minimal separating sets, some of which are non-separating, and some of which are separating but not minimal separating. Consider the examples in Figure 1.

[1] $\mathcal{M}$ for *minimal separating*, $\mathcal{L}$ for *least separating genus*, and $\mathcal{C}$ for *connected*

[2] The name 'map' emphasizes that these look like maps in the sense of cartography. The components of $X \setminus \eta(G)$ are the countries and the edges of G are the borders between them.

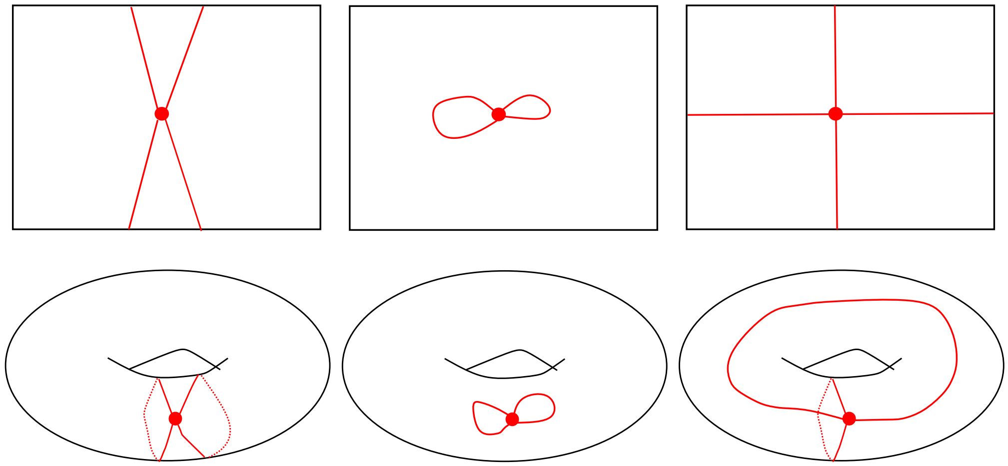

Figure 1: Three embeddings of the figure 8 in the torus. Each is shown drawn on the standard square torus with sides identified (upper) and on the torus as a donut. The left is minimal separating and the other two are not.

Because of this, we aim first to classify minimal separating sets in terms of equivalent embeddings. Then, a topological graph will be realizable as a minimal separating set in a surface S if it has such an embeding. Obtaining an actual list of such graphs can then be done as follows: Use the classification to obtain a complete list of equivalent minimal separating embeddings, then construct a list of minimal separating sets up to graph homeomorphism by testing the graphs underlying the embeddings for graph isomorphism on a computer. This notion of equivalence is formalized with the following defintion.

Definition 2.3. *Let G be a topological graph, X be a surface, and $\eta : G \to X$, $\psi : G \to X$ be embeddings of G into X. We say η and ψ are equivalent embeddings if there is a homeomorphism $f : X \to X$ such that $\psi = f \circ \eta$. We say that two minimal separating sets $M_1, M_2 \subset X$ are equivalent minimal separating sets if they can be realized as equivalent embeddings of the same topological graph.*

From now on, we'll freely use the phrase *minimal separating set* in place of the more precise but tediously long *equivalence class of minimal separating sets*.

Since the underlying surface of an embedding is oriented, given an edge e of G, it makes sense to speak of the two sides of $\eta(e)$. Thinking of starting at a point in $\eta(e)$ and walking along $\eta(e)$ towards one of the two endpoints of $\eta(e)$, the two sides are the left and right of the walker.

We can now immediately state the following result about when a graph embedding in a surface is minimal separating:

Lemma 2.1. *An embedding $\eta(G)$ of a graph G in a surface X is minimal separating in X if and only if $X \setminus \eta(G)$ has two connected components, A and B, and for every edge e of G, $\eta(e)$ has A on one side and B on the other.*

Proof. If $\eta(G)$ is minimal separating, then certainly $X \setminus \eta(G)$ has at least two components, which we call $A, B, \ldots$. By minimality, for any edge e of G, $(X \setminus \eta(G)) \cup \eta(e)$ is not separated. Thus $\eta(e)$ is incident to all the components of $X \setminus \eta(G)$. Thus there are exactly two components A and B of $X \setminus \eta(G)$ with A lying to one side of $\eta(e)$ and B lying to the other.

Now assume $\eta(G)$ is an embedding with $X \setminus \eta(G)$ having two disjoint components A and B, and for all edges e of G, $\eta(e)$ is incident to A on one side and B on the other. $\eta(G)$ is clearly separating so we just need to address minimality. Since each edge e satisfies $\eta(e)$ is incident to A on one side and B on the other, for any point $x \in \eta(e)$, there is a path in A from any point in A to x, and a path in B from any point in B to x. Thus $X \setminus (\eta(G) \setminus x)$ is connected and $\eta(G)$ is minimal separating. □

3 Disconnected Graphs and Non-Cellular Embeddings

We wish to use powerful existing tools for working with cellular embeddings of connected graphs, but in general, the graph embeddings which give us minimal separating sets are neiher connected nor cellular. As an example, consider the left embedding from Figure 1.

Definition 3.1. *Given a graph G and an embedding $\eta : G \to X$.*
i) The ribbon graph $R(\eta)$ is a collared neighborhood of $\eta(G)$ that is small enough to retract to $\eta(G)$.
ii) X_η is the surface obtained by gluing discs onto the boundary components of $R(\eta)$.

We obtain the following results relating $\mathcal{L}_g$ with the sets $\mathcal{C}_h$ for $h \leq g$.

