TOTALLY SYMMETRIC SETS

NOAH CAPLINGER AND DAN MARGALIT

Dedicated to Mike Mihalik on the occasion of his retirement

Abstract. We survey the theory of totally symmetric sets, with applications to homomorphisms of symmetric groups, braid groups, linear groups, and mapping class groups.

1. Introduction

The theory of totally symmetric sets is a tool that has been proven to be useful in classifying homomorphisms between certain types of groups. The basic definitions were introduced by Kordek and the second author in their study of homomorphisms between braid groups [19]. Since that work, the theory has been used in the study of homomorphisms between symmetric groups, braid groups, linear groups, and mapping class groups.

Here is the definition. A totally symmetric set in a group G is a subset

$$X = \{x_1, \ldots, x_k\} \subset G$$

with the following property: for every $\sigma \in \Sigma_k$, there is a $g_\sigma \in G$ such that

$$g_\sigma x_i g_\sigma^{-1} = x_{\sigma(i)}$$

for all i. It follows from the definition that the elements of a totally symmetric set lie in a single conjugacy class.

Evidently, if $X \subseteq G$ is a totally symmetric set and $f : G \to H$ is a homomorphism, then $f(X)$ is a totally symmetric set in H. As we show in Section 2.1, a much stronger condition is true: $f(X)$ is either a totally symmetric set of size $|X|$ or it is a singleton. In the phrasing of Salter and the first author: collision implies collapse. This is the fundamental property of totally symmetric sets.

For groups G and H, the collision-implies-collapse property yields an (unreasonably effective) blueprint for classifying homomorphisms $f : G \to H$, as follows:

Step 1. Find a large totally symmetric set $X \subset G$.
Step 2. Classify the large totally symmetric sets in H.
Step 3. Deduce properties of $f(X)$ and draw conclusions about f.

For instance, if G has a totally symmetric set X with $|X| = k$, and H has no totally symmetric set of cardinality k, then any homomorphism $f : G \to H$ must collapse X. Moreover, for any $x_i, x_j \in X$ the normal closure of $x_i x_j^{-1}$ lies in the kernel of f. In particular, if the $x_i x_j^{-1}$ are normal generators for G, then f is the trivial map.

Step 2 of the blueprint is generally the most challenging (and interesting). For a given group H, the approach is to choose a space Y on which H acts. As explained in Section 2.2, totally symmetric sets in H correspond to totally symmetric geometric configurations in Y. These configurations can be, for example, configurations of eigenspaces for linear maps or

---

This material is based upon work supported by the National Science Foundation under Grant No. DMS-2203431.

canonical reduction systems for mapping classes. The main strategy is to classify the totally symmetric configurations, and then use this to classify the totally symmetric sets.

Overview. In Section 2 we introduce the basic notions and examples in the theory of totally symmetric sets. We also prove the collision-implies-collapse property. We then use the blueprint to give an intuitive explanation for why the outer automorphism group of the symmetric group is (usually) trivial.

In Sections 3 and 4 we explain how the blueprint is applied in the cases of the general linear group and the braid group. The former case was addressed in a paper by the first author and Salter [9] and the latter in a paper by the second author and Kordek [19]. As per the blueprint, the strategy in both cases is to classify large totally symmetric configurations and then to promote this to a classification of large totally symmetric sets. For the general linear group, the configurations are configurations of subspaces. We then use the theory of Jordan normal form to do the promotion. For the braid group the configurations are configurations of multicurves. In this case we use Nielsen–Thurston theory to do the analogous promotion.

In Section 5 we prove a theorem of Kolay, which says that the standard map $B_n \to \Sigma_n$ gives the smallest non-cyclic quotient of the braid group. To streamline Kolay's argument, we first introduce a variation on totally symmetric sets, namely, collapsing sets. These are exactly the sets that satisfy the collision-implies-collapse property. We then present Kolay's proof of the theorem.

Finally, in Section 6 we make an explicit analogy between the collision-implies-collapse property and Schur's lemma from representation theory. We discuss several questions that arise from this analogy.

Acknowledgments. We are grateful to Michael Griffin, Koichi Oyakawa, Larry Rolen, Shuxian Song, Isabelle Steinmann, and an anonymous referee for helpful comments and conversations. We would also like to thank Lei Chen, Kevin Kordek, Justin Lanier, Dan Minahan, and Nick Salter for enlightening conversations about totally symmetric sets.

2. Totally symmetric sets and the blueprint

The three subsections in this section correspond to the three steps of the blueprint for totally symmetric sets. In Section 2.1 we give some basic examples of totally sets, and state and prove the collision-implies-collapse lemma. In Section 2.2, we define totally symmetric configurations, and use them to give upper bounds on the sizes of totally symmetric sets in certain groups. Finally, in Section 2.3 we use the results of Sections 2.1 and 2.2 to give a conceptual explanation for the classification of automorphisms of the symmetric group Σ_n.

