

Effective finite generation for $[\mathrm{IA}_n, \mathrm{IA}_n]$ and the Johnson kernel

Mikhail Ershov and Daniel Franz

Abstract. Let IA_n denote the group of IA -automorphisms of a free group of rank n , and let \mathcal{I}_n^b denote the Torelli subgroup of the mapping class group of an orientable surface of genus n with b boundary components, $b = 0, 1$. In 1935, Magnus proved that IA_n is finitely generated for all n , and in 1983, Johnson proved that \mathcal{I}_n^b is finitely generated for $n \geq 3$. It was recently shown that for each $k \in \mathbb{N}$, the k -th terms of the lower central series $\gamma_k \mathrm{IA}_n$ and $\gamma_k \mathcal{I}_n^b$ are finitely generated when $n \gg k$; however, no information about finite generating sets was known for $k > 1$. The main goal of this paper is to construct an explicit finite generating set for $\gamma_2 \mathrm{IA}_n = [\mathrm{IA}_n, \mathrm{IA}_n]$ and almost explicit finite generating sets for $\gamma_2 \mathcal{I}_n^b$ and the Johnson kernel, which contains $\gamma_2 \mathcal{I}_n^b$ as a finite index subgroup.

1. Introduction

1.1. Discussion of the problem

Given non-negative integers n and b , let Σ_n^b be an orientable surface of genus n with b boundary components, and let $\mathrm{Mod}_n^b = \mathrm{Mod}(\Sigma_n^b)$ be its mapping class group. The corresponding Torelli group \mathcal{I}_n^b is the subgroup of Mod_n^b consisting of elements acting trivially on $H_1(\Sigma_n^b, \mathbb{Z})$. In this paper, we will only consider the cases $b = 0, 1$. Given $n \in \mathbb{N}$, let F_n denote a free group on n generators, and let IA_n be the subgroup of $\mathrm{Aut}(F_n)$ consisting of automorphisms acting trivially on the abelianization $F_n^{ab} = F_n/[F_n, F_n]$. The group IA_n is often called the Torelli subgroup of $\mathrm{Aut}(F_n)$ and is known to behave similarly to \mathcal{I}_n^1 in many ways.

In 1935, Magnus [18] proved that IA_n is finitely generated for all $n \geq 2$; in fact, he found an explicit and simple-to-describe generating set of smallest possible cardinality $n \binom{n}{2}$. In 1983, Johnson [14] proved that \mathcal{I}_n^b is finitely generated for $n \geq 3$; his generating set is also explicit and also of optimal size for $n = 3$, but of considerably larger size in general (growing exponentially in n). More recently, Putman [22] found a smaller

generating set whose size grows cubically with n , which is known to be asymptotically optimal.¹

It was a very interesting question whether the commutator subgroup $[G, G]$ is finitely generated for $G = \text{IA}_n$ or \mathcal{I}_n^b . In both cases, G has large abelianization, so there was no a priori reason to expect $[G, G]$ to be finitely generated. On the other hand, G possesses a generating set in which many pairs of generators commute, which can be seen as positive evidence for finite generation of $[G, G]$. An additional motivation for the finite generation question in the mapping class group case is given by the fact that $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ is a finite index subgroup of the Johnson kernel \mathcal{K}_n^b , a group of major interest in topology.

In [9], it was proved that the commutator subgroups of the Torelli groups are indeed finitely generated in sufficiently large rank: $[\text{IA}_n, \text{IA}_n]$ for $n \geq 4$ and $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ for $n \geq 12$. In [5], $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ was shown to be finitely generated for all $n \geq 4$, and finite generation was also extended to some higher terms of the lower central series, namely $\gamma_k \text{IA}_n$ for $n \geq 4k - 3$ and $\gamma_k \mathcal{I}_n^b$ for $n \geq 2k + 1$ (see Section 1.4 for an additional discussion). However, neither [9] nor [5] dealt directly with the finite generation question, instead reducing the problem to an analysis of Bieri–Neumann–Strebel (BNS) invariants.

Given a finitely generated group G , its character sphere $\mathbb{S}(G)$ is the set of nonzero homomorphisms from G to $(\mathbb{R}, +)$ modulo the equivalence given by multiplication by positive scalars. The BNS-invariant of G , introduced by Bieri, Neumann and Strebel in [2] and denoted by $\Sigma(G)$, is a subset of $\mathbb{S}(G)$ which determines which subgroups of G containing $[G, G]$ are finitely generated (see Theorem 2.1, often called the BNS criterion). In particular, $[G, G]$ itself is finitely generated if and only if $\Sigma(G) = \mathbb{S}(G)$, and it is the latter equality that was established in [5, 9] for $G = \text{IA}_n$ and \mathcal{I}_n^b for $n \geq 4$. Finite generation of higher terms of the lower central series was established in [5] by an inductive application of the BNS criterion.

Since the proof of the BNS criterion is not effective, [5, 9] did not yield an actual construction of finite generating sets for $[G, G]$ (or higher terms) for $G = \text{IA}_n$ and \mathcal{I}_n^b . The main goal of the present paper is to give an effective proof of finite generation for $[G, G]$ when $n \geq 8$. In Section 1.4, we will discuss the main obstacle to extending this method to $\gamma_k G$ for $k > 2$.

The proof of finite generation that we will provide does not make a formal reference to the BNS invariant; however, it relies on the proof of the BNS criterion in a substantial way. Essentially, we follow the proof of the BNS criterion given in [24] (which is considerably simpler than the original argument from [2]), replace all the non-effective steps with explicit constructions and make some simplifications which are not possible in general. In addition to providing explicit generating sets for $[G, G]$, the proof in this paper is algorithmic in the following sense: given a sufficiently nice generating set S of G and an

¹The smallest size of a generating set for IA_n is indeed $n \binom{n}{2}$ since its abelianization IA_n^{ab} is free of rank $n \binom{n}{2}$; see Section 3 for details. Likewise, the asymptotic optimality of the generating set from Putman [22] follows from the fact that the torsion-free rank of $(\mathcal{I}_n^b)^{\text{ab}}$ is cubic in n which was proved by Johnson [13, 15]; see Section 4 for details.

element $g \in [G, G]$ expressed in terms of S , our proof yields a procedure for writing g in terms of a finite generating set for $[G, G]$ (which is explicitly derived from S). Our general method for proving effective finite generation of $[G, G]$ that will be developed in Section 2 is sufficiently flexible and could be applicable to other groups.

1.2. Main results

We proceed with stating our main result for IA_n . Throughout the paper, for a group G and elements $x, y \in G$, we set $x^y = y^{-1}xy$ and $[x, y] = x^{-1}y^{-1}xy$.

Theorem 1.1. *Let $n \geq 8$. Let $N = n \binom{n}{2}$, and let $S = \{s_1, \dots, s_N\}$ be the standard generating set for IA_n constructed by Magnus (see the beginning of Section 3). Then $[\mathrm{IA}_n, \mathrm{IA}_n]$ is generated by elements of the form*

$$[s_i, s_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_N^{a_N}}, \quad \text{where } 1 \leq i < j \leq N \text{ and } 0 \leq |a_m| < 5 \cdot 10^{12} \text{ for each } m.$$

In particular, the minimal number of generators of $[\mathrm{IA}_n, \mathrm{IA}_n]$ is at most $n \binom{n}{2} \cdot (10^{13})^{n \binom{n}{2}}$.

Remark. (1) It was previously known [26]² that if G is any group generated by a finite set $\{x_1, \dots, x_k\}$, then $[G, G]$ is generated by elements of the form $[x_i, x_j]^{x_i^{a_i} x_{i+1}^{a_{i+1}} \dots x_k^{a_k}}$ with $1 \leq i < j \leq k$ and $a_m \in \mathbb{Z}$. Thus, the novel part of Theorem 1.1 is that in the case $G = \mathrm{IA}_n$ it suffices to take only such elements where a_m are bounded by an explicit constant independent of n .

(2) In Section 3, we will show that the number of generators of $[\mathrm{IA}_n, \mathrm{IA}_n]$ is actually bounded by a function of the form C^{n^2} (see Theorem 3.5) which is slightly better than a bound of the form C^{n^3} given by Theorem 1.1.

(3) Our method of proof is in principle applicable to all $n \geq 4$, but would yield a constant larger than $5 \cdot 10^{12}$ for $n = 6, 7$ and a much larger constant for $n = 4, 5$.

We now turn to the mapping class groups. In this case, we will describe generating sets for two different subgroups of \mathcal{I}_n^b – the commutator subgroup $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ and the Johnson kernel \mathcal{K}_n^b . One can define the Johnson kernel algebraically, as the second term of the Johnson filtration of \mathcal{I}_n^b (see Section 4.1 for the definition of the Johnson filtration in the case $b = 1$; the definition in the case $b = 0$ is similar) or topologically, as the subgroup of Mod_n^b generated by Dehn twists about separating curves. It also follows from work of Johnson that $\mathcal{I}_n^b / \mathcal{K}_n^b$ is the largest quotient of \mathcal{I}_n^b which is abelian and torsion-free. This, together with finite generation of \mathcal{I}_n^b , immediately implies that \mathcal{K}_n^b contains $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ as a finite-index subgroup for $n \geq 3$.

For simplicity, we state Theorems 1.2 and 1.3 below for $b = 1$. It is well known that there is a natural surjective map $\mathrm{Mod}_n^1 \rightarrow \mathrm{Mod}_n^0$ which sends \mathcal{I}_n^1 to \mathcal{I}_n^0 and \mathcal{K}_n^1 to \mathcal{K}_n^0 , so any finite generating set of \mathcal{K}_n^1 yields the corresponding finite generating set for \mathcal{K}_n^0 .

²It suffices to prove this result when G is free, in which case the given generating set for $[G, G]$ freely generates $[G, G]$. The latter is the main result of [26]; recently, a more geometric proof was given in [23].

Theorem 1.2. *Let $n \geq 4$, let $G = \mathcal{I}_n^1$ and $K = \mathcal{K}_n^1$. Let $N = \binom{2n}{3}$ and $M = \binom{2n}{2} + \binom{2n}{1}$. The following hold:*

- (1) *G has a generating set $S^{(1)} \sqcup S^{(2)} \sqcup S^{(3)}$ with the following properties:*
 - (i) *$|S^{(1)}| = N$ and the elements s_1, \dots, s_N of $S^{(1)}$ project to a basis of G/K (which is free abelian of rank N).*
 - (ii) *$|S^{(2)}| = M$, the elements t_1, \dots, t_M of $S^{(2)}$ lie in K , and $S^{(2)}$ projects to a basis of $K/[G, G]$ (which is a vector space over \mathbb{F}_2 of dimension M).*
 - (iii) *$S^{(3)}$ is contained in $[G, G]$ and $|S^{(3)}| = 42\binom{n}{3} - |S^{(1)}| - |S^{(2)}|$.*
- (2) *There exists an absolute constant R such that K is generated by elements of the form*
 - (a) $[s_i, s_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_N^{a_N}}$, where $1 \leq i < j \leq N$ and $|a_m| \leq R$ for all m ;
 - (b) $x^{s_1^{a_1} s_2^{a_2} \dots s_N^{a_N}}$, where $x \in S^{(2)} \cup S^{(3)}$ and $|a_m| \leq R$ for all m .
- (3) *$[G, G]$ is generated by elements of the form*
 - (a) $[s_i, s_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_N^{a_N} t_1^{\varepsilon_1} \dots t_M^{\varepsilon_M}}$ with $1 \leq i < j \leq N$;
 - (b) $[s_i, t_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_N^{a_N} t_1^{\varepsilon_1} \dots t_M^{\varepsilon_M}}$ with $1 \leq i \leq N$ and $1 \leq j \leq M$;
 - (c) $[t_i, t_j]^{t_i^{\varepsilon_i} t_{i+1}^{\varepsilon_{i+1}} \dots t_M^{\varepsilon_M}}$ with $1 \leq i < j \leq M$;
 - (d) $x^{s_1^{a_1} s_2^{a_2} \dots s_N^{a_N} t_1^{\varepsilon_1} \dots t_M^{\varepsilon_M}}$ with $x \in S^{(3)}$,

where in each part $|a_m| \leq R$, $\varepsilon_m \in \{0, 1\}$ for all m and R is the same as in (2).

Remark. We will construct explicit generating sets for $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ and \mathcal{K}_n^b only for $n \geq 8$. Theorem 1.2 is still valid for all $n \geq 4$ since $[\mathcal{I}_n^b, \mathcal{I}_n^b]$ and \mathcal{K}_n^b are known to be finitely generated for all $n \geq 4$ by [5], and we are not making any assertions about the constant R above.

The problem of explicitly estimating R , at least for $n \geq 8$, reduces to a certain computation in the Torelli group \mathcal{I}_5^1 . We did not compute a precise upper bound, but we believe that this can be achieved by carefully examining the proofs of Johnson [14] and Stylianakis [25] (see the end of Section 4.8 for a detailed discussion).

1.3. Generating the Johnson kernel by finitely many Dehn twists

One drawback of the generating sets from Theorem 1.2 is that they do not seem to have any natural geometric or topological interpretation. We will now address this issue in the case of the Johnson kernel \mathcal{K}_n^1 . Recall that \mathcal{K}_n^1 is generated by the Dehn twists about separating curves. Thus, one way to produce an explicit geometrically meaningful finite generating set for \mathcal{K}_n^1 is to show that \mathcal{K}_n^1 is generated by the Dehn twists about separating curves of explicitly bounded word length.

Let us now make our task more precise. Fix a point p_0 on the boundary of $\Sigma = \Sigma_n^1$. The fundamental group $\pi_1(\Sigma, p_0)$ is free of rank $2n$ and admits a basis $\alpha_1, \beta_1, \dots, \alpha_n, \beta_n$ such that $\prod_{i=1}^n [\alpha_i, \beta_i]$ is represented by $\partial\Sigma$; below we will refer to such a basis as *natural*.

Fix a natural basis ω of $\pi_1(\Sigma, p_0)$. Given $m \in \mathbb{N}$, let $\mathrm{SC}(m)$ be the set of all elements of $\pi_1(\Sigma, p_0)$ which have word length at most m with respect to ω and which are represented by a separating simple curve on Σ . Let $T_{\mathrm{sc}}(m) \subset \mathrm{Mod}(\Sigma_n^1)$ be the set of Dehn twists about the elements of $\mathrm{SC}(m)$. As we will explain in Section 4, any two natural bases of $\pi_1(\Sigma, p_0)$ lie in the same orbit under the action of $\mathrm{Mod}(\Sigma, p_0)$ on $\pi_1(\Sigma, p_0)$ (see the remark after Theorem 4.1). This easily implies that $T_{\mathrm{sc}}(m)$ is independent of ω up to conjugation; in particular, the smallest m for which $T_{\mathrm{sc}}(m)$ generates \mathcal{K}_n^1 does not depend on the choice of ω .

We can now formulate our theorem describing an explicit finite generating set for \mathcal{K}_n^1 consisting of Dehn twists.

Theorem 1.3. *Assume that $n \geq 4$. Then there exists an absolute constant D such that \mathcal{K}_n^1 is generated by the set $T_{\mathrm{sc}}(D^{n^3})$ defined above.*

Theorem 1.3 will be obtained as a relatively easy consequence of Theorem 1.2 and some auxiliary results established in Section 4. The constant D in Theorem 1.3 can be expressed in terms of the constant R from Theorem 1.2, and two other absolute constants which we believe can be estimated explicitly for $n \geq 8$.

1.4. Some questions and remarks

As we already stated in Section 1.1, finite generation results from [9] were extended to higher terms of the lower central series in [5], where it was shown that $\gamma_k \mathrm{IA}_n$ is finitely generated whenever $n \geq 4k - 3$ and $\gamma_k \mathcal{I}_n^b$ is finitely generated whenever $n \geq 2k + 1$. Thus it is natural to ask if the method of the current paper can also provide explicit generating sets for higher terms. We did not succeed in doing this.

The proof in [5] was ineffective for two reasons: similarly to [9], it exploited the BNS invariant. In addition, a combinatorial calculation from [9] was replaced by an ineffective Zariski density argument in [5]. The latter is not a real obstacle to constructing explicit generating sets, and one can show that algebraic geometry can be eliminated from the proof in [5] at the expense of increasing the lower bound on n in terms of k (for which we are claiming that $\gamma_k G$ is finitely generated), with the new bound being quadratic in k . What does cause a problem is the fact that for $k \geq 2$, $G = \mathrm{IA}_n$ or \mathcal{I}_n^b , very little seems to be known about the torsion in $\gamma_k G / \gamma_{k+1} G$ or the presentation of $\gamma_k G / \gamma_{k+1} G$ by generators and relations (as an abelian group).

We conclude this section with some speculations on the asymptotic growth of the number of generators of $[\mathrm{IA}_n, \mathrm{IA}_n]$ and \mathcal{K}_n^b as $n \rightarrow \infty$. Below, for a group Γ we will denote by $d(\Gamma)$ the minimal number of generators of Γ . The following inequalities are obvious:

$$d(\Gamma) \geq d(\Gamma^{\mathrm{ab}}) = d(H_1(\Gamma, \mathbb{Z})) \geq \dim H_1(\Gamma, \mathbb{Q}).$$

It is known that $\dim H_1(\mathcal{K}_n^b, \mathbb{Q})$ grows polynomially with n , and in fact a precise formula for this dimension for $n \geq 6$ can be immediately extracted from a recent paper of Morita,

Sakasai and Suzuki [20, Theorem 1.4] which, in turn, makes essential use of an earlier work of Dimca, Hain and Papadima [8]. This provides at least some evidence that $d(\mathcal{K}_n^b)$ might grow polynomially as well. We are not aware of analogous results dealing with $[\mathrm{IA}_n, \mathrm{IA}_n]$.

Finally, it is natural to ask if the assertions of Theorems 1.1 and 1.2(2) would remain true if the condition $|a_m| \leq C$ on the exponents is replaced by the much more restrictive condition $\sum |a_m| \leq C$ for some absolute constant C . Clearly, if this stronger version of Theorem 1.1 (resp. Theorem 1.2(2)) holds, it would immediately imply polynomial growth for $d([\mathrm{IA}_n, \mathrm{IA}_n])$ (resp. $d(\mathcal{K}_n^b)$).

2. BNS invariant and effective finite generation

We begin this section with some basic terminology. Let G be a group and let S be a subset of G .

The *Cayley graph* of G with respect to S , denoted by $\mathrm{Cay}(G, S)$, is the graph whose vertex set is G and where $g, h \in G$ are connected by an edge if and only if $h = gs^{\pm 1}$ for some $s \in S$. It is clear that $\mathrm{Cay}(G, S)$ is connected if and only if S generates G .

By an *S -word*, we will mean a formal expression $s_1 \dots s_k$ with $s_i \in S \cup S^{-1}$. Thus each S -word naturally represents an element of G , and every element of G is represented by some S -word if and only if S generates G . For each $g \in G$, there is a natural bijection between S -words representing g and paths in $\mathrm{Cay}(G, S)$ from 1 to g .

If $w = s_1 \dots s_k$ is an S -word representing $g \in G$, by a *prefix* of w we will mean a subword of the form $s_1 \dots s_l$ with $l \leq k$. Thus, geometrically, prefixes of w correspond to initial segments of the corresponding path in $\mathrm{Cay}(G, S)$ from 1 to g .

If S generates G , for each $g \in G$ we denote by $\|g\|_S$ the *word length* of g with respect to S , that is, the smallest $k \in \mathbb{Z}_{\geq 0}$ such that g is represented by an S -word $s_1 \dots s_k$. Geometrically, $\|g\|_S$ is the distance from 1 to g in $\mathrm{Cay}(G, S)$.

2.1. Review of the BNS invariant

We start by recalling the definition of the BNS invariant. By a character of a group G , we will mean a homomorphism from G to the additive group of \mathbb{R} . Two characters χ and χ' will be considered equivalent if they are positive multiples of each other, and the equivalence class of a character χ will be denoted by $[\chi]$. The character sphere $\mathbb{S}(G)$ is the set of equivalence classes of nonzero characters of G .

Assume now that G is generated by a finite set S . Given a character χ of G , denote by $\mathrm{Cay}(G, S)_\chi$ the full subgraph of $\mathrm{Cay}(G, S)$ with vertex set $\{g \in G: \chi(g) \geq 0\}$. Note that $\mathrm{Cay}(G, S)_\chi$ is completely determined by the equivalence class of χ . The BNS invariant of G , denoted by $\Sigma(G)$, is defined by

$$\Sigma(G) = \{[\chi] \in \mathbb{S}(G): \mathrm{Cay}(G, S)_\chi \text{ is connected}\}.$$

It is not hard to show that $\Sigma(G)$ does not depend on the choice of S although this is not obvious from definition.