Theorem 3.1. *Let G be a topological graph which is the disjoint union of graphs G_1 and G_2. Then*

$$G \in \mathcal{L}_g \quad \Longleftrightarrow \quad G_i \in \mathcal{L}_{g_i} \text{ with } g = g_1 + g_2 + 1\,.$$

Proof. Let $G \in \mathcal{L}_g$ and η be a minimal separating embedding of G into a surface X with genus g. We'll let η_1 and η_2 be the restriction of η to G_1 and G_2 respectively. Consider the decomposition of X into $R(\eta) = R(\eta_1) \cup R(\eta_2)$ and the two connected components of $X \setminus R(\eta)$ as shown in Figure 2. As in the figure, we color one component white and the other black. The white punctured sphere has a hole for each white boundary component of $R(\eta)$ similar for the black punctured sphere. By cutting both the black cylinder along the red curve C_b, and cutting the white cylinder along the red curve C_w, and "plugging" the

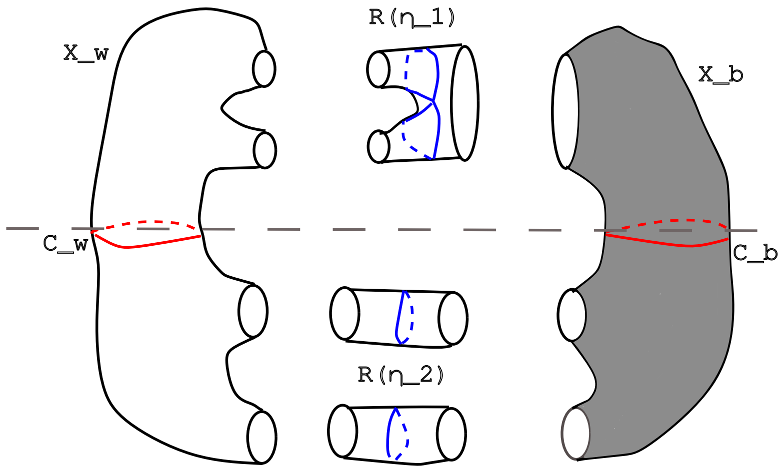

Figure 2: An example where G_1 is a single loop, G_2 is a pair of loops, and gluing the components together recovers a minimal separating embedding of $G = G_1 \cup G_2$ in a surface.

new holes by gluing on discs, when we glue the spheres back onto $R(\eta)$, we see that we now have two surfaces, X_1 and X_2 with η_i minimally separating in X_i. Define g_i and $\chi(X_i)$ to be the genus and Euler characteristic of X_i.

Now suppose we reassemble X from X_1 and X_2 by cutting out an open disc from the black components of X_1 and X_2 and identifying the boundaries (along C_b). Similarly for the white components (identifying along C_w). Then we have (taking into account that we took out two open "white" disks and two "black" ones.)

$$\begin{aligned}\chi(X) &= (\chi(X_1) - 2) + (\chi(X_2) - 2) - 0\\ 2 - 2g &= 2 - 2g_1 + 2 - 2g_2 - 4\,.\end{aligned}$$

from which the condition on the g_i follows. If, for example, g_1 is *not* the least separating genera of G_1, then it can be embedded in a surface of genus $h_1 < g_1$. That surface can be joined to X_2 to get a minimal separating embedding of G in a surface of genus

$$h_1 + g_2 + 1 < g$$

which is impossible by hypothesis. The same reasoning works for g_2. This proves "$\Longrightarrow$".

Now we prove "$\Longleftarrow$". Let both G_i are in $\mathcal{L}_{g_i}$. The same method of puncturing the surfaces and gluing them together produces a minimal separating embedding of G of genus $g = g_1 + g_2 + 1$. Assume g there is an embedding of G in a surface of genus $g' < g$. Then cutting up the surface again — exactly as before — we will obtain minimal separating embeddings in two surfaces X_i' with genera g_i' such that at least one of these satisfies $g_i' < g_i$, contradicting the assumption of the g_i. $\square$

Corollary 3.1.1. *Let $G \in \mathcal{L}_g$ and let $G_1, \ldots, G_k$ be the connected components of G. Then there exist $g_1, \ldots, g_k$ such that for each $i \in 1, 2, \ldots, k$ we have $G_i \in \mathcal{C}_{g_i}$ and $g = (k-1) + \sum_{i=1}^{k} g_i$.*

Proof. The proof follows from Theorem 3.1 by induction on the number of connected components of G. □

With this result, we now know that the contents of the sets $\mathcal{C}_h$ for all $h \leq g$ determine the contents of $\mathcal{L}_g$ and in Section 5 we'll use this to compute $|\mathcal{L}_g|$ in terms of the $|\mathcal{C}_h|$ for $h \leq g$. So, from here on we restrict our investigation to the sets $\mathcal{C}_g$, the connected graphs.

4 Topological and Combinatorial Maps

Given a topological map or equivalently a cellular embedding, there is discrete data called a combinatorial map which allows us to recover G, X, and η up to equivalence of embeddings. These objects are equivalently known as rotation systems, fat graphs, ribbon graphs[3], or dessins d'enfants ([4], [12],[10]). Precise definitions vary a bit from source to source based on the types of embeddings of interest to the author (particularly if the author is interested in non-orientable surfaces). We use the following definition.

Definition 4.1. *[10] A* combinatorial map *is an ordered triple*[4] $(\sigma, \alpha, \varphi)$[5] *of permutations in S_{2n} (the symmetric group on $2n$ elements), such that*

i) α *is an involution with no fixed points.*

ii) *The permutation group* $\langle \sigma, \alpha \rangle = \langle \sigma, \alpha, \varphi \rangle$ *acts transitively on the set* $\{1, 2, \ldots, 2n\}$.

iii) $\sigma\alpha\varphi = 1$.

The second condition is equivalent to the underlying graph being connected.

We give an example of the correspondence between combinatorial maps and connected topological maps (see Definition 2.2). For a proof of this, see [10]. Given a connected graph G and a cellular embedding of G into an orientable surface X, the following four-step reasoning associates a combinatorial map to our embedding of G (see Figure 3).

1. Assign labels to the edge ends and mark them on the surface X as described above.

2. Define α as the involution that swaps the two ends of each edge. The number $c(\alpha)$ of cycles of α gives the number of edges.

[3]However, we use the term *ribbon graph* exclusively as in Definition 3.1.

[4]Or, equivalently, an ordered pair (σ, α).

[5]The letters σ, α, φ come from the French *sommet* for vertex, *arc* for edge and *face* for face[5].