2.1. Totally symmetric sets and collision implies collapse.

Examples of totally symmetric sets. Among the most basic examples of totally symmetric sets are:

$$\{(1\ 2),(3\ 4),\ldots\} \subseteq \Sigma_n, \quad \{(1\ i),\ldots,(1\ n)\} \subseteq \Sigma_n, \quad \{E_{1,2},\ldots,E_{1,n}\} \subseteq \mathrm{GL}_n(\mathbb{Z}),$$
$$\{E_1,\ldots,E_n\} \subseteq \mathrm{GL}_n(\mathbb{Z}), \quad \text{and } \{\sigma_1,\sigma_3,\ldots\} \subseteq B_n.$$

where the $(i\ j)$ are transpositions in the symmetric group Σ_n, the σ_i are the standard half-twist generators for the braid group B_n, the $E_{i,j}$ are elementary matrices in the general linear group $\mathrm{GL}_n(\mathbb{Z})$, and the E_i are the elements of $\mathrm{GL}_n(\mathbb{Z})$ obtained from the identity by negating the ith diagonal entry. We leave it as an exercise to verify that these are all totally symmetric sets.

Collision implies collapse. The following lemma, mentioned in the introduction, encapsulates the fundamental property of totally symmetric sets. The lemma originally appears in the work of Kordek and the second author of this paper [19, Lemma 2.1].

Lemma 2.1. *Let $f : G \to H$ be a homomorphism of groups. If $X \subseteq G$ is a totally symmetric set then $f(X)$ is either a totally symmetric set of cardinality $|X|$ or a singleton.*

Proof. Let $x, y, z \in X$ and suppose $f(x) = f(y)$. Total symmetry guarantees some $g \in G$ such that $g(x, y, z)g^{-1} = (x, z, y)$. We have:

$$f(xz^{-1}) = f(gxy^{-1}g^{-1}) = f(g)f(xy^{-1})f(g)^{-1} = 1.$$

Thus $f(z) = f(x)$ and the lemma follows. □

In the original paper by Kordek and the second author, totally symmetric sets were assumed to have the additional property that the elements commute pairwise. So in that paper, the set $\{(1\ 2),\ (3\ 4), \ldots\}$ would be considered as a totally symmetric set in Σ_n, whereas $\{(1\ i), \ldots, (1\ n)\}$ would not. The commutativity condition was included because it simplifies the classification of totally symmetric configurations for braid groups. Since the more general totally symmetric sets considered here still satisfy Lemma 2.1, we will henceforth use the term "commuting totally symmetric set" to refer to a totally symmetric set with the additional property that the elements commute pairwise.

In defining totally symmetric sets, Kordek and the second author were directly inspired by the work of Aramayona–Souto, who used a symmetric group action on a collection of Dehn twists to similar effect [1, Section 5].

2.2. **Totally symmetric configurations and upper bounds on totally symmetric sets.** Recall that Step 2 of the blueprint concerns the classification of large totally symmetric sets in a given group G (in the blueprint the group is called H). After the fashion in geometric group theory, we approach this problem by considering a suitable action of G on a space Y. Given such an action, we often have at least one natural choice of function:

$$G \to \text{subsets of } Y$$
$$g \mapsto Y_g$$

For a given $g \in G$, the subset Y_g might be the fixed set, or an eigenspace, or an invariant axis, etc. As long as the association $g \mapsto Y_g$ is natural, it will satisfy the equation

$$Y_{hgh^{-1}} = h \cdot Y_g$$

for all $h \in G$ (we can take this equivariance condition as the definition of naturality). In particular, if $\{x_1, \ldots, x_k\}$ is a totally symmetric subset of G, then $\{Y_{x_1}, \ldots, Y_{x_k}\}$ is a totally symmetric configuration in Y in the sense that any permutation of the Y_{x_i} can be realized by the action of G.

Totally symmetric configurations. Motivated by this discussion, we can give the definition of a totally symmetric configuration. Suppose that a group G acts on a space Y. Let

$$\{Y_1, \ldots, Y_k\}$$

be a collection of subspaces of Y. We say that $\{Y_i\}$ is a totally symmetric configuration if for each $\sigma \in \Sigma_k$, there is a $g_\sigma \in G$ so that

$$g_\sigma \cdot (Y_i) = Y_{\sigma(i)}$$

for all i. Again, the point of the definition is that, as long as the association of a subspace to a group element satisfies the naturality property $Y_{hgh^{-1}} = h \cdot Y_g$, the configuration associated to a totally symmetric set is a totally symmetric configuration.

