

Rigidity of the Strongly Separating Curve Graph

BRIAN H. BOWDITCH

ABSTRACT. We define the strongly separating curve graph to be the full subgraph of the curve graph of a compact orientable surface, where the vertex set consists of all separating curves that do not bound a three-holed sphere. We show that, for all but finitely many surfaces, any automorphism of the strongly separating curve graph is induced by an element of the mapping class group.

1. Introduction

The main aim of this paper is to prove a rigidity result (Theorem 1.1) for certain curve graphs associated to compact orientable surfaces. It is a variation on some well-known results in this direction. Our main motivation for this particular statement is its application to the quasi-isometric rigidity of the Weil–Petersson metric.

Let Σ be a compact orientable surface. We write $g(\Sigma)$ for its genus, and $p(\Sigma)$ for the number of boundary components. The *complexity* $\xi(\Sigma)$ of Σ is defined by $\xi(\Sigma) = 3g(\Sigma) + p(\Sigma) - 3$. (It equals the number of disjoint simple closed curves needed to cut Σ into a collection of 3-holed spheres.)

Let $\mathcal{G}(\Sigma)$ be the curve graph associated to Σ , that is, the 1-skeleton of the curve complex as defined in [H]. It has vertex set $C(\Sigma)$, the set of nontrivial nonperipheral simple closed curves in Σ , defined up to homotopy. Two elements of $C(\Sigma)$ are deemed adjacent if they can be homotoped to be disjoint. Note that the mapping class group $\text{Map}(\Sigma)$ acts cofinitely on $\mathcal{G}(\Sigma)$. The rigidity theorems of [Iv; Ko; L] tell us (in particular) that if $\xi(\Sigma) \geq 2$, then any automorphism of $\mathcal{G}(\Sigma)$ is induced by an element of $\text{Map}(\Sigma)$. (Note that since the curve complex is a flag complex, this is equivalent to the same statement for the curve complex.)

There are a number of variations of this. Given a subset $A \subseteq C(\Sigma)$ we write $\mathcal{G}(\Sigma, A)$ for the full subgraph of $\mathcal{G}(\Sigma)$ with vertex set A . If A is $\text{Map}(\Sigma)$ -invariant, then $\text{Map}(\Sigma)$ also acts on $\mathcal{G}(\Sigma, A)$. We say that $\mathcal{G}(\Sigma, A)$ is *rigid* if every automorphism is induced by an element of $\text{Map}(\Sigma)$.

For example, if $C_s(\Sigma)$ is the set of separating curves, then we refer to $\mathcal{G}_s(\Sigma) = \mathcal{G}(\Sigma, C_s(\Sigma))$ as the *separating curve graph*. (Note that if $g = 0$, then this is the same as $\mathcal{G}(\Sigma)$.) The results of [BrM; Ki], together with that cited for planar surfaces, tell us $\mathcal{G}_s(\Sigma)$ is rigid if $g(\Sigma) \geq 3$ or ($g(\Sigma) = 2$ and $p(\Sigma) \geq 2$) or ($g(\Sigma) = 1$ and $p(\Sigma) \geq 2$) or ($g(\Sigma) = 0$ and $p(\Sigma) \geq 5$).

We remark that the nonseparating curve graphs $\mathcal{G}(\Sigma, C(\Sigma) \setminus C_s(\Sigma))$ of a large class of surfaces of genus at least 2 are also rigid [If], though this is not directly relevant to the present paper.

Let $C_0(\Sigma) \subseteq C_s(\Sigma)$ be the set of curves that bound some three-holed sphere in Σ . Let $C_{ss}(\Sigma) = C_s(\Sigma) \setminus C_0(\Sigma)$, and let $\mathcal{G}_{ss}(\Sigma) = \mathcal{G}(\Sigma, C_{ss}(\Sigma))$. We refer to elements of $C_{ss}(\Sigma)$ as *strongly separating curves* and to $\mathcal{G}_{ss}(\Sigma)$ as the *strongly separating curve graph*.

We will show here that $\mathcal{G}_{ss}(\Sigma)$ is rigid in all but finitely many cases:

THEOREM 1.1. *If $g(\Sigma) + p(\Sigma) \geq 7$, then $\mathcal{G}_{ss}(\Sigma)$ is rigid.*

Note that if $p(\Sigma) \leq 1$, then $\mathcal{G}_{ss}(\Sigma) = \mathcal{G}_s(\Sigma)$, and so this is covered by the results of [BrM; Ki].

This still leaves unresolved about a dozen cases, which I suspect are also rigid. We can probably deal with a few more cases with some elaboration on the arguments here, though a complete answer may require new ideas.

It is natural to ask more generally for what classes of subsets $A \subseteq C_s(\Sigma)$ is $\mathcal{G}(\Sigma, A)$ rigid. (Note there are only finitely many possibilities for A for any given topological type.)

The motivation for studying this particular case is the application given in [Bo2] to the Weil–Petersson metric on a Teichmüller space. There it was shown that the rigidity of $\mathcal{G}_{ss}(\Sigma)$ implies the quasi-isometric rigidity of the Weil–Petersson metric associated to Σ . In view of Theorem 1.1, this holds whenever $g(\Sigma) + p(\Sigma) \geq 7$. In particular, Theorem 1.1, together with the results of that paper, shows that in all but at most finitely many cases, the Weil–Petersson space is quasi-isometrically rigid.

We remark that the quasi-isometric rigidity of the Teichmüller metric on a Teichmüller space has been proven independently in [EMR] and [Bo1] (by different methods). For this, the rigidity of the curve graphs [Iv; Ko; L] was used (in place of the strongly separating curve graphs as needed for the Weil–Petersson metric).

In the course of proving the main result of this paper, we also show that most strongly separating curve graphs are distinct (see Proposition 5.2).

An earlier (rather different) draft of this paper was written at the Tokyo Institute of Technology, and I am grateful for the hospitality of that institution and for the invitation of Sadayoshi Kojima. I thank Javier Aramayona for his interest and suggestions.

2. Outline of Proof

We begin by introducing some terminology and notation used throughout the paper.

Let Σ be a compact surface. We assume that $\xi(\Sigma) \geq 2$. Given a curve $\gamma \in C_s(\Sigma)$, we will (by slight abuse of terminology) use the term *complementary component* to refer to the closure of a connected component of $\Sigma \setminus \gamma$. We write $B(\gamma)$ for the complementary component that has smaller complexity. (We will only use this notation when it is unambiguous.) We write $C_i(\Sigma) \subseteq C_s(\Sigma)$ for the

set of $\gamma \in C_s(\Sigma)$ for which $B(\gamma)$ has complexity at most i . When $i = 0$, $B(\gamma)$ is an $S_{0,3}$. (Recall that $C_0(\Sigma)$ was defined in this way in Section 1.) Given distinct $\alpha, \beta \in C_s(\Sigma)$, we write $\alpha < \beta$ to mean that $B(\alpha) \subseteq B(\beta)$. We will always assume curves in $C_s(\Sigma)$ to be realised in general position and with minimal intersection. Given $\alpha, \beta \in C_{ss}(\Sigma)$, we write $\iota(\alpha, \beta) = |\alpha \cap \beta|$ for the geometric intersection number. The following notion will be central to the proof.

DEFINITION. (When $p(\Sigma) \geq 5$.) We say that $\alpha, \beta \in C_{ss}(\Sigma)$ form a *surrounding pair* if $B(\alpha)$ and $B(\beta)$ are both $S_{0,4}$ and if $B(\alpha) \cap B(\beta)$ is an $S_{0,3}$.

(We will need to modify this definition slightly when $p(\Sigma) \leq 4$, as we discuss in Section 6.)

Note that $\partial(B(\alpha) \cap B(\beta))$ gives a curve $\omega \in C_0(\Sigma)$ satisfying $\omega < \alpha$ and $\omega < \beta$. Indeed, ω is uniquely determined by this property, and we say that the pair α, β *surrounds* ω .

DEFINITION. We say that $\alpha, \beta, \gamma \in C_{ss}(\Sigma)$ form a *surrounding triple* if any pair of them form a surrounding pair and there is some (necessarily unique) $\omega \in C_0(\Sigma)$, with $\omega < \alpha$, $\omega < \beta$, and $\omega < \gamma$.

(Modulo the definition of “surrounding pair”, the definition of “surrounding triple” will remain unchanged when $p(\Sigma) \leq 4$.)

We will say that we can *recognise* a given property of a collection of curves in $C_{ss}(\Sigma)$ if it is preserved under any automorphism of $\mathcal{G}_{ss}(\Sigma)$. Thus, by the definition of $\mathcal{G}_{ss}(\Sigma)$ we can recognise disjointness of curves. The ultimate goal will be to show that we can recognise all combinatorial (i.e., $\text{Map}(\Sigma)$ -invariant) properties.

In particular, an intermediate goal is to recognise surrounding pairs and surrounding triples. This will allow us to “reconstruct” the graph $\mathcal{G}_s(\Sigma)$ from $\mathcal{G}_{ss}(\Sigma)$ in a canonical way. (We have already observed that a surrounding pair determines an element of $C_0(\Sigma)$, which is a step in that direction.) Thus, any automorphism of $\mathcal{G}_{ss}(\Sigma)$ extends to an automorphism to $\mathcal{G}_s(\Sigma)$, and we can apply the results of [Ko; BrM; Ki] to see that it is induced by an element of $\text{Map}(\Sigma)$.