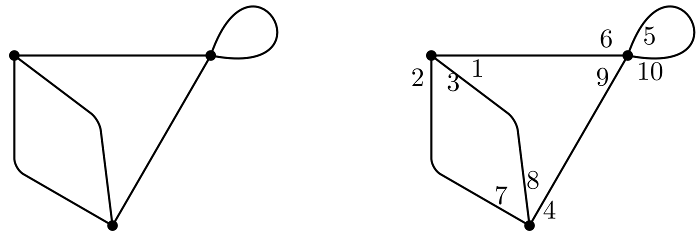

Figure 3: An embedding of a graph in the plane is shown on the left, and the same embedding equipped with a labelling of the edge ends is shown on the right. Each label is immediately clockwise of its corresponding edge end. The corresponding combinatorial map is $\sigma = (1,2,3)(4,8,7)(5,6,9,10), \alpha = (1,6)(2,7)(3,8)(4,9)(5,10), \varphi = (1,8,9)(2,6,10,4)(3,7)(5)$.

3. Define σ as the permutation that assigns the next (counterclockwise) edge at each vertex. The number $c(\sigma)$ of cycles of σ gives the number of vertices.

4. Then $\varphi = (\sigma\alpha)^{-1}$ turns out to be (see below) the 'boundary walk'. That is: the cycles of φ are precisely the edges that form the boundary of each component (or cell) of the complement of the graph. The number $c(\varphi)$ of cycles of φ gives the number of faces.

The procedure above certainly gives a way to associate a triple of permutations in S_{2n} to any connected graph with n edges. Since G is connected it's possible to walk from any edge end to any other one by travelling to the opposite end of the edge (applying α), then moving to another edge end incident to the current vertex (applying σ) and repeating. This shows that $\langle \sigma, \alpha \rangle$ acts transitively on $\{1, \ldots, 2n\}$. It remains to verify that φ in Definition 4.1 is really a boundary walk. We illustrate this with a simple picture looking at two vertices with an edge between them, see Figure 4.

Notice that given any connected graph G and an embedding $\eta : G \to X$ (cellular or not) we can construct a combinatorial map with the same ribbon graph as η by steps 1-3 of the procedure described above, but if the embedding isn't cellular then point 4 will not be true. Thus for any embedding of a connected graph we can talk about *the combinatorial map associated to the embedding.*

Lemma 4.1. *Given a combinatorial map $\eta := (\sigma, \alpha, \varphi)$ of a cellular embedding η, we have that the genus g_η of the cellular surface is given by*

$$2 - 2g_\eta = V - E + F = c(\sigma) - c(\alpha) + c(\varphi).$$

Proof. This follows directly from the four points listed above. □

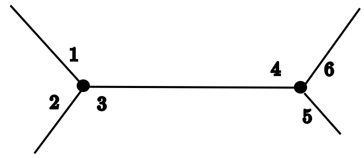

Figure 4: With counterclockwise as the positive orientation we have $\alpha(4) = 3$ and $\sigma(3) = 1$, so $(\sigma\alpha)(4) = 1$. Then $(\sigma\alpha)^{-1}(1) = 4$, and the edge with label 4 is the edge following the edge with label 1 in counterclockwise order about their shared face.

To deal with the matter of a choice of labelling, there is the notion of *isomorphism* of combinatorial maps, again following [10]. It captures the idea that two combinatorial maps arising from different choices of labelling on the same topological map are really *the same object*. In terms of the S_{2n} action on $\{1, \ldots, 2n\}$, relabeling can be thought of as conjugation by the relabelling permutation ρ, and we arrive at the following definition:

Definition 4.2. *We say $\rho \in S_{2n}$ is an isomorphism between combinatorial maps $(\sigma_1, \alpha_1, \varphi_1)$ and $(\sigma_2, \alpha_2, \varphi_2)$ if*

$$\rho^{-1}\sigma_1\rho = \sigma_2, \qquad \rho^{-1}\alpha_1\rho = \alpha_2.$$

And this, of course, implies that $\rho^{-1}\varphi_1\rho = \varphi_2$.

There is one more attribute of combinatorial and topological maps which we must introduce, the dual. This is generalization of the dual of a graph embedded in the plane (or sphere) to a graph cellularly embedded in an arbitrary surface.

Definition 4.3. *The geometric dual of a topological map M is the map M^* constructed as follows: We place a single vertex of M^* inside each face of M. Then, for every edge e in M, we place an edge in M^* running between the vertices of M^* corresponding to the faces of M incident to e, with new new edge passing through the midpoint of e.*

Corollary 4.1.1. *If $(\sigma, \alpha, \varphi)$ is combinatorial map with M the corresponding topological map, then $(\varphi, \alpha, \sigma)$ is the combinatorial map corresponding to M^* [10].*

Proof. Definition 4.3 immediately implies that the faces of the dual of M correspond to the vertices of M and so forth. □

So far, all of the sets of minimal separating sets we've discussed are very difficult to enumerate directly. We now introduce a set that will be easier to directly enumerate, and which we will use to determine the contents of $\mathcal{C}_g, \mathcal{L}_g$, and $\mathcal{M}_g$.

Definition 4.4. *We define the set*[6] $\mathcal{E}_g$ *to be the set of all isomorphism classes of combinatorial maps (Definition 4.1)* $(\sigma, \alpha, \varphi)$ *such that* $(\sigma, \alpha, \varphi)$ *is the combinatorial map of a minimal separating set in a surface of genus* g*, but not in any surface of genus less than* g*.*

We'll abuse language in the future and refer to a combinatorial map $(\sigma, \alpha, \varphi)$ as an element of $\mathcal{E}_g$ when we really mean the isomorphism class of combinatorial maps with representative $(\sigma, \alpha, \varphi)$. We now aim to enumerate the elements of the sets $\mathcal{E}_g$.