Unifying the definitions. Our definitions of totally symmetric sets and totally symmetric configurations are almost identical. As observed by the first author and Salter [9, Definition 2.1], they can be combined into one definition as follows:

> *Let G act on a set Z. A subset $X = \{x_1, \ldots, x_k\} \subset Z$ is totally symmetric if for all $\sigma \in \Sigma_k$, there is some $g_\sigma \in G$ such that $g_\sigma \cdot x_i = x_{\sigma(i)}$*

The definition of a totally symmetric set is recovered by considering the action of G on itself by conjugation, and the definition of a totally symmetric configuration is recovered by considering the action of G on a set of subsets of a space Y that carries an action of G. Even in this general setting, totally symmetric sets obey the collision-implies-collapse principle where the homomorphism in Lemma 2.1 is replaced by a G-equivariant map.

Example: Dihedral groups. We will use the notion of totally symmetric configurations to prove the following fact:

> *A totally symmetric set $X \subseteq D_n$ has $|X| \leq 3$.*

The first step is to prove that a totally symmetric set of rotations has cardinality at most two (exercise). Suppose then that X is a totally symmetric set consisting of reflections. To each reflection in X we can associate the corresponding line of reflection in the plane. As above, this gives a totally symmetric configuration of lines in the plane. The largest such configuration has three lines (another exercise). Since reflections are determined by the corresponding lines, the desired statement follows.

This argument shows more:

> *If X is a totally symmetric set of D_n with $|X| = 3$ then $3 \mid n$ and X consists of reflections about lines that pairwise form an angle of $\pi/3$.*

As a sample consequence, we have the following fact:

> *If $n \geq 8$ and $m \geq 3$, then every homomorphism $B_n \to D_m$ has cyclic image.*

While this fact is not difficult to prove directly, the theory of totally symmetric sets gives a natural, conceptual explanation.

Example: Free groups. We now use totally symmetric configurations to prove the following:

> *A totally symmetric set in $X \subseteq F_2$ has $|X| \leq 1$.*

Consider the action of F_2 on its Cayley graph, the regular four-valent tree T_4. Each element $x \in F_2$ acts on T_4 by translating along an axis, which is a bi-infinite geodesic in T_4. Again, this gives a totally symmetric configuration of geodesics.

Within a given conjugacy class in F_2, an element is determined by its axis. Therefore, it suffices to show that there is no totally symmetric configuration consisting of two bi-infinite geodesics in T_4.

Let $Y \subset T_4$ be a totally symmetric configuration of bi-infinite geodesics. If Y_1 and Y_2 are distinct elements of Y, then by total symmetry there is an element of F_2 interchanging Y_1 and Y_2. This is impossible, since (using the usual embedding of T_4 in the plane) the elements of F_2 act on T_4 by orientation-preserving planar automorphisms.

An upper bound for all groups. The first author proved the following result [7, Theorem 1], which gives an upper bound for the cardinality of a totally symmetric set in an arbitrary group.

Theorem 2.2 (Caplinger)**.** *Let X be a totally symmetric set in a group G. If $|X| \geq 4$, then*

$$|G| \geq (|X| + 1)!$$

Equality is attained only when $G = \Sigma_n$.

As a sample application, any totally symmetric set in the monster group M has cardinality less than 44, since $44! > |M|$.

Analogous (but not sharp) upper bounds on the cardinalities of commuting totally symmetric sets were proved by Chudnovsky–Kordek–Li–Partin [12, Proposition 2.2] and by Scherich–Verberne [22, Theorem A].

Other upper bounds on commuting totally symmetric sets. Kordek–Li–Partin [18] provide a suite of upper bounds for the cardinality of a commutative totally symmetric set. For instance they show that the largest cardinality of a commutative totally symmetric set in the dihedral group is 2 [18, Theorem 3.3] and that the largest cardinality of a commutative totally symmetric set in the Baumslag–Solitar group $BS(1, n)$ with $n \neq -1$ is 1. They also prove that the largest cardinality of a commutative totally symmetric set in a product $G \times H$ or $G * H$ is the supremum of the cardinalities for totally symmetric sets in a single factor. They also prove that a solvable group cannot have a commutative totally symmetric set with 5 elements. We refer the reader to their paper for a full accounting of their results.

2.3. **Application to the symmetric group.** Using the basic theory of totally symmetric sets already established, we can give a conceptual explanation of the following classical theorem. This argument originally appeared in the work of the first author [7].