In Sections 3 and 4, we begin by considering the case where Σ is an $S_{0,7}$. The key point here is that there is only one heptagon (that is 7-cycle) in $\mathcal{G}_{ss}(S_{0,7})$ up to the action of $\text{Map}(S_{0,7})$ (see Proposition 3.1). A surrounding pair can now be recognised as a pair of vertices at distance 2 apart in some heptagon. In Section 4, we proceed to show that $\mathcal{G}_{ss}(S_{0,7})$ is rigid.

For the general case, we need to recognise the topological type of a multicurve in the surface Σ . This argument is largely independent of the $S_{0,7}$ case and is discussed in Section 5. In Section 6, we then combine this with what we know about $S_{0,7}$ to prove Theorem 1.1 in general (for $g(\Sigma) + p(\Sigma) \geq 7$)—except, that is, when Σ is an $S_{0,8}$, which we will treat as a special case in Section 7.

3. Heptagons in the 7-Holed Sphere

We begin with a description of 7-cycles (or “heptagons”) in the separating curve graph of $S_{0,7}$. In general, by an n -cycle in $\mathcal{G}_{ss}(\Sigma)$ we mean a cyclically ordered sequence of n vertices, where consecutive vertices are adjacent. We refer to it as an *odd* or *even* cycle depending on whether n is odd or even. Note that any shortest odd cycle is necessarily isometrically embedded.

For the purposes of this and the next section, it will be convenient to view $S_{0,7}$ as (the complement of) the 2-sphere S with a set of seven preferred points $\Pi \subseteq S$. (In other words, S is obtained by collapsing each boundary component of $S_{0,7}$ to a point of Π . We can recover $S_{0,7}$ by removing a small open disc about each point of Π .)

Note that if $\gamma \in C_{ss}$, then γ bounds a disc $B(\gamma)$ with $|B(\gamma) \cap \Pi| = 3$. We write $\pi(\gamma) = B(\gamma) \cap \Pi$. Note that $\alpha, \beta \in C_{ss}$ are adjacent if and only if we can realise α, β so that $B(\alpha) \cap B(\beta) = \emptyset$. We say that two curves are n -distant if they are a distance exactly n apart in \mathcal{G}_{ss} .

Note that if $\iota(\alpha, \beta) = 2$, then $\alpha \cup \beta$ cuts S into four discs. In particular, $B(\alpha) \cap B(\beta)$, $B(\alpha) \cup B(\beta)$, $B(\alpha) \setminus B(\beta)$, and $B(\beta) \setminus B(\alpha)$ are all discs.

It is easily seen that $\mathcal{G}_{ss}(S_{0,7})$ has no 3-cycles. Also it has no 4-cycles. (For if $\gamma_1, \gamma_2, \gamma_3, \gamma_4$ were a 4-cycle, $B(\gamma_1) \cup B(\gamma_3)$ and $B(\gamma_2) \cup B(\gamma_4)$ would be disjoint connected subsurfaces of S , each containing at least four points of Π , which clearly is not possible.) We will also see (Lemma 3.3) that $\mathcal{G}_{ss}(S_{0,7})$ has no 5-cycles. This implies that any 7-cycle must be isometrically embedded.

In fact, $\mathcal{G}_{ss}(S_{0,7})$ does contain 6-cycles. (For example, take disjoint discs D_1, D_2, F_1, F_2, F_3 in S so that each $|D_i \cap \Pi| = 2$ and each $|F_j \cap \Pi| = 1$. Now connect each D_i to each F_j by a set of six disjoint arcs a_{ij} . Let β_{ij} be the boundary of a regular neighbourhood of $D_i \cup a_{ij} \cup F_j$. Then $\beta_{11}, \beta_{23}, \beta_{12}, \beta_{21}, \beta_{13}, \beta_{22}$ is a 6-cycle.) We will however focus on the 7-cycles since these are more symmetrical and will serve for our purposes.

Here is a description of a 7-cycle. Let $\lambda \subseteq \Sigma$ be an embedded circle with $\Pi \subseteq \lambda$. This determines a cyclic ordering on Π where we index the punctures as $p_1, p_3, p_5, p_7, p_2, p_4, p_6$. Let l_{13} be the segment between p_1 and p_3 , and so on. Thus, $\lambda = l_{13} \cup l_{35} \cup l_{57} \cup l_{72} \cup l_{24} \cup l_{46} \cup l_{61}$. Let B_i be a regular neighbourhood of $l_{i-2,i} \cup l_{i,i+2}$ with $B_i \cap \Pi = \{p_{i-2}, p_i, p_{i+2}\}$, and let $\gamma_i = \partial B_i$. Thus, $B_i = B(\gamma_i)$. Note that $B_i \cap B_{i+1} = \emptyset$, and so $\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6, \gamma_7$ is a 7-cycle. Any nonadjacent pairs of curves intersect exactly twice (Figure 1).

Now $\text{Map}(S_{0,7})$ acts on $\mathcal{G}_{ss}(S_{0,7})$. The main aim of this section is to show the following:

PROPOSITION 3.1. *There is exactly one 7-cycle in $\mathcal{G}_{ss}(S_{0,7})$ up to the action of $\text{Map}(S_{0,7})$.*

We begin with an analogous statement for 3-sets in a 7-set. Given any set Ψ , let $\Theta = \Theta(\Psi)$ be the graph whose vertex set $V(\Theta)$ consists of subsets of cardinality 3 in Ψ and whose edge set $E(\Theta)$ consists of pairs of disjoint such 3-sets.




Figure 1 A 7-cycle in $G_{SS}(S_{0,7})$

For the remainder of this section, we will assume that $|\Psi| = 7$. In this case, it is a connected 4-regular graph on thirty-five vertices. Note also that there is an edge colouring $\chi : E(\Theta) \rightarrow \Psi$ given by $\Psi \setminus (P \cup Q) = \{\chi(e)\}$, where e is the edge from P to Q .

The following simple observation will be useful:

LEMMA 3.2. *If $P, Q \in V(\Theta)$ are 2-distant, then $|P \cap Q| = 2$. If P, Q are 3-distant, then $|P \cap Q| = 1$.*

The following is the analogue, in Θ , of the main result of this section:

LEMMA 3.3. *If $|\Psi| = 7$, there are no 3-cycles or 5-cycles in $\Theta(\Psi)$. There is exactly one 7-cycle up to the action of $\text{Sym}(\Psi)$.*

Proof. The nonexistence of 3-cycles is trivial.

Suppose P_1, P_2, P_3, P_4, P_5 were a 5-cycle. Then P_3 and P_4 are both 2-distant from P_1 , so $|P_1 \cap P_3| = |P_1 \cap P_4| = 2$, and we get the contradiction that $P_3 \cap P_4 \neq \emptyset$.

Writing $\Psi = \{1, 2, 3, 4, 5, 6, 7\}$, there is a 7-cycle in Ψ given by 613–724–135–246–357–461–572. We want to show that this is the only one up to the action of $\text{Sym}(\Psi)$.

Suppose then that $P_1, P_2, P_3, P_4, P_5, P_6, P_7$ is a 7-cycle. Since there is no smaller odd cycle, this must be isometrically embedded in Ψ . Suppose that two edges have the same colour, p , say. Since the P_i are all distinct, these edges cannot meet or contain adjacent vertices. Thus, up to cyclic reordering, the only possibility for this pair of edges is P_1, P_2 and P_4, P_5 . Now P_2 and P_4 are 2-distant, and so $|P_2 \cap P_4| = 2$. Now, P_1 and P_5 are the complements of these sets in $\Psi \setminus \{p\}$,




Figure 2 a : α -type arc, b : β -type arc, c : crossing arc

and so we also have $|P_1 \cap P_5| = 2$. But these are 3-distant, so this contradicts Lemma 3.2.

Therefore, each colour occurs exactly once around the cycle. Up to $\text{Sym}(\Psi)$, we can assume that they occur in the cyclic order 1234567, starting with the edge P_1, P_2 . Now consider the sequence P_1, P_3, P_5, P_7 . We must proceed by replacing 2 by 1, then 4 by 3, then 6 by 5. So we must have started with P_1 being 246 (and ended with P_7 as 135). But now the whole 7-cycle starting with P_1 is completely determined by the colours on the edges. In fact, it must be precisely the cyclic sequence given above. This proves the result. \square

We now move on to the proof of Proposition 3.1. We set $\Psi = \Pi$. Recall that we have defined $\pi : C_{ss} \rightarrow V(\Pi)$ by $\pi(\gamma) = \Pi \cap B(\gamma)$. This extends to a map $\pi : \mathcal{G}_{ss} \rightarrow \Theta$, sending edges to edges. Composing with χ , we also get a colouring of the edges of \mathcal{G}_{ss} , which we also denote by $\chi : E(\mathcal{G}_{ss}) \rightarrow \Pi$. Note that it is now an immediate consequence that there are no 5-cycles in \mathcal{G}_{ss} , as stated earlier.