The following remarkable result was proved by Bieri, Neumann and Strebel in [2].

Theorem 2.1 (BNS criterion). *Let K be a normal subgroup of G such that G/K is abelian. Then K is finitely generated if and only if $\Sigma(G)$ contains $[\chi]$ for every character χ which vanishes on K . In particular, $[G, G]$ is finitely generated if and only if $\Sigma(G) = \mathbb{S}(G)$.*

The original proof of Theorem 2.1 given in [2] was quite involved. A much simpler and more transparent proof appears in an unpublished manuscript of Strebel [24] who attributes the argument to Bieri. While still ineffective, the proof in [24] is almost entirely algorithmic apart from one step, as we will explain later in this section.

2.2. On the proof of the BNS criterion

In this subsection, we will give a brief outline of the proof of the “if” part of Theorem 2.1 from [24]. With the exception of Lemma 2.3 below, the results discussed in this subsection will not be used in the rest of the paper, and the main purpose of providing this outline is to help the reader follow the proofs later in this section where we will establish an effective version of (the “if” part of) Theorem 2.1 under some additional hypotheses.

The following theorem (Theorem 2.2) should be well known, although we are not aware of a reference in the literature where it is stated exactly in this form. We are grateful to Andrew Putman for pointing out the formulation below.

Theorem 2.2. *Let G be a group generated by a finite set S , let K be a subgroup of G (not necessary normal), and let $\theta: G \rightarrow G/K$ be the natural projection. Then K is finitely generated if and only if there is a finite subset A of G/K such that $\theta^{-1}(A)$ is connected in $\mathrm{Cay}(G, S)$.*

The “only if” part (which is not essential for our purposes) is a straightforward exercise. The “if” part of Theorem 2.2 is an immediate consequence of [24, Theorem A4.7]. Later in this section we will prove Theorem 2.14 which is an effective version of the “if” part of Theorem 2.2.

We now begin a sketch of proof of the “if” direction of Theorem 2.1. Keeping all the notations from Theorem 2.2, suppose now that G/K is abelian and $\Sigma(G) \supseteq \mathbb{S}(G/K)$. We wish to show that K is finitely generated. Since G is finitely generated, after replacing K by a finite index overgroup (which does not affect finite generation), we can assume that G/K is torsion-free. In addition, we want to impose an extra condition on the generating set S given by (2.1) below.

Lemma 2.3. *Let G be a finitely generated group and let K be a normal subgroup of G such that G/K is abelian and torsion-free. Let $\theta: G \rightarrow G/K$ be the natural projection, and choose a basis E of G/K . Then G has a generating set S such that*

$$\theta(S) = E \quad \text{or} \quad \theta(S) = E \cup \{0\}. \quad (2.1)$$

Moreover, if $S_0 = \{s_1, \dots, s_n\}$ is any finite generating set of G , one can obtain a generating set S satisfying (2.1) from S_0 by applying a sequence of right Nielsen transformations, that is, transformations of the form $(g_1, \dots, g_i, \dots, g_n) \mapsto (g_1, \dots, g_i g_j^{\pm 1}, \dots, g_n)$ for some $i \neq j$ and possibly inverting one of the generators.

Proof. Write $E = \{e_1, \dots, e_m\}$. Since the sequence $(\theta(s_1), \dots, \theta(s_n))$ generates G/K , by basic linear algebra,³ using transformations of the form $(v_1, \dots, v_i, \dots, v_n) \mapsto (v_1, \dots, v_i \pm v_j, \dots, v_n)$ for some $i \neq j$, we can reduce $(\theta(s_1), \dots, \theta(s_n))$ to a sequence of the form $(\pm e_1, e_2, \dots, e_m, 0, \dots, 0)$. Since each such linear transformation is induced by a right Nielsen transformation, applying the corresponding sequence of right Nielsen transformations to S_0 and then replacing one of the generators by its inverse if needed, we obtain a generating set S with desired properties. ■

From now on, assume that S satisfies the conclusion of Lemma 2.3 (with respect to some fixed basis E). Choose an isomorphism $G/K \cong \mathbb{Z}^m$ which maps E onto the standard basis of \mathbb{Z}^m . Let $\|\cdot\|$ denote the corresponding l^2 -norm on G/K , and let $B(R)$ denote the l^2 -ball of radius R with respect to this norm (centered at 0). The goal now is to show that $\theta^{-1}(B(R))$ is connected for sufficiently large R (this would imply that K is finitely generated by Theorem 2.2). To do this, one chooses an arbitrary path p in $\text{Cay}(G, S)$ whose end vertices a and b lie in $\theta^{-1}(B(R))$ and then applies a (finite) sequence of modifications to p , so that the resulting path lies entirely in $\theta^{-1}(B(R))$.

At each step the modification is as follows. Choose a vertex g on the current path such that $\|\theta(g)\|$ is maximal. If $\|\theta(g)\| \leq R$, there is nothing to do, so assume that $\|\theta(g)\| > R$. Let gy_1 and gy_2 be the vertices preceding and succeeding g on this path. Define the character χ_g of G by $\chi_g(x) = -(\theta(g), \theta(x))$, and choose $t \in S^{\pm 1}$ such that $\chi_g(t) > 0$. Since by assumption $\|\theta(gy_i)\| \leq \|\theta(g)\|$ for $i = 1, 2$ and

$$\begin{aligned} \|\theta(gy_i)\|^2 &= \|\theta(g) + \theta(y_i)\|^2 = \|\theta(g)\|^2 + \|\theta(y_i)\|^2 + 2(\theta(g), \theta(y_i)) \\ &= \|\theta(g)\|^2 + \|\theta(y_i)\|^2 - 2\chi_g(y_i), \end{aligned}$$

we deduce that $\chi_g(y_i) \geq 0$ for $i = 1, 2$.

Since χ_g vanishes on K and we assume that $\Sigma(G) \supseteq \mathbb{S}(G/K)$, there exist paths $p_{t, y_2, g}$ from t to y_2t and $q_{t, y_1, g}$ from y_1t to t such that χ_g is positive on any vertex of those paths. Now we replace the segment (gy_1, g, gy_2) of the current path by a new subpath passing through $gy_1, gy_1t, gt, gy_2t, gy_2$ where one moves from gy_1t to gt using the path $g \circ q_{t, y_1, g}$ and from gt to gy_2t using the path $g \circ p_{t, y_2, g}$ (see Figure 1).

A key step of the proof is a compactness argument⁴ which shows that there is a finite set of paths Ω and $\varepsilon > 0$ such that for any character χ_g arising above we can find desired

³This can be done by writing $\theta(s_i) = \sum_{j=1}^m a_{ij}e_j$ with $a_{ij} \in \mathbb{Z}$ for each $1 \leq i \leq n$ and turning the matrix (a_{ij}) into Smith normal form using a standard algorithm.

⁴In order to use compactness, one has to replace χ_g by the normalized character $\chi'_g = \frac{1}{\|\theta(g)\|} \chi_g$ (note that χ_g and χ'_g are equivalent). If we extend the chosen inner product from G/K to $G/K \otimes \mathbb{R}$ and consider characters of G as elements of the dual space $\text{Hom}(G/K \otimes \mathbb{R}, \mathbb{R})$, then characters of the form χ'_g always have norm 1.




Figure 1. Parts of the original and new paths joining gy_1 and gy_2 .

paths $p_{t,y_2,g}$ and $q_{t,y_1,g}$ in Ω and moreover $\chi_g(v) \geq \varepsilon \|\theta(g)\|$ for any vertex v of $p_{t,y_2,g}$ or $q_{t,y_1,g}$. Now let r be the maximum of $\|\theta(v)\|$ where v ranges over the vertices of all paths from Ω . The direct computation below shows that if $R \geq \frac{r^2}{2\varepsilon}$, then $\|\theta(z)\| < \|\theta(g)\|$ for any vertex z on the newly added segment. This concludes the (sketch of) proof of the “if” direction of Theorem 2.1.

Indeed, any new vertex z has the form gv where v lies on $q_{t,y_1,g}$ or $p_{t,y_2,g}$. Therefore,

$$\begin{aligned} \|\theta(z)\|^2 &= \|\theta(g) + \theta(v)\|^2 = \|\theta(g)\|^2 + \|\theta(v)\|^2 + 2(\theta(g), \theta(v)) \\ &= \|\theta(g)\|^2 + \|\theta(v)\|^2 - 2\chi_g(v) \leq \|\theta(g)\|^2 + r^2 - 2\varepsilon \|\theta(g)\| \\ &\leq \|\theta(g)\|^2 + 2\varepsilon(R - \|\theta(g)\|) < \|\theta(g)\|^2. \end{aligned}$$

As the above outline suggests, in order to turn this proof into an actual algorithm (for a specific group), one needs to have an explicit procedure for constructing the paths $p_{t,y_2,g}$ and $q_{t,y_1,g}$. We will now discuss some additional conditions which make this possible.

2.3. Explicitly finding non-negative forms

The following notion of a non-negative form of a group element provides a convenient way to reformulate the definition of $\Sigma(G)$.

Definition 2.4. Let S be a generating set of a group G , let χ be a nonzero character of G and assume that $\chi(g) \geq 0$ for some $g \in G$. By a (χ, S) -non-negative form of g , we will mean an S -word w which represents g such that $\chi(v) \geq 0$ for every prefix v of w .

It is clear that a given $g \in G$ with $\chi(g) \geq 0$ admits a (χ, S) -non-negative form if and only if there is a path in $\text{Cay}(G, S)_\chi$ connecting 1 with g (recall that $\text{Cay}(G, S)_\chi$ is the full subgraph of $\text{Cay}(G, S)$ with the vertex set $\{x: \chi(x) \geq 0\}$). Thus, $[\chi] \in \Sigma(G)$ if and only if every g with $\chi(g) \geq 0$ admits a (χ, S) -non-negative form.

There is a well-known sufficient condition for a character to lie in the BNS invariant.

Lemma 2.5. *Suppose G is generated by $S = \{s_1, \dots, s_n\}$ and χ is a character of G such that*

- (i) $\chi(s_1) > 0$;
- (ii) *for every $i \geq 2$, there exists $j < i$ such that $[s_j, s_i] = 1$ and $\chi(s_j) \neq 0$.*

Then $[\chi] \in \Sigma(G)$.

Lemma 2.5 was proved in [17, Lemma 1.9], although indirectly it appeared already in [19]; see also [9, Lemma 2.4] for a generalization. Although the condition in Lemma 2.5 may appear very special, there are many important classes of groups G for which every character in $\Sigma(G)$ does satisfy this condition for some S – for instance, this is the case for right-angled Artin groups [19], groups of pure symmetric automorphisms of free groups [21], pure braid groups [17], IA_n for $n \geq 5$ and \mathcal{I}_n^b for $n \geq 5$, $b = 0, 1$ (see [5, 9]).

If a generating set S and a character χ satisfy the hypotheses of Lemma 2.5, it is not difficult to describe an algorithm which computes an (S, χ) -non-negative form for a given $g \in G$ with $\chi(g) \geq 0$. In fact, such an algorithm implicitly appears in the proof of [19, Theorem 4.1].

In this paper, it will be more convenient to work with a slightly more restrictive condition, which still holds for IA_n and \mathcal{I}_n^b for sufficiently large n and leads to a very simple formula for a (χ, S) -non-negative form.

We will need some technical definitions.

Definition 2.6. Let S be a finite generating set for a group G and Z a subset of S .

(a) We will say that the pair (S, Z) is *chain-centralizing* if for every $s \in S$ and $z \in Z$ there exists $z' \in Z$ which commutes with both s and z .

(b) If χ is a character of G , we will say that χ is *regular* for (S, Z) if $\chi(z) \neq 0$ for all $z \in Z$.

Remark. If χ is regular for a chain-centralizing pair (S, Z) , it is easy to show that the pair (χ, S) satisfies the hypothesis of Lemma 2.5 (for a suitable ordering of S), but we will not use this fact in the proofs.

The name *chain-centralizing* is motivated by the following property which is an obvious consequence of the definition.

Observation 2.7. *Suppose that (S, Z) is chain-centralizing. Then for any finite sequence $s_1, \dots, s_k \in S$ there exists a sequence $z_1, \dots, z_k \in Z$ such that*

- (i) z_i commutes with s_i for each $1 \leq i \leq k$;
- (ii) z_i commutes with z_{i-1} for each $2 \leq i \leq k$.

The following lemma shows that if χ is regular for a chain-centralizing pair (S, Z) , it is very easy to construct (χ, S) -non-negative forms.

Lemma 2.8. *Suppose that (S, Z) is chain-centralizing and χ is regular for (S, Z) . Let $g \in G$ with $\chi(g) \geq 0$, and write $g = s_1 \dots s_k$ with $s_i \in S^{\pm 1}$. Choose z_1, \dots, z_k satisfying conditions (i) and (ii) of Observation 2.7, and choose $n_i \in \mathbb{Z}$ such that*

- (a) $n_i \chi(z_i) \geq 0$ for all $1 \leq i \leq k$;
 - (b) $n_i \chi(z_i) \geq -\sum_{j=1}^i \chi(s_j)$ for all $1 \leq i \leq k$;
 - (c) $n_i \chi(z_i) \geq -\sum_{j=1}^{i-1} \chi(s_j)$ for all $2 \leq i \leq k$
- (such n_i exist since χ is regular). Then the S -word

$$w_\chi = z_1^{n_1} s_1 z_2^{n_2} z_1^{-n_1} s_2 z_3^{n_3} z_2^{-n_2} s_3 \dots z_k^{n_k} z_{k-1}^{-n_{k-1}} s_k z_k^{-n_k}$$

is a (χ, S) -non-negative form of g .

Proof. The word w_χ represents g by conditions (i) and (ii) of Observation 2.7: we first move $z_1^{-n_1}$ past s_1 and $z_2^{n_2}$ and cancel it with $z_1^{n_1}$, then we move $z_2^{-n_2}$ past s_2 and $z_3^{n_3}$ and cancel it with $z_2^{n_2}$ etc.

Let us now prove that $\chi(v) \geq 0$ for every S -prefix v of w_χ . Without loss of generality, we can assume that $\chi(z_i) > 0$ for all i , in which case $n_i \geq 0$ by (a). If v does not end with $z_i^{-n_i}$ or s_i for some i , we can produce another S -prefix v' of w_χ with $\chi(v') \leq \chi(v)$ by either removing the last letter of v or adding the next letter of w_χ to the end of v . Thus, it suffices to prove that $\chi(v) \geq 0$ when v ends with $z_i^{-n_i}$ or s_i .

Case 1: $v = z_1^{n_1} \prod_{j=1}^i (s_j z_{j+1}^{n_{j+1}} z_j^{-n_j})$ for some $1 \leq i \leq k-1$. In this case, by (c) $\chi(v) = n_1 \chi(z_1) + \sum_{j=1}^i (\chi(s_j) + n_{j+1} \chi(z_{j+1}) - n_j \chi(z_j)) = n_{i+1} \chi(z_{i+1}) + \sum_{j=1}^i \chi(s_j) \geq 0$.

Case 2: $v = z_1^{n_1} \prod_{j=1}^i (s_j z_{j+1}^{n_{j+1}} z_j^{-n_j}) s_{i+1}$ for some $0 \leq i \leq k-1$ (the case $i = 0$ corresponds to the empty product). Then $\chi(v) = n_{i+1} \chi(z_{i+1}) + \sum_{j=1}^{i+1} \chi(s_j) \geq 0$ by (b).

Case 3: v is the full word w_χ . In this case, $\chi(v) = \chi(g) \geq 0$. ■

2.4. Extra hypothesis

In this subsection, we introduce the additional condition that will allow us to turn the proof of the BNS criterion into an algorithm. As before, we will assume that G, K, E, θ and S satisfy the hypotheses and conclusion of Lemma 2.3: G/K is abelian and torsion-free, $\theta: G \rightarrow G/K$ is the natural projection, E is a fixed basis of G/K and S is a generating set of G such that $\theta(S) = E$ or $E \cup \{0\}$.

In order to make use of Lemma 2.8, we need to know that every nonzero character is regular with respect to some chain-centralizing pair. Note that a single chain-centralizing pair would rarely work for all the characters (apart from rather trivial examples). Also observe that if we have one chain-centralizing pair (S, Z) , then for any $\phi \in \text{Aut}(G)$, the pair $(\phi(S), \phi(Z))$ is also chain-centralizing.

This motivates our new hypothesis. We would like to assume that there is a finite subset $\Phi \subseteq \text{Aut}(G)$ with the following property: for every nonzero character χ of G which vanishes on K , there is some $Z \subseteq S$ and $\phi \in \Phi$ such that (S, Z) is chain-centralizing

and χ is regular for $(\phi(S), \phi(Z))$. In fact, we will need to assume a bit more (see the regularity hypothesis below), but first we will introduce some additional notations involving automorphisms of G .

Constants A and B . Let $\phi \in \text{Aut}(G)$. Define

$$B(\phi, S) = \max\{\|\theta(\phi(s))\|_1 : s \in S\}$$

(where $\|\cdot\|_1$ denotes the l^1 -norm with respect to E).

Now define $A = A(\phi, S)$ to be the smallest integer with the following property: for every $s \in S$, there is an S -word $w_{\phi,s}$ representing $\phi(s)$ such that $\|\theta(v)\|_1 \leq A$ for every S -prefix v of $w_{\phi,s}$. Since $w_{\phi,s}$ is its own S -prefix, we have the obvious inequality $B(\phi, S) \leq A(\phi, S)$.

If Φ is a finite subset of $\text{Aut}(G)$, we define

$$A(\Phi, S) = \max\{A(\phi, S) : \phi \in \Phi\} \quad \text{and} \quad B(\Phi, S) = \max\{B(\phi, S) : \phi \in \Phi\}.$$

When S is fixed or clear from the context, we will usually suppress it from the notation and write $A(\phi)$ for $A(\phi, S)$ etc.

Constant M . For each character χ define

$$M(\chi) = M(\chi, S) = \max\{|\chi(s)| : s \in S\}.$$

We are now ready to state our additional hypothesis. Let $\text{Aut}(G, K)$ denote the subgroup of $\text{Aut}(G)$ consisting of automorphisms which leave K invariant.

Regularity hypothesis. *There exist a finite subset $\Phi \subseteq \text{Aut}(G, K)$ and a constant $C > 0$ with the following property:*

(***) *For every nonzero character $\chi \in \text{Hom}(G/K, \mathbb{R})$, there exist a subset $Z \subseteq S$ and $\phi \in \Phi$ such that (S, Z) is chain-centralizing and $|\chi(\phi(z))| \geq \frac{M(\chi)}{C}$ for all $z \in Z$.*

Remark. The inequality in (***) implies, in particular, that χ is regular for $(\phi(S), \phi(Z))$.

Before proceeding, we establish a few simple inequalities involving the constants A , B and M .

Observation 2.9. *Let $g \in G$. Then $\|\theta(g)\|_1$ is the smallest integer m for which there exists $h \in G$ with $\|h\|_S = m$ and $\theta(h) = \theta(g)$.*

Proof. This immediately follows from the assumption that $\theta(S) = E$ or $E \cup \{0\}$. ■

Claim 2.10. *Let χ be a character of G which vanishes on K . Then*

$$|\chi(g)| \leq \|\theta(g)\|_1 \cdot M(\chi)$$

for all $g \in G$.

Proof. Let $m = \|\theta(g)\|_1$. By Observation 2.9, there exist $s_1, \dots, s_m \in S^{\pm 1}$ such that $\theta(g) = \theta(s_1 \dots s_m)$. Since the character χ vanishes on K , it factors through θ and thus

$$|\chi(g)| = |\chi(s_1 \dots s_m)| \leq \sum_{i=1}^m |\chi(s_i)| \leq m \cdot \max\{|\chi(s_i)|\} \leq mM(\chi). \quad \blacksquare$$

Claim 2.11. *Let $\phi \in \text{Aut}(G, K)$, and let $A = A(\phi^{-1})$ and $B = B(\phi)$. Then for every $s \in S^{\pm 1}$, there exists a $\phi(S)$ -word \tilde{s} representing s such that $\|\theta(w)\|_1 \leq AB$ for every $\phi(S)$ -prefix w of \tilde{s} .*

Proof. Fix $s \in S^{\pm 1}$. By definition of the constant A , there exists an S -word $x_1 \dots x_k$ (with $x_i \in S^{\pm 1}$) representing $\phi^{-1}(s)$ such that $\|\theta(\prod_{j=1}^l x_j)\|_1 \leq A$ for all $1 \leq l \leq k$.