Theorem 2.3. *For $n \geq 7$, the outer automorphism group of Σ_n is trivial.*

Let Z_n denote the totally symmetric set

$$Z_n = \{(1\ i) \mid i \geq 2\} \subseteq \Sigma_n.$$

To prove Theorem 2.3 we use the following auxiliary result [7, Theorem 2], which is a classification of large totally symmetric sets in Σ_n:

> *Let $n \geq 7$. If $X \subset \Sigma_n$ is a totally symmetric set with $|X| \geq n - 1$. Then X is conjugate to Z_n.*

From this fact, the proof of Theorem 2.3 proceeds as follows. Let $f : \Sigma_n \to \Sigma_n$ be an automorphism. Then $f(Z_n)$ is equivalent to Z_n in the following sense: there exists $\tau \in \Sigma_n$ with $\tilde{\tau} f(Z_n) = Z_n$, where $\tilde{\tau}$ is the inner automorphism corresponding to τ. Then $\tilde{\tau} \circ f$ permutes Z_n, so total symmetry gives some $\sigma \in \Sigma_n$ so that

$$\tilde{\sigma} \circ \tilde{\tau} \circ f = \widetilde{\sigma\tau} \circ f$$

is the identity on Z_n. Since Z_n generates Σ_n, we conclude that $\widetilde{\sigma\tau} \circ f = \mathrm{id}$, that is, $f = (\widetilde{\sigma\tau})^{-1}$. In particular, f is an inner automorphism, completing the proof of the theorem.

A slight modification of the above argument yields the following (well-known) generalization of Theorem 2.3.

Theorem 2.4. *Let $7 \leq n \leq m$, and let $f : \Sigma_m \to \Sigma_n$ be a homomorphism whose image is not cyclic. Then $m = n$ and f is an inner automorphism.*

To obtain this stronger theorem, the only additional observation required is that—by the classification of large totally symmetric sets in Σ_n—the restriction $f|Z_n$ cannot be injective when $n < m$. Thus it must be trivial, and so the image is cyclic (of order at most 2).

We would be remiss not to describe the situation for Σ_6, which does have a nontrivial outer automorphism. From the perspective of totally symmetric sets, the reason why this outer automorphism exists is that Σ_6 has two conjugacy classes of totally symmetric sets with five elements: the standard one and its image under the nontrivial outer automorphism.

The proof of Theorem 2.3 given here is not simpler than the classical proof. However, it gives a conceptually simple, structural explanation. Also, the classification of large totally symmetric sets creates a broad tool for studying any homomorphism to or from Σ_n. We will return to this theme several times in what follows.

3. Totally symmetric sets in the general linear group

In this section we turn our attention to the general linear group. The first author and Salter give a classification of large totally symmetric sets in $\mathrm{GL}_n(\mathbb{C})$, Theorem 3.1 below. They used this classification to give a new, conceptual proof of the following classical fact:

Any non-abelian representation of Σ_n has dimension at least $n-1$.

We will start by describing the largest totally symmetric sets in $\mathrm{GL}_n(\mathbb{C})$, then state the classification theorem, and then explain the applications to representation theory. Following the blueprint, we then classify large totally symmetric configurations in $\mathbb{C}^n$, before using this classification to prove the classification of totally symmetric sets in $\mathrm{GL}_n(\mathbb{C})$.

Standard totally symmetric sets. Consider a regular n-simplex $\Delta \subset \mathbb{R}^n$ centered at the origin. The set of vertices $v_1, \ldots, v_{n+1}$ of Δ is a totally symmetric configuration in the sense that for all $\sigma \in \Sigma_{n+1}$, there is an A_σ in $\mathrm{GL}_n(\mathbb{C})$ (in fact in $O(n)$) such that $A_\sigma v_i = v_{\sigma(i)}$. The hyperplanes $v_i^\perp$ also form a totally symmetric configuration with the same choice of A_σ. This allows us to form a totally symmetric set in $\mathrm{GL}_n(\mathbb{C})$ as follows. Pick $\lambda, \mu \in \mathbb{C}$ distinct and non-zero, and define $A_i \in \mathrm{GL}_n(\mathbb{C})$ by declaring $\mathbb{C}v_i$ and $v_i^\perp$ to be λ- and μ-eigenspaces respectively. The set $\mathcal{A}_n = \{A_1, \ldots, A_{n+1}\}$ is then totally symmetric. We refer to (any conjugate of) any such $\mathcal{A}_n$ as a standard totally symmetric set in $\mathrm{GL}_n(\mathbb{C})$.

The classification of large totally symmetric sets. The following theorem says that the above construction is the only construction of totally symmetric sets in $\mathrm{GL}_n(\mathbb{C})$ of cardinality $n+1$.

Theorem 3.1 (Caplinger–Salter)**.** *Let $n \neq 5$ and let $X \subset \mathrm{GL}_n(\mathbb{C})$ be a totally symmetric set. Then $|X| \leq n+1$, and equality is achieved exactly when X is standard.*

This theorem immediately applies to bound the dimension of a faithful representation of any group: if G contains a totally symmetric subset of size n, then G has no faithful representations in dimension less than $n-1$. In the case of Σ_n, we can say more.

Application to representations of the symmetric group. Building on the last idea, let $\rho : \Sigma_n \to \mathrm{GL}_m(\mathbb{C})$ be a non-abelian representation of Σ_n. We would like to show $m \geq n-1$.