Lemma 4.2. *Let* G *be a connected graph and* $\eta : G \to X$ *be an embedding of* G *in a surface* X *such that* $X \setminus \eta(G)$ *has two connected components. Let* X_w, X_b, X_η *be the surfaces constructed as described in the previous section. Let* g, g_w, g_b, g_η *be the genera of* $X, X_w, X_b,$ *and* X_η *respectively (with the holes plugged). Then*

$$g = g_\eta + g_w + g_b + c(\varphi) - 2 = \frac{c(\varphi) + c(\alpha) - c(\sigma)}{2} + g_w + g_b$$

where $c(\varphi)$ *is the number of (white plus black) faces of* X_η *and* g_η *is given by Lemma 4.1*

Proof. We obtain X from X_η by gluing X_w and X_b along the boundary of the ribbon graph $R(\eta)$ corresponding to the embedding η. $R(\eta)$ is a surface of genus g_η from which F_η faces have been removed. By Definition 4.1, $F_\eta = c(\varphi)$. So its Euler characteristic equals $2 - 2g_\eta - c(\varphi)$. Similarly, X_w and X_b had n_w, resp. n_b disks removed and so they have Euler characteristic $2 - 2g_w - n_w$ and $2 - 2g_b - n_b$, where, of course, $n_w + n_b$ equals $c(\varphi)$, the number of holes in $R(\eta)$. Since the Euler characteristic is additive, we obtain

$$2 - 2g = 2 - 2g_\eta - c(\varphi) + 2 - 2g_w - n_w + 2 - 2g_b - n_b = 2 - 2(g_\eta + g_w + g_b + c(\varphi) - 2)\,,$$

which implies the lemma. □

Clearly, the least separating genus is obtained from Lemma 4.2 by setting $g_w = g_b = 0$. This yields the following corollary.

Corollary 4.2.1. *For* $g > 0$*, a combinatorial map* $(\sigma, \alpha, \varphi)$ *is in* $\mathcal{E}_g$ *if and only if it has no degree 2 vertices, the dual map* $(\varphi, \alpha, \sigma)$ *is a combinatorial map of a bipartite graph, and*

$$g = \frac{c(\varphi) + c(\alpha) - c(\sigma)}{2} - 1\,.$$

Proof. The equation for g is obtained from Lemma 4.2 by setting $g_w = g_b = 0$. The no degree 2 vertices condition is a consequence of our choice in section 2 to identify connected minimal separating sets with a representative without degree 2 vertices (for $g > 0$). The condition that $(\varphi, \alpha, \sigma)$ is bipartite is equivalent to Lemma 2.1. □

[6] $\mathcal{E}$ for embedding

Lemma 4.3. *If $(\sigma, \alpha, \varphi) \in \mathcal{E}_g$, and the underlying graph has E edges, then $1 + g \leq E \leq 4g$ for all $g > 0$.*

Proof. Recall (from section 2) that for $g > 0$, we can assume there are no degree 2 vertices in a minimal separating set.

Additionally, since each edge is incident to both components of the separated surface the degree of each vertex must be even. Thus every vertex has degree at least 4 and we get $E \geq 2V$ or $V \leq \frac{E}{2}$. By separating there must at least be two faces, $F \geq 2$. So using the language of Definition 4.1 and Corollary 4.2.1 we have

$$g = \frac{E + c(\varphi) - V}{2} - 1 \geq \frac{E + 2 - \frac{E}{2}}{2} - 1$$

which gives the second inequality of the lemma.

From Corollary 4.2.1 we obtain that

$$c(\alpha) = 2g - c(\eta) + c(\sigma) + 2\,.$$

By Lemma 4.1, $g = g_\varphi + c(\eta) - 2$. Substitute this in the above expression to get

$$c(\alpha) = g + (g_\varphi + c(\eta) - 2) - c(\eta) + c(\sigma) + 2 = g + g_\varphi + c(\sigma)\,.$$

Since $g_\varphi \geq 0$ and $c(\sigma) \geq 1$, the first inequality follows. □

As we'll see in the next section, we can construct an algorithm to enumerate the elements of $\mathcal{E}_g$ with a computer, but we can make a much faster algorithm if we introduce the following generalization of combinatorial maps:

Definition 4.5. *A hypermap is a triple of permutations (ψ, ρ, θ)[7] in S_n such that*
i) $\langle \psi, \rho, \theta \rangle$ acts transitively on $\{1, \ldots, n\}$
ii) $\psi\rho\theta = 1$.

Note that this definition is effectively a relaxation of Definition 4.1, since it no longer requires that one of the permutations be a fixed-point-free involution. Dropping that requirement means we need no longer require the set of labels has an even number of elements. Hence the other change: considering permutations in S_n for any n, rather than only even n.

Given a 2-colored bipartite combinatorial map $(\sigma, \alpha, \varphi)$, we construct the corresponding hypermap by considering the colored labelled topological map corresponding to $(\sigma, \alpha, \varphi)$. We remove the label from one end of each edge so that now we have one label per edge (see

[7]Traditionally, the letters $(\sigma, \alpha, \varphi)$ are used for hypermaps, just as for combinatorial maps, but since we'll be discussing the two together, we use (ψ, ρ, θ) for hypermaps and $(\sigma, \alpha, \varphi)$ for combinatorial maps.

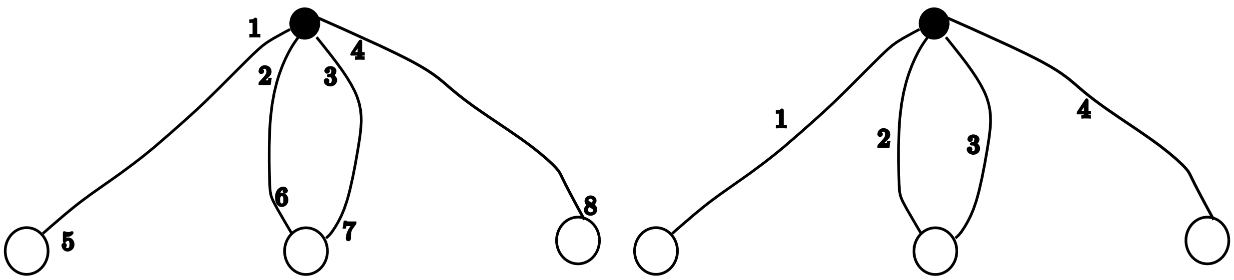

Figure 5: On the left is a drawing of the combinatorial map $\sigma = (1,2,3,4)(6,7), \alpha = (1,5)(2,6)(3,7)(4,8), \varphi = (1,8,4,7,2,5)(3,6)$, and on the right is the corresponding hypermap $\psi = (1,2,3,4), \rho = (2,3), \theta = (1,4,2)$.