We need to make a few observations about configurations of curves separated by given distances in \mathcal{G}_{ss} .

If $\alpha, \beta \in C_{ss}$ are adjacent, then we set $A = A(\alpha, \beta)$ to be (the closure of) $S \setminus (B(\alpha) \cup B(\beta))$. This is an annulus with $\partial A = \alpha \cup \beta$ and with $A(\alpha, \beta) \cap \Pi = \{p\}$, where $p = \chi(\alpha, \beta)$.

By an *arc* in A , we mean an arc $a \subseteq A \setminus \Pi$ with endpoints $\partial a = a \cap \partial A$. We generally regard such an arc as defined up to homotopy in $A \setminus \Pi$, allowing ourselves to slide an endpoint of a in ∂A . Up to homotopy, there are exactly three types of arc, depending on whether a meets only α , only β , or both α and β . We refer to these classes as α -type, β -type, or *crossing* arcs, respectively. Note that an α -type arc and a β -type arc meet (minimally) in exactly two points (Figure 2).

Suppose now that $\beta, \delta \in C_{ss}$ are 2-distant. Then $|\pi(\beta) \cap \pi(\delta)| = 2$, so $\pi(\delta) \setminus \pi(\beta) = \{q\}$ for some $q \in \Pi$. Let $D = D(\delta, \beta)$ be (the closure of) the component of $B(\delta) \setminus B(\beta)$ containing q . We claim that D is a bigon:




Figure 3 The bigon D

LEMMA 3.4. *If $\beta, \delta \in C_{SS}$ are 2-distant, and $D = D(\beta, \delta)$, then $D \cap B(\beta)$ consists of a single arc in β (Figure 3).*

Proof. Let $\gamma \in C_{SS}$ be adjacent to both β and δ . Note that $B(\beta), B(\delta) \subseteq S \setminus B(\gamma)$, and so $D \subseteq A(\beta, \gamma)$ and $D \cap \gamma = \emptyset$. Thus, ∂D can contain only β -type arcs in $A(\beta, \gamma)$. There can only be one of these, so the statement follows easily. \square

Suppose now that α, β are adjacent and that $\varepsilon \in C_{SS}$ is 3-distant from both α and β . Let $A = A(\alpha, \beta)$ and $p = \chi(\alpha, \beta)$ as before. By Lemma 3.2, $|\pi(\varepsilon) \cap \pi(\alpha)| = |\pi(\varepsilon) \cap \pi(\beta)| = 1$, and so it follows that $p \in \pi(\varepsilon)$. Let $F = F(\alpha, \beta; \varepsilon)$ be (the closure of) the component of $A \cap B(\varepsilon)$ containing p . Note that F must intersect either α or β , possibly both. (In fact, given that there are only three classes of arc in A , we can easily see that F can meet each of $B(\alpha)$ and $B(\beta)$ in at most a single arc, though we will not explicitly need this.)

Now let σ be any 7-cycle in \mathcal{G}_{SS} .

Suppose that α, β is an edge of σ . Let ε be the vertex of σ opposite this edge. Thus, ε is 3-distant from both α and β , as before. Let $F = F(\alpha, \beta; \varepsilon)$, as before.

Suppose that $\beta \cap F \neq \emptyset$. This implies that any α -type arc in $A = A(\alpha, \beta)$ must intersect F . Let γ be the vertex of σ adjacent to β and distinct from α (so that γ and ε are 2-distant). We claim:

LEMMA 3.5. *Let α, β, γ be consecutive vertices of σ . Let ε be the vertex of σ opposite the edge α, β , and let $F = F(\alpha, \beta; \varepsilon)$. If $F \cap \beta \neq \emptyset$, then $\iota(\alpha, \gamma) = 2$.*

Proof. Let δ be the vertex of σ between γ and ε (so that $\alpha, \beta, \gamma, \delta, \varepsilon$ are consecutive vertices of σ). Let $D = D(\delta, \beta)$. By Lemma 3.4 this is a bigon, that is, $\partial D = b \cup d$, where b and d are respectively arcs of β and δ . Now $F \subseteq B(\varepsilon)$, $D \subseteq B(\delta)$, and $B(\varepsilon) \cap B(\delta) = \emptyset$, so $F \cap D = \emptyset$. Let $\pi(\delta) \setminus \pi(\beta) = \{q\}$. Note that $q \neq p$, and so $q \in \pi(\alpha) \subseteq B(\alpha)$. Now $d \cap F \subseteq D \cap F = \emptyset$, and so $d \cap A$ contains no α -type arcs. Since $d \cap \beta$ are the endpoints of d , it cannot contain any β -type arcs either. Thus, $d \cap A$ consists only of crossing arcs, of which there must be




Figure 4 $D = D(\delta, \beta)$, $F = F(\alpha, \beta; \varepsilon)$




Figure 5 The curves $\alpha, \delta, \varepsilon$

exactly two. This means that $|d \cap \alpha| = 2$ (with $d \cap B(\alpha)$ consisting of a single arc); see Figure 4.

Let $R = B(\beta) \cup D$. Now R is a disc with $R \cap \Pi = \pi(\beta) \cup \{q\} = \pi(\beta) \cup \pi(\alpha)$. Also, $R \subseteq B(\beta) \cup B(\delta)$, so $R \cap B(\gamma) = \emptyset$. It follows that γ and ∂R are homotopic in $S \setminus \Pi$. In other words, they represent the same element of C_{SS} . But now, $\partial R \subseteq \beta \cup d$, so $|\partial R \cap \alpha| = 2$. Thus, $\iota(\alpha, \gamma) \leq 2$, and so, in fact, $\iota(\alpha, \gamma) = 2$ as required (Figure 5). □

Now, as already observed, at most one of $F \cap \alpha$ or $F \cap \beta$ can be empty. If $F \cap \beta = \emptyset$, we refer to β as a *bad* endpoint of the edge α, β . We say that a vertex of σ is *bad* if it is a bad endpoint of both incident edges of σ . Thus, Lemma 3.5 tells us that if β is not bad, then the two vertices adjacent to β in σ correspond to curves which intersect twice. (In fact, Proposition 3.1 will retrospectively rule out bad vertices altogether.) In Figure 4, β would be a bad vertex of the edge α, β .

Proof of Proposition 3.1. Let Σ be a 7-cycle.




Figure 6 The curves $\gamma_2, \gamma_3, \gamma_4, \gamma_5$

Now no two bad vertices of σ can be adjacent. It follows that there can be at most three bad vertices in total. In fact, we can index the vertices of σ consecutively (mod 7) as $\gamma_1, \gamma_2, \gamma_3, \gamma_4, \gamma_5, \gamma_6, \gamma_7$, so that none of $\gamma_1, \gamma_3, \gamma_4$, or γ_6 are bad. It then follows that $\iota(\gamma_7, \gamma_2) = \iota(\gamma_2, \gamma_4) = \iota(\gamma_3, \gamma_5) = \iota(\gamma_5, \gamma_7) = 2$.

We write $B_i = B(\gamma_i)$ and label the points of Π as p_i , so that $\pi(\gamma_i) = B_i \cap \Pi = \{p_{i-2}, p_i, p_{i+2}\}$.

Consider first the discs B_2, B_3, B_4, B_5 . We have $B_2 \cap B_3 = B_3 \cap B_4 = B_4 \cap B_5 = \emptyset$. Let $A = A(\gamma_3, \gamma_4)$. Since $\iota(\gamma_3, \gamma_5) = 2$, we see that $\gamma_5 \cap A$ consists of a single γ_3 -type arc. Similarly, $\gamma_2 \cap A$ consists of a single γ_4 -type arc. It follows that $\iota(\gamma_2, \gamma_5) = 2$. Thus, $B_2 \cap B_5$ is a disc with $B_4 \cap B_5 \cap \Pi = \{p_7\}$. Now $B_3 \cap B_5$ and $B_2 \cap B_4$ are also discs with $B_3 \cap B_5 \cap \Pi = \{p_3, p_5\}$ and $B_2 \cap B_4 \cap \Pi = \{p_2, p_4\}$. We can therefore find an arc l with $\Pi \subseteq l$, with endpoints p_1 and p_6 and with the points $p_1, p_3, p_5, p_7, p_2, p_4, p_6$ occurring in this order along l , and so that B_3, B_5, B_2, B_4 are respectively regular neighbourhoods of $l_{13} \cup l_{35}, l_{35} \cup l_{57}, l_{72} \cup l_{24}$ and $l_{24} \cup l_{46}$, where we have cut l into six arcs, $l = l_{13} \cup l_{35} \cup l_{57} \cup l_{72} \cup l_{24} \cup l_{46}$ connecting the points of Π . Note in particular, that $R = B_2 \cup B_5$ is a disc with $R \cap \Pi = \{p_3, p_5, p_7, p_2, p_4\}$. Also, $B_2 \cap B_5$ is a disc with $B_2 \cap B_5 \cap \Pi = \{p_7\}$ (Figure 6).