For the standard totally symmetric set Z_n in Σ_n we have that $\rho(Z_n) \subset \mathrm{GL}_m(\mathbb{C})$ is a totally symmetric set. We show that if $m < n-1$ then $\rho(Z_n)$ is a singleton, implying that ρ has cyclic image. We treat the cases $m < n-2$ and $m = n-2$ in turn.

If $m < n-2$ then Theorem 3.1 gives that there is no totally symmetric subset of $\mathrm{GL}_m(\mathbb{C})$ of size $n-1$. Thus by Lemma 2.1, $\rho(Z_n)$ is a singleton, as desired.

If $m = n-2$, Theorem 3.1 gives that $\rho(Z_n)$ is a singleton or is standard. But no two distinct elements of the standard totally symmetric set $\mathcal{A}_n$ satisfy the braid relation, so $\rho(Z_n)$ must be a singleton, as desired.

Application to algebraic geometry. Let $\mathcal{U}_{n,d}$ denote the space of smooth degree d hypersurfaces in $\mathbb{CP}^n$. We can continue the above line of reasoning in order to constrain the dimension of a representation of $\pi_1(\mathcal{U}_{n,d})$:

If $\rho : \pi_1(\mathcal{U}_{n,d}) \to \mathrm{GL}_m(\mathbb{C})$ is a non-cyclic representation with $d \geq 5$, then

$$m \geq \left\lceil \frac{d-1}{2} \right\rceil^n - 1.$$

We now outline the proof of this fact. Lönne gave a presentation of $\pi_1(\mathcal{U}_{n,d})$ generalizing the standard presentation of the braid group [21]. From this presentation, we can see that $\pi_1(\mathcal{U}_{n,d})$ has a totally symmetric set that is analogous to the standard totally symmetric set in the braid group (see Section 4), and has cardinality $\lceil \frac{d-1}{2} \rceil^n$. Just like in the braid group, when $d \geq 5$ this set collapses if and only if ρ has cyclic image, so Theorem 3.1 gives the desired bound. (In fact, a slightly better bound can be obtained from a related result in the paper of the first author and Salter [9, Theorem A] that classifies commutative totally symmetric sets in $\mathrm{GL}_n(\mathbb{C})$.)

Totally symmetric configurations. As per Step 2 of the blueprint, we now explain the classification of large totally symmetric configurations used in the classification of large totally symmetric sets in $\mathrm{GL}_n(\mathbb{C})$.

The first author and Salter give the following classification [9, Theorem A]. In the statement, a standard totally symmetric configuration is the collection of 1-dimensional eigenspaces of the elements of the standard totally symmetric set, or the image of this configuration under any element of $\mathrm{GL}_n(\mathbb{C})$. The dual of such a configuration is the set of orthogonal complements (in the paper by the first author and Salter these are referred to as simplex configurations and their duals).

Proposition 3.2 (Caplinger–Salter). *Let $\mathcal{W}$ be a totally symmetric configuration of subspaces in $\mathbb{C}^n$. Then $|\mathcal{W}| \leq n+1$, and when $n \neq 5$, equality is realized only by a standard configuration or the dual of such.*

In the work of the first author and Salter, Proposition 3.2 is proved inductively, in tandem with Theorem 3.1. Here, we assume the proposition without proof, and show how the inductive step for Theorem 3.1 proceeds.

Proof of Theorem 3.1 assuming Proposition 3.2. We discuss the two statements in turn, namely, the upper bound on the size of a totally symmetric set and the classification of large totally symmetric sets.

We proceed by induction on n, with base case $n = 1$. In this base case we are considering $\mathrm{GL}_1(\mathbb{C}) \cong \mathbb{C}^*$. Since the latter is abelian, the largest totally symmetric set is a singleton. As $1 \leq 2$ the first statement of the proposition is verified. The second statement is vacuous since the upper bound $n+1$ is not realized in this case.

Let $X = \{A_1, \ldots, A_k\} \subset \mathrm{GL}_n(\mathbb{C})$ be a totally symmetric set. Consider the generalized eigenspaces

$$E^i_{\lambda,j} = \ker\,(A_i - \lambda I)^j.$$

If for any λ and j the arrangement of subspaces

$$\{E^1_{\lambda,j}, \dots E^k_{\lambda,j}\}$$

is non-degenerate (i.e. not a singleton), then Lemma 2.1 (really, the version for configurations) and Proposition 3.2 give the bound $k \leq n+1$. Thus, we may henceforth assume that all arrangements $\{E^i_{\lambda,j}\}_{i=1}^k$ are degenerate. In other words, the A_i share a common Jordan filtration

$$E^i_{\lambda,1} \subset E^i_{\lambda,2} \subset \cdots$$

for every eigenvalue λ. We may therefore drop the superscript and write $E_{\lambda,j}$ for $E^i_{\lambda,j}$.