Figure 5). If σ, α, φ were elements of S_{2n}, we now have n labels remaining. If needed, change the labels to be the numbers $\{1, 2, \ldots, n\}$. Now the cycles of ψ correspond to the black vertices of the map, and the cycle corresponding to a vertex is the list of incident edges in order as we travel around the vertex. The cycles of ρ are defined similarly, but for the white vertices instead of the black vertices. This defines θ, since $\theta = (\psi\rho)^{-1}$, but the cycles of θ correspond to the faces of the map, but each cycle has half the length of the corresponding cycle in φ. The genus of (ψ, ρ, θ) is equal to the genus of $(\sigma, \alpha, \varphi)$ and also satisfies the formula:

$$2 - 2g = c(\psi) + c(\rho) - n + c(\theta)$$

where $\psi, \rho, \theta \in S_n$. For proofs of these properties and more details, see [10] and [20]. It's a result of [20] that there is a one to one correspondence between isomorphism classes of 2-colored combinatorial maps of bipartite graphs on n edges and isomorphism classes of hypermaps $(\psi, \rho, \theta) \in S_n^3$.

Isomorphisms of hypermaps are defined analogously to isomorphisms for combinatorial maps:

Definition 4.6. *An isomorphism between two hypermaps $(\psi_1, \rho_1, \theta_1)$ and $(\psi_2, \rho_2, \theta_2)$ in S_n, is a permuation $\gamma \in S_n$ with*

$$\gamma^{-1}\psi_1\gamma = \psi_2, \quad \gamma^{-1}\rho_1\gamma = \rho_2, \quad \gamma^{-1}\theta_1\gamma = \theta_2.$$

Note that as an immediate consequence, any automorphism γ of a hypermap (ψ, ρ, θ) must commute with each of ψ, ρ, and θ.

Translating Corollary 4.2.1 into terms of hypermaps we have

Theorem 4.4. *A combinatorial map $(\sigma, \alpha, \varphi)$ is an element of $\mathcal{E}_g$ if and only if its dual map $(\varphi, \alpha, \sigma)$ corresponds to a hypermap (ψ, ρ, θ) satisfying*

$$g = \frac{c(\psi) + c(\rho) + n - c(\theta)}{2} - 1$$

where n is the number of edges of the combinatorial map and θ has no fixed points.

Proof. From the construction of the hypermap corresponding to a colored bipartite map, we have $c(\varphi) = c(\psi) + c(\rho)$, $c(\alpha) = n$, and $c(\sigma) = c(\theta)$. Similarly, θ has a fixed point (a 1-cycle) if and only if σ has a 2-cycle (a degree 2 vertex). Then the proof follows immediately from Corollary 4.2.1. $\square$

5 Computation and Results

Here we introduce an algorithm to enumerate the elements of $\mathcal{E}_g$. Then we use the results of the previous section to determine the sizes of $|\mathcal{C}_g|, |\mathcal{L}_g|$, and $|\mathcal{M}_g|$.

At this point we've bounded the number of edges by $g + 1 \leq E \leq 4g$ so we could search all combinatorial maps in S_{2E} to try and find all of the minimal separating ones.

Here's a possible approach. For each E with $g + 1 \leq E \leq 4g$, we construct all fixed-point-free involutions in S_{2E} as candidate αs, and then for every $\sigma \in S_{2E}$ we check if $(\sigma, \alpha, (\sigma\alpha)^{-1})$ is in $\mathcal{C}_g$ using Corollary 4.2.1. We can fix a choice of α (fixing the choice of labels attached to each edge) and can use the fact that all cycles of σ have even length of at least 4 to reduce our search space.

From Definition 4.4, we can determine the elements of $\mathcal{C}_g$ if we know the contents of $\mathcal{E}_h$ for all $h \leq g$. Each element of $\mathcal{E}_h$ has an underlying graph G. So, for each element of $\mathcal{E}_g$ we check the sets $\mathcal{E}_h$ for $h < g$, and see if any contain another combinatorial map with the same underlying graph. If not, then g is the least separating genus of G and $G \in \mathcal{C}_g$.

Just for genus 3, this would involve an enormous search, (S_{24} has $24! \sim 6.2 \times 10^{24}$ elements and $\frac{24!}{6!} \sim 8.6 \times 10^{20}$ candidate σs once we fix α). But we would like to eliminate as much computational work as possible, so we take advantage of the correspondence between 2-colored bipartite combinatorial maps and hypermaps due to Walsh[20]. Translating to the language of hypermaps allows us to effectively trade the "all cycles of σ have even length" condition for smaller symmetric groups to search, saving a great deal of computing time.