Let $\tilde{s} = \phi(x_1) \dots \phi(x_k)$. Then \tilde{s} is a $\phi(S)$ -word representing $\phi(\phi^{-1}(s)) = s$. Thus, to prove the claim, it suffices to show that

$$\left\| \theta\left(\prod_{j=1}^l \phi(x_j)\right) \right\|_1 = \left\| \theta\left(\phi\left(\prod_{j=1}^l x_j\right)\right) \right\|_1 \leq AB \quad \text{for all } 1 \leq l \leq k. \quad (2.2)$$

Fix such an index l . By Observation 2.9 and the choice of the word $x_1 \dots x_k$, there exists an S -word $v = y_1 \dots y_m$ with $m \leq A$ such that $\theta(\prod_{j=1}^l x_j) = \theta(v)$ and hence $\prod_{j=1}^l x_j = vz$ with $z \in K$.

Then $\phi(\prod_{j=1}^l x_j) = \phi(v)\phi(z)$, and since K is ϕ -invariant, we have $\theta(\phi(\prod_{j=1}^l x_j)) = \theta(\phi(v))$. Note that

$$\|\theta(\phi(v))\|_1 = \|\theta(\phi(y_1 \dots y_m))\|_1 \leq \sum_{i=1}^m \|\theta(\phi(y_i))\|_1.$$

Since $m \leq A$ and $\|\theta(\phi(y_i))\|_1 \leq B$ by definition of B , we get that $\|\theta(\phi(v))\|_1 \leq AB$ which proves (2.2). \blacksquare

2.5. The main result

In this subsection, we will prove the main result of this section, Theorem 2.13. We start with a key proposition which shows that under the regularity hypothesis, we can control the “size” of (χ, S) -non-negative forms with respect to any character χ of G/K .

Proposition 2.12. *Let G, K and S be as in Lemma 2.3, and assume that the regularity hypothesis holds for some $\Phi \subseteq \text{Aut}(G, K)$ and a constant $C > 0$. Let $A = A(\Phi \cup \Phi^{-1})$ and $B = B(\Phi \cup \Phi^{-1})$. Let χ be a nonzero character of G/K , let $g \in G$ with $\chi(g) \geq 0$, and write $g = s_1 \dots s_r$ with $s_i \in S^{\pm 1}$. Then there exists a (χ, S) -non-negative form w of g such that for any prefix v of w we have*

$$\|\theta(v)\|_1 \leq (2BC + 1)(AB + A + r) + 2A.$$

Proof. The basic idea is very simple. Choose $\phi \in \Phi$ such that (***) in the regularity hypothesis holds for χ , express g as a $\phi(S)$ -word and then use Lemma 2.8 to construct a $(\chi, \phi(S))$ -non-negative form of g . This almost works – the issue is that when we rewrite the obtained $(\chi, \phi(S))$ -non-negative form as an S -word, prefixes of this S -word may have negative χ -values. However, we have a lower bound for those χ -values: $\chi(v) \geq -MA$ for every such prefix v , where $M = M(\chi)$.

To resolve this problem, we choose $t \in S^{\pm 1}$ with $\chi(t) = M$ and apply the same argument to the element $t^{-A}gt^A = t^{-A} \prod_{i=1}^r s_i t^A$. It is easy to see that if w is an S -word representing $t^{-A}gt^A$ such that $\chi(v) \geq -MA$ for every such prefix v of w , then t^Awt^{-A} is a (χ, S) -non-negative form of g .

We now present the full argument. For convenience, we break the construction into three steps.

Step 1. Choose $t \in S^{\pm 1}$ with $\chi(t) = M$. Recall that $s_1 \dots s_r$ is an S -word representing g and hence $p = t^{-A}s_1 \dots s_r t^A$ is an S -word representing $t^{-A}gt^A$. Replacing each factor $s \in \{s_1, \dots, s_r, t^{\pm 1}\}$ in the word p by the corresponding \tilde{s} given by Claim 2.11, we obtain a $\phi(S)$ -word \tilde{p} representing $t^{-A}gt^A$.

Let v be an arbitrary $\phi(S)$ -prefix of \tilde{p} . We claim that

- (1) $\|\theta(v)\|_1 \leq AB + A + r$;
- (2) $|\chi(v)| \leq (AB + A + r)M$.

For convenience, let us write $t^{-A}gt^A = \prod_{i=1}^{r+A} s_i$, where $s_i = t^{-1}$ for $i \leq 0$ and $s_i = t$ for $i > r$. Thus, $\tilde{p} = \prod_{i=1}^{r+A} \tilde{s}_i$, and there exist some $-A \leq j < r + A$ and a $\phi(S)$ -prefix w of \tilde{s}_{j+1} such that $v = (\prod_{i \leq j} \tilde{s}_i)w$. Then

$$\|\theta(v)\|_1 \leq \left\| \theta \left(\prod_{i \leq j} \tilde{s}_i \right) \right\|_1 + \|\theta(w)\|_1 = \left\| \sum_{i \leq j} \theta(s_i) \right\|_1 + \|\theta(w)\|_1.$$

The first summand is bounded above by $A + r$ (this upper bound may occur for $j = r$ since for larger j , we start getting cancellations of $\theta(t)$ with $\theta(t^{-1})$). Also $\|\theta(w)\|_1 \leq AB$ by Claim 2.11, so we proved (1). Inequality (2) immediately follows from item (1) and Claim 2.10.

Step 2. Next we use Lemma 2.8 to construct a $(\chi, \phi(S))$ non-negative form of \tilde{p} , call it \tilde{q} (the elements z_i in Lemma 2.8 will be chosen from $\phi(Z)$). For each i , we shall choose n_i to be smallest in absolute value satisfying the inequality in Lemma 2.8. Since $|\chi(z_i)| \geq \frac{M}{C}$ by the regularity hypothesis and $|\chi(v)| \leq (AB + A + r)M$ for every prefix v of \tilde{p} , we have $|n_i| \leq C(AB + A + r)$.

Now we need to establish the analogue of (1) for prefixes of \tilde{q} . Let v be a $\phi(S)$ -prefix of \tilde{q} . It is clear from the construction that there is a $\phi(S)$ -prefix w of \tilde{p} such that $\theta(v)$ is equal to $\theta(w)$ or $\theta(wz_i^m)$ with $|m| \leq |n_i|$ for some i or $\theta(wz_i^{n_i}z_{i+1}^m)$ with $|m| \leq |n_{i+1}|$ for some i . In any case, we have

$$\|\theta(v)\|_1 \leq \|\theta(w)\|_1 + |n_i| \cdot \|\theta(z_i)\|_1 + |n_{i+1}| \cdot \|\theta(z_{i+1})\|_1.$$

Since $\|\theta(w)\|_1 \leq AB + A + r$ by (1), $\|\theta(z_j)\|_1 \leq B$ for all j by definition of B and $|n_j| \leq C(AB + A + r)$ for all j as established above, we get

$$(3) \quad \|\theta(v)\|_1 \leq (AB + A + r) + 2BC(AB + A + r) = (2BC + 1)(AB + A + r).$$

Step 3. Next we rewrite \tilde{q} as an S -word, call it q . Let v be an S -prefix of q . We claim that

$$(4) \quad \|\theta(v)\|_1 \leq (2BC + 1)(AB + A + r) + A;$$

$$(5) \quad \chi(v) \geq -MA.$$

Indeed, similarly to step 1, we can write $v = v_1 v_2$, where v_1 is a $\phi(S)$ -prefix of \tilde{q} and v_2 is an S -prefix of some $x \in \phi(S^{\pm 1})$. Now $\|\theta(v_1)\|_1 \leq (2BC + 1)(AB + A + r)$ by (3) and $\|\theta(v_2)\|_1 \leq A$ by definition of A , so (4) holds.

Since $\|\theta(v_2)\|_1 \leq A$, Claim 2.10 yields $|\chi(v_2)| \leq MA$, so in particular $\chi(v_2) \geq -MA$. Since \tilde{q} is a $(\chi, \phi(S))$ -non-negative form, we must have $\chi(v_1) \geq 0$, so $\chi(v) = \chi(v_1) + \chi(v_2) \geq -MA$. Thus we proved (5).

Recall now that q represents $t^{-A}gt^A$ in G and hence t^Aqt^{-A} represents g . We claim that t^Aqt^{-A} is the desired (χ, S) -non-negative form. Take any S -prefix w of t^Aqt^{-A} . We need to show that

$$(6) \quad \|\theta(w)\|_1 \leq (2BC + 1)(AB + A + r) + 2A;$$

$$(7) \quad \chi(w) \geq 0.$$

Clearly, there are three cases:

Case 1: $w = t^m$ for some $0 \leq m \leq A$. In this case, both (6) and (7) are obvious.

Case 2: $w = t^A v$ for some prefix v of q . In this case, $\|\theta(v)\|_1 \leq (2BC + 1)(AB + A + r) + A$ by (4), so (6) holds, and (7) follows from (5) and the fact that $\chi(t^A) = MA$ (by the choice of t).

Case 3: $w = t^A q t^{-m}$ for some $0 \leq m \leq A$. In this case, (6) again follows from (4). Finally, $\chi(t^A q t^{-A}) \geq 0$ since $t^A q t^{-A}$ represents g and hence $\chi(w) = \chi(t^A q t^{-A} t^{A-m}) \geq (A - m)\chi(t) \geq 0$. ■

We can now construct an explicit finite subset of G/K whose preimage in $\text{Cay}(G, S)$ is connected. We will show that the l^∞ -ball of a certain radius has this property.

Theorem 2.13. *We keep all the hypotheses and notations from Proposition 2.12 and let $R = 16(BC + 1)^2(AB + 3A + 3)$. Let $B_\infty(R)$ be the l^∞ -ball of radius R in G/K . Then $\theta^{-1}(B_\infty(R))$ is connected in $\text{Cay}(G, S)$.*

To simplify terminology, in the proof of Theorem 2.13, we will occasionally talk about the l^2 -norm or l^∞ -norm for an element $g \in G$, by which we mean the corresponding norm of $\theta(g)$.

Proof. Recall that E is a fixed basis of G/K such that $\theta(S) = E$ or $E \cup \{0\}$. As before, we choose an isomorphism $G/K \rightarrow \mathbb{Z}^m$ which maps E to the standard basis of \mathbb{Z}^m .

We will generally follow the outline of the proof of Theorem 2.1 given earlier in this section. Here is the summary of the new ingredients that we will use:

- (i) The compactness argument will be replaced by a reference to Proposition 2.12, which enables us to get an explicit formula for R .
- (ii) In the proof, we will use not just l^2 -norm, but also l^∞ -norm and l^1 -norm (where all the norms are taken with respect to the chosen identification of G/K with \mathbb{Z}^m).
- (iii) When modifying the path at each step, we will make slightly more complicated “detours”.

Note that modification (i) is essential, while (ii) and (iii) will only be used to obtain a better estimate.

So take any $a, b \in G$ with $\|\theta(a)\|_\infty, \|\theta(b)\|_\infty \leq R$. Our goal is to show that there is a path from a to b in $\text{Cay}(G, S)$ whose projection lies in $B_\infty(R)$. We start by choosing some path p from a to b in $\text{Cay}(G, S)$ and then apply a sequence of modifications, eventually pushing it inside $\theta^{-1}(B_\infty(R))$.

If $\theta(p)$ lies inside $B_\infty(R)$, we are done, so assume that $\theta(p)$ has at least one vertex outside $B_\infty(R)$. Among all the vertices of p outside $B_\infty(R)$ we choose one with the largest l^2 -norm, call it g . Our goal is to replace p by another path p' from a to b which does not have any vertices with l^2 -norm larger than $\|g\|_2$ and has fewer vertices with l^2 -norm equal to $\|g\|_2$ than p . Clearly, after applying such a modification finitely many times, we will obtain a path inside $\theta^{-1}(B_\infty(R))$, as desired. Note that the maximal l^∞ -norm may increase during some initial steps – this is not a problem.

Define the character χ of G/K by $\chi(x) = -(\theta(g), \theta(x))$. Let gy_1 and gy_2 , where $y_1, y_2 \in S \cup S^{-1}$, be the vertices preceding and succeeding g in the path p , respectively. We claim that $\|\theta(gy_i)\|_2 \leq \|\theta(g)\|_2$ for $i = 1, 2$.

To prove the claim recall that by assumption g has the largest l^2 -norm among the vertices of p which lie outside of $\theta^{-1}(B_\infty(R))$. Thus, if $\theta(gy_i) \notin B_\infty(R)$, it is automatic that $\|\theta(gy_i)\|_2 \leq \|\theta(g)\|_2$. On the other hand, if $\theta(gy_i) \in B_\infty(R)$, we have $\|\theta(gy_i)\|_\infty \leq R < \|\theta(g)\|_\infty$. But by assumption, $\theta(y_i) \in \{\pm e_j\} \cup \{0\}$. Hence $\theta(g)$ and $\theta(gy_i)$ differ in at most one coordinate, so $\|\theta(gy_i)\|_\infty < \|\theta(g)\|_\infty$ forces $\|\theta(gy_i)\|_2 < \|\theta(g)\|_2$ as well.

Thus, we proved that $\|\theta(gy_i)\|_2 \leq \|\theta(g)\|_2$ for $i = 1, 2$. Since

$$\|\theta(gy_i)\|_2^2 = \|\theta(g)\|_2^2 + \|\theta(y_i)\|_2^2 + 2(\theta(g), \theta(y_i)) = \|\theta(g)\|_2^2 + \|\theta(y_i)\|_2^2 - 2\chi(y_i),$$

it follows that $\chi(y_i) \geq 0$ for $i = 1, 2$.

Let $\theta(g)_i$ denote the i -th coordinate of $\theta(g)$, and choose i such that $\theta(g)_i$ is maximal in absolute value. Since by assumption $\theta(g) \notin B_\infty(R)$, we have $|\theta(g)_i| > R$. Choose $t \in S^{\pm 1}$ with $\theta(t) = \pm e_i$ where the sign is chosen to be the same as the sign of $\theta(g)_i$. Then $(\theta(g), \theta(t)) = |\theta(g)_i| > R$.

We proceed with the construction of p' (see Figure 2 for an illustration). The path p' will coincide with p prior to gy_1 and after gy_2 , but we will connect gy_1 and gy_2 by




Figure 2. Parts of the original and new paths joining gy_1 and gy_2 .

a new subpath which passes through the vertices gy_1t^{-r} , gt^{-r} and gy_2t^{-r} , in this order, where r satisfying $1 \leq r \leq R$ will be chosen later. We connect gy_1 with gy_1t^{-r} in the natural way (multiplying r times by t^{-1}) and similarly gy_2t^{-r} with gy_2 . To connect gt^{-r} with gy_2t^{-r} , we write $gy_2t^{-r} = (gt^{-r})(t^r y_2t^{-r})$ and then replace the suffix $t^r y_2t^{-r}$ by its (χ, S) -non-negative form constructed in Proposition 2.12 (the latter is applicable since $\chi(t^r y_2t^{-r}) = \chi(y_2) \geq 0$). Similarly, we connect gy_1t^{-r} with gt^{-r} .

We now need to show that if v is any vertex on this new subpath different from the end vertices gy_1 and gy_2 , then $\|\theta(v)\|_2 < \|\theta(g)\|_2$. Recall that $\theta(t) = \pm e_i$, so all the vertices on the segment between gy_1 to gy_1t^{-r} differ only in the i -th coordinate. Moreover, by assumption $\theta(gy_1)_i \geq \theta(g)_i - 1 \geq R$ and $\theta(t)_i$ and $\theta(gy_1)_i$ have the same sign. Thus, if we require that $r \leq R$, then $|\theta(gy_1t^{-j})_i| = |\theta(gy_1)_i| - j$ for $0 \leq j \leq r$, so the l^2 -norm strictly goes down as we move from gy_1 to gy_1t^{-r} . Similarly, the l^2 -norm strictly goes down when we move from gy_2 to gy_2t^{-r} .

Let us now consider the vertices of p' between gt^{-r} and gy_2t^{-r} (the vertices between gy_1t^{-r} and gt^{-r} are treated similarly). Any such vertex is of the form $gt^{-r}v$ where v is a prefix of the (χ, S) -non-negative form of $t^r y_2t^{-r}$ chosen above. Since $\|t^r y_2t^{-r}\|_1 \leq 2r + 1$, by construction we get

$$\|\theta(v)\|_1 \leq (2BC + 1)(AB + A + 2r + 1) + 2A \leq (2BC + 1)(AB + 3A + 2r + 1)$$

and $\chi(v) \geq 0$, so $(\theta(g), \theta(v)) = -\chi(v) \leq 0$. Also recall that $(\theta(g), \theta(t)) \geq R$. Therefore, we have

$$\begin{aligned} \|\theta(gt^{-r}v)\|_2^2 &= \|\theta(g)\|_2^2 + \|\theta(t^{-r}v)\|_2^2 + 2(\theta(g), \theta(t^{-r}v)) + 2(\theta(g), \theta(v)) \\ &\leq \|\theta(g)\|_2^2 + \|\theta(t^{-r}v)\|_1^2 - 2r(\theta(g), \theta(t)) + 2(\theta(g), \theta(v)) \\ &\leq \|\theta(g)\|_2^2 + (\|\theta(v)\|_1 + r)^2 - 2rR \\ &\leq \|\theta(g)\|_2^2 + ((2BC + 1)(AB + 3A + 2r + 1) + r)^2 - 2rR \\ &< \|\theta(g)\|_2^2 + ((2BC + 2)(AB + 3A + 2r + 1))^2 - 2rR. \end{aligned}$$

The first inequality above holds since for any $(a_1, \dots, a_n) \in \mathbb{Z}^n$ we have $\|(a_1, \dots, a_n)\|_2^2 = \sum |a_i|^2 \leq (\sum |a_i|)^2 = \|(a_1, \dots, a_n)\|_1^2$.

Recall that $R = 16(BC + 1)^2(AB + 3A + 3)$. For simplicity of notation, write $P = 16(BC + 1)^2$ and $Q = \frac{AB+3A+1}{2}$, so that $R = 2P(Q + 1)$. If we set $r = \lceil Q \rceil$ (which satisfies the condition $r \leq R$ imposed earlier in the proof), then $Q \leq r < Q + 1$ and hence $Q + r \leq 2r < 2(Q + 1)$. Therefore,

$$((2BC + 2)(AB + 3A + 2r + 1))^2 - 2rR = P(Q + r)^2 - P \cdot 2r \cdot 2(Q + 1) < 0$$

which completes the proof. \blacksquare

2.6. An explicit generating set for K

To finish the constructive proof of finite generation given in this section, we need to establish an effective version of Theorem 2.2.

For a technical reason, in the following two results it will be convenient to work with left Cayley graphs (note that earlier in this section, we worked with the commonly used right Cayley graphs). The left Cayley graph of a group G with respect to S , denoted by $\text{Cay}_{\text{left}}(G, S)$, is defined in the same way as $\text{Cay}(G, S)$ except that edges have the form $(g, s^{\pm 1}g)$ with $s \in S$. Note that $\text{Cay}_{\text{left}}(G, S)$ and $\text{Cay}(G, S)$ are isomorphic as graphs via the inversion map $g \mapsto g^{-1}$.

Theorem 2.14 below is a variation of the Reidemeister–Schreier rewriting process. This result is undoubtedly well known, but we are not aware of a specific reference in the literature, so we will provide a proof.

Recall that if K is a subgroup of a group G , a left transversal of K in G is a subset T of G which contains exactly one element from each left coset of K .

Theorem 2.14. *Let G be a group generated by a set S , let K be a (not necessary normal) subgroup of G , and let $\theta: G \rightarrow G/K$ be the natural projection. Let $\mathcal{F} \subseteq G/K$ be such that $\theta^{-1}(\mathcal{F})$ is connected in $\text{Cay}_{\text{left}}(G, S)$. Let T be a left transversal for K in G , let $T_{\mathcal{F}} = T \cap \theta^{-1}(\mathcal{F})$ (note that $|T_{\mathcal{F}}| = |\mathcal{F}|$) and let*

$$U = \{(s, t): t \in T_{\mathcal{F}}, s \in S, \theta(st) \in \mathcal{F}\}.$$

Then K is generated by the set

$$S_K = \{\overline{st}^{-1}st: (s, t) \in U\},$$

where for every $g \in G$ by $\overline{g} \in T$ we denote the unique element of T such that $\theta(g) = \theta(\overline{g})$. In particular, if \mathcal{F} and S are finite, then K can be generated by (at most) $|\mathcal{F}| |S|$ elements.

Remark. If we take $\mathcal{F} = G/K$, then $\theta^{-1}(\mathcal{F}) = G$ is automatically connected. In this case, S_K is the usual Reidemeister–Schreier generating set for K .