Restricting the A_i to any $E_{\lambda,j}$ gives a totally symmetric set in some $\mathrm{GL}_d(\mathbb{Z})$ with $d < n$. If this restricted totally symmetric set is non-degenerate, then by induction we have $k \leq d+1 < n+1$, as desired. Similarly, the maps induced by A_i on the quotients $\mathbb{C}^n/E_{\lambda,j}$ are totally symmetric, and if they are nondegenerate, induction applies.

We are thus left with the case where all restrictions and quotients associated to each $E_{\lambda,j}$ are identical. This is a strong condition. From here, the first author and Salter use a variety of techniques to coax out totally symmetric sets of smaller dimension. They then apply induction to obtain the bound $k \leq n$.

We now turn to the second statement of the theorem, the classification of totally symmetric sets of size $n+1$. The above argument shows that if $k = n+1$, there must be an eigenvalue λ and an index j so that the eigenspace arrangement $\{E^i_{\lambda,j}\}$ corresponding to $X = \{A_i\}$ is nondegenerate. Moreover, Proposition 3.2 implies that this arrangement must be the standard totally symmetric configuration or the dual to such.

Let $\phi : \Sigma_{n+1} \to \mathrm{GL}_n(\mathbb{C})$ be a realization map for $X = \{A_i\}$, by which we mean that

$$\phi(\sigma) A_i \phi(\sigma)^{-1} = A_{\sigma(i)}$$

for all $\sigma \in \Sigma_{n+1}$ and all $i \in \{1, \dots, k\}$. The elements of the image of ϕ must permute the set $\{E^i_{\lambda,j}\}_{i=1}^k$ accordingly.

The first author and Salter show that, up to scaling, the only realization map for the standard totally symmetric subspace configuration (or its dual) is the standard representation $\Sigma_{n+1} \to \mathrm{GL}_n(\mathbb{C})$. Since the standard representation is irreducible, no (proper) arrangement of eigenspaces can be degenerate, as such an arrangement would give an invariant subspace.

The next step is to show that each A_i has two eigenvalues. There are two cases, namely, where $\{E^i_{\lambda,j}\}_{i=1}^k$ is the standard totally symmetric configuration and its dual. To illustrate the idea, we treat the former case. Assume for the purposes of contradiction that λ is the only eigenvalue for the A_i. In this case, $\mathbb{C}^n$ is a single Jordan block, and $\{E_{\lambda,j}\}$ does not stabilize until $j = n$. But then if $n > 2$, the generalized eigenspace $\{E_{\lambda,2}\}$ would be a non-degenerate totally symmetric collection of 2-dimensional subspaces with $n+1$ elements. This is impossible by Proposition 3.2.

By the previous paragraph, we may assume that each A_i has two distinct eigenvalues, λ and μ. The standard representation $\phi : \Sigma_{n+1} \to \mathrm{GL}_n(\mathbb{C})$ must be a realization map for both $\{E_{\lambda,1}\}$ and $\{E_{\mu,1}\}$. Then both the λ- and μ-eigenspaces of A_i must be stabilized by $\phi(\mathrm{Stab}(i))$, whose fixed subspaces are D_i and $\mathbb{C}v_i$. Since $\mathrm{Stab}(i) \to D_i$ has no subrepresentations, the eigenspaces must be exactly D_i and $\mathbb{C}v_i$. The theorem follows.

4. Totally symmetric sets in braid groups

In this section we use the theory of totally symmetric sets to outline a proof of the following result originally due to Dyer–Grossman [14]. Our argument is a modification of the one used by Kordek and the second author [19], simplified for this special case.

For the statement, let $\sigma_1, \ldots, \sigma_{n-1}$ denote the standard half-twist generators for the braid group B_n, and let ϵ be the automorphism of B_n given by $\sigma_i \mapsto \sigma_i^{-1}$ for all i.

Theorem 4.1 (Dyer–Grossman)**.** *For $n \geq 3$ the automorphism ϵ represents the unique nontrivial outer automorphism of B_n. In particular,*

$$\mathrm{Aut}(B_n) \cong B_n/Z(B_n) \rtimes \mathbb{Z}/2.$$

The original proof of this theorem by Dyer–Grossman has an algebraic flavor. The proof we outline here, while only valid as stated for $n \geq 5$, uses combinatorial topology and the theory of mapping class groups.

In what follows we consider B_n as the mapping class group of the disk D_n with n marked points in the interior (not to be confused with the dihedral group!). We will use a number of aspects of the theory of mapping class groups, including the Nielsen–Thurston classification theorem, the canonical reduction systems of Birman–Lubotzky–McCarthy, and the change of coordinates principle. We refer the reader to the book by Farb and the second author of this article for background on these topics [15]. From the point of view of mapping class groups, the Dyer–Grossman theorem can be stated as: every automorphism of B_n is induced by a homeomorphism of D_n.