Here is the algorithm we use: For each E with $g + 1 \leq E \leq 4g$, we construct list of triples (P, R, T) of possible cycle types for a hypermap (ψ, ρ, θ) satisfying the conditions of Theorem 4.4. In the correspondence between 2-colored maps and hypermaps, ψ and ρ correspond to the two colors of vertex, and the actual choice of coloring doesn't matter to us, so we always pick P and R so that ρ has no fewer cycles than ψ. In all the most computationally demanding cases this results in there being more permutations with cycle type P than with either R or T. Since our assignment of labels is free, so for each triple we fix a permutation of cycle type P to be ψ. Since $\psi\rho\theta = 1$, the choice of any two elements of (ψ, ρ, θ) determines the third uniquley. Then from R and T, we select whichever cycle type corresponds to the smallest number of permutations. Again, without loss of generality, assume that it is R. Rather than look at every permutation with cycle type R, we note that we're only interested in isomorphism classes, and whenever a permutation γ commutes with ψ, we can see from Definition 4.6 that (ψ, ρ, θ) is isomorphic to $(\psi, \gamma^{-1}\rho\gamma, \gamma^{-1}\theta\gamma)$ for any ρ, θ. So, we find a representative from each orbit of the permutations with cycle type R under conjugation by the centralizer of ψ, and check each representative ρ to see of $(\psi, \rho, (\psi\rho)^{-1})$ forms a hypermap satisfying the conditions of Theorem 4.4. This generates a complete list of isomorphism classes of hypermaps corresponding to the elements of $\mathcal{E}_g$. Unfortunately, for the purposes of finding $|\mathcal{E}_g|$, we've overcounted, since Walsh's correspondence is between *2-colored* bipartite maps and hypermaps, and the choice of how to assign the colors means that most elements of $\mathcal{E}_g$ correspond to two distinct hypermaps. Since we always choose cycle-type R to have at least as many cycles as P, we only run the risk of double-counting when P and R have the same number of cycles. In that case, the only time an element of $\mathcal{E}_g$ doesn't correspond to two isomorphism classes of hypermaps in our count is when switching the coloring gives an isomorphic hypermap. I.e., (ψ, ρ, θ) is isomorphic to (ρ, ψ, θ). So for each hypermap (ψ, ρ, θ) that our search returns, if it has the potential to be double counted, we check if (ρ, ψ, θ). If it isn't, then it's been double counted and we subtract $1/2$ from out count ($1/2$ from each time it's counted to cancel the out the double count).

From Definition 4.4, we can similarly determine the elements of $\mathcal{C}_g$ if we know the elements of $\mathcal{E}_h$ for all $h \leq g$. Once $\mathcal{E}_h$ has been computed for all $h \leq g$, we can inductively find the elements of $\mathcal{C}_g$ by taking the underlying graphs of each map in $\mathcal{E}_g$ and graph isomorphism testing them against the elements of $\mathcal{C}_h$ for $h \leq g$. A computer program to compute the values of $\mathcal{E}_g$ and $\mathcal{C}_g$ using this method can be found at this GitHb link.

Then, by Corollary 3.1.1, we know that the elements of $\mathcal{L}_g$ are unions of elements in the sets $\mathcal{C}_h$ (for $h \leq g$) satisfying the genus sum condition from the corollary. The following lemma gives an explicit formula for $|\mathcal{L}_g|$ when $g \leq 4$.

Lemma 5.1. *The values of* $|\mathcal{L}_1|, |\mathcal{L}_2|, |\mathcal{L}_3|$, *and* $|\mathcal{L}_4|$ *are given by the following:*

$$\begin{aligned}
|\mathcal{L}_1| &= |\mathcal{C}_1| + \binom{|\mathcal{C}_0|+2-1}{2} \\
|\mathcal{L}_2| &= |\mathcal{C}_2| + |\mathcal{C}_1|\cdot|\mathcal{C}_0| + \binom{|\mathcal{C}_0|+3-1}{3} \\
|\mathcal{L}_3| &= |\mathcal{C}_3| + |\mathcal{C}_2|\cdot|\mathcal{C}_0| + \binom{|\mathcal{C}_1|+2-1}{2} + |\mathcal{C}_1|\binom{|\mathcal{C}_0|}{2} + \binom{|\mathcal{C}_0|+4-1}{4} \\
|\mathcal{L}_4| &= |\mathcal{C}_4| + |\mathcal{C}_3|\cdot|\mathcal{C}_0| + |\mathcal{C}_2|\cdot|\mathcal{C}_1| + |\mathcal{C}_2|\cdot\binom{|\mathcal{C}_0|+2-1}{2} + \binom{|\mathcal{C}_1|+2-1}{2}\cdot|\mathcal{C}_0| \\
&\quad + |\mathcal{C}_1|\cdot\binom{|\mathcal{C}_0|+3-1}{3} + \binom{|\mathcal{C}_0|+5-1}{5}
\end{aligned}$$

Proof. An immediate consequence of Corollary 3.1.1 is that any graph in $|\mathcal{L}_g|$ has at most $g+1$ connected components. Then every graph in $\mathcal{L}_1$ is either connected or has 2 connected components. The connected elements of $\mathcal{L}_1$ are by definition $\mathcal{C}_1$. If $G \in \mathcal{L}_1$ has 2 components, then by Corollary 3.1.1 they must both be elements of $\mathcal{C}_0$. The number of graphs with two components, each in $\mathcal{C}_0$ is equal to the number of ways to choose, with replacement, 2 elements of $\mathcal{C}_0$. By [18], the number of ways to choose k elements with replacement from a set of size n is $\binom{n+k-1}{k}$, so we have the formula

$$|\mathcal{L}_1| = |\mathcal{C}_1| + \binom{|\mathcal{C}_0|+2-1}{2}$$

For $\mathcal{L}_2$, we need to consider graphs with up to 3 components. Again, the subset of $\mathcal{L}_2$ with one component form $\mathcal{C}_2$. By Corollary 3.1.1, if $G \in \mathcal{L}_2$ has 2 connected components, then they are elements of some $\mathcal{C}_i, \mathcal{C}_j$ where $i+j+1=2$. This means $\{i,j\}=\{0,1\}$, so G is the union of a graph in $\mathcal{C}_1$ and a graph in $|\mathcal{C}_0|$. There are exactly $|\mathcal{C}_1|\cdot|\mathcal{C}_0|$ ways to pick the components of G.