Proof. Let K' be the subgroup generated by S_K . Clearly, $S_K \subseteq K$ (since $\overline{g}K = gK$ for all $g \in G$ by definition of \overline{g}), so $K' \subseteq K$. Let us now prove that $K \subseteq K'$.

Take any $k \in K$. As $K = \theta^{-1}(1) \subseteq \theta^{-1}(\mathcal{F})$ and $\theta^{-1}(\mathcal{F})$ is connected in $\text{Cay}_{\text{left}}(G, S)$, we can find a path $1 = y_0, y_1, \dots, y_m = k$ in $\text{Cay}_{\text{left}}(G, S)$ with $\theta(y_i) \in \mathcal{F}$ for each i . Since T is a transversal for K , for each $0 \leq i < m$ we can uniquely write $y_i = t_i k_i$, where $t_i \in T$ and $k_i \in K$. Note that $k_0 = 1$ and $k_m = k$. Thus, to prove that $k \in K'$, it suffices to show that $k_{i+1}k_i^{-1} \in K'$ for each $0 \leq i < m$.

Let us fix $0 \leq i < m$. By assumption, y_i and y_{i+1} are connected by an edge in $\text{Cay}_{\text{left}}(G, S)$, so there exists $s \in S$ such that $y_{i+1} = sy_i$ or $y_i = sy_{i+1}$.

First consider the case $y_{i+1} = sy_i$. We have $\theta(t_i) = \theta(t_i k_i) = \theta(y_i) \in \mathcal{F}$, so $t_i \in T_{\mathcal{F}}$ and $\theta(st_i) = \theta(st_i k_i) = \theta(sy_i) = \theta(y_{i+1}) \in \mathcal{F}$, whence $(s, t_i) \in U$.

Also $t_{i+1}k_{i+1} = y_{i+1} = sy_i = st_i k_i$ and hence $t_{i+1} = t_{i+1}k_{i+1} = \overline{st_i k_i} = \overline{st_i}$. Finally, the equality $t_{i+1}k_{i+1} = st_i k_i$ implies that

$$k_{i+1}k_i^{-1} = t_{i+1}^{-1}st_i = \overline{st_i}^{-1}st_i \in S_K,$$

as desired.

In the case $y_i = sy_{i+1}$, we can repeat the above argument swapping i and $i+1$ in every expression and observe that in the end $k_i k_{i+1}^{-1} \in K'$ forces

$$k_{i+1}k_i^{-1} = (k_i k_{i+1}^{-1})^{-1} \in K'. \quad \blacksquare$$

Making the generating set for K more explicit. Let us now consider the special case where G/K is abelian and torsion-free. Let $\theta: G \rightarrow G/K$ be the natural projection. Choose an ordered basis $E = \{e_1, \dots, e_n\}$ of G/K and use it to identify G/K with \mathbb{Z}^n . Recall that by Lemma 2.3, G has a finite generating set S such that $\theta(S) = E$ or $\theta(S) = E \cup \{0\}$. Note that θ restricted to S need not be injective.

We will show that if \mathcal{F} satisfies a certain technical condition (see the definition of a Schreier set below), one can obtain an even more explicit finite generating set for K (by slightly modifying the set S_K from Theorem 2.14).

Choose $s_1, \dots, s_n \in S$ with $\theta(s_i) = e_i$, and let $S_1 = \{s_1, \dots, s_n\}$. Let S_2 be the elements of $S \setminus S_1$ which lie outside of K , and let S_3 be the elements of $S \setminus S_1$ which lie in K . Thus $S = S_1 \sqcup S_2 \sqcup S_3$. For each $s \in S_2$, let $d(s)$ be the unique integer such that $\theta(s) = e_{d(s)}$.

Definition 2.15. In the above notations, a subset \mathcal{F} of G/K will be called *Schreier* if the following hold:

- (i) $(0, \dots, 0) \in \mathcal{F}$;
- (ii) for every $(a_1, \dots, a_n) \in \mathcal{F}$ and any path p in $\text{Cay}(G/K, E)$ from $(0, \dots, 0)$ to (a_1, \dots, a_n) of minimal length, all vertices of p lie in \mathcal{F} .

Equivalently, a subset \mathcal{F} of G/K is Schreier if for every $(a_1, \dots, a_n) \in \mathcal{F}$, any element $(b_1, \dots, b_n) \in G/K$ such that $|b_i| \leq |a_i|$ and $a_i b_i \geq 0$ for all i also lies in \mathcal{F} .

Remark. It is clear that for any $p \in [1, \infty]$, the l^p -ball centered at 0 is a Schreier set.

Recall that for group elements x, y , we set $[x, y] = x^{-1}y^{-1}xy$ and $x^y = y^{-1}xy$.

Theorem 2.16. Assume that G , K and S satisfy the above conditions. Let $\mathcal{F} \subseteq G/K$ be a Schreier subset of G/K such that $\theta^{-1}(\mathcal{F})$ is connected. Then K is generated by the following three types of elements:

- (a) $[s_i, s_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_n^{a_n}}$, where $1 \leq i < j \leq n$ and $(\underbrace{0, \dots, 0}_{i-1 \text{ times}}, a_i, a_{i+1}, \dots, a_n) \in \mathcal{F}$;
- (b) $(s^{-1} s_{d(s)})^{s_1^{a_1} s_2^{a_2} \dots s_n^{a_n}}$, where $s \in S_2$ and $(a_1, \dots, a_n) \in \mathcal{F}$;
- (c) $s^{s_1^{a_1} s_2^{a_2} \dots s_n^{a_n}}$, where $s \in S_3$ and $(a_1, \dots, a_n) \in \mathcal{F}$.

Proof. Clearly, all elements in (a)–(c) above lie in K . Let $T = \{s_1^{a_1} \dots s_n^{a_n} : a_i \in \mathbb{Z} \text{ for all } i\}$. By construction, T is a transversal for K in G , and let S_K be the corresponding set from Theorem 2.14. We need to show that every element of S_K can be expressed in terms of elements of type (a)–(c). Let us denote the subgroup generated by those elements by K' .

So take any $s \in S$, $t \in T$ with $\theta(t) \in \mathcal{F}$. Thus $t = \prod_{i=1}^n s_i^{a_i}$ for some $(a_1, \dots, a_n) \in \mathcal{F}$. We are also allowed to assume that $\theta(st) \in \mathcal{F}$, but this extra condition will not be needed for the argument. We will consider three cases depending on which of the subsets S_1 , S_2 and S_3 the generator s lies in.

Case 1: $s \in S_1$, so $s = s_j$ for some $1 \leq j \leq n$. Let us write $t = uv$ where $u = \prod_{i=1}^j s_i^{a_i}$ and $v = \prod_{i=j+1}^n s_i^{a_i}$. Then $\overline{st} = \prod_{i=1}^{j-1} s_i^{a_i} \cdot s_j^{a_j+1} \prod_{i=j+1}^n s_i^{a_i} = us_j v$, so

$$\overline{st}^{-1} st = v^{-1} s_j^{-1} u^{-1} s_j uv = [s_j, u]^v = ([u, s_j]^v)^{-1}.$$

Using the formulas $[xy, z] = [x, z]^y [y, z]$ and $[x^{-1}, y] = ([x, y]^{x^{-1}})^{-1}$ and the fact that \mathcal{F} is Schreier, it is straightforward to express $[u, s_j]^v$ in terms of elements of type (a). Let us illustrate this in the case $j = 3$, $u = s_1^2 s_2^{-1}$ (so that $v = \prod_{i=3}^n s_i^{a_i}$ for some a_i). We have

$$\begin{aligned} [u, s_j]^v &= [s_1^2 s_2^{-1}, s_3]^v = [s_1^2, s_3]^{s_2^{-1} v} [s_2^{-1}, s_3]^v \\ &= [s_1, s_3]^{s_1 s_2^{-1} v} [s_1, s_3]^{s_2^{-1} v} ([s_2, s_3]^{s_2^{-1} v})^{-1}. \end{aligned} \quad (2.3)$$

By assumption, $(2, -1, a_3, \dots, a_n) \in \mathcal{F}$. Since \mathcal{F} is Schreier, $(1, -1, a_3, \dots, a_n) \in \mathcal{F}$ and $(0, -1, a_3, \dots, a_n) \in \mathcal{F}$ as well. Thus, all the factors in the last expression in (2.3) are generators of type (a), as desired.

Case 2: $s \in S_2$. Let $j = d(s)$. As in case 1, we write $t = uv$, where $u = \prod_{i=1}^j s_i^{a_i}$ and $v = \prod_{i=j+1}^n s_i^{a_i}$, so $\overline{st} = us_j v$. Hence $\overline{st}^{-1} st = (us_j v)^{-1} suv$ and therefore

$$(\overline{st}^{-1} st)^{-1} (s_j t^{-1} s_j t) = (suv)^{-1} us_j v (us_j v)^{-1} s_j uv = (s^{-1} s_j)^{uv} = (s^{-1} s_{d(s)})^t. \quad (2.4)$$

The last expression in (2.4) is an element of type (b). Since $\overline{s_j t}^{-1} s_j t$ lies in K' by case 1, it follows from (2.4) that $\overline{st}^{-1} st \in K'$ as well.

Case 3: $s \in S_3$. In this case $s \in K$, so $\overline{st} = t$ and hence $\overline{st}^{-1} st = s^t$, which is an element of type (c). ■

3. Effective finite generation of $[IA_n, IA_n]$

In this section, we will prove Theorem 1.1. Throughout this section we fix an integer $n \geq 2$ and let $[n] = \{1, 2, \dots, n\}$.

Magnus [18] proved that IA_n is generated by the elements K_{ij} with $i \neq j \in [n]$ and K_{ijk} with $i, j, k \in [n]$, i, j, k distinct defined by

$$K_{ij}: \begin{cases} x_i \mapsto x_j^{-1} x_i x_j, \\ x_k \mapsto x_k \quad \text{for } k \neq i, \end{cases} \quad K_{ijk}: \begin{cases} x_i \mapsto x_i [x_j, x_k], \\ x_l \mapsto x_l \quad \text{for } l \neq i. \end{cases}$$

Clearly, $K_{ikj} = K_{ijk}^{-1}$, so IA_n is generated by the set $\{K_{ij}\} \cup \{K_{ijk}: j < k\}$.

Throughout this section, we set $S = \{K_{ij}\} \cup \{K_{ijk}: j < k\}$ and will refer to S as the *Magnus generating set* for IA_n .⁵ An easy computation shows that $|S| = n \binom{n}{2}$.

The commutation relations below between the Magnus generators of IA_n are straightforward to check.

Lemma 3.1. *The following hold:*

- (a) $[K_{ij}, K_{kl}] = 1$ if $i \notin \{k, l\}$ and $k \notin \{i, j\}$.
- (b) $[K_{ij}, K_{klm}] = 1$ if $i \notin \{k, l, m\}$ and $k \notin \{i, j\}$.

In particular, two Magnus generators commute if their sets of indices are disjoint.

Since IA_n is normal in $\text{Aut}(F_n)$ and $\text{Aut}(F_n)/IA_n \cong \text{GL}_n(\mathbb{Z})$, the abelianization $IA_n^{\text{ab}} \cong IA_n/[IA_n, IA_n]$ has the natural structure of a $\text{GL}_n(\mathbb{Z})$ -module. As a $\text{GL}_n(\mathbb{Z})$ -module, IA_n^{ab} is canonically isomorphic to $V^* \otimes (V \wedge V)$, where $V = \mathbb{Z}^n$, considered as a standard $\text{GL}_n(\mathbb{Z})$ -module, and $V^* = \text{Hom}(V, \mathbb{Z})$ is the dual module. This isomorphism was first established by Formanek [11], but there are several alternative proofs in the literature (e.g., see [7]).

Let e_1, \dots, e_n be the standard basis of V , and let e_1^*, \dots, e_n^* be the dual basis. Given $x \in IA_n$, let $[x]$ denote the image of x in IA_n^{ab} . The above isomorphism $IA_n^{\text{ab}} \cong V^* \otimes (V \wedge V)$ is given by

$$[K_{ij}] \mapsto e_i^* \otimes (e_i \wedge e_j) \quad \text{and} \quad [K_{ijk}] \mapsto e_i^* \otimes (e_j \wedge e_k). \quad (3.1)$$

From now on, we will identify IA_n^{ab} with $V^* \otimes (V \wedge V)$ via map (3.1).

Let $N = n \binom{n}{2} = \frac{n^2(n-1)}{2}$. By the above discussion, $IA_n^{\text{ab}} \cong \mathbb{Z}^N$ as abelian groups, and moreover the natural projection $IA_n \rightarrow IA_n^{\text{ab}}$ is injective on $S = \{K_{ij}\} \cup \{K_{ijk}: j < k\}$ and maps S to $E = \{e_i^* \otimes (e_j \wedge e_k): j < k\}$, which is a basis of IA_n^{ab} . In particular, S satisfies the conclusion of Lemma 2.3 for $G = IA_n$ and $K = [IA_n, IA_n]$.

⁵New, more geometric, proofs of the fact that S generates IA_n were given in [1, 7]. The proof given in [18] has two parts: one first shows that S generates IA_n as a normal subgroup of $\text{Aut}(F_n)$ and then shows that the subgroup generated by S is normal in $\text{Aut}(F_n)$. The articles [1, 7] gave very different proofs for the first part, but followed the original argument of Magnus for the second part.

Now let

$$Z = \{K_{12}, K_{34}, K_{56}, K_{78}\}.$$

By Lemma 3.1, elements of Z commute with each other, and it is easy to check that every element of S commutes with an element of Z . These two properties immediately imply that the pair (S, Z) is chain-centralizing. We now need to construct Φ satisfying the regularity hypothesis, but first we make some general observations.

Action on the space of characters. For any group G , we have a natural action of $\text{Aut}(G)$ on the space of characters $\text{Hom}(G, \mathbb{R})$ given by

$$(\phi\chi)(x) = \chi(\phi^{-1}(x)) \quad \text{for any } x \in G \text{ and } \phi \in \text{Aut}(G).$$

Clearly, $\text{Inn}(G)$ acts trivially, so we get an action of $\text{Out}(G) = \text{Aut}(G)/\text{Inn}(G)$.

Next note the centralizer of IA_n in $\text{Aut}(F_n)$ is trivial (since already the centralizer of $\text{Inn}(F_n)$ in $\text{Aut}(F_n)$ is trivial) and hence the conjugation action of $\text{Aut}(F_n)$ on IA_n yields an embedding of $\text{Aut}(F_n)$ into $\text{Aut}(\text{IA}_n)$.⁶

This, in turn, induces an embedding of $\text{Aut}(F_n)/\text{IA}_n \cong \text{GL}_n(\mathbb{Z})$ into $\text{Out}(\text{IA}_n)$ and thereby an action of $\text{GL}_n(\mathbb{Z})$ on $\text{Hom}(\text{IA}_n, \mathbb{R})$. It is easy to check that this is a “standard” action, dual to the action of $\text{GL}_n(\mathbb{Z})$ on IA_n^{ab} discussed above, but it is important for us that it comes from an action of $\text{Aut}(F_n)$ on IA_n .

The key technical result that we will prove in this section is the following lemma.

Lemma 3.2. *There exists a finite subset Ω of $\text{GL}_n(\mathbb{Z})$ with the following properties:*

- (a) *For any nonzero character χ of $G = \text{IA}_n$, there is $g \in \Omega$ such that $|g\chi(z)| \geq \frac{M(\chi)}{3}$ for all $z \in Z$.*
- (b) *Every $g \in \Omega$ admits a lift $\phi \in \text{Aut}(F_n)$ with $B(\phi^{\pm 1}) \leq 150$ and $A(\phi^{\pm 1}) \leq 8100$.*

Remark. Recall that $M(\chi) = \max\{|\chi(s)| : s \in S\}$ and that throughout this section S is the Magnus generating set of IA_n .

As a consequence of Lemma 3.2, we deduce that the pair $(G, K) = (\text{IA}_n, [\text{IA}_n, \text{IA}_n])$ satisfies the regularity hypothesis with explicit constants.

Corollary 3.3. *Let $G = \text{IA}_n$ and $K = [\text{IA}_n, \text{IA}_n]$. Then the regularity hypothesis holds for $C = 3$ and some finite Φ satisfying $B(\Phi \cup \Phi^{-1}) \leq 150$ and $A(\Phi \cup \Phi^{-1}) \leq 8100$.*

Proof. Let Φ be the set of all elements ϕ^{-1} , where ϕ ranges over all lifts from the conclusion of Lemma 3.2 (b). Then Φ is finite (since Ω is finite) and $B(\Phi \cup \Phi^{-1}) \leq 150$ and $A(\Phi \cup \Phi^{-1}) \leq 8100$ by Lemma 3.2 (b).

Now let χ be any nonzero character of G . By Lemma 3.2 (a), there exists $g \in \Omega$ such that $|g\chi(z)| \geq \frac{M(\chi)}{3}$ for all $z \in Z$. Let $\phi \in \text{Aut}(F_n)$ be the lift of g from Lemma 3.2 (b).

⁶In fact, it has been recently proved in [3] that this embedding is an isomorphism for $n \geq 3$, but this is not essential for our purposes.

Since $\chi(\phi^{-1}(x)) = (\phi\chi)(x) = g\chi(x)$ for all $x \in G$ and $\phi^{-1} \in \Phi$ by construction, regularity hypothesis holds for this Φ and $C = 3$. ■

The proof of Lemma 3.2 will consist of two parts. First we will construct g satisfying (a). This will be done in several steps, and g will be constructed as a product of at most 9 unit transvections and at most 2 permutation matrices (this ensures that there are only finitely many possibilities for g). Then we will prove (b) using the specific form of g constructed in the proof of (a).

In the computations below, it will be convenient to use the following notation: for a character λ of G and $i, j, k \in [n]$, we set

$$c_{ijk}(\lambda) = \lambda(e_i^* \otimes (e_j \wedge e_k)).$$

Note that we can reformulate the condition on g in Lemma 3.2 (a) in terms of the coefficients c_{ijk} as follows:

$$|c_{iij}(g\chi)| \geq \frac{M(\chi)}{3} \quad \text{for } (i, j) \in \{(1, 2), (3, 4), (5, 6), (7, 8)\}. \quad (3.2)$$

Given $i, j \in [n]$ with $i \neq j$ and a permutation σ of $[n]$, define $E_{ij}, F_\sigma \in \mathrm{GL}_n(\mathbb{Z})$ by

$$E_{ij}: \begin{cases} e_j \mapsto e_j + e_i, \\ e_k \mapsto e_k \quad \text{for } k \neq j, \end{cases} \quad F_\sigma: e_k \mapsto e_{\sigma(k)} \quad \text{for all } k. \quad (3.3)$$

Proof of Lemma 3.2 (a). In each step below, M will denote a positive real number, and λ will denote an arbitrary character of G (which will vary from step to step).

Step 1. If $M(\lambda) \geq M$, there is a permutation matrix g_1 such that $|c_{112}(g_1\lambda)| \geq M$ or $|c_{132}(g_1\lambda)| \geq M$.

By definition of $M(\lambda)$, we have $|c_{iij}(\lambda)| \geq M$ or $|c_{ijk}(\lambda)| \geq M$ for some distinct i, j, k . The result now follows from the obvious fact that

$$c_{ijk}(F_\sigma\lambda) = c_{\sigma^{-1}(i)\sigma^{-1}(j)\sigma^{-1}(k)}(\lambda).$$

In the next step, we give different arguments depending on which case occurred in step 1.

Step 2A. If $|c_{132}(\lambda)| \geq M$ and $|c_{112}(\lambda)| < M$, there exists $g_2 = E_{31}^{\pm 2}$ such that

$$|c_{112}(g_2\lambda)| \geq M \quad \text{and} \quad |c_{332}(g_2\lambda)| \geq M.$$

We have

$$c_{112}(E_{31}^{\pm 2}\lambda) = c_{112}(\lambda) \mp 2c_{132}(\lambda), \quad c_{332}(E_{31}^{\pm 2}\lambda) = c_{332}(\lambda) \pm 2c_{132}(\lambda).$$

By assumption, $|c_{112}(\lambda) \mp 2c_{132}(\lambda)| \geq M$ for any choice of sign and $|c_{332}(\lambda) \pm 2c_{132}(\lambda)| \geq M$ for some choice of sign, so either E_{31}^2 or E_{31}^{-2} can be used as g_2 .

Step 2B. If $|c_{112}(\lambda)| \geq M$, there exists $g \in \text{GL}_n(\mathbb{Z})$ which is either the identity matrix or $E_{13}^\varepsilon E_{31}^\varepsilon$ with $\varepsilon = \pm 1$ such that $|c_{112}(g_2\lambda)| \geq M$ and $|c_{332}(g_2\lambda)| \geq \frac{M}{3}$.