Now consider B_7 . We have $B_7 \cap \Pi = \{p_5, p_7, p_2\}$. Now $\iota(\gamma_5, \gamma_7) = 2$, so $\gamma_7 \cap B_5$ consists of a single arc separating p_3 from p_5 and p_7 in B_5 . We can therefore realise it so that it is disjoint from $B_2 \cap B_5$. In fact, we can take this arc to meet l just once, in a point of the segment l_{35} . Similarly, we can realise $\gamma_7 \cap B_2$ as a single arc, also disjoint from $B_2 \cap B_5$, and meeting l in a single point of l_{24} . In this way, $\gamma_7 \setminus R$ consists of exactly two arcs. Since $(B_7 \setminus R) \cap \Pi = \emptyset$, each of these arcs can be homotoped into ∂R in $S \setminus \Pi$, fixing their endpoints. We can therefore realise B_7 as a regular neighbourhood of $l_{57} \cup l_{72}$.

Now let l_{61} be any arc in S meeting l exactly at their common endpoints. Thus, $\lambda = l \cup l_{61}$ is a circle containing Π . Now the homotopy classes of γ_1 and γ_6 are determined as $\partial(B_7 \cup B_2)$ and $\partial(B_7 \cup B_4)$. We can therefore realise B_1 and B_6 respectively as regular neighbourhoods of $l_{46} \cup l_{61}$ and of $l_{61} \cup l_{13}$.

We are therefore exactly in the situation of the example of the 7-cycle described earlier (as in Figure 1).

This proves Proposition 3.1. □

We note the following immediate consequence.




Figure 7 A surrounding pair α, β

LEMMA 3.6. Any automorphism of $\mathcal{G}_{SS}(S_{0,7})$ preserves the set of surrounding pairs.

Proof. In view of Proposition 3.1, we see that α, β form a surrounding pair if and only if they are 2-distant vertices in some 7-cycle in \mathcal{G}_{SS} . Note that this is determined just by the structure of \mathcal{G}_{SS} . □

4. Rigidity for the 7-Holed Sphere

Recall that $C_0 = C_0(S_{0,7})$ is the set of curves in $S \setminus \Pi$ that bound a disc containing exactly two points of Π . If $\omega \in C_0$, then this disc is $B(\omega)$, and $\pi(\omega) = B(\omega) \cap \Pi$. If $\omega \in C_0$ and $\alpha \in C_{SS}$, then $\omega < \alpha$ means that $B(\omega) \subseteq B(\alpha)$.

By an ω -arc we mean an arc a in S meeting $B(\omega)$ at one endpoint (the *initial* endpoint) and Π at the other (*terminal* endpoint). We regard a as being defined up to homotopy relative to Π , fixing the terminal endpoint, and allowing the initial endpoint to slide along ω . Note that a determines an element $\alpha \in C_{SS}$ with $\omega < \alpha$, so that $B(\alpha)$ is a regular neighbourhood of $B(\omega) \cup a$. In fact, every $\alpha \in C_{SS}$ with $\omega < \alpha$ arises in this way.

If $\alpha, \beta \in C_{SS}$ is a surrounding pair, then there is a unique $\omega \in C_0$ with $\omega < \alpha$ and $\omega < \beta$. In fact, $\omega = \partial(B(\alpha) \cap B(\beta))$. Note that α, β correspond to disjoint ω -arcs, a, b (i.e. we can realise a, b to be disjoint). We say that α, β surround ω (Figure 7).

Recall that a *surrounding triple* consists of three curves $\alpha, \beta, \gamma \in C_{SS}$ such that $\{\alpha, \beta\}$, $\{\beta, \gamma\}$, and $\{\gamma, \alpha\}$ are all surrounding pairs and such that there is some $\omega \in C_0$ with $\omega < \alpha, \beta, \gamma$. In this case, α, β, γ correspond to three pairwise disjoint ω -arcs, a, b, c .

Suppose we just know that $\alpha, \beta, \gamma \in C_{SS}$ are such that any pair forms a surrounding pair. Then $\pi(\alpha), \pi(\beta), \pi(\gamma)$ pairwise intersect in sets of two elements. It follows easily that $\pi(\alpha) \cup \pi(\beta) \cup \pi(\gamma)$ has either four or five elements. In the former case, $|\pi(\alpha) \cap \pi(\beta) \cap \pi(\gamma)| = 1$, and we see easily that there is a curve in C_{SS} (namely $\partial(B(\alpha) \cup B(\beta) \cup B(\gamma))$), disjoint from each of α, β, γ . In the latter case, $|\pi(\alpha) \cap \pi(\beta) \cap \pi(\gamma)| = 2$, and we see easily that α, β, γ form a surrounding triple. We deduce:

LEMMA 4.1. *Suppose $\alpha, \beta, \gamma \in C_{SS}$. Then α, β, γ form a surrounding triple if and only if any pair of them form a surrounding pair and there is no curve in C_{SS} that is disjoint from each of α, β, γ .*

In view of Lemma 3.6, we see that we can recognise when three elements of C_{SS} form a surrounding triple.

Let \mathcal{H} be the graph whose vertex set is the set of all surrounding pairs in C_{SS} , and where two vertices are deemed adjacent if the union of the two pairs is a surrounding triple. Given any $\omega \in C_0$, let $\mathcal{H}(\omega)$ be the full subgraph whose vertex set consists of those surrounding pairs that surround ω . Now adjacent vertices of \mathcal{H} determine the same element of C_0 , and so $\mathcal{H}(\omega)$ is a union of components of \mathcal{H} . In fact:

LEMMA 4.2. *If $\omega \in C_0$, then the graph $\mathcal{H}(\omega)$ is connected.*

Proof. This is best seen in terms of ω -arcs. Recall that a vertex of $\mathcal{H}(\omega)$ corresponds to a pair of disjoint ω -arcs and that an edge of $\mathcal{H}(\omega)$ corresponds to a triple of pairwise disjoint ω -arcs.

Suppose that a, b, c, d are ω -arcs with $a \cap b = \emptyset$ and $c \cap d = \emptyset$. We realise them in general position in S . (They do not need to have minimal intersection in their homotopy classes.) We aim to connect the vertices a, b and c, d by a path in $\mathcal{H}(\omega)$. Write $I(a, b; c, d) = (a \cup b) \cap (c \cup d) \setminus \Pi$ for the set of interior intersection points. We proceed by induction on $|I(a, b; c, d)|$.

The case where $I(a, b; c, d) = \emptyset$ is elementary; so we assume that there is some $x \in I(a, b; c, d)$. After permuting a, b and c, d , we can assume that $x \in a \cup c$ and that x is the first intersection point along c ; that is, the initial segment e of c ending at x meets $a \cup b$ only at x . Let f be the initial segment of a ending at x , and set $a' = (a \setminus f) \cup e$. We can move a' slightly so that $a \cap a'$ meet precisely in their terminal points while retaining disjointness from b . Now, we can easily find an ω -arc e disjoint from each of b, a, a' , so a, b, e , and a', b, e correspond to surrounding triples. It follows that a, b and a', b correspond to surrounding pairs in the same component of $\mathcal{H}(\omega)$. Note that $|I(a', b; c, d)| < |I(a, b; c, d)|$. We therefore replace a, b by a', b and proceed inductively. □

We deduce:

LEMMA 4.3. *There is a natural bijective correspondence between the elements of C_0 and the connected components of \mathcal{H} such that if $\omega \in C_0$ and $\alpha \in C_{SS}$, then $\omega < \alpha$ if and only if α occurs as a curve in a surrounding pair of some vertex of the corresponding component of \mathcal{H} .*

Since \mathcal{H} can be constructed out of \mathcal{G}_{SS} , we see that we can also reconstruct C_0 and the relation $<$ between C_0 and C_{SS} out of \mathcal{G}_{SS} .

We can also recognise disjointness. If $\omega \in C_0$ and $\alpha \in C_{SS}$, then ω and α are disjoint if and only if there is some $\beta \in C_{SS}$ such that $\omega < \beta$ and α, β are either equal or disjoint. Similarly, if $\omega, \omega' \in C_0$, then ω and ω' are disjoint if and only if there are disjoint curves $\alpha, \alpha' \in C_{SS}$ with $\omega < \alpha$ and $\omega' < \alpha'$.

We can therefore reconstruct the graph $\mathcal{G}(S_{0,7}) = \mathcal{G}_s(S_{0,7})$ from the graph $\mathcal{G}_{ss}(S_{0,7})$. It follows that any automorphism of the former extends to an automorphism of the latter. But we know that $\mathcal{G}(S_{0,7})$ is rigid by Theorem 1 of [Ko], so we have shown the following:

PROPOSITION 4.4. *The graph $\mathcal{G}_{ss}(S_{0,7})$ is rigid.*

5. Multicurves

In this section, we explain how to identify classes of multicurves from the structure of the strongly separating curve graph. We assume that Σ is a compact orientable surface with boundary $\partial\Sigma$, which we view as a (possibly empty) set of curves. (Thus, $S_{0,7}$ reverts to being a bona fide 7-holed sphere.) We will assume that $\xi(\Sigma) \geq 4$.

As before, we say that we can *recognise* a property of a collection of curves in $C_{ss} = C_{ss}(\Sigma)$ if this property is preserved under any automorphism of $\mathcal{G}_{ss}(\Sigma)$. In other words, it can be seen just in terms of the graph structure. We similarly say that we can “tell” if a given property holds. We also say that another graph can be “constructed” (from \mathcal{G}_{ss}), and so on.