Finally, we consider graphs with 3 components. By the corollary, such a graph must have all 3 components in $\mathcal{C}_0$, so there are $\binom{|\mathcal{C}_0|+3-1}{3}$ such graphs. This gives

$$|\mathcal{L}_2| = |\mathcal{C}_2| + |\mathcal{C}_1|\cdot|\mathcal{C}_0| + \binom{|\mathcal{C}_0|+3-1}{3}.$$

For $\mathcal{L}_3$, we can use the same reasoning but need to consider graphs with up to 4 components. The graphs with one component form $\mathcal{C}_3$. If $G \in \mathcal{L}_3$ has 2 components, they're elements of $\mathcal{C}_i$ and $\mathcal{C}_j$ with $i+j+1=3$, so $i+j=2$. This means either G has 1 component in $\mathcal{C}_2$ and 1 in $\mathcal{C}_0$, or both components from $\mathcal{C}_1$. The are $|\mathcal{C}_2|\cdot|\mathcal{C}_0| + \binom{|\mathcal{C}_1|+2-1}{2}$ ways to construct such

a graph.
Next we consider graphs $G \in \mathcal{L}_3$ with 3 connected components. By Corollary 3.1.1 these are elements of $\mathcal{C}_i, \mathcal{C}_j, \mathcal{C}_k$ with $i+j+k+2=3$, so $i+j+k=1$. This means one component comes from $\mathcal{C}_1$, and the other two come from $\mathcal{C}_0$. There are $|\mathcal{C}_1| \cdot \binom{|\mathcal{C}_0|+2-1}{2}$ such graphs.
Finally, we consider graphs in $\mathcal{L}_4$ with 4 connected components. By the corollary, all these components are in $\mathcal{C}_0$, so there are $\binom{|\mathcal{C}_0|+4-1}{4}$ such graphs. We now have

$$|\mathcal{L}_3| = |\mathcal{C}_3| + |\mathcal{C}_2| \cdot |\mathcal{C}_0| + \binom{|\mathcal{C}_1|+2-1}{2} + |\mathcal{C}_1| \cdot \binom{|\mathcal{C}_0|+4-1}{4}.$$

Only $\mathcal{L}_4$ remains. The connected elements form $\mathcal{C}_4$. By the corollary, any $G \in \mathcal{L}_4$ with two components has one component each from $\mathcal{C}_i, \mathcal{C}_j$ where $i+j+1=4$. This means either $\{i,j\} = \{0,3\}$ or $\{i,j\} = \{1,2\}$, so there are $|\mathcal{C}_3| \cdot |\mathcal{C}_0| + |\mathcal{C}_2| \cdot |\mathcal{C}_1|$ such graphs.
By the same type of reasoning the graphs with 3 connected components either have one component from $\mathcal{C}_2$ and two components from $\mathcal{C}_0$ or two components from $\mathcal{C}_1$ and one component from $\mathcal{C}_0$. This means there are $|\mathcal{C}_2| \cdot \binom{|\mathcal{C}_0|+2-1}{2} + \binom{|\mathcal{C}_1|+2-1}{2} \cdot |\mathcal{C}_0|$ such graphs.
For graphs in $\mathcal{L}_4$ with 4 components, the only way to satisfy Corollary 3.1.1 is by having one component from $\mathcal{C}_1$ and 3 from $\mathcal{C}_0$. So there are $|\mathcal{C}_1| \cdot \binom{|\mathcal{C}_0|+3-1}{3}$ of these.
Finally, if a graph in $\mathcal{L}_4$ has 5 connected components, then all 5 are from $\mathcal{C}_0$, so there are $\binom{|\mathcal{C}_0|+5-1}{5}$ of these. Together this gives

$$\begin{aligned} |\mathcal{L}_4| = |\mathcal{C}_4| + |\mathcal{C}_3| \cdot |\mathcal{C}_0| + |\mathcal{C}_2| \cdot |\mathcal{C}_1| + |\mathcal{C}_2| \binom{|\mathcal{C}_0|+2-1}{2} + \binom{|\mathcal{C}_1|+2-1}{2} \cdot |\mathcal{C}_0| \\ + |\mathcal{C}_1| \cdot \binom{|\mathcal{C}_0+3-1}{3} + \binom{|\mathcal{C}_0|+5-1}{5}. \end{aligned}$$

□

From Definition 2.1, knowing the elements of $\mathcal{L}_h$ for all $h \leq g$ is enough to know the elements of $\mathcal{M}_g$, and $|\mathcal{M}_g| = |\mathcal{M}_{g-1}| + |\mathcal{L}_g|$.

Here are the results of running this algorithm for $g \leq 4$:

g	$\lvert\mathcal{E}_g\rvert$	$\lvert\mathcal{C}_g\rvert$	$\lvert\mathcal{L}_g\rvert$	$\lvert\mathcal{M}_g\rvert$
0	1	1	1	1
1	3	3	4	5
2	31	17	21	26
3	1831	164	191	217
4	462645	3096	3338	3555

6 Concluding Remarks

A search of the OEIS does not find any sequences matching any of the sequences $\mathcal{E}_g$, $\mathcal{C}_g, \mathcal{L}_g$, or $\mathcal{M}_g$. It thus appears unlikely that any currently known formula will yield correct values for the cardinalities of these sets.

At this point, we do not see any way forward towards a method to directly determine the values of $|\mathcal{C}_g|, |\mathcal{L}_g|$, or $|\mathcal{M}_g|$ without the computationally intensive process of directly them. It would be interesting to attempt to find a recursive formula for the values of $|\mathcal{E}_g|$ using some of the techniques for enumerating isomorphism classes of similar flavors of nonplanar hypermaps and combinatorial maps developed in [14] and [15]. The general method used there is to use the fact that it's easier to counted rooted maps and hypermaps (maps/hypermaps with a distinguished edge) directly than to count unrooted ones, and the authors develop formulae to convert counts of rooted maps/hypermaps into counts of unrooted maps and hypermaps. This is done by considering rooted objects which can be realized as quotients under an automorphism of the unrooted objects in question. For formal definitions and more details on quotients of maps see [6] or [11] or the unpublished [16].