If $|c_{332}(\lambda)| \geq \frac{M}{3}$, then $g_2 = 1$ obviously works, so assume that $|c_{332}(\lambda)| < \frac{M}{3}$. For $\varepsilon = \pm 1$ by direct computation, we have

$$\begin{aligned} c_{112}(E_{13}^\varepsilon E_{31}^\varepsilon \lambda) &= 2c_{112}(\lambda) - c_{332}(\lambda) + \varepsilon(c_{312}(\lambda) - 2c_{132}(\lambda)), \\ c_{332}(E_{13}^\varepsilon E_{31}^\varepsilon \lambda) &= 2c_{332}(\lambda) - c_{112}(\lambda) - \varepsilon(c_{312}(\lambda) - 2c_{132}(\lambda)). \end{aligned}$$

Since $|c_{332}(\lambda)| < \frac{M}{3}$ and $|c_{112}(\lambda)| > M$, we have $|2c_{112}(\lambda) - c_{332}(\lambda)| > \frac{5M}{3}$ and $|2c_{332}(\lambda) - c_{112}(\lambda)| > \frac{M}{3}$; moreover, $2c_{112}(\lambda) - c_{332}(\lambda)$ and $2c_{332}(\lambda) - c_{112}(\lambda)$ have different signs. Therefore, choosing $\varepsilon = \pm 1$ such that $2c_{112}(\lambda) - c_{332}(\lambda)$ and $\varepsilon(c_{312}(\lambda) - 2c_{132}(\lambda))$ have the same sign, we obtain the desired g_2 .

Step 3. If $|c_{112}(\lambda)| \geq M$ and $|c_{332}(\lambda)| \geq \frac{M}{3}$, then there exists $g_3 = E_{24}^{\pm 1}$ such that $|c_{112}(g_3\lambda)| \geq M$ and $|c_{334}(g_3\lambda)| \geq \frac{M}{3}$.

This is clear since $c_{112}(E_{24}^{\pm 1}\lambda) = c_{112}(\lambda)$ and $c_{334}(E_{24}^{\pm 1}\lambda) = c_{334}(\lambda) \mp c_{332}(\lambda)$.

Step 4. If $|c_{112}(\lambda)| \geq M$ and $|c_{334}(\lambda)| \geq \frac{M}{3}$, there exists g_4 equal to either $E_{26}^{\pm 1}$ or $E_{26}^{\pm 1} E_{15}^\varepsilon E_{51}^\varepsilon$ for $\varepsilon = \pm 1$ such that $|c_{112}(g_4\lambda)| \geq M$ and $|c_{ijj}(g_4\lambda)| \geq \frac{M}{3}$ for $(i, j) = (3, 4)$ and $(5, 6)$.

To do this, we apply steps 2 and 3 with indices 3 and 4 replaced by 5 and 6. Since we will be acting by matrices E_{ij} with $i, j \notin \{3, 4\}$, the value of c_{334} will not change.

Step 5. If $|c_{112}(\lambda)| \geq M$ and $|c_{ijj}(\lambda)| \geq \frac{M}{3}$ for $(i, j) = (3, 4)$ and $(5, 6)$, there exists g_5 equal to either $E_{28}^{\pm 1}$ or $E_{28}^{\pm 1} E_{17}^\varepsilon E_{71}^\varepsilon$ for $\varepsilon = \pm 1$ such that $|c_{112}(g_5\lambda)| \geq M$ and $|c_{ijj}(g_5\lambda)| \geq \frac{M}{3}$ for $(i, j) = (3, 4), (5, 6)$ and $(7, 8)$.

The argument here is identical to step 4.

Putting all the steps together, we obtain the desired $g \in \text{GL}_n(\mathbb{Z})$ given by the product $g_5 g_4 g_3 g_2 g_1$ (recall that g_i is the matrix we acted by in step i). ■

Before turning to the proof of Lemma 3.2(b), we will first explain how the lifts from the conclusion of Lemma 3.2(b) will be constructed and derive some general bounds on the constants A and B for certain maps.

Constructing lifts. Since g in the proof of Lemma 3.2(a) is explicitly constructed as a product, we can obtain a lift of g by simply lifting each factor. The natural lift of a transposition matrix F_{ij} is $\widetilde{F}_{ij} \in \text{Aut}(F_n)$ which swaps x_i and x_j and fixes other generators. A transvection matrix E_{ij} can be lifted to either left or right Nielsen map R_{ji} or L_{ji} defined by

$$R_{ji}: \begin{cases} x_j \mapsto x_j x_i, \\ x_k \mapsto x_k \text{ for } k \neq j, \end{cases} \quad L_{ji}: \begin{cases} x_j \mapsto x_i x_j, \\ x_l \mapsto x_l \text{ for } l \neq i. \end{cases}$$

Note that we can use different lifts for different occurrences of E_{ij} but this does not seem to matter for the resulting bound.

It is clear that $A(\widetilde{F_{ij}}) = B(\widetilde{F_{ij}}) = 1$, and an explicit computation (see [7, Appendix A, Table 1]) shows that

$$A(R_{ij}^{\pm 1}) = A(L_{ij}^{\pm 1}) = 6 \quad \text{and} \quad B(R_{ij}^{\pm 1}) = L(R_{ij}^{\pm 1}) = 4.$$

Note that g in the proof of Lemma 3.2 (a) is a product of at most 2 transpositions and at most 9 unit transvections. Combining these facts with the easy observation that

$$A(\phi\psi) \leq A(\phi)A(\psi) \quad \text{and} \quad B(\phi\psi) \leq B(\phi)B(\psi) \quad (3.4)$$

for any $\phi, \psi \in \text{Aut}(F_n)$, we already deduce that g has a lift ϕ with $A(\phi^{\pm 1}) \leq 6^9$ and $B(\phi^{\pm 1}) \leq 4^9$.

To improve those bounds, we will prove the following general lemma.

Lemma 3.4. *Let G, K and S be as in Lemma 2.3. The following hold:*

- (a) *For any $\phi, \psi \in \text{Aut}(G)$, we have $A(\phi\psi) \leq A(\psi)B(\phi) + A(\phi)$.*
- (b) *Suppose that $\phi_1, \dots, \phi_k \in \text{Aut}(G)$ with $A(\phi_i) \leq A$ for each i . For $1 \leq m \leq k$, define $\psi_m = \prod_{i=1}^{m-1} \phi_i$ (by convention $\psi_1 = 1$). Then*

$$A(\phi_1 \dots \phi_k) \leq A \sum_{m=1}^k B(\psi_m).$$

Proof. Item (b) follows from (a) by straightforward induction, so we will only prove (a). Fix $s \in S^{\pm 1}$. By definition of $A(\psi)$ and Observation 2.9, there exists an S -word $s_1 \dots s_r$ representing $\psi(s)$ such that for all $1 \leq j \leq r$, one can write $\theta(s_1 \dots s_j)$ as a product of at most $A(\psi)$ elements $\theta(s)$, $s \in S^{\pm 1}$. Next for each $s \in S^{\pm 1}$ choose an S -word w_s representing $\phi(s)$ such that $\|\theta(u)\|_1 \leq A(\phi)$ for every S -prefix u of w_s .

Consider the S -word $w = w_{s_1} \dots w_{s_r}$. Then w represents $\phi\psi(s)$, and any S -prefix of w is equal to $w_{s_1} \dots w_{s_{j-1}}u$ for some $1 \leq j \leq r$ and S -prefix u of w_{s_j} . By assumption, $\theta(w_{s_1} \dots w_{s_{j-1}})$ is the product of at most $A(\psi)$ elements $\theta(w_s)$, $s \in S^{\pm 1}$. Since $\|\theta(w_s)\|_1 \leq B(\phi)$ by definition of $B(\phi)$ and $\|\theta(u)\|_1 \leq A(\phi)$, we have

$$\|\theta(w_{s_1} \dots w_{s_{j-1}}u)\|_1 \leq \|\theta(w_{s_1} \dots w_{s_{j-1}})\|_1 + \|\theta(u)\|_1 \leq A(\psi)B(\phi) + A(\phi).$$

Thus, $A(\phi\psi) \leq A(\psi)B(\phi) + A(\phi)$, as desired. ■

We can now finish the proof of Lemma 3.2.

Proof of Lemma 3.2 (b). First note that for $\phi \in \text{Aut}(F_n)$, the constant $B(\phi)$ depends only on the image of ϕ in $\text{GL}_n(\mathbb{Z})$, so we can talk about $B(g)$ for $g \in \text{GL}_n(\mathbb{Z})$.

We will consider the case where g constructed in the proof of part (a) is equal to $E_{28}E_{17}E_{71}E_{26}E_{15}E_{51}E_{24}E_{13}E_{31}$. It is not hard to check that this g represents the worst-case scenario. It is also easy to see that $B(g) \geq B(h)$ for any prefix h of g and the same is true for g^{-1} .

To estimate $B(g^{-1})$, we first compute the action of g on the basis elements of V and V^* . We have

$$\begin{aligned}
 ge_1 &= 8e_1 + 4e_7 + 2e_5 + e_3, & ge_1^* &= e_1^* - e_3^* - e_5^* - e_7^*, \\
 ge_2 &= e_2, & ge_2^* &= e_2^* - e_4^* - e_6^* - e_8^*, \\
 ge_3 &= 4e_1 + 2e_7 + e_5 + e_3, & ge_3^* &= 2e_3^* - e_1^* + e_5^* + e_7^*, \\
 ge_4 &= e_4 + e_2, & ge_4^* &= e_4^*, \\
 ge_5 &= 2e_1 + e_7 + e_5, & ge_5^* &= 2e_5^* - e_1^* + e_7^*, \\
 ge_6 &= e_6 + e_2, & ge_6^* &= e_6^*, \\
 ge_7 &= e_1 + e_7, & ge_7^* &= 2e_7^* - e_1^* + e_7^*, \\
 ge_8 &= e_8 + e_2, & ge_8^* &= e_8^*.
 \end{aligned}$$

Now it is easy to see that $\|g(e_i^* \otimes (e_j \wedge e_k))\|_1$ takes on its maximal value when $i = 3$, $j = 1$ and $k = 4, 6$ or 8 , and the maximum is equal to $5 \cdot 15 \cdot 2 = 150$. It follows that $B(g^{-1}) \leq 150$ and similarly $B(g) \leq 150$.

Now let ϕ be the lift of g or g^{-1} defined above. By construction, we can write $\phi = \prod_{i=1}^k \phi_i$, where $k \leq 9$ and $A(\phi_i) \leq 6$ for all i . We just argued that $B(\prod_{i=1}^j \phi_i) \leq 150$ for all $1 \leq j \leq k$. Hence Lemma 3.4 (b) yields $A(\phi) = A(\prod_{i=1}^k \phi_i) \leq 6k \cdot 150 \leq 6 \cdot 9 \cdot 150 = 8100$, as desired. ■

We are finally ready to prove Theorem 1.1. In fact, we will prove a slightly stronger statement.

Theorem 3.5. *Let $n \geq 8$. Let $N = n \binom{n}{2}$, and let $S = \{s_1, \dots, s_N\}$ be the Magnus generating set for IA_n . Then $[\text{IA}_n, \text{IA}_n]$ is generated by elements of the form*

$$[s_i, s_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_N^{a_N}}, \quad \text{where } 1 \leq i < j \leq N \text{ and } 0 \leq |a_m| < 5 \cdot 10^{12} \text{ for each } m. \quad (***)$$

In fact, $[\text{IA}_n, \text{IA}_n]$ is generated by elements of this form with the additional property that $|\{m: a_m \neq 0\}| \leq 8n^2$.

Proof. We start by proving that $[\text{IA}_n, \text{IA}_n]$ is generated by all elements of form (***), which is precisely the assertion of Theorem 1.1.

First we compute the radius R in Theorem 2.13 applied to $G = \text{IA}_n$ and $K = [\text{IA}_n, \text{IA}_n]$. Corollary 3.3 shows that for a suitable Φ , we have $B \leq 150$, $A \leq 8100$ and can take $C = 3$. Hence the preimage of the l^∞ -ball of radius

$$R = 16(BC + 1)^2(AB + 3A + 3) \leq 16 \cdot 451^2 \cdot (8100 \cdot 153 + 3) < 5 \cdot 10^{12}$$

is connected. Now Theorem 1.1 follows directly from Theorem 2.16 (note for $G = \text{IA}_n$, we have $S = S_1$ in the notations from that theorem).

Let us now prove the full statement of Theorem 3.5. As before, we identify IA_n^{ab} with \mathbb{Z}^N via the map

$$\prod_{i=1}^N s_i^{a_i} [\text{IA}_n, \text{IA}_n] \mapsto (a_1, \dots, a_N).$$

For a point $(a_1, \dots, a_N) \in \mathbb{Z}^N$, define its support as

$$\text{supp}((a_1, \dots, a_N)) = |\{m: a_m \neq 0\}|.$$

Let \mathcal{F} be the set of all $(a_1, \dots, a_N) \in \mathbb{Z}^N$ with $|a_i| \leq R$ for each i and

$$|\text{supp}((a_1, \dots, a_N))| \leq 8n^2.$$

Clearly, \mathcal{F} is a Schreier set. Thus, Theorem 2.16 reduces Theorem 3.5 to showing that $\theta^{-1}(\mathcal{F})$ is connected (where $\theta: \text{IA}_n \rightarrow \mathbb{Z}^N$ is the natural projection).

So take any $x, y \in \text{IA}_n$ with $\theta(x), \theta(y) \in \mathcal{F}$. We already know that x and y can be connected by a path p which lies in the θ -preimage of the l^∞ -ball of radius R . If $\theta(g) \in \mathcal{F}$ for every vertex g of p , we are done; otherwise, consider all vertices g on p such that $|\text{supp}(\theta(g))|$ is largest possible, and among these vertices (if there is more than one), choose one where $\|\theta(g)\|_1$ is maximal. In particular, by assumption $|\text{supp}(\theta(g))| > 8n^2$ and thus $\theta(g) \neq \theta(x), \theta(y)$.

The vertices of p which precede and succeed g have the form gy_1 and gy_2 for some $y_1, y_2 \in S^{\pm 1}$. Write $\theta(g) = (a_1, \dots, a_N)$. Since $|\text{supp}(\theta(g))| > 8n^2$, there exist more than $8n^2$ indices m such that $a_m \neq 0$, and an easy calculation using Lemma 3.1 shows that there exists $1 \leq m \leq N$ such that $a_m \neq 0$, $s_m \neq y_1^{\pm 1}, y_2^{\pm 1}$ and s_m commutes with both y_1 and y_2 . Without loss of generality, we can assume that $a_m > 0$. Now modify p replacing the segment (gy_1, g, gy_2) of p by

$$(gy_1, gy_1s_m^{-1} = gs_m^{-1}y_1, gs_m^{-1}, gs_m^{-1}y_2 = gy_2s_m^{-1}, gy_2).$$

Another easy calculation shows that for every vertex v on this segment, we have

$$|\text{supp}(\theta(v))| \leq |\text{supp}(\theta(g))| \quad \text{and} \quad \|\theta(v)\|_1 < \|\theta(g)\|_1.$$

Thus, after applying such modification finitely many times, we will obtain a path connecting x to y which lies in $\theta^{-1}(\mathcal{F})$, as desired. \blacksquare

4. Effective finite generation of $[\mathcal{I}_n^1, \mathcal{I}_n^1]$ and the Johnson kernel

4.1. Preliminaries

Throughout this section, we fix an integer $n \geq 0$ and let $\Sigma = \Sigma_n^1$ be an orientable surface of genus n with 1 boundary component. We start by introducing some curves and subsurfaces on Σ that will be used throughout the proof.

First, it will be convenient to think of Σ as a (closed) disk with n handles attached; let us number the handles from 1 to n . We also fix a point p_0 on the boundary $\partial\Sigma$ (it will serve as the base point for all the fundamental groups considered below).

For each $1 \leq i \leq n$, choose a point p_i on the i -th handle and curves a_i and b_i passing through p_i as shown on Figure 3. Also choose (oriented) paths γ_i from p_0 to p_i as in Figure 3. In particular, we require the sets $\gamma_i \setminus \{p_0\}$ to be disjoint.




Figure 3. The curves a_i and b_i and the paths γ_i .

Next define the curves α_i and β_i by

$$\alpha_i = \gamma_i a_i \gamma_i^{-1} \quad \text{and} \quad \beta_i = \gamma_i b_i \gamma_i^{-1};$$

note that these are closed curves based at p_0 (see Figure 3). For simplicity, we will also use the notations α_i and β_i for the corresponding classes in the fundamental group $\pi_1(\Sigma, p_0)$.

Let $(\partial\Sigma)_{p_0}$ denote the boundary of Σ considered as a closed path from p_0 to itself oriented clockwise. It is easy to check that $(\partial\Sigma)_{p_0}$ is homotopic to $\prod_{i=1}^n [\alpha_i, \beta_i]$.

The *mapping class group* $\text{Mod}_n^1 = \text{Mod}(\Sigma)$ is defined as the subgroup of orientation preserving homeomorphisms of Σ which fix the boundary $\partial\Sigma$ pointwise modulo the isotopies which fix $\partial\Sigma$ pointwise. The action of $\text{Mod}(\Sigma)$ on Σ induces an action on $\pi = \pi_1(\Sigma, p_0)$ and thus we obtain a homomorphism $\iota: \text{Mod}(\Sigma) \rightarrow \text{Aut}(\pi)$. The group π is free of rank $2n$ with generators $\alpha_1, \beta_1, \dots, \alpha_n, \beta_n$, and since $(\partial\Sigma)_{p_0}$ is fixed under the action, the image of ι stabilizes $\prod_{i=1}^n [\alpha_i, \beta_i]$ by property (ii) above. In fact, a stronger statement holds.

Theorem 4.1. *The map ι is injective and $\text{Im } \iota$ is equal to the full stabilizer of $\prod_{i=1}^n [\alpha_i, \beta_i]$ in $\text{Aut}(\pi)$.*

Theorem 4.1 is proved, for example, in [28]: the above map ι is injective by [28, Theorem 5.13.2] and surjective (that is, $\text{Im } \iota$ is the full stabilizer) by [28, Theorem 5.7.1]. The surjectivity part is originally due to Zieschang [27].

Remark. (a) The surjectivity part of Theorem 4.1 can be rephrased by saying that $\text{Mod}(\Sigma)$ acts transitively on the set of (ordered) bases $\delta_1, \delta'_1, \dots, \delta_n, \delta'_n$ of π satisfying $\prod_{i=1}^n [\delta_i, \delta'_i] = (\partial\Sigma)_{p_0}$. Recall that such bases were called *natural* in the introduction.

(b) Theorem 4.1 is a variation of the classical Dehn–Nielsen–Baer theorem which asserts that for a closed surface Σ of genus n , the mapping class group $\text{Mod}(\Sigma)$ is isomorphic to an index 2 subgroup of the outer automorphism group of a surface group on $2n$ generators.

Johnson filtration. For each $k \in \mathbb{N}$, define $\mathcal{I}_n^1(k)$ to be the kernel of the induced map $\mathrm{Mod}_n^1 \rightarrow \mathrm{Aut}(\pi/\gamma_{k+1}\pi)$. The filtration $\{\mathcal{I}_n^1(k)\}_{k=1}^\infty$ is called the *Johnson filtration*. The first term of this filtration $\mathcal{I}_n^1 = \mathcal{I}_n^1(1)$ is the *Torelli subgroup* of Mod_n^1 . It can also be defined as the set of elements of Mod_n^1 acting trivially on $H_1(\Sigma_n^1)$. The second term of the Johnson filtration $\mathcal{K}_n^1 = \mathcal{I}_n^1(2)$ is known as the *Johnson kernel*. One can characterize \mathcal{K}_n^1 purely topologically as the subgroup generated by Dehn twists about separating curves. The equivalence of these two definitions of \mathcal{K}_n^1 is a deep theorem of Johnson [15].

Recall that we consider n as being fixed, and for the rest of the section, we will use the simplified notations $\mathcal{M} = \mathrm{Mod}_n^1$, $\mathcal{I} = \mathcal{I}_n^1$ and $\mathcal{K} = \mathcal{K}_n^1$. Occasionally, we will also use the notations $\mathcal{I}(\Omega)$ and $\mathcal{K}(\Omega)$ for the Torelli subgroup (resp. Johnson kernel) of the mapping class group of a surface Ω .