Recall that $C_1 \subseteq C_{ss}$ is the set of $\alpha \in C_{ss}$ for which $B(\alpha)$ has complexity 1, in other words, it is either an $S_{0,4}$ or an $S_{1,1}$.

Suppose that $\alpha, \beta, \gamma \in C_{ss}$. We can tell if α separates β from γ . (It is equivalent to saying that any curve in C_{ss} that is disjoint from α must also be disjoint from either β or γ .) Therefore, we can recognise elements of C_1 : a curve in C_{ss} lies in C_1 if and only if it does not separate any two other elements of C_{ss} .

By a *multicurve* τ in C_{ss} we mean a nonempty set of pairwise disjoint curves in C_{ss} . We will sometimes abuse notation by regarding τ as a subset of Σ . We claim that we can identify τ up to the action of $\text{Map}(\Sigma)$. This is equivalent to saying that we can recognise the topological types of each of the components of $\Sigma \setminus \tau$ together with the elements of τ that bound a given component.

Let X be a component of $\Sigma \setminus \tau$. Write $\partial X \subseteq \tau \cup \partial\Sigma$ for its intrinsic boundary, and $\partial_\Sigma X = \tau \cap \partial X$ for the relative boundary. We write $p(X) = |\partial X|$, $q(X) = |\partial_\Sigma X|$, and $g(X)$ for the genus of X . Write $C_{ss}(\Sigma, X)$ for the set of elements of $C_{ss}(\Sigma) \setminus \tau$ contained in X . Note that $C_{ss}(\Sigma, X)$ is either empty or infinite.

DEFINITION. We say that a complementary component X is *large* if $C_{ss}(\Sigma, X)$ is infinite.

Thus, X is not large if and only if it is either an $S_{0,3}$ with $q(X) \geq 2$ or else has the form $B(\alpha)$ for some $\alpha \in C_1(\Sigma)$.

Note that if $\alpha \in C_{ss}$, then $\alpha \in C_{ss}(\Sigma, X)$ for some large component X if and only if it does not lie in τ and is disjoint from τ (i.e. is disjoint from each element of τ). Given two such curves $\alpha, \beta \in C_{ss}$, then α, β lie in the same set $C_{ss}(\Sigma, X)$ if and only if they are not separated by any element of τ . Thus, from τ we can identify the collection of sets $C_{ss}(\Sigma, X)$ that arise from the large components X of $\Sigma \setminus \tau$. We next want to recognise the topological type of such X .




Figure 8 Here, X is the subsurface bounded by $\alpha \cup \beta$

To this end, we define a *chain* in C_{ss} to be a sequence, $\gamma_0, \gamma_1, \dots, \gamma_n$ of disjoint curves such that γ_j separates γ_i from γ_k whenever $i < j < k$.

Suppose that X is a component of $\Sigma \setminus \tau$ with $q(X) \geq 2$. Choose any distinct $\alpha, \beta \in \partial_\Sigma X$, and let n be maximal so that there is a chain $\alpha = \gamma_0, \dots, \gamma_n = \beta$ in C_{ss} (so that $\gamma_i \in C_{ss}(\Sigma, X)$ for all $i \neq 0, n$). Each component of $X \setminus \bigcup_i \gamma_i$ is either an $S_{0,3}$ or an $S_{1,2}$ (Figure 8). Moreover, we can tell if the component between γ_i and γ_{i+1} is an $S_{1,2}$, since in that case, there will be some $\delta \in C_{ss}(\Sigma, X)$ such that γ_i separates δ from α and γ_{i+1} separates δ from β . Thus, we know the number m of such $S_{1,2}$ components. We see that $g(X) = m$ and $p(X) = n - m + 2$, so we have determined the type of X in this case.

Now suppose that $q(X) = 1$ and that X is large (equivalently, not an $S_{0,4}$ nor an $S_{1,1}$). Write $\partial_\Sigma X = \{\alpha\}$. Suppose that $\beta \in C_{ss}(\Sigma, X) \cap C_1(\Sigma)$ (which we can recognise). Let $Y = X \setminus B(\beta)$, so that Y is a component of $\Sigma \setminus (\tau \cup \beta)$, and $\partial_\Sigma Y = \{\alpha, \beta\}$. By the previous paragraph we know the type of Y (given β). Therefore, we know the collection of types of all such Y that can arise in this way. Given that $B(\beta)$ is either an $S_{0,4}$ or an $S_{1,1}$, there are at most two such types. If $\xi(X) \geq 4$, then we now see easily that this data determines the topological type of X . (Note that if there is only one type for Y , then either $g(X) = 0$ or $p(X) \leq 3$. If there are two types, then $p(X) \geq 4$.) However, if $\xi(X) \leq 3$, then the data does not allow us to distinguish the pairs $\{S_{0,5}, S_{1,2}\}$ or $\{S_{0,6}, S_{1,3}\}$. (For if X is an $S_{0,5}$ or $S_{1,2}$, then Y must be an $S_{0,3}$; and if X is an $S_{0,6}$ or $S_{1,3}$, then Y must be an $S_{0,4}$.)

Now suppose that $\delta \in C_1(\Sigma)$. Let $Z = \Sigma \setminus B(\delta)$. Suppose that $\xi(\Sigma) \geq 6$ and that $\Sigma \neq S_{0,7}$. We see that $Z \neq S_{0,5}, S_{0,6}, S_{1,2}, S_{1,3}$. Thus, by the previous paragraph we can determine the type of Z . Since this holds for all elements of $C_1(\Sigma)$, we can now easily determine the topological type of Σ and also tell whether an element of $C_1(\Sigma)$ bounds an $S_{0,4}$ or an $S_{1,1}$.

Retrospectively, we can now go back to the earlier setup and distinguish an $S_{0,5}$ from an $S_{1,2}$, or an $S_{0,6}$ from an $S_{1,3}$ in the complement of τ . We therefore now know the types of all large components. From this we can easily determine τ up to the action of $\text{Map}(\Sigma)$.

In summary, we have shown the following:

LEMMA 5.1. *Suppose $\xi(\Sigma) \geq 6$. Suppose that $\tau, \tau' \subseteq C_{ss}(\Sigma)$ are two multicurves and that there is an automorphism of $\mathcal{G}_{ss}(\Sigma)$ taking τ to τ' . Then there is an element of $\text{Map}(\Sigma)$ taking τ to τ' .*

Given a component X of $\Sigma \setminus \tau$, set $C_{ss}(X)$ and $\mathcal{G}_{ss}(X)$ as defined intrinsically to X . Thus, $\mathcal{G}_{ss}(X)$ is the full subgraph of $\mathcal{G}_{ss}(\Sigma, X)$ with vertex set $C_{ss}(X)$. Note that we can tell whether a curve $\gamma \in C_{ss}(\Sigma, X)$ lies in $C_{ss}(X)$ (since it does not bound an $S_{0,3}$ component of $\Sigma \setminus (\tau \cup \gamma)$). Thus, we can construct $\mathcal{G}_{ss}(X)$ out of $\mathcal{G}_{ss}(\Sigma)$, given τ .

Note that the above encompasses all cases where $g(\Sigma) + p(\Sigma) \geq 7$.

We remark that we have also proven the following:

PROPOSITION 5.2. *Suppose that Σ, Σ' are compact surfaces with $\mathcal{G}_{ss}(\Sigma)$ isomorphic to $\mathcal{G}_{ss}(\Sigma')$. If $\xi(\Sigma) \geq 6$, then $\Sigma' = \Sigma$.*

Of course, this leaves open a number of cases, which we will not address here.

6. Rigidity of Other Surfaces

In this section, we prove Theorem 1.1, except in the case where $\Sigma = S_{0,8}$. We assume that $g(\Sigma) + p(\Sigma) \geq 7$. We will split into three cases: First, $p(\Sigma) \geq 5$, $\Sigma \neq S_{0,8}$; second, $p(\Sigma) \leq 4$, $g(\Sigma) \geq 4$; and finally $\Sigma \in \{S_{3,4}, S_{4,3}, S_{4,4}\}$.

Recall that for $\alpha, \beta \in C_{ss}(\Sigma)$, $\alpha < \beta$ means that $\alpha \neq \beta$ and $B(\alpha) \subseteq B(\beta)$.

First, consider the case where $p(\Sigma) \geq 5$. In this case, we will also assume that $\Sigma \neq S_{0,8}$.

We begin by giving a criterion for recognising such pairs (as defined in Section 2).

LEMMA 6.1. *Suppose that $\alpha, \beta \in C_{ss}(\Sigma)$. Then α, β is a surrounding pair if and only each of the following three conditions holds:*

- (S1): $B(\alpha)$ and $B(\beta)$ are both $S_{0,4}$,
- (S2): There is some multicurve τ in Σ such that α, β both lie in some component X of $\Sigma \setminus \tau$ of type $S_{0,7}$, and
- (S3): α, β form a surrounding pair intrinsically in X .

(To make sense of (S3), note that necessarily $\alpha, \beta \in C_{ss}(X)$.)