There are, however, two important distinctions between the sets considered in those papers and the $\mathcal{E}_g$ we consider. The first is that the families of maps and hypermaps considered in those papers are closed under quotients by map automorphisms. An example of the set of maps hypermaps considered in [15] may help to clarify this. The authors enumerate the number of hypermaps of genus g with n edges. If M is a hypermap of with genus g with n edges, and f is an automorphism of M, then the quotient M/f is a hypermap with genus at most g and at most n edges (strictly fewer unless f is the identity). So, if we proceed inductively with respect to genus and number of edges, we will have already considered these maps. As observed in [11] this property is extremely useful for such enumeration problems. Without this property one needs to determine precisely which rooted objects arise as quotients of the unrooted ones being counted and then enumerate those. The difficulty with using those techniques to enumerates elements of $\mathcal{E}_g$ arises here. The quotient of a map in $\mathcal{E}_g$ is not generally another map in $\mathcal{E}_g$, or even a map which be embedded as a minimal separating set in any surface. As an example, consider the map and its quotient shown in Figure 6 (for details on quotients of maps see [11] or [16]).

The map on the left is an element of $\mathcal{E}_2$, but the quotient map cannot be minimal separating in any genus. Additionally, such a method would require a formula enumerating the numer of rooted versions of the relevant hypermaps (which is as yet unknown). To the best of our knowledge, in all existing work enumerating unrooted maps and hypermaps using the tools developed in [14], [15], and [11] the familes of objects being considered tend to either count all maps/hypermaps with a given genus and edge count or to heavily restrict to counting only regular maps/hypermaps ([7],[8],[9]). The case of all maps/hypermaps puts

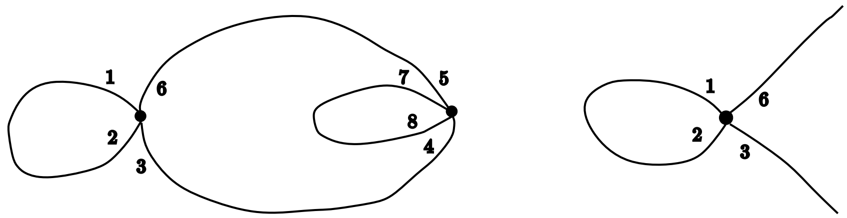

Figure 6: The combinatorial map $\sigma = (1, 2, 3, 6)(4, 5, 7, 8), \alpha = (1, 2)(3, 4)(5, 6)(7, 8)$ (left) and its quotient under the automorphism $\rho = (1, 7)(2, 8)(3, 4)(5, 6)$. The unlabelled edge ends that occur are called *singular edges* in [11] and the desire to admit singular edges in combinatorial maps is why the definition of map used in [14] only requires that α is an involution, rather than a fixed-point free involution as in the definition we use

no restrictions on cycle types of the permutations, while a requiring a map to be k-regular for some k fully determines the cycle types of σ and α. In our case, there are some restrictions on cycle types, but nothing nearly as strong as regularity. It remains to be seen whether these obstruction can be worked around and the techniques developed in [14] and [15] can the be applied to enumerate $\mathcal{E}_g$ but suggests they may form an interesting class of maps for further study of these methods.

References

[1] James Bernhard and J.J.P. Veerman. The topology of surface mediatrices. *Topology and its Applications*, 154(1):54–68, January 2007.

[2] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.

[3] Logan S Fox and J J P Veerman. Equidistant sets on alexandrov surfaces. *arXiv Preprint*, May 2022. arXiv: 2205.09155.

[4] Ruben A. Hidalgo. Bipartite graphs and their dessins d'enfants. *arXiv:1611.02901 [math]*, August 2017. arXiv: 1611.02901.

[5] Alain Jacques. *Constellations et propriétés algébriques des graphes topologiques*. Theses, Faculté des Sciences, Université de Paris, January 1969.

[6] Gareth A. Jones and David Singerman. Theory of Maps on Orientable Surfaces. *Proceedings of the London Mathematical Society*, s3-37(2):273–307, September 1978.

[7] Evgeniy Krasko, Igor Labutin, and Alexander Omelchenko. Enumeration of 3-regular one-face maps on orientable or non-orientable surface up to all symmetries. *ArXiv Preprint*, 2019.

[8] Evgeniy Krasko and Alexander Omelchenko. Enumeration of r-regular maps on the torus. part i: Rooted maps on the torus, the projective plane and the klein bottle. sensed maps on the torus. *Discrete Mathematics*, 342(2):584–599, 2019.

[9] Evgeniy Krasko and Alexander Omelchenko. Enumeration of r-regular maps on the torus. part ii: Unsensed maps. *Discrete Mathematics*, 342(2):600–614, 2019.

[10] S. K. Lando and A. K. Zvonkin. *Graphs on surfaces and their applications*. Number v. 141. 2 in Encyclopaedia of mathematical sciences, Low-dimensional topology. Springer, Berlin ; New York, 2004.

[11] V. A. Liskovets. Enumeration of nonisomorphic planar maps: Nonisomorphic Planar Graphs. *Journal of Graph Theory*, 5(1):115–117, March 1981.

[12] Yanpei Liu. *Enumerative theory of maps*. Science Press, Beijing, 1999.

[13] Mario Ponce and Patricio Santibáñez. On Equidistant Sets and Generalized Conics: The Old and the New. *The American Mathematical Monthly*, 121(1):18, 2014.

[14] Alexander Mednykh and Roman Nedela. Enumeration of unrooted maps of a given genus. *Journal of Combinatorial Theory, Series B*, 96(5):706–729, September 2006.

[15] Alexander Mednykh and Roman Nedela. Enumeration of unrooted hypermaps of a given genus. *Discrete Mathematics*, 310(3):518–526, February 2010.

[16] Roman Nedela. Maps, Hypermaps and Related Topics. *Unpublished*, page 92.

[17] Oscar – open source computer algebra research system, version 0.12.2-dev, 2023.

[18] Richard P. Stanley. *Enumerative combinatorics*. The Wadsworth & Brooks/Cole mathematics series. Wadsworth & Brooks/Cole Advanced Books & Software, Monterey, Calif, 1986.

[19] J.J.P. Veerman and James Bernhard. Minimally separating sets, mediatrices, and Brillouin spaces. *Topology and its Applications*, 153(9):1421–1433, March 2006.

[20] T.R.S Walsh. Hypermaps versus bipartite maps. *Journal of Combinatorial Theory, Series B*, 18(2):155–163, April 1975.