4.2. Generators for the mapping class group

It is well known that the mapping class group Mod_n^1 is generated by Dehn twists. The minimal number of Dehn twists needed to generate Mod_n^b for $b = 0, 1$ is $2n + 1$ as proved by Humphries [12] for $b = 0$ and by Johnson [14] for $b = 1$ (the generating set in [14] is a natural analogue of the one in [12]). More specifically, Mod_n^1 is generated by the Dehn twists about the curves $c_1, c_2, \dots, c_{2n}, b$ defined in [14, p. 428, Figure 5].

Usually, in the definition of the Dehn twist T_γ one assumes that γ is an essential simple closed curve, but for the discussion below it will be convenient to introduce the following convention: If γ is a closed curve on Σ which is not simple, but freely homotopic to some essential simple closed curve γ' , we set $T_\gamma = T_{\gamma'}$. The right-hand side is well defined since two freely homotopic curves on a surface are isotopic, and the Dehn twist $T_{\gamma'}$ is determined by the isotopy class of γ' .

With this convention, we can relate the Humphries-Johnson generating set to the curves α_i, β_i introduced in Section 4.1. It is not hard to see that c_{2i} is freely homotopic to β_i for $1 \leq i \leq n$, that c_{2i-1} is freely homotopic to $\alpha_i \alpha_{i-1}^{-1}$ for $2 \leq i \leq n$ and that c_1 and b are freely homotopic to α_1 and α_2 , respectively. Thus, [14, Theorem 3] can be restated as follows.

Theorem 4.2. *The mapping class group Mod_n^1 is generated by the following Dehn twists: $\{T_{\beta_i} : 1 \leq i \leq n\}$, $\{T_{\alpha_i \alpha_{i-1}^{-1}} : 2 \leq i \leq n\}$, T_{α_1} and T_{α_2} .*

We will not explicitly refer to Theorem 4.2 in this paper, but we will use several results whose proof relies on Theorem 4.2.

4.3. Generators for the Torelli subgroup and subsurfaces Σ_I

It is a celebrated theorem of Johnson [14] that the Torelli group $\mathcal{I} = \mathcal{I}_n^1$ is finitely generated for $n \geq 3$. Johnson's generating set from [14] is explicit, but it lacks a key feature of Magnus' generating set for IA_n – the fact that most generating pairs commute – that is essential for our purposes. A generating set for \mathcal{I} which has the latter property was constructed by Church and Putman [6] using an earlier work of Putman [22].




Figure 4. The surface $\Sigma_{\{1,3,4\}}$.

Recall the curves α_i and β_i , $1 \leq i \leq n$, defined in Section 4.1. For each $I \subseteq [n]$, choose a subsurface $\Sigma_I \subseteq \Sigma$ satisfying the following properties:

- (i) The curves α_i and β_i lie on Σ_I for all $i \in I$.
- (ii) Σ_I has genus $|I|$ and 1 boundary component.
- (iii) The boundary of Σ_I is homotopic to $\prod_{i \in I} [\alpha_i, \beta_i]$ (where the product is taken in increasing order).
- (iv) $\Sigma_I \cap \partial\Sigma$ is an interval which contains p_0 and does not depend on I .

For a longer but more transparent definition of Σ_I , see [6, §4] or [9, §7]. For an illustration, see Figure 4.

The subsurfaces Σ_I are uniquely defined up to isotopy and satisfy the following properties (1)–(4). Properties (1)–(3) follow immediately from the definitions and (4) can be proved by a standard application of the *change of coordinates principle* [10, §1.1.3].

Observation 4.3. *The following hold:*

- (1) $\Sigma_{[n]}$ is isotopic to Σ .
- (2) If $I \subseteq J$, there exist subsurfaces Σ'_I isotopic to Σ_I and Σ'_J isotopic to Σ_J such that $\Sigma'_I \subseteq \Sigma'_J$.
- (3) If I and J are disjoint and uncrossed (as defined below), there exist subsurfaces Σ'_I isotopic to Σ_I and Σ'_J isotopic to Σ_J such that Σ'_I and Σ'_J are disjoint.
- (4) If $|I| = |J|$, there exists an orientation-preserving homeomorphism g of Σ acting trivially on $\partial\Sigma$ such that $g(\Sigma_I) = \Sigma_J$.

Definition 4.4. Let I and J be disjoint subsets of $[n]$. We will say that I and J are *crossed* if there exist $i_1, i_2 \in I$ and $j_1, j_2 \in J$ such that $i_1 < j_1 < i_2 < j_2$ or $j_1 < i_1 < j_2 < i_2$. Otherwise, I and J will be called *uncrossed*. Clearly, if I consists of consecutive integers, then I is uncrossed with any subset J disjoint from it.

Remark. The technical condition (iv) in the definition of Σ_I (which was not imposed in [6, 9]) is needed to ensure that property (4) in Observation 4.3 holds. Note that an easier way to achieve (4) would be to require that $\Sigma_I \cap \partial\Sigma = \emptyset$. However, the latter would prevent us from considering the fundamental groups of Σ_I as subgroups of $\pi_1(\Sigma, p_0)$, something that is essential for our purposes.

For each $I \subseteq [n]$, define Mod_I to be the subgroup of $\mathrm{Mod}(\Sigma)$ consisting of mapping classes which have a representative supported on Σ_I , and let $\mathcal{I}_I = \mathrm{Mod}_I \cap \mathcal{I}$. Parts (1)–(3) of Observation 4.3 have obvious group-theoretic consequences.

Observation 4.5. *The following hold:*

- (i) $\mathrm{Mod}_{[n]} = \mathrm{Mod}(\Sigma)$.
- (ii) If $I \subseteq J$, then $\mathrm{Mod}_I \subseteq \mathrm{Mod}_J$.
- (iii) If I and J are disjoint and uncrossed, then Mod_I and Mod_J commute.

A deep result of Church and Putman [6, Proposition 4.5] asserts that if $n \geq 3$, then $\mathcal{I} = \langle \mathcal{I}_I : |I| = 3 \rangle$. This fact played a key role in the proof of finite generation of \mathcal{K}_n^1 in [9], but in [9] there was no need to work with a specific finite generating set inside $\bigcup_{|I|=3} \mathcal{I}_I$. In this paper, we will need to choose such a generating set S as follows: we will start with some finite generating set $S_{\{1,2,3\}}$ of $\mathcal{I}_{\{1,2,3\}}$ and then add to it the images of $S_{\{1,2,3\}}$ under carefully chosen isomorphisms between $\mathcal{I}_{\{1,2,3\}}$ and \mathcal{I}_I for every 3-element subset I of $[n]$. The details of this construction will be given later in this section.

4.4. Abelian quotients of the Torelli subgroup

Let $V = H_1(\Sigma)$. Then V is a free abelian group of rank $2n$, and it is well known that the algebraic intersection form on V is symplectic. Clearly, \mathcal{M}/\mathcal{I} acts on V preserving this form, so there is a canonical group homomorphism $\mathcal{M}/\mathcal{I} \rightarrow \mathrm{Sp}(V)$ where $\mathrm{Sp}(V)$ is the group of automorphisms of V preserving this form. It is also well known that this homomorphism is an isomorphism, which enables us to identify \mathcal{M}/\mathcal{I} with $\mathrm{Sp}(V)$.

From the definition of the Johnson filtration, it is easy to see that the quotient $\mathcal{I}/\mathcal{K} = \mathcal{I}_n^1(1)/\mathcal{I}_n^1(2)$ is abelian and torsion-free. A complete description of the abelianization of \mathcal{I} for $n \geq 3$ was obtained in a series of Johnson's papers [13, 15, 16]. Below we collect some specific results about abelian quotients of \mathcal{I} that will be needed in this paper.

Theorem 4.6. *Assume that $n \geq 3$. The following hold:*

- (a) \mathcal{I}/\mathcal{K} is the largest torsion-free abelian quotient of \mathcal{I} .
- (b) There is a canonical isomorphism of $\mathrm{Sp}(V)$ -modules $\mathcal{I}/\mathcal{K} \cong \bigwedge^3 V$. In particular, $\mathcal{I}/\mathcal{K} \cong \mathbb{Z}^{\binom{2n}{3}}$ as a group.
- (c) $\mathcal{K}/[\mathcal{I}, \mathcal{I}]$ (which by (a) is the torsion part of $\mathcal{I}/[\mathcal{I}, \mathcal{I}]$) has exponent 2 and rank $\binom{2n}{2} + \binom{2n}{1} + \binom{2n}{0}$.

Remark. We briefly comment on how (a), (b) and (c) follow from the results of [13, 15, 16]. In [13], Johnson constructed

- an epimorphism $\tau: \mathcal{I} \rightarrow \bigwedge^3 V$ such that $\mathcal{K} \subseteq \mathrm{Ker} \tau$ and the induced map $\mathcal{I}/\mathcal{K} \rightarrow \bigwedge^3 V$ is a homomorphism of $\mathrm{Sp}(V)$ -modules;
- a group epimorphism $\sigma: \mathcal{I} \rightarrow B_3$, where B_3 is an elementary abelian 2-group of rank $\sum_{i=0}^3 \binom{2n}{i}$ (there is also a natural $\mathrm{Sp}(V)$ -module structure on B_3 , but it is not essential for our purposes).

In [15], it was proved that $\mathcal{K} = \text{Ker } \tau$ which yields (b). One of the main results of [16] is that $[\mathcal{I}, \mathcal{I}] = \mathcal{K} \cap \text{Ker } \sigma$. This implies that $\mathcal{I}/[\mathcal{I}, \mathcal{I}]$ embeds into $\mathcal{I}/\mathcal{K} \oplus B_3$ which yields (a). The other main result of [16] is that B_3 is the largest exponent 2 quotient of \mathcal{I} . Combined with (a) and (b), this implies (c).

We proceed with the description of \mathcal{I}/\mathcal{K} . Recall the curves a_i and b_i on Σ introduced in Section 4.1. By slight abuse of notation below, we will use the notations a_i and b_i for the corresponding homology classes in $H_1(\Sigma)$. It is clear that $a_1, b_1, \dots, a_n, b_n$ is a symplectic basis of $H_1(\Sigma)$, that is, $a_i \cdot a_j = b_i \cdot b_j = 0$ for all i, j and $a_i \cdot b_j = \delta_{ij}$. Also note that a_i and b_i span $H_1(\Sigma_{\{i\}})$.

Let $V = H_1(\Sigma)$, for $i \in [n]$ let $V_i = H_1(\Sigma_{\{i\}}) = \text{Span}\{a_i, b_i\}$, and for $I \subseteq [n]$, let

$$V_I = \bigoplus_{i \in I} V_i = \text{Span}\{a_i, b_i : i \in I\} \subseteq V.$$

Thus, $\bigwedge^3 V = \sum_{|I|=3} \bigwedge^3 V_I$. Recall that we have an isomorphism of $Sp(V)$ -modules $\mathcal{I}/\mathcal{K} \rightarrow \bigwedge^3 V$. For every $I \subseteq [n]$, let $\phi_I: \mathcal{I}_I \rightarrow \bigwedge^3 V$ be the map obtained by precomposing the isomorphism $\mathcal{I}/\mathcal{K} \rightarrow \bigwedge^3 V$ with the natural projection $\mathcal{I} \rightarrow \mathcal{I}/\mathcal{K}$ and the inclusion $\mathcal{I}_I \rightarrow \mathcal{I}$.

The following result is an immediate consequence of Johnson's paper [16].

Lemma 4.7. *For every $I \subseteq [n]$, we have $\phi_I(\mathcal{I}_I) = V_I$.*

We can now construct a generating set for \mathcal{I} with certain nice properties. First order the chosen basis of V as follows: $a_1 < b_1 < a_2 < \dots < b_n$, and let \mathcal{B} be the set of all wedges $x \wedge y \wedge z$ with $x, y, z \in \{a_i, b_i\}$ and $x < y < z$. Clearly, \mathcal{B} is a basis for $\bigwedge^3 V$.

Lemma 4.8. *Suppose that $n \geq 3$. Then for every $I \subseteq [n]$ with $|I| = 3$, there exists a generating set S_I for \mathcal{I}_I with $|S_I| = 42$ whose image in \mathcal{I}/\mathcal{K} is equal to $(\mathcal{B} \cap V_I) \sqcup \{0\}$. Moreover, if we let $S = \bigcup_{|I|=3} S_I$, then S is generating set for \mathcal{I} which satisfies the conclusion of Lemma 2.3 for $G = \mathcal{I}$, $K = \mathcal{K}$ and $E = \mathcal{B}$.*

Proof. By [14], the group $\mathcal{I}_I \cong \mathcal{I}_3^1$ has a generating set S_I^0 with 42 elements. By Lemma 4.7, the image of \mathcal{I}_I in V is equal to V_I , which is a free abelian group of rank 20 with basis $\mathcal{B} \cap V_I$. As in the proof of Lemma 2.3, applying suitable Nielsen moves to S_I^0 (replacing x by $xy^{\pm 1}$, where x and y are distinct generators), we obtain another generating set S_I of \mathcal{I}_I with 42 elements whose image in V is equal to $(\mathcal{B} \cap V_I) \sqcup \{0\}$.

The set $S = \bigcup_{|I|=3} S_I$ generates \mathcal{I} by [6, Proposition 4.5]. Also the image of S in $\bigwedge^3 V$ is equal to $\bigcup_{|I|=3} (\mathcal{B} \cap V_I) \sqcup \{0\} = \mathcal{B} \sqcup \{0\}$, so S indeed satisfies the conclusion of Lemma 2.3 (for the desired G , K and E). ■

Our effective generation procedure described in Section 2 can be applied to any generating set S constructed in Lemma 4.8. However, some additional compatibility assumptions on $\{S_I\}$ will be needed in order to explicitly estimate the constants A and B from Theorem 2.13. We postpone this discussion until the end of this section.

4.5. Some elements of $\text{Mod}(\Sigma)$

In this subsection, we will prove the existence of certain elements of $\text{Mod}(\Sigma)$ with a prescribed induced action on $\pi_1(\Sigma, p_0)$. These results will play a key role in estimating the constant A later this section.

Lemma 4.9. *Let I be a non-empty subset of $[n]$, and let $k = |I|$. Suppose we are given closed curves $\{\delta_j, \delta'_j\}_{j=1}^k$ which are based at p_0 , lie on Σ_I , form a basis of $\pi_1(\Sigma_I, p_0)$ and such that $\prod_{j=1}^k [\delta_j, \delta'_j] = (\partial \Sigma_I)_{p_0}$ in $\pi_1(\Sigma, p_0)$. Then there exists an element $f \in \text{Mod}(\Sigma)$ such that*

- (a) *some representative of f maps $\Sigma_{[k]}$ onto Σ_I ;*
- (b) *the induced action of f on $\pi_1(\Sigma, p_0)$ satisfies $f^*(\alpha_j) = \delta_j$ and $f^*(\beta_j) = \delta'_j$ for all $1 \leq j \leq k$.*

Moreover, suppose that $f' \in \text{Mod}(\Sigma)$ is another element satisfying (a) and (b). Then f' coincides with f on $\Sigma_{[k]}$, and therefore, $fsf^{-1} = f's(f')^{-1}$ for all $s \in \text{Mod}_{[k]}$.

Proof. We first construct an element f satisfying (a) and (b). By Observation 4.3 (4), there exists $g \in \text{Mod}(\Sigma)$ such that some representative of g maps $\Sigma_{[k]}$ to Σ_I .

Since $(\partial \Sigma_{[k]})_{p_0} = \prod_{j=1}^k [\alpha_j, \beta_j]$ and $(\partial \Sigma_I)_{p_0} = \prod_{j=1}^k [\delta_j, \delta'_j]$ in $\pi_1(\Sigma, p_0)$, the induced action of g on $\pi_1(\Sigma, p_0)$ satisfies

$$g^*\left(\prod_{j=1}^k [\alpha_j, \beta_j]\right) = \prod_{j=1}^k [\delta_j, \delta'_j]. \quad (4.1)$$

Since a representative of g maps $\Sigma_{[k]}$ homeomorphically onto Σ_I and $\{\alpha_j, \beta_j\}_{j=1}^k$ generate $\pi_1(\Sigma_{[k]}, p_0)$, the elements $\{g^*(\alpha_j), g^*(\beta_j)\}_{j=1}^k$ generate $\pi_1(\Sigma_I, p_0)$.

Since $\pi_1(\Sigma_I, p_0)$ is free of rank $2k$ and the automorphism group of a free group F acts transitively on the sets of bases of F , there exists $\phi \in \text{Aut}(\pi_1(\Sigma_I, p_0))$ such that $\phi(\delta_j) = g^*(\alpha_j)$ and $\phi(\delta'_j) = g^*(\beta_j)$ for all $j \in [k]$. Then

$$\begin{aligned} \phi\left(\prod_{j=1}^k [\delta_j, \delta'_j]\right) &= \prod_{j=1}^k [\phi(\delta_j), \phi(\delta'_j)] = \prod_{j=1}^k [g^*(\alpha_j), g^*(\beta_j)] \\ &= g^*\left(\prod_{j=1}^k [\alpha_j, \beta_j]\right) = \prod_{j=1}^k [\delta_j, \delta'_j] \end{aligned}$$

(where the last equality holds by (4.1)), so by Theorem 4.1 applied to the surface Σ_I , there exists $h \in \text{Mod}(\Sigma_I)$ such that $h^* = \phi$. If we extend h to $\text{Mod}(\Sigma)$ by letting it act trivially on $\Sigma \setminus \Sigma_I$, then clearly $f = h^{-1}g$ satisfies both (a) and (b).

Let us now prove the ‘moreover’ part. Suppose that $f' \in \text{Mod}(\Sigma)$ satisfies both (a) and (b), and let $h = (f')^{-1}f$. By (a), some representative of h stabilizes $\Sigma_{[k]}$. By (b), h^* acts trivially on $\pi_1(\Sigma_{[k]}, p_0)$, and hence by the injectivity part of Theorem 4.1, h acts

trivially on $\Sigma_{[k]}$, so f' coincides with f on $\Sigma_{[k]}$. Finally, any $s \in \text{Mod}_{[k]}$ is supported on $\Sigma_{[k]}$, so $hs = sh$, which yields the last assertion. ■

The following corollary describes a key special case of Lemma 4.9.

Corollary 4.10. *Let I be a subset of $[n]$, let $k = |I|$, and let $i_1 < i_2 < \dots < i_k$ be the elements of I listed in increasing order. Then there exists an element $f_I \in \text{Mod}(\Sigma)$ such that*

- (a) *some representative of f_I maps $\Sigma_{[k]}$ onto Σ_I ;*
- (b) *the induced action of f_I on $\pi_1(\Sigma)$ satisfies $f_I^*(\alpha_j) = \alpha_{i_j}$ and $f_I^*(\beta_j) = \beta_{i_j}$ for all $1 \leq j \leq k$.*

Moreover, if $f_I \in \text{Mod}(\Sigma)$ is any element satisfying (a) and (b), then for any subset J of $[k]$, some representative of f_I maps Σ_J onto Σ_{I_J} where $I_J = \{i_j : j \in J\}$.

Proof. The first assertion is a special case of Lemma 4.9, so we only need to prove the ‘moreover’ part. So take any $f_I \in \text{Mod}(\Sigma)$ satisfying (a) and (b). Then

$$f_I^*((\partial\Sigma_J)_{p_0}) = f_I^*\left(\prod_{j \in J} [\alpha_j, \beta_j]\right) = \prod_{i \in I_J} [\alpha_i, \beta_i] = (\partial\Sigma_{I_J})_{p_0}. \quad (4.2)$$

Thus, if \tilde{f}_I is any orientation-preserving homeomorphism of Σ which fixes $\partial\Sigma$ pointwise and represents f_I , then $\tilde{f}_I(\partial\Sigma_J)$ is homotopic (and hence isotopic) to $\partial\Sigma_{I_J}$. By [10, Proposition 1.11], an isotopy between $\tilde{f}_I(\partial\Sigma_J)$ and $\partial\Sigma_{I_J}$ can be extended to an isotopy of Σ (acting trivially on $\partial\Sigma$). Therefore, after perturbing \tilde{f}_I by an isotopy of Σ , we can assume that $\tilde{f}_I(\partial\Sigma_J) = \partial\Sigma_{I_J}$ and hence $\tilde{f}_I(\Sigma \setminus \partial\Sigma_J) = \Sigma \setminus \partial\Sigma_{I_J}$.