Proof of Lemma 6.1. First, suppose that α, β is a surrounding pair in Σ . Note that $Y = \Sigma \setminus (B(\alpha) \cup B(\beta))$ satisfies $g(Y) = g(\Sigma)$ and $p(Y) = p(\Sigma) - 3$. In particular, $g(Y) + p(Y) \geq 4$, and Y is not an $S_{0,5}$ (since $\Sigma \neq S_{0,8}$). Thus, we can find a multicurve $\tau \subseteq C_{ss}(Y)$ such that $X = Y \setminus \bigcup_{\gamma \in \tau} B'(\gamma)$ is an $S_{0,7}$. Here, $B'(\gamma)$ is the component of $\Sigma \setminus \gamma$ not containing α, β . (It might not be the lower-complexity component.) We can, of course, assume the $B'(\gamma)$ to be disjoint. Now $B(\alpha), B(\beta) \subseteq X$, and we see that α, β is a surrounding pair in X . Thus, α, β satisfies (S1)–(S3).

Conversely, if α, β satisfies (S1)–(S3), then again we must have $B(\alpha), B(\beta) \subseteq X$, and so $B(\alpha) \cap B(\beta)$ is an $S_{0,3}$, so α, β is a surrounding pair in Σ . □

LEMMA 6.2. *If $g(\Sigma) + p(\Sigma) \geq 7$, $p(\Sigma) \geq 5$, and $\Sigma \neq S_{0,8}$, then the collection of surrounding pairs is invariant under any automorphism of $\mathcal{G}_{ss}(\Sigma)$.*

Proof. By Lemma 3.1 it is enough to show that properties (S1)–(S3) are all recognisable in terms of $\mathcal{G}_{ss}(\Sigma)$. We can certainly recognise a multicurve τ in $C_{ss}(\Sigma)$, and by Lemma 5.1 we can tell if α, β lie in a component X of $\Sigma \setminus \tau$ of type $S_{0,7}$. As explained at the end of Section 5, we can also construct $\mathcal{G}_{ss}(X)$. By Lemma 3.6 we can tell if α, β form a surrounding pair intrinsically to X . \square

The class of surrounding triples (as defined in Section 2) is also recognisable in $\mathcal{G}_{ss}(\Sigma)$:

LEMMA 6.3. *If $g(\Sigma) + p(\Sigma) \geq 7$, $p(\Sigma) \geq 5$, and $\Sigma \neq S_{0,8}$, then the collection of surrounding triples is invariant under any automorphism of $\mathcal{G}_{ss}(\Sigma)$.*

Proof. We first make the following general observation (which holds for any surface Σ). Suppose that $B_1, B_2, B_3 \subseteq \Sigma$ are connected subsurfaces (in general position) such that the ∂B_i are all connected and such that $|\partial B_i \cap \partial B_j| = 2$ whenever $i \neq j$. If $B_1 \cap B_2 \cap B_3$ and $\Sigma \setminus (B_1 \cup B_2 \cup B_3)$ are both nonempty, then they are also both connected with connected boundary. This is a simple exercise on noting that a regular neighbourhood of $\partial B_1 \cup \partial B_2 \cup \partial B_3$ is an $S_{0,8}$.

Given this, we can now recognise a surrounding triple as a triple α, β, γ with each pair forming a surrounding pair and such that there is a fourth curve δ disjoint from α, β, γ and such that α, β, γ all lie in an $S_{0,6}$ component of $\Sigma \setminus \delta$. Note that this implies that $B(\alpha) \cap B(\beta) \cap B(\gamma)$ must be nonempty (it must contain a boundary component of Σ). It now follows easily from the previous paragraph that $B(\alpha) \cap B(\beta) \cap B(\gamma)$ is in fact an $S_{0,3}$, and so we can set ω to be its relative boundary in Σ . \square

Proof of Theorem 1.1 when $p(\Sigma) \geq 5$. The remainder of the proof now follows exactly as for $S_{0,7}$. We define the graphs \mathcal{H} and $\mathcal{H}(\omega)$ in the same way. It is sufficient to show that $\mathcal{H}(\omega)$ is connected. The argument follows exactly as with that of Lemma 4.2. We can define an ω -arc to be an arc connecting $B(\omega)$ to $\partial \Sigma \setminus B(\omega)$. Taking a regular neighbourhood of $B(\omega) \cup a \cup \varepsilon$, where ε is the boundary component, we get an $S_{0,4}$, namely $B(\alpha)$, where $\alpha = \partial_\Sigma B(\alpha) \in C_{ss}$, so that $\omega < \alpha$. A surrounding pair corresponds to a disjoint pair of ω -arcs, terminating in distinct boundary components. Similarly, a surrounding triple corresponds to a disjoint triple of ω -arcs to distinct boundary components. We can now copy the argument of Lemma 4.2 (after collapsing each component of $\partial \Sigma \setminus B(\omega)$ to a point). This allows us to reconstruct $\mathcal{G}_s(\Sigma)$, and so, using [Ko; BrM; Ki], we see that $\mathcal{G}_{ss}(\Sigma)$ is rigid in these cases. \square

We now move on to the cases where $p(\Sigma) \leq 4$. We can assume that $p(\Sigma) \geq 2$ (otherwise, $\mathcal{G}_{ss}(\Sigma) = \mathcal{G}_s(\Sigma)$, and we are covered by [Ko; BrM; Ki]). Note that $g(\Sigma) \geq 3$ in these cases.

We will use the following construction. Let $\mathcal{T}(\Sigma)$ be the graph whose vertex set consists of those $\alpha \in C_1(\Sigma)$ for which $B(\alpha)$ is an $S_{1,1}$ and where α, β are deemed adjacent if $B(\alpha) \cap B(\beta) = \emptyset$. (This is a full subcomplex of $\mathcal{G}_{ss}(\Sigma)$.) We note the following:

LEMMA 6.4. *If $g(\Sigma) \geq 3$, then $\mathcal{T}(\Sigma)$ is connected.*

Proof. Suppose α, β are vertices of $\mathcal{T}(\Sigma)$. Since the separating curve graph of a closed surface of genus at least 3 is connected, we can connect α, β by a vertex path $\alpha = \gamma_0, \gamma_1, \dots, \gamma_n = \beta$ in $\mathcal{G}_s(\Sigma)$, so that no complementary component of any γ_i is planar. (To see this, just cap off the boundary components of Σ by discs.) Taking n to be minimal, we see that γ_{i-1} must intersect γ_{i+1} for all $i \neq 0, n$. Let δ_i be a vertex of $\mathcal{T}(\Sigma)$ contained in the component of $\Sigma \setminus \gamma_i$ not containing $\gamma_{i-1}, \gamma_{i+1}$. We see that $\alpha, \delta_1, \dots, \delta_{n-1}, \beta$ is a path in $\mathcal{T}(\Sigma)$ connecting α to β . □

Let $\hat{\mathcal{T}}(\Sigma)$ be the flag simplicial complex with 1-skeleton $\mathcal{T}(\Sigma)$, so that every complete subgraph of $\mathcal{T}(\Sigma)$ is contained in a simplex of $\hat{\mathcal{T}}(\Sigma)$. Given $n \geq 1$, let $\mathcal{S}_n(\Sigma)$ be the graph whose vertex set consists of n -simplices of $\hat{\mathcal{T}}(\Sigma)$ and where two such simplices are deemed adjacent if they have a common $(n - 1)$ -face.

LEMMA 6.5. *If $g(\Sigma) \geq n + 2$, then $\mathcal{S}_n(\Sigma)$ is connected.*

Proof. If $0 \leq m \leq n - 2$, then the link of any m -simplex in $\hat{\mathcal{T}}(\Sigma)$ is isomorphic to $\mathcal{T}(\Sigma')$, where Σ' is obtained by removing $m + 1$ disjoint copies of $S_{1,1}$ from Σ . Thus, $g(\Sigma') = g(\Sigma) - m - 1 \geq 3$, and so this is connected by Lemma 6.4. Since $\hat{\mathcal{T}}(\Sigma)$ is itself connected, the statement now follows easily. □

We now consider the case where $g(\Sigma) \geq 5$. (This will cover all cases with $p(\Sigma) \leq 4$, except $S_{3,4}, S_{4,3}$, and $S_{4,4}$, which we discuss later.) For this, we need to modify the definition of a surrounding pair.

DEFINITION. (When $p(\Sigma) \leq 4$ and $g(\Sigma) \geq 5$.) A *surrounding pair* is a pair of curves $\alpha, \beta \in C_{ss}(\Sigma)$ such that $B(\alpha), B(\beta)$ are both $S_{1,3}$ and such that $B(\alpha) \cap B(\beta)$ is an $S_{0,3}$.