The outline for this section mirrors the one for Section 3. This stands to reason, as we will be following the same blueprint.

A large commutative totally symmetric set. As per Step 1 of the blueprint, we will require the services of a large (commutative) totally symmetric set in B_n. The desired set is:

$$\{\sigma_1, \sigma_3, \sigma_5, \ldots, \sigma_m\}$$

where m is $n-1$ or $n-2$, according to whether n is even or odd, respectively. As such, the cardinality of this set is $\lfloor n/2 \rfloor$. That this set is a totally symmetric set is an application of the change of coordinates principle from the theory of mapping class groups. We refer to any B_n-conjugate of this totally symmetric set as a standard totally symmetric set in B_n.

Crash course in Nielsen–Thurston theory. Briefly, the Nielsen–Thurston classification gives that every braid is periodic, pseudo-Anosov, or reducible. Periodic braids have powers that are central in B_n; they correspond to rotations of D_n. Each reducible braid preserves a multicurve, that is, the isotopy class of a collection of pairwise disjoint and pairwise non-homotopic simple closed curves in D_n. Any such multicurve is called a reduction system for the braid. Pseudo-Anosov braids do not preserve any multicurve.

For a reducible braid b, we may restrict b to the complementary components, and inductively apply the classification. Thus, there exists a reduction system with the property that the associated restrictions are all periodic or pseudo-Anosov. There is in fact a unique minimal such reduction system, called the canonical reduction system $\mathrm{CRS}(b)$. We will use the following properties of canonical reduction systems:

1. for $a, b \in B_n$ we have $\mathrm{CRS}(aba^{-1}) = a\,\mathrm{CRS}(b)$, and
2. if a and b commute then $\mathrm{CRS}(a)$ and $\mathrm{CRS}(b)$ have trivial geometric intersection.

For the proof, the main points to keep in mind are that there is a map

$$\mathrm{CRS} : B_n \to \{\text{multicurves in } D_n\},$$

and that this map satisfies the above two properties.

Totally symmetric configurations of multicurves. The configurations we will use to address Step 2 of the blueprint are collections of multicurves in the disk D_n. To each element g of B_n we associate its canonical reduction system $\mathrm{CRS}(g)$. Then, to a totally symmetric set $X = \{g_1, \dots, g_k\}$ we can associate the collection of multicurves

$$\{\mathrm{CRS}(g_1), \dots, \mathrm{CRS}(g_k)\}.$$

Because X is totally symmetric, this multicurve configuration is totally symmetric in the sense that for any $\sigma \in \Sigma_k$ there is a braid g_σ so that

$$g_\sigma \cdot (\mathrm{CRS}(g_1), \dots, \mathrm{CRS}(g_k)) = (\mathrm{CRS}(g_{\sigma(1)}), \dots, \mathrm{CRS}(g_{\sigma(k)}))$$

(using the first property of CRS above). When X is a commutative totally symmetric set, the multicurves $\mathrm{CRS}(g_i)$ have trivial intersection pairwise (using the second property); in what follows, we refer to such a collection as a noncrossing multicurve configuration.

In the original work of Kordek and the second author, they associate a single labeled multicurve to X instead of a configuration of multicurves. This is equivalent, but we take this point of view to make the analogy with $\mathrm{GL}_n(\mathbb{C})$ more clear.

Large noncrossing totally symmetric multicurve configurations. We now turn towards Step 2 of the blueprint. Let us realize D_n at the closed unit disk in the complex plane, with all marked points on the real axis. For $1 \le i \le n-1$, let c_i be the isotopy class of curves corresponding to a round circle surrounding the ith and $(i+1)$st marked points. The curves are chosen precisely so that $\mathrm{CRS}(\sigma_i)$ is equal to c_i. We define the following noncrossing totally symmetric multicurve in D_n:

$$M_n = \{c_1, c_3, \dots\}.$$

This is the noncrossing totally symmetric multicurve configuration associated to the totally symmetric set $\{\sigma_1, \sigma_3, \dots\}$ in B_n.

There are two variations on M_n that we will need to consider. First, we have the dual totally symmetric configuration

$$M_n^* = \{\{c_1\}^{\mathsf{C}}, \{c_3\}^{\mathsf{C}}, \dots\}$$

where $\{c_i\}^{\mathsf{C}}$ is the complementary configuration to c_i in M_n. In other words, $\{c_i\}^{\mathsf{C}}$ is the multicurve whose components are all the curves appearing in M_n except for c_i.