The topological spaces $\Sigma \setminus \partial\Sigma_J$ and $\Sigma \setminus \partial\Sigma_{I_J}$ both have two connected components: $\Sigma_J \setminus \partial\Sigma_J$ and $\Sigma \setminus \Sigma_J$ (resp. $\Sigma_{I_J} \setminus \partial\Sigma_{I_J}$ and $\Sigma \setminus \Sigma_{I_J}$). Since \tilde{f}_I acts trivially on $\partial\Sigma$ and since the intersection $(\Sigma \setminus \Sigma_J) \cap (\Sigma \setminus \Sigma_{I_J}) \cap \partial\Sigma$ is non-trivial by construction, \tilde{f}_I must map $\Sigma \setminus \Sigma_J$ to $\Sigma \setminus \Sigma_{I_J}$. Thus

$$\tilde{f}_I(\Sigma_J \setminus \partial\Sigma_J) = \Sigma_{I_J} \setminus \partial\Sigma_{I_J}$$

and hence $\tilde{f}_I(\Sigma_J) = \Sigma_{I_J}$. ■

4.6. An analogue of Lemma 3.2

In this subsection, we will establish an analogue of Lemma 3.2 for mapping class groups (see Lemma 4.11 below). The proof of the second part of Lemma 4.11 will be postponed till the next subsection, as estimation of the constants A and B in the mapping class group case requires more work.

Since \mathcal{I}/\mathcal{K} is the maximal torsion-free abelian quotient of \mathcal{I} , we have $\text{Hom}(\mathcal{I}, \mathbb{R}) = \text{Hom}(\mathcal{I}/\mathcal{K}, \mathbb{R})$. By the same logic as in Section 3, there is a natural embedding of $\mathcal{M} = \text{Mod}(\Sigma)$ into $\text{Aut}(\mathcal{I})$ and a natural action of $\text{Sp}(V) = \mathcal{M}/\mathcal{I}$ on $\text{Hom}(\mathcal{I}/\mathcal{K}, \mathbb{R})$.

The symplectic group $\text{Sp}(V) \cong \text{Sp}_{2n}(\mathbb{Z})$ is generated by the elements w_i, τ_i for $i \in [n]$ and τ_{ij} for $i \neq j \in [n]$ defined as follows (all basis elements whose image is not specified are fixed):

$$w_i: \begin{cases} a_i \mapsto b_i, \\ b_i \mapsto -a_i, \end{cases} \quad \tau_i: a_i \mapsto a_i + b_i, \quad \tau_{ij}: \begin{cases} a_i \mapsto a_i + a_j, \\ b_j \mapsto b_j - b_i. \end{cases}$$

In addition, for each $1 \leq i \neq j \leq n$, let $f_{ij} \in \text{Sp}(V)$ be the element which swaps a_i and a_j and swaps b_i and b_j . We will refer to the elements f_{ij} as transpositions.

Lemma 4.11. *Let S be as in Lemma 4.8, let χ be a nonzero character of $G = \mathcal{I}$ and let $M = M(\chi)$. The following hold:*

- (a) *There exists $g \in \text{Sp}_{2n}(\mathbb{Z})$ and $Z \subset S$ with the following properties:*
 - (i) *(S, Z) is chain-centralizing.*
 - (ii) *$|g\chi(z)| \geq M$ for all $z \in Z$.*
 - (iii) *g is a product of at most 15 elements of the form τ_{ij} and at most 3 transpositions f_{ij} .*
- (b) *There exists an absolute constant C_0 (independent of n) such that any g in part (a) admits a lift $\phi \in \text{Mod}_n^1$ with $A(\phi^{\pm 1}) \leq C_0$.*

Remark. Arguing exactly as in the proof of Corollary 3.3, we deduce from Lemma 4.11 that the pair $(G, K) = (\mathcal{I}, \mathcal{K})$ satisfies the regularity hypothesis for $C = 1$ and some finite Φ with $B(\Phi) \leq A(\Phi) \leq C_0$.

Proof. Unlike the case of $\text{Aut}(F_n)$, the set Z will depend on χ , but there will be only boundedly many possibilities, and in each case Z will contain 4 elements, exactly one element from each of the sets $S \cap G_I$ for $I = \{1, 2\}, \{3, 4\}, \{5, 6\}$ and $\{7, 8\}$. For each such Z , the pair (S, Z) is chain-centralizing – this follows immediately from Observation 4.5 (iii).

The proof of Lemma 4.11 (a) is very similar to that of Lemma 3.2 (a), so we will just list the main steps and skip the details of the computations. As in the proof of Lemma 3.2, in each step M denotes a positive real number and λ is an arbitrary character of G .

Step 1. If $M(\lambda) \geq M$, there exists $g \in \text{Sp}_{2n}(\mathbb{Z})$ which is a product of at most 3 transpositions f_{ij} such that $gz \in \mathcal{B} \cap V_{\{1,2\}}$ or $gz \in \mathcal{B} \cap V_{\{1,2,3\}}$.

Below we will consider the case where z in step 1 lies in $V_{\{1,2\}}$; the other case is similar. Without loss of generality we can assume that $z = a_1 \wedge b_1 \wedge b_2$.

Step 2. If $|\lambda(a_1 \wedge b_1 \wedge b_2)| \geq M$, there exists $g = \tau_{13}^e$ with $|e| \leq 2$ such that $|g\lambda(a_1 \wedge b_1 \wedge b_2)| \geq M$ and $|g\lambda(a_1 \wedge b_3 \wedge b_2)| \geq M$.

Step 3. If $|\lambda(a_1 \wedge b_1 \wedge b_2)| \geq M$ and $|\lambda(a_1 \wedge b_3 \wedge b_2)| \geq M$, there exists $g = \tau_{31}^e$ with $|e| \leq 2$ such that $|g\lambda(a_1 \wedge b_1 \wedge b_2)| \geq M$ and $|g\lambda(a_3 \wedge b_3 \wedge b_2)| \geq M$.

Step 4. $|\lambda(a_1 \wedge b_1 \wedge b_2)| \geq M$ and $|\lambda(a_3 \wedge b_3 \wedge b_2)| \geq M$, there exists $g = \tau_{24}^e$ with $|e| \leq 1$ such that $|g\lambda(a_1 \wedge b_1 \wedge b_2)| \geq M$ and $|g\lambda(a_3 \wedge b_3 \wedge b_4)| \geq M$.

Combining steps 1–4, we conclude that there exists $g_1 \in \text{Sp}_{2n}(\mathbb{Z})$ equal to the product of at most 3 transpositions and at most 5 elements of the form $\tau_{13}^{\pm 1}, \tau_{31}^{\pm 1}, \tau_{24}^{\pm 1}$ such that $|g_1\chi(z)| \geq M$ for some $z \in V_{\{1,2\}} \cap \mathcal{B}$ and also for some $z \in V_{\{3,4\}} \cap \mathcal{B}$. Repeating steps 2–4 two more times, first replacing indices 3 and 4 by 5 and 6, respectively, and then by 7 and 8, we obtain an element g satisfying the conclusion of Lemma 4.11.

(b) For an essential simple closed curve α , the action of T_α on $H_1(\Sigma)$ is given by the following formula (see, e.g., [10, Proposition 6.3]):⁷

$$T_\alpha([\beta]) = [\beta] + ([\beta] \cdot [\alpha])[\alpha]. \quad (4.3)$$

A direct computation using this formula shows that

- (i) the Dehn twist T_{β_i} is a lift of τ_i ;
- (ii) the element $W_i = T_{\beta_i} T_{\alpha_i}^{-1} T_{\beta_i}$ is a lift of w_i ;
- (iii) the element $W_i T_{\alpha_i \alpha_j^{-1}}^{-1} T_{\alpha_i} T_{\alpha_j} W_i^{-1}$ is a lift of $\tau_{i,j}$.

Below we exhibit the calculation for (iii) (assuming the result for (ii)). Recall that $a_i = [\alpha_i]$ and $b_i = [\beta_i]$. We have

$$\begin{aligned} a_i &\xrightarrow{W_i^{-1}} -b_i \xrightarrow{T_{\alpha_j}} -b_i \xrightarrow{T_{\alpha_i}} a_i - b_i \xrightarrow{T_{\alpha_i \alpha_j^{-1}}^{-1}} a_i - b_i - (a_i - a_j) = a_j - b_i \xrightarrow{W_i} a_j + a_i, \\ b_i &\xrightarrow{W_i^{-1}} a_i \xrightarrow{T_{\alpha_j}} a_i \xrightarrow{T_{\alpha_i}} a_i \xrightarrow{T_{\alpha_i \alpha_j^{-1}}^{-1}} a_i \xrightarrow{W_i} b_i, \\ a_j &\xrightarrow{W_i^{-1}} a_j \xrightarrow{T_{\alpha_j}} a_j \xrightarrow{T_{\alpha_i}} a_j \xrightarrow{T_{\alpha_i \alpha_j^{-1}}^{-1}} a_j \xrightarrow{W_i} a_j, \\ b_j &\xrightarrow{W_i^{-1}} b_j \xrightarrow{T_{\alpha_j}} b_j - a_j \xrightarrow{T_{\alpha_i}} b_j - a_j \xrightarrow{T_{\alpha_i \alpha_j^{-1}}^{-1}} b_j - a_j - (a_i - a_j) = b_j - a_i \xrightarrow{W_i} b_j - b_i. \end{aligned}$$

In the case of transpositions f_{ij} , it will be more convenient to define lifts directly instead of expressing them as products of Dehn twists. Given distinct $i < j \in [n]$, let F_{ij} be the unique element of $\text{Mod}(\Sigma)$ which is supported on $\Sigma_{\{i,j\}}$ and acts on $\pi_1(\Sigma_{\{i,j\}})$ as follows:

$$\alpha_i \mapsto \alpha_j, \quad \beta_i \mapsto \beta_j, \quad \alpha_j \mapsto \alpha_i^{[\alpha_j, \beta_j]}, \quad \beta_j \mapsto \beta_i^{[\alpha_j, \beta_j]}.$$

Such an element F_{ij} exists (and is unique) by Lemma 4.9. It is clear that F_{ij} is a lift of f_{ij} .

Now any g in part (a) has a lift ϕ which can be written as a product of at most 3 transposition-lifts F_{ij} and at most 135 Dehn twists $T_{\alpha_i}, T_{\beta_i}$ or $T_{\alpha_i \alpha_j^{-1}}$ or their inverses. By (3.4), in order to get an absolute bound for $A(\phi^{\pm 1}, S)$, it suffices to prove the following proposition.

⁷Proposition 6.3 in [10] makes the additional assumption that the curve β is also simple. However, since $H_1(\Sigma)$ is generated by the homology classes of simple closed curves, by linearity (4.3) holds for all β .

Proposition 4.12. *For a suitable choice of S , there exists an absolute constant C (independent of n) such that $A(g, S) \leq C$ for $g = T_{\alpha_i}^{\pm 1}, T_{\beta_i}^{\pm 1}, T_{\alpha_i \alpha_j^{-1}}^{\pm 1}$ or $F_{ij}^{\pm 1}$.*

Remark. The assertion of Proposition 4.12 does not appear to be obvious even if we restrict ourselves to, say, $g = T_{\alpha_i}$ with i fixed but n tending to infinity.

Proposition 4.12 will be proved in the next subsection. ■

4.7. Estimating the A constants

Recall that by Corollary 4.10, for every non-empty subset $I = \{i_1 < i_2 < \dots < i_k\}$ of $[n]$, there exists an element $f_I \in \text{Mod}(\Sigma)$ which maps $\Sigma_{[k]}$ to Σ_I and satisfies $f_I^*(\alpha_j) = \alpha_{i_j}$ and $f_I^*(\beta_j) = \beta_{i_j}$ for all $1 \leq j \leq k$. From now on, we will fix such an element f_I for every I .

Let us record one more simple observation, which is an immediate consequence of the moreover parts of Lemma 4.9 and Corollary 4.10.

Observation 4.13. *Let $I = \{i_1 < i_2 < \dots < i_k\}$ be a subset of $[n]$, let J be a subset of $[k]$, and let $I_J = \{i_j : j \in J\}$. Then the elements f_{I_J} and $f_I f_J$ coincide on $\Sigma_{[J]}$ and hence $f_{I_J} s f_{I_J}^{-1} = f_I f_J s (f_I f_J)^{-1}$ for all $s \in \text{Mod}_{[J]}$.*

Proof. Let $t = |J|$, and write $J = \{j_1 < \dots < j_t\}$. Recall that some representative of f_J maps $\Sigma_{[k]}$ to Σ_J , and by Corollary 4.10 some representative of f_I maps Σ_J to Σ_{I_J} . Thus, if $f = f_I f_J$, the following hold:

- (i) Some representative of f maps $\Sigma_{[k]}$ to Σ_{I_J} .
- (ii) For all $1 \leq m \leq t$, we have $f^*(\alpha_m) = f_I^*(\alpha_{j_m}) = \alpha_{i_{j_m}}$ and similarly $f^*(\beta_m) = \beta_{i_{j_m}}$.

If $f = f_{I_J}$, then both (i) and (ii) also hold by construction. Hence, the assertion of Observation 4.13 follows from the moreover part of Lemma 4.9. ■

We are now ready to prove Proposition 4.12.

Proof of Proposition 4.12. Define a generating set S for \mathcal{I} as follows: Choose any generating set $S_{[3]}$ of $\mathcal{I}_{[3]}$ satisfying the requirement of Lemma 4.8. For every $I \subseteq [n]$ with $|I| = 3$, define $S_I = f_I S_{[3]} f_I^{-1}$ and let $S = \bigcup_{|I|=3} S_I$. It is clear that each S_I (and hence S) also satisfy the requirement of Lemma 4.8.

By definition, $A(g, S) = A(g, \bigcup_{|L|=3} S_L) = \max_{|L|=3} A(g, S_L)$, so it is enough to bound $A(g, S_L)$ with $|L| = 3$. From now on, we fix L with $|L| = 3$.

Case 1: $g = T_{\alpha_u \alpha_v^{-1}}^{\pm 1}$ for some $u \neq v$.

Since the Dehn twist does not depend on the orientation of the curve, without loss of generality we can assume that $u < v$. Below we will consider the subcase $g = T_{\alpha_u \alpha_v^{-1}}$; the subcase $g = T_{\alpha_u \alpha_v^{-1}}^{-1}$ is analogous.

Let $I = L \cup \{u, v\}$ (thus $3 \leq |I| \leq 5$). Let $i_1 < i_2 < \dots < i_{|I|}$ be the elements of I listed in increasing order. Then, in the notations of Observation 4.13, we have $L = I_J$ where J is a 3-element subset of $[|I|] \subseteq [5]$.

Now take any element $x \in S_L$. By definition, $x = f_L s f_L^{-1}$ where $s \in S_{[3]}$. Since $f_L = f_I f_J$ on $\Sigma_{[3]}$ by Observation 4.13 and $f_J s f_J^{-1} \in S_J$ (by definition of S_J), we have $x = f_I t f_I^{-1}$ for some $t \in S_J$. Hence

$$g x g^{-1} = g(f_I t f_I^{-1})g^{-1} = f_I(f_I^{-1} g f_I) t (f_I^{-1} g f_I)^{-1} f_I^{-1}. \quad (4.4)$$

By the conjugation formula for Dehn twists (see, e.g., [10, Fact 3.7]), we have

$$f_I^{-1} g f_I = f_I^{-1} T_{\alpha_u \alpha_v^{-1}} f_I = T_{(f_I^{-1})^*(\alpha_u \alpha_v^{-1})} = T_{(f_I^{-1})^*(\alpha_u)} (f_I^{-1})^*(\alpha_v)^{-1} = T_{\alpha_r \alpha_s^{-1}},$$

where $r, s \in [|I|]$ are the unique indices such that $i_r = u$ and $i_s = v$.

It follows that $f_I^{-1} g f_I$ is an element of $\text{Mod}_{[|I|]}$, for which we have only boundedly many (in fact, at most $\binom{5}{2} = 10$) possibilities. Hence we also have boundedly many possibilities for $y = (f_I^{-1} g f_I) t (f_I^{-1} g f_I)^{-1}$ (since $t \in S_J$ and $J \subset [5]$ with $|J| = 3$, we have at most $\binom{5}{3} |S_{[3]}| = 420$ possibilities for t and hence at most 4200 possibilities for y). Note that y also lies in $\mathcal{I}_{[|I|]}$ and thus we can write $y = t_1^{\pm 1} \dots t_m^{\pm 1}$, where each $t_j \in S_{[|I|]}$ and $m \leq C_1$ for some absolute constant C_1 . By (4.4), we have $g x g^{-1} = f_I y f_I^{-1} = \prod_{k=1}^m f_I t_k f_I^{-1}$. We claim that each factor in the latter product lies in S . This would imply that $\|g x g^{-1}\|_S \leq m \leq C$ and thus finish the proof.

Indeed, for each k as above, we have $t_k \in S_{J_k}$ with $|J_k| = 3$ and $J_k \subseteq [|I|]$, so $t_k = f_{J_k} s_k f_{J_k}^{-1}$ for some $s_k \in S_{[3]}$. Hence, using Observation 4.13 again, we have

$$f_I t_k f_I^{-1} = f_I f_{J_k} s_k f_{J_k}^{-1} f_I^{-1} = f_{I_{J_k}} s_k f_{I_{J_k}}^{-1} \in S_{I_{J_k}} \subset S.$$

Case 2: $g = T_{\alpha_u}^{\pm 1}$ or $T_{\beta_u}^{\pm 1}$ for some u . This case is similar to (and easier than) case 1.

Case 3: $g = F_{uv}^{\pm 1}$ for some $u < v$. The argument in this case is mostly similar to case 1, so we will just outline the differences. First as in case 1, without loss of generality we can assume that $g = F_{uv}$.

Fix $L \subseteq [n]$, and let $I = L \cup \{u, v\}$ (thus $3 \leq |I| \leq 5$). Take any $x \in S_L$. As in case 1, $x = f_I t f_I^{-1}$ with $t \in S_{[|I|]}$, and we just need to find an appropriate expression for the conjugate $f_I^{-1} F_{uv} f_I$.

Let $i_1 < i_2 < \dots < i_{|I|}$ be the elements of I listed in increasing order, and let $r, s \in [|I|]$ be such that $u = i_r$ and $v = i_s$. By definition of f_I , some representative of f_I sends $\Sigma_{\{r,s\}}$ to $\Sigma_{\{u,v\}}$. Since F_{uv} is trivial on the complement of $\Sigma_{\{u,v\}}$, we conclude that $f_I^{-1} F_{uv} f_I$ is trivial on the complement of $\Sigma_{\{r,s\}}$. The action on $\Sigma_{\{r,s\}}$ is determined by direct computation below. We have

$$(f_I^{-1} F_{uv} f_I)^*(\alpha_r) = (f_I^{-1} F_{uv})^*(\alpha_u) = (f_I^{-1})^*(\alpha_v) = \alpha_s, \quad (4.5)$$

$$(f_I^{-1} F_{uv} f_I)^*(\alpha_s) = (f_I^{-1} F_{uv})^*(\alpha_v) = (f_I^{-1})^*(\alpha_u^{\alpha_v}) = \alpha_r^{\alpha_s}. \quad (4.6)$$

Similarly, $(f_I^{-1}F_{uv}f_I)^*(\beta_r) = \beta_s$ and $(f_I^{-1}F_{uv}f_I)^*(\beta_s) = \beta_r^{\beta_s}$. Hence

$$f_I^{-1}gf_I = f_I^{-1}F_{uv}f_I = F_{rs}.$$

Thus, $f_I^{-1}gf_I$ is supported on $\Sigma_{[I]}$, and we have boundedly many possibilities for $f_I^{-1}gf_I$, and we can finish the proof as in case 1. ■

4.8. Conclusion of the proof of Theorem 1.2

Proof of Theorem 1.2. Let us first prove Theorem 1.2 (1). Let S be a generating set for the group \mathcal{I} constructed in Proposition 4.12. This set need not satisfy the conclusion of Theorem 1.2 (1). First we decompose $S = S_1 \sqcup S_2 \sqcup S_3$ as in the setup introduced before Theorem 2.16. Thus, $S_3 = S \cap \mathcal{K}$, $S_1 = \{s_1, \dots, s_N\}$ is a subset of $S \setminus S_3$ such that the natural projection $\theta: \mathcal{I} \rightarrow \mathcal{I}/\mathcal{K}$ maps S_1 bijectively onto \mathcal{B} (and $S_2 = S \setminus (S_1 \sqcup S_3)$). Also recall that for $s \in S_2$, we denote by $d(s)$ the unique integer such that $\theta(s) = \theta(s_{d(s)})$.