This implies that $|\partial B(\alpha) \cap \partial B(\beta)| = 2$. Again, there is a unique $\omega \in C_0(\Sigma)$ with $\omega < \alpha, \beta$, namely $\omega = \partial(B(\alpha) \cap B(\beta))$. This property is also recognisable:

LEMMA 6.6. *If $p(\Sigma) \leq 4$ and $g(\Sigma) \geq 5$, then the collection of surrounding pairs is invariant under any automorphism of $\mathcal{G}_{ss}(\Sigma)$.*

Proof. The argument follows exactly as with Lemma 4.2. The criterion of Lemma 6.1 still holds, except that $B(\alpha), B(\beta)$ are now both $S_{1,3}$ instead of $S_{0,4}$. Also, in verifying (S1)–(S3), the multicurve τ will include curves in $B(\alpha)$ and $B(\beta)$ that bound $S_{1,1}$. □

We also define a *surrounding triple* the same way as before (modulo the definition of a surrounding pair). Given Lemma 6.6, we see easily that α, β, γ form a surrounding triple if and only if they pairwise form surrounding pairs and if there is some $\delta \in C_{ss}(\Sigma)$ such that α, β, γ all lie in an $S_{3,3}$ component of $\Sigma \setminus \delta$.

Proof of Theorem 1.1 when $p(\Sigma) \leq 4$ and $g(\Sigma) \geq 5$. We now proceed, as usual, to define the graphs \mathcal{H} and $\mathcal{H}(\omega)$ for $\omega \in C_0$. We claim that $\mathcal{H}(\omega)$ is connected. This is a bit more involved in this case.

Suppose that $\alpha \in C_{ss}(\Sigma)$ with $B(\alpha)$ an $S_{1,3}$ and with $B(\omega) \subseteq B(\alpha)$. Let $\varepsilon \in C_{ss}(\Sigma)$ be any curve with $B(\varepsilon)$ an $S_{1,1}$ with $B(\varepsilon) \subseteq B(\alpha) \setminus B(\omega)$. Thus, $B(\alpha) \setminus (B(\omega) \cup B(\varepsilon))$ is an $S_{0,3}$, so there is (up to homotopy) a unique arc a in $B(\alpha)$ from $B(\omega)$ to $B(\varepsilon)$, meeting $B(\omega)$ and $B(\varepsilon)$ precisely at its endpoints. Note that $B(\alpha)$ is a regular neighbourhood of $B(\omega) \cup a \cup B(\varepsilon)$. Conversely, given any $\varepsilon \in C_{ss}(\Sigma)$ with $B(\varepsilon)$ an $S_{1,1}$ disjoint from $B(\omega)$, we can obtain such an α as the boundary of a regular neighbourhood of $B(\omega) \cup a \cup B(\varepsilon)$. (Of course, such a representation of α is not unique, but that will not matter.)

Now, a vertex of $\mathcal{H}(\omega)$ arises from a pair of disjoint such curves ε, η and disjoint arcs a, b connecting $B(\omega)$ respectively to $B(\varepsilon)$ and $B(\eta)$ in $\Sigma \setminus (B(\omega) \cup B(\varepsilon) \cup B(\eta))$. Similarly, a surrounding triple arises from three disjoint such curves ε, η, ζ and three disjoint arcs a, b, c in the complement of $B(\omega) \cup B(\varepsilon) \cup B(\eta) \cup B(\zeta)$.

Suppose we fix ε, η, ζ , and let $\mathcal{H}(\omega; \varepsilon, \eta, \zeta)$ be the full subgraph of $\mathcal{H}(\omega)$, where all the vertices arise (as before) from some pair of curves in $\{\varepsilon, \eta, \zeta\}$. Now $\mathcal{H}(\omega; \varepsilon, \eta, \zeta)$ is connected. This can be seen by the same argument as in the previous case (applied to the surface $\Sigma \setminus (B(\varepsilon) \cup B(\eta) \cup B(\zeta))$, where the notion of “surrounding pair” reverts to the previous case, as defined in Section 2).

Now if $\varepsilon, \eta, \zeta, \theta$ are all disjoint such curves, we can connect $B(\omega)$ to each of $B(\varepsilon), B(\eta), B(\zeta), B(\theta)$ by disjoint arcs in the complement of $B(\omega) \cup B(\varepsilon) \cup B(\eta) \cup B(\zeta) \cup B(\theta)$. It then follows that $\mathcal{H}(\omega; \varepsilon, \eta, \zeta) \cap \mathcal{H}(\omega; \varepsilon, \eta, \theta) \neq \emptyset$. But now, by Lemma 6.5 we can get between any two triples ε, η, ζ and $\varepsilon', \eta', \zeta'$ by a sequence of such moves, replacing one curve at a time. Since $\mathcal{H}(\omega)$ is a union of such $\mathcal{H}(\omega; \varepsilon, \eta, \zeta)$, it follows that $\mathcal{H}(\omega)$ is connected as claimed.

This allows us to construct $\mathcal{G}_s(\Sigma)$ out of $\mathcal{G}_{ss}(\Sigma)$, and so rigidity follows as before. □

Finally, we should discuss the cases where $\Sigma \in \{S_{3,4}, S_{4,3}, S_{4,4}\}$. Again, we need to redefine a “surrounding pair”:

DEFINITION. (When $\Sigma \in \{S_{3,4}, S_{4,3}, S_{4,4}\}$.) A *surrounding pair* is a pair $\alpha, \beta \in C_{ss}(\Sigma)$ with $B(\alpha), B(\beta)$, each either an $S_{0,4}$ or an $S_{1,3}$.

Given this, we define “surrounding triple” as usual.

Note that $B(\alpha)$ is now determined by a curve which is either a component of $\partial\Sigma \setminus B(\omega)$ or a curve $\varepsilon \in C_{ss}(\Sigma)$ with $B(\varepsilon)$ an $S_{1,1}$ disjoint from $B(\omega)$, together (in either case) with an arc a from $B(\omega)$ to ε . Thus, $B(\alpha)$ is a regular neighbourhood of $B(\omega) \cup a \cup \varepsilon$ or of $B(\omega) \cup a \cup B(\varepsilon)$, respectively. Surrounding pairs and triples then arise from disjoint curves and arcs similarly as before.

We can define $\mathcal{H}(\omega; \varepsilon, \eta, \zeta)$ similarly as in the previous case, where ε, η, ζ are disjoint curves, each either a boundary curve or in $C_{ss}(\Sigma)$ as before. The same argument shows that the graph is connected.

Moreover, we claim that we can get between any two such triples ε, η, ζ and $\varepsilon', \eta', \zeta'$, replacing each curve at time by a disjoint curve. In the case where $p(\Sigma) \geq 4$, there are at least two components of $\partial\Sigma \setminus B(\omega)$, so this follows easily applying Lemma 6.4 to $\Sigma \setminus B(\omega)$. If $p(\Sigma) = 3$, then $g(\Sigma \setminus B(\omega)) = g(\Sigma) \geq 4$, and so the statement follows since $\mathcal{S}_2(\Sigma \setminus B(\omega))$ is connected by Lemma 6.5.

The argument can now be completed as before, proving Theorem 1.1 in all cases except $S_{0,8}$.

7. The 8-Holed Sphere

In this section, we outline the proof of rigidity for $\mathcal{G}_{ss}(S_{0,8})$. The argument is essentially the same as for $S_{0,7}$, except we need to start by finding a different rigid subgraph in order to recognise pairs of curves that intersect exactly twice. We will revert to thinking of $S_{0,8}$ as $S \setminus \Pi$, where Π is a subset of the 2-sphere S with $|\Pi| = 8$.

Let Δ be the graph obtained by adding the four longest diagonal edges to an octagon. More formally, we write $V(\Delta) = \{v_1, \dots, v_8\}$, where v_i is deemed adjacent to v_j whenever $|i - j| = 1$ or $|i - j| = 4$ (taking indices mod 8).

We can realise Δ as follows. Recall from Section 3 that $\Theta(\Psi)$ is the graph whose vertex set consists 3-sets in Ψ and where two such sets are deemed adjacent if they are disjoint subsets of Ψ . If $\Psi = \{1, \dots, 8\}$ and $P_i = \{i - 3, i, i + 3\}$, then the full subgraph of $\Theta(\Psi)$ with vertex set $\{P_1, \dots, P_8\}$ is isomorphic to Δ .

In fact, we claim that all copies of Δ in $\Theta(\Psi)$ arise in this way.

First note that there are no 3-cycles in $\Theta(\Psi)$, and so any 5-cycle in $\Theta(\Psi)$ is isometrically embedded. Now any two vertices of Δ lie in a 5-cycle in Δ . It follows that any map of Δ into $\Theta(\Psi)$ sending edges to edges must be injective.