Define $S^{(1)} = S_1$ and $S^{(4)} = S_3 \cup \{s^{-1}s_{d(s)} : s \in S_2\}$. Then $\langle S^{(1)} \cup S^{(4)} \rangle = \langle S_1 \sqcup S_2 \sqcup S_3 \rangle = \mathcal{I}$. Also note that $S^{(4)}$ lies in \mathcal{K} and the image of $S^{(4)}$ in \mathcal{I}^{ab} generates $\mathcal{K}/[\mathcal{I}, \mathcal{I}]$ (the latter holds since the image of $S^{(1)} \cup S^{(4)}$ generates \mathcal{I}^{ab} and the images of elements of $S^{(1)}$ are linearly independent modulo \mathcal{K}). Since $\mathcal{K}/[\mathcal{I}, \mathcal{I}]$ is a vector space over \mathbb{F}_2 , we can choose a subset $S^{(2)}$ of $S^{(4)}$ whose image in $\mathcal{K}/[\mathcal{I}, \mathcal{I}]$ is a basis of $\mathcal{K}/[\mathcal{I}, \mathcal{I}]$. Finally, we can multiply each element of $S^{(4)} \setminus S^{(2)}$ by a product of elements of $S^{(2)}$ (on either side) so that the obtained element lies in $[\mathcal{I}, \mathcal{I}]$, and let $S^{(3)}$ be the set of all such elements.

Since $S^{(1)} \sqcup S^{(2)} \sqcup S^{(3)}$ is obtained from S by Nielsen transformations, it clearly generates \mathcal{I} , and the additional properties asserted in Theorem 1.2 (1) hold by construction.

Now let $\theta: \mathcal{I} \rightarrow \mathcal{I}/\mathcal{K}$ be the natural projection. By Theorem 2.13 and the remark after Lemma 4.11, the set $\theta^{-1}(B_\infty(R))$ is connected where R is bounded by an absolute constant. By Theorem 2.16, \mathcal{K} is generated by elements of the form

- (a) $[s_i, s_j]^{s_i^{a_i} s_{i+1}^{a_{i+1}} \dots s_n^{a_n}}$, where $1 \leq i < j \leq n$ and $|a_j| \leq R$ for all j ;
- (b) $x^{s_1^{a_1} s_2^{a_2} \dots s_n^{a_n}}$, where $x \in S^{(4)}$ and $|a_j| \leq R$ for all j .

This generating set is almost the same as the set in Theorem 1.2 (2) – the only difference is that the condition $x \in S^{(4)}$ above is replaced by $x \in S^{(2)} \sqcup S^{(3)}$. But by construction the subgroup generated by $S^{(4)}$ is equal to the subgroup generated by $S^{(2)}$ and $S^{(3)}$, so the set described in Theorem 1.2 (2) also generates G . This proves Theorem 1.2 (2).

Finally, Theorem 1.2 (3) follows from Theorem 1.2 (2) by a standard application of the Reidemeister–Schreier process and straightforward computations involving basic commutator identities. ■

Estimating the constant R . Let us now briefly address the problem of explicitly estimating the constant R in Theorem 1.2. Let $S_{[3]}^0$ be the generating set of $\mathcal{I}_{[3]}^1$ with 42 elements constructed in Johnson’s paper [14]. We can algorithmically construct another generating set $S_{[3]}$ of $\mathcal{I}_{[3]}^1$, also with 42 elements, satisfying the conclusion of Lemma 4.8. Next for

each $I \subseteq [5]$ with $|I| = 3$, we choose f_I satisfying the conclusion of Corollary 4.10 (it is easy to do this explicitly) and define $S_I = f_I S_{[3]} f_I^{-1}$. Also define $S_{[5]} = \bigcup_{|I|=3, I \subset [5]} S_I$.

Now the key step is estimating the constant C_1 from the proof of Proposition 4.12. To do this, we need to explicitly express each conjugate yty^{-1} , where $t \in S_{[5]}$ and $y = T_{\alpha_i \alpha_j}^{\pm 1}, T_{\alpha_i}^{\pm 1}$ or $T_{\beta_i}^{\pm 1}$, $1 \leq i \neq j \leq 5$, as a product of elements of $S_{[5]}$. An algorithm for obtaining such expressions is given in the Ph.D. thesis of Stylianakis [25]; however, some of the arguments in [25] rely on certain computations in [14].

Finally, it is easy to express the constant R in terms of C_1 following the steps of the proof of Theorem 1.2 (2).

4.9. Proof of Theorem 1.3

Throughout this subsection, S will denote the generating set for \mathcal{I} constructed in Proposition 4.12 (not the modified set from the proof of Theorem 1.2). Let us recall some notations from the introduction. We let $\omega = \{\alpha_i, \beta_i\}_{i=1}^n$; also for $1 \leq l \leq n$ we set $\omega_l = \{\alpha_i, \beta_i\}_{i=1}^l$. Given $m \in \mathbb{N}$, we denote by $T_{sc}(m)$ the set of all Dehn twists T_γ , where $\gamma \in \pi_1(\Sigma, p_0)$ can be represented by a separating curve and $\|\gamma\|_\omega \leq m$.

Theorem 1.3 will be deduced from Theorem 1.2 (or rather its proof) and the following technical lemma.

Lemma 4.14. *The following hold:*

- (a) *Let w be any word in the free group on 2 generators. There exists a constant C_w depending only on w with the following property: for every $s, t \in S$ such that $w(s, t) \in \mathcal{K}$, the element $w(s, t)$ lies in $\langle T_{sc}(C_w) \rangle$, the subgroup generated by $T_{sc}(C_w)$.*
- (b) *There exists an absolute constant C_2 such that $\|s^*(\alpha)\|_\omega \leq C_2 \|\alpha\|_\omega$ for all $\alpha \in \pi_1(\Sigma, p_0)$ and $s \in S^{\pm 1}$.*

Before proving Lemma 4.14, we establish a simple auxiliary result.

Claim 4.15. *For every $I \subseteq [n]$, we have $\mathcal{K}(\Sigma_I) = \mathcal{K} \cap \text{Mod}_I$.*

Proof. Let $\pi = \pi_1(\Sigma, p_0)$ and $\pi_I = \pi_1(\Sigma_I, p_0)$. Using Theorem 4.1, we can identify $\text{Mod}(\Sigma)$ and $\text{Mod}_I = \text{Mod}(\Sigma_I)$ with subgroups of $\text{Aut}(\pi)$ and $\text{Aut}(\pi_I)$, respectively. Under this identification, $\mathcal{K}(\Sigma_I)$ consists of all elements of $\text{Aut}(\pi_I)$ which act trivially modulo $\gamma_3 \pi_I$ and lie in Mod_I , while $\mathcal{K} \cap \text{Mod}_I$ consists of all elements of $\text{Aut}(\pi_I)$ which act trivially modulo $\gamma_3 \pi \cap \pi_I$ and lie in Mod_I .

Thus, proving the equality $\mathcal{K}(\Sigma_I) = \mathcal{K} \cap \text{Mod}_I$ reduces to showing that $\gamma_3 \pi_I = \gamma_3 \pi \cap \pi_I$. The latter holds since π is free and π_I is a free factor of π and hence a retract of π . ■

Remark. Claim 4.15 is a special case of [4, Theorem 4.6]; however, we decided to give a direct proof since we are dealing with a much more specific situation compared to [4].

Proof of Lemma 4.14. (a) Take any $s, t \in S$ and a word w such that $w(s, t) \in \mathcal{K}$. By definition of S , there exist $I, J \subseteq [n]$ with $|I| = |J| = 3$ and $u, v \in S_{[3]}$ such that $s = f_I u f_I^{-1}$ and $t = f_J v f_J^{-1}$. Let $L = I \cup J$ and $l = |L|$ (thus, $l \leq 6$). Then, using the notations of Observation 4.13, $I = L_A$ and $J = L_B$ for some $A, B \subseteq [l] \subseteq [6]$. Hence by Observation 4.13, $s = f_{L_A} u f_{L_A}^{-1} = f_L s' f_L^{-1}$, where $s' = f_A u f_A^{-1}$. Similarly, $t = f_L t' f_L^{-1}$, where $t' = f_B v f_B^{-1}$.

Note that $s', t' \in S_{[l]}$. Also note that s' and t' are obtained from s and t by conjugation by the same element of $\text{Mod}(\Sigma)$. Since $w(s, t) \in \mathcal{K}$ and \mathcal{K} is normal in $\text{Mod}(\Sigma)$, it follows that $w(s', t') \in \mathcal{K} \cap \text{Mod}_{[l]} = \mathcal{K}(\Sigma_{[l]})$ where the last equality holds by Claim 4.15.

Since $s', t' \in S_{[l]}$ and $l \leq 6$, there are boundedly many possibilities for the pair (s', t') . Let us enumerate all such pairs $(s'_1, t'_1), \dots, (s'_l, t'_l)$ satisfying the additional restriction $w(s'_j, t'_j) \in \mathcal{K}$ (and hence $w(s'_j, t'_j) \in \mathcal{K}(\Sigma_{[l]})$). Since $\mathcal{K}(\Sigma_{[l]})$ is generated by Dehn twists T_γ with $\gamma \in \pi_1(\Sigma_{[l]}, p_0)$ separating, there exists a constant C_w depending only on w such that each element $w(s'_j, t'_j)$ (and in particular, $w(s', t')$) lies in the subgroup generated by all T_γ , where $\gamma \in \pi_1(\Sigma_{[l]}, p_0)$ is separating and $\|\gamma\|_{\omega_l} \leq C_w$.

Thus we can write

$$w(s', t') = \prod_{i=1}^k T_{\gamma_i}^{\pm 1},$$

where each $\gamma_i \in \pi_1(\Sigma_{[l]}, p_0)$ is separating with $\|\gamma_i\|_{\omega_l} \leq C_w$. Then

$$\begin{aligned} w(s, t) &= w(f_L s' f_L^{-1}, f_L t' f_L^{-1}) = f_L w(s', t') f_L^{-1} \\ &= \prod_{i=1}^k f_L T_{\gamma_i}^{\pm 1} f_L^{-1} = \prod_{i=1}^k T_{f_L^*(\gamma_i)}^{\pm 1}. \end{aligned} \quad (4.7)$$

Finally, by definition of f_L , we have $f_L^*(\omega_l) \subseteq \omega$ and therefore for any γ_i in (4.7), we have $\|f_L^*(\gamma_i)\|_\omega \leq \|\gamma_i\|_{\omega_l} \leq C_w$, and hence $w(s, t)$ lies in $T_{\text{sc}}(C_w)$ by (4.7).

(b) Take any $1 \neq \alpha \in \pi_1(\Sigma, p_0)$, let $k = \|\alpha\|_\omega$, and write $\alpha = \prod_{j=1}^k \lambda_j^{\pm 1}$ with $\lambda_j \in \omega$. Since $\|s^*(\alpha)\|_\omega = \|\prod_{j=1}^k s^*(\lambda_j)^{\pm 1}\|_\omega \leq \sum_{j=1}^k \|s^*(\lambda_j)\|_\omega$, to prove (b) it suffices to show that there exists an absolute constant C_2 such that

$$\|s^*(\lambda)\|_\omega \leq C_2 \quad \text{for all } \lambda \in \omega, \text{ that is, for } \lambda = \alpha_i \text{ or } \beta_i.$$

By symmetry, it suffices to consider the case $\lambda = \alpha_i$ and $s \in S$. Write $s = f_J v f_J^{-1}$, where $|J| = 3$ and $v \in S_{[3]}$. Let $I = \{i\}$, $L = I \cup J$ and $l = |L| \leq 4$. As in (a), we have $s = f_L s' f_L^{-1}$, where $s' \in S_{[l]}$ and $\lambda = f_L^*(\alpha_j)$ for some $j \in [l]$.

Then $s^* \lambda = (f_L s' f_L^{-1})^* f_L^*(\alpha_j) = f_L^*((s')^* \alpha_j)$. Since we have boundedly many possibilities for j and s' , we have $\|(s')^* \alpha_j\|_{\omega_l} \leq C_2$ for some absolute constant C_2 , and as in the proof of (a), we have

$$\|f_L^*((s')^* \alpha_j)\|_\omega \leq \|(s')^* \alpha_j\|_{\omega_l}.$$

Thus, $\|s^* \lambda\|_\omega \leq C_2$, as desired. ■

Proof of Theorem 1.3. Recall that S is the generating set for \mathcal{I} from Proposition 4.12. In the proof of Theorem 1.2, we showed that the Johnson kernel \mathcal{K} has a generating set consisting of elements of the form x^y , where $x = [s, t]$ for some $s, t \in S$ or $x \in S^{(4)}$ (in the notations from that proof) and $y = \prod_{j=1}^m s_j$, where each $s_j \in S$ and $m \leq R(N + M) < 2Rn^3$, where R, N and M are as in Theorem 1.2 (b). Also recall that every element of $S^{(4)}$ either lies in S or has the form $s^{-1}t$ with $s, t \in S$.

Consider the following three words in the free group on two generators a, b :

$$w_1(a, b) = a, \quad w_2(a, b) = a^{-1}b \quad \text{and} \quad w_3(a, b) = [a, b].$$

Let C_1 be the maximum of the constants $C_{w_1}, C_{w_2}, C_{w_3}$ from the conclusion of Lemma 4.14 (a). Now take any x and y as in the previous paragraph. By Lemma 4.14 (a), we can write $x = \prod_{i=1}^k T_{\gamma_i}$, where $\gamma_i \in \pi_1(\Sigma, p_0)$ is separating and $\|\gamma_i\|_\omega \leq C_1$. Hence

$$x^y = \prod_{i=1}^k T_{\gamma_i}^{\prod_{j=1}^m s_j} = \prod_{i=1}^k T_{(\prod_{j=1}^m s_j^{-1})^*(\gamma_i)}.$$

By Lemma 4.14 (b) for each i , we have

$$\left\| \left(\prod_{j=m}^1 s_j^{-1} \right)^*(\gamma_i) \right\|_\omega \leq C_2^m \|\gamma_i\|_\omega \leq C_2^{2Rn^3} C_1 \leq (C_1 C_2)^{2Rn^3}.$$

Thus, Theorem 1.3 holds with $D = (C_1 C_2)^{2R}$. ■

Acknowledgements. We are extremely grateful to Andrew Putman for explaining to us the proof of the BNS criterion given in [24] and to the anonymous referee who made a number of suggestions that helped improve the exposition. We also thank Thomas Church, Thomas Koberda and Andrew Putman for useful discussions related to the subject of this paper. After this paper was completed, the authors learned that results similar to those in this paper were independently obtained by Church and Putman (unpublished).

References

- [1] M. Bestvina, K.-U. Bux, and D. Margalit, [Dimension of the Torelli group for \$\text{Out}\(F_n\)\$](#) . *Invent. Math.* **170** (2007), no. 1, 1–32 Zbl [1135.20026](#) MR [2336078](#)
- [2] R. Bieri, W. D. Neumann, and R. Strebel, [A geometric invariant of discrete groups](#). *Invent. Math.* **90** (1987), no. 3, 451–477 Zbl [0642.57002](#) MR [914846](#)
- [3] M. R. Bridson and R. D. Wade, [Commensurations of \$\text{Aut}\(F_n\)\$ and its Torelli subgroup](#). 2023, arXiv:[2306.13437](#)
- [4] T. Church, [Orbits of curves under the Johnson kernel](#). *Amer. J. Math.* **136** (2014), no. 4, 943–994 Zbl [1306.57016](#) MR [3245184](#)
- [5] T. Church, M. Ershov, and A. Putman, [On finite generation of the Johnson filtrations](#). *J. Eur. Math. Soc. (JEMS)* **24** (2022), no. 8, 2875–2914 Zbl [07523092](#) MR [4416592](#)

- [6] T. Church and A. Putman, [Generating the Johnson filtration](#). *Geom. Topol.* **19** (2015), no. 4, 2217–2255 Zbl [1364.20025](#) MR [3375526](#)
- [7] M. Day and A. Putman, [The complex of partial bases for \$F_n\$ and finite generation of the Torelli subgroup of \$\text{Aut}\(F_n\)\$](#) . *Geom. Dedicata* **164** (2013), 139–153 Zbl [1294.20041](#) MR [3054621](#)
- [8] A. Dimca, R. Hain, and S. Papadima, [The abelianization of the Johnson kernel](#). *J. Eur. Math. Soc. (JEMS)* **16** (2014), no. 4, 805–822 Zbl [1344.57001](#) MR [3191977](#)
- [9] M. Ershov and S. He, [On finiteness properties of the Johnson filtrations](#). *Duke Math. J.* **167** (2018), no. 9, 1713–1759 Zbl [1498.20082](#) MR [3813595](#)
- [10] B. Farb and D. Margalit, *A primer on mapping class groups*. Princeton Math. Ser. 49, Princeton University Press, Princeton, NJ, 2012 Zbl [1245.57002](#) MR [2850125](#)
- [11] E. Formanek, [Characterizing a free group in its automorphism group](#). *J. Algebra* **133** (1990), no. 2, 424–432 Zbl [0715.20022](#) MR [1067415](#)
- [12] S. P. Humphries, [Generators for the mapping class group](#). In *Topology of low-dimensional manifolds*, pp. 44–47, Lecture Notes in Math. 722, Springer, Berlin, 1979 Zbl [0732.57004](#) MR [547453](#)
- [13] D. Johnson, [An abelian quotient of the mapping class group \$\mathcal{I}_g\$](#) . *Math. Ann.* **249** (1980), no. 3, 225–242 Zbl [0409.57009](#) MR [579103](#)
- [14] D. Johnson, [The structure of the Torelli group I: A finite set of generators for \$\mathcal{I}\$](#) . *Ann. of Math. (2)* **118** (1983), no. 3, 423–442 Zbl [0549.57006](#) MR [727699](#)
- [15] D. Johnson, [The structure of the Torelli group–II: A characterization of the group generated by twists on bounding curves](#). *Topology* **24** (1985), no. 2, 113–126 Zbl [0571.57009](#) MR [793178](#)
- [16] D. Johnson, [The structure of the Torelli group–III: The abelianization of \$\mathcal{T}\$](#) . *Topology* **24** (1985), no. 2, 127–144 Zbl [0571.57010](#) MR [793179](#)
- [17] N. Koban, J. McCammond, and J. Meier, [The BNS-invariant for the pure braid groups](#). *Groups Geom. Dyn.* **9** (2015), no. 3, 665–682 Zbl [1326.20045](#) MR [3420539](#)
- [18] W. Magnus, [Über \$n\$ -dimensionale Gittertransformationen](#). *Acta Math.* **64** (1935), no. 1, 353–367 Zbl [0012.05405](#) MR [1555401](#)
- [19] J. Meier and L. VanWyk, [The Bieri–Neumann–Strebel invariants for graph groups](#). *Proc. Lond. Math. Soc. (3)* **71** (1995), no. 2, 263–280 Zbl [0835.20037](#) MR [1337468](#)
- [20] S. Morita, T. Sakasai, and M. Suzuki, [Torelli group, Johnson kernel, and invariants of homology spheres](#). *Quantum Topol.* **11** (2020), no. 2, 379–410 Zbl [1462.57017](#) MR [4118638](#)
- [21] L. A. Orlandi-Korner, [The Bieri–Neumann–Strebel invariant for basis-conjugating automorphisms of free groups](#). *Proc. Amer. Math. Soc.* **128** (2000), no. 5, 1257–1262 Zbl [0963.20016](#) MR [1712889](#)
- [22] A. Putman, [Small generating sets for the Torelli group](#). *Geom. Topol.* **16** (2012), no. 1, 111–125 Zbl [1296.57008](#) MR [2872579](#)
- [23] A. Putman, [The commutator subgroups of free groups and surface groups](#). *Enseign. Math.* **68** (2022), no. 3–4, 389–408 Zbl [1506.20064](#) MR [4452428](#)
- [24] R. Strebel, [Notes on the Sigma invariants](#). 2013, arXiv:1204.0214
- [25] C. Stylianakis, *Braid groups, mapping class groups, and torelli groups*. Ph.D. thesis, 2016, University of Glasgow, <https://theses.gla.ac.uk/7466/>
- [26] W. Tomaszewski, [A basis of Bachmuth type in the commutator subgroup of a free group](#). *Canad. Math. Bull.* **46** (2003), no. 2, 299–303 Zbl [1031.20015](#) MR [1981684](#)
- [27] H. Zieschang, [Über Automorphismen ebener diskontinuierlicher Gruppen](#). *Math. Ann.* **166** (1966), 148–167 Zbl [0151.33102](#) MR [201521](#)

- [28] H. Zieschang, E. Vogt, and H.-D. Coldewey, *Surfaces and planar discontinuous groups*. Lecture Notes in Math. 835, Springer, Berlin, 1980 Zbl [0438.57001](#) MR [606743](#)

Received 8 January 2021.

Mikhail Ershov

Department of Mathematics, University of Virginia, 141 Cabell Drive, Charlottesville, VA 22904, USA; ershov@virginia.edu

Daniel Franz

Department of Mathematics, Jacksonville University, 2800 University Blvd N, Jacksonville, FL 32211, USA; df Franz1@ju.edu