Suppose that $[v_i \mapsto P_i]$ is any embedding of Δ into $\Theta(\Psi)$. Now v_1, v_2, v_6, v_5 is a 4-cycle in Δ , and so P_1, P_2, P_6, P_5 is a 4-cycle in $\Theta(\Psi)$. Let $P_{16} = P_1 \cup P_6$ and $P_{52} = P_5 \cup P_2$. Then $P_{16} \cap P_{52} = \emptyset$, and so $|P_{16}| = |P_{52}| = 4$. In other words, $\{P_{16}, P_{52}\}$ is a partition of Ψ into two 4-sets. The same holds for $\{P_{63}, P_{27}\}$, $\{P_{38}, P_{74}\}$, and $\{P_{85}, P_{41}\}$. Now, $P_{41} \neq P_{16}$. (Otherwise, we would have $P_{41} = P_{16}$ and $P_{85} = P_{52}$, so $P_1 \cap P_2 \subseteq P_{41} \cap P_{52} = \emptyset$. But also $P_3 \cap P_2 = P_3 \cap P_1 = \emptyset$, giving a contradiction.) Similarly, since the P_i are all distinct, we easily see that $\{P_{16}, P_{52}\} \neq \{P_{38}, P_{74}\}$. It follows that $\{P_{16} \cap P_{38}, P_{16} \cap P_{74}, P_{52} \cap P_{38}, P_{52} \cap P_{74}\}$ is a partition of Ψ into four 2-sets. The same holds for $\{P_{63} \cap P_{85}, P_{63} \cap P_{41}, P_{27} \cap P_{85}, P_{27} \cap P_{41}\}$. From this information we can easily find a permutation of Ψ so that each $P_i = \{i - 3, i, i + 3\}$ as in our example. This shows the following:

LEMMA 7.1. *There is exactly one embedded copy of Δ in $\Theta(\Psi)$ up to the action of $\text{Sym}(\Psi)$.*

We now move on to consider $\mathcal{G}_{ss} = \mathcal{G}_{ss}(S_{0,8})$. We set $\Pi = \Psi$. If $\gamma \in C_1 = C_1(S_{0,8})$, then write $B(\gamma)$ for the disc with $|\Pi \cap B(\gamma)| = 3$. Let $\pi(\gamma) = \Pi \cap B(\gamma)$. Thus, π maps $\mathcal{G}(S_{0,8}, C_1)$ to $\Theta(\Pi)$ sending edges to edges.

We note that $\mathcal{G}(S_{0,8}, C_1) \subseteq \mathcal{G}_{S,S}(\Sigma)$ contains an embedded copy of Δ constructed as follows. Let $\lambda \subseteq S$ be an embedded circle with $\Pi \subseteq \lambda$. We label the elements of Π so that the cyclic order induced from λ is given by $p_1, p_4, p_7, p_2, p_5, p_8, p_3, p_6$. Let $l_{i,i+3} \subseteq \lambda$ be the segment from p_i to p_{i+3} . Let B_i be a regular neighbourhood of the arc $l_{i-3,i} \cup l_{i,i+3}$, and let $\gamma_i = \partial B_i$. The map $[v_i \mapsto \gamma_i]$ now gives an embedding of Δ into $\mathcal{G}(S_{0,8}, C_1)$.

We claim the following:

LEMMA 7.2. *There is exactly one embedded copy of Δ in $\mathcal{G}(S_{0,8}, C_1)$ up to the action of $\text{Map}(S_{0,8})$.*

Let $\Delta \subseteq \mathcal{G}(S_{0,8}, C_1)$ be such a copy. After composing with π , we get a map of Δ into $\Theta(\Pi)$ sending edges to edges, which as we have already noted, must also be an embedding. By Lemma 7.1 we can now label the elements of Π as $\{p_1, \dots, p_8\}$ so that $\pi(\gamma_i) = \{p_{i-3}, p_i, p_{i+3}\}$ for all i . (In other words, as in the given example.) Write $B_i = B(\gamma_i)$.

Now $\gamma_1, \gamma_2, \gamma_6, \gamma_5$ is a square in Δ , and so $(B_1 \cup B_6) \cap (B_5 \cup B_2) = \emptyset$. Also, $|(B_1 \cup B_6) \cap \Pi| = |(B_5 \cup B_2) \cap \Pi| = 4$. It follows that we can find disjoint discs $B_{16}, B_{52} \subseteq S$ with $B_1 \cup B_6 \subseteq B_{16}$ and $B_5 \cup B_2 \subseteq B_{52}$. Note that ∂B_{16} and ∂B_{52} are homotopic in $S \setminus \Pi$ and so determine a curve $\gamma_{16} = \gamma_{52}$. We have similar pairs of discs $\{B_{63}, B_{27}\}, \{B_{38}, B_{74}\}$, and $\{B_{85}, B_{41}\}$. Note that we have similar curves γ_{ij} defined whenever $|i - j|$ is 3 or 5.

Now consider the curves $\gamma_1, \gamma_2, \gamma_3, \gamma_4$. Let $A = S \setminus (B_1 \cup B_2)$. This is an annulus with $A \cap \Pi = \{p_3, p_8\}$. Now $B_3 \cap B_2 = \emptyset$, $B_3 \cap B_1 \neq \emptyset$ (since $p_6 \in B_3 \cap B_1$), and $p_3, p_8 \in B_3 \cap A$. Thus, we can find (disjoint) arcs e, f in $B_3 \cap A$, respectively connecting p_3 and p_8 to ∂B_1 .

Let D be (the closure of) the component of $B_4 \setminus B_1$ containing p_7 . Now $B_4 \cup B_1 \subseteq B_{41}$ and $(B_{41} \setminus B_4) \cap \Pi = \{p_7\}$. It follows (cf. Lemma 3.4) that $D \cap B_1$ is a single arc, say, b . Thus, $\partial D = b \cup d$, where $d \subseteq \gamma_4$ is an arc with endpoints $d \cap B_1 = d \cap b$.

Now, $d \cap (e \cup f) \subseteq B_4 \cap B_3 = \emptyset$. It follows that any arc of $d \cap A$ that does not include an endpoint could be homotoped into B_2 . It follows that $d \cap B_2$ must consist of a single arc, and so $|d \cap B_2| = 2$. Now $R = B_1 \cup D$ is a disc with $R \cap \Pi = B_{41} \cap \Pi$, and so ∂R is homotopic in $S \setminus \Pi$ to $\gamma_{41} = \gamma_{52}$. Also, $|\partial R \cap \gamma_2| = 2$, and so $\iota(\gamma_{41}, \gamma_2) = 2$.

By symmetry it follows that each curve γ_i intersects each curve of the form γ_{jk} either 0 or 2 times.

It is now fairly straightforward to see that $\{\gamma_1, \dots, \gamma_8\}$ must be precisely the set of curves described in our example. For example, note that $B_{41} \cap \Pi = \{p_7, p_4, p_1, p_6\}$ and that $B_{41} \cap \gamma$ is an arc that cuts off a disc, namely $B_{41} \cap B_2$, containing p_7 . Now $B_1 \cap \Pi = \{p_4, p_1, p_6\}$, and so $\gamma_1 = \partial B_1$ is homotopic to $\partial(B_{41} \setminus B_2)$. Similarly, γ_4 is homotopic to $\partial(B_{41} \setminus B_3)$. It follows that $B_1 \cap B_4$ is a disc containing p_1, p_4 . We can now proceed to show that any pair of the curves γ_i intersect at most twice and that this determines them completely.

This proves Lemma 7.2.

The remainder of the proof of rigidity is now essentially identical to that for $S_{0,7}$. We define surrounding pairs in the same way. A surrounding pair α, β (as defined in Section 2) can be recognised by the fact that $\alpha, \beta \in C_1(S_{0,8})$, there is a curve $\gamma \in C_{ss}(S_{0,8}) \setminus C_1(S_{0,8})$ disjoint from both α and β , and α, β are vertices of some embedded copy of Δ in $\mathcal{G}(S_{0,8}, C_1)$. (Note that $\mathcal{G}(S_{0,8}, C_1)$ can be constructed from $\mathcal{G}_{ss}(S_{0,8})$.) We now proceed as before.

This shows that $\mathcal{G}_{ss}(S_{0,8})$ is rigid, completing the proof of Theorem 1.1.

References

- [Bo1] B. H. Bowditch, *Large-scale rank and rigidity of the Teichmüller metric*, preprint, Tokyo/Warwick, 2015.
- [Bo2] ———, *Large-scale rank and rigidity of the Weil–Petersson metric*, preprint, 2015, *J. Topol.* (to appear).
- [BrM] T. E. Brendle and D. Margalit, *Commensurations of the Johnson kernel*, *Geom. Topol.* 8 (2004), 1361–1384.
- [EMR] A. Eskin, H. Masur, and K. Rafi, *Rigidity of Teichmüller space*, preprint, 2015, [arXiv:1506.04774](https://arxiv.org/abs/1506.04774).
- [H] W. J. Harvey, *Boundary structure of the modular group*, Riemann surfaces and related topics: proceedings of the 1978 Stony Brook conference (I. Kra, B. Maskit eds.), *Ann. of Math. Stud.*, 97, pp. 245–251, Princeton University Press, 1981.
- [Ir] E. Irmak, *Complexes of nonseparating curves and mapping class groups*, *Michigan Math. J.* 54 (2006), 81–110.
- [Iv] N. V. Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces*, *Int. Math. Res. Not. IMRN* 14 (1997), 651–666.
- [Ki] Y. Kida, *Automorphisms of the Torelli complex and the complex of separating curves*, *J. Math. Soc. Japan* 62 (2011), 363–417.
- [Ko] M. Korkmaz, *Automorphisms of complexes of curves on punctured spheres and on punctured tori*, *Topology Appl.* 95 (1999), 85–111.
- [L] F. Luo, *Automorphisms of the complex of curves*, *Topology* 39 (2000), 283–298.

Mathematics Institute
University of Warwick
Coventry, CV4 7AL
Great Britain