Second, when n is odd, we have the totally symmetric labeled multicurve

$$\widehat{M_n} = \{\{c_1, d\}, \{c_3, d\}, \dots\},$$

where d is represented by the round curve surrounding the first $n-1$ marked points. Finally, we can combine these two variations in order to obtain

$$\widehat{M_n^*} = \{\{c_1\}^{\mathsf{C}} \cup \{d\}, \{c_3\}^{\mathsf{C}} \cup \{d\}, \dots\}.$$

That all of these multicurve configurations are totally symmetric follows again from the change of coordinates principle.

Kordek and the second author prove [19, Lemma 2.3] that in fact these are the only examples of large noncrossing totally symmetric multicurve configurations in D_n. In the statement, a configuration of multicurves $\{m_1, \dots, m_k\}$ is degenerate if two m_i are equal.

Proposition 4.2 (Kordek–Margalit)**.** *Let $M = \{m_1, \dots, m_k\}$ be a nondegenerate, noncrossing, totally symmetric multicurve configuration in D_n with $k = \lfloor n/2 \rfloor$. Then M is B_n-equivalent to one of M_n, M_n^*, $\widehat{M}_n$, or $\widehat{M}_n^*$.*

The main idea of the proof is as follows. Suppose that some m_i contains a curve c_i^p surrounding p marked points. By total symmetry each of the m_i contains such a curve. Then if $p > 2$ it must be that these c_i^p are not distinct (otherwise the noncrossing condition would be violated). We are then led to consider the case that some curve c surrounds exactly p marked points and lies in exactly d of the m_i. Again applying total symmetry, there must be $\binom{k}{d}$ such curves, all with pairwise trivial geometric intersection, and all surrounding p marked points. But for $d < k$, the quantity $\binom{k}{d}$ is quadratic in $k = \lfloor n/2 \rfloor$, hence quadratic in n. Again, this violates the noncrossing condition (there are in fact at most $n-2$ pairwise non-isotopic curves in D_n with pairwise trivial geometric intersection). It follows that the only possibilities are that each curve appearing in an m_i surrounds exactly two marked points, or it lies in all the m_i. From here the proof is straightforward.

Large commutative totally symmetric sets. Continuing with Step 2 of the blueprint, we now explain how Proposition 4.2 is used to classify large commutative totally symmetric sets in B_n.

To each noncrossing totally symmetric multicurve configuration M_n, M_n^*, $\widehat{M}_n$ and $\widehat{M}_n^*$, there is an associated commutative totally symmetric set in B_n, namely:

$$Z_n = \{\sigma_1, \sigma_3, \dots\},\ Z_n^* = \{\sigma_1^*, \sigma_3^*, \dots\},\ \widehat{Z}_n = \{\sigma_1 T_d, \sigma_3 T_d, \dots\},\ \text{and } \widehat{Z}_n^* = \{\sigma_1^* T_d, \sigma_3^* T_d, \dots\}.$$

Here σ_i^* is the product of the elements of Z_n not equal to σ_i and T_d is the Dehn twist about the curve used in the definitions of $\widehat{M}_n$ and $\widehat{M}_n^*$. We refer to any B_n-conjugate of any of these as a standard commutative totally symmetric set in B_n.

We can modify any of the standard totally symmetric sets by raising all elements to the same nonzero power. We can also modify them by multiplying all elements of the set by the same power of z, a generator for the (cyclic) center of B_n. We refer to any totally symmetric set obtained in this way as a modification of a standard commutative totally symmetric set in B_n. Kordek and the second author prove the following [19, Lemma 2.6].

Proposition 4.3. *Let $n \geq 3$ and let $k = \lfloor n/2 \rfloor$. Any commutative totally symmetric set X in B_n with $|X| = k$ is a modification of a standard one.*

To prove the proposition, we of course first use the fact that, through canonical reduction systems, X gives rise to a noncrossing totally symmetric multicurve configuration. We then use Proposition 4.2 to reduce to the four cases of multicurve configurations given there.

We then treat the four cases in turn. For the case $\mathrm{CRS}(X) = M_n$, the idea is as follows. Say that X is $\{g_1, \dots, g_k\}$. Up to conjugation in B_n we may assume that $\mathrm{CRS}(g_i)$ is equal to c_{2i-1}. Note that this is the canonical reduction system of σ_{2i-1}, the ith element of Z_n.

We would like to show that X is equal to Z_n. On the exterior of c_1 the element g_1 is either the identity, periodic, or pseudo-Anosov. But since g_1 commutes with the other g_i, it must fix the curves $c_3, c_5, \dots$ There is no (nontrivial) periodic or pseudo-Anosov map that can fix these curves. It follows that this exterior component of g_1 is trivial, and hence that g_1 is equal to $\sigma_1^\ell z^s$ for some nonzero ℓ and some s. The other three cases are similar.

Proof of Theorem 4.1 assuming Proposition 4.3. Let $\rho : B_n \to B_n$ be an automorphism. By Lemma 2.1, the image of $Z_n = \{\sigma_1, \sigma_3, \dots\}$ is either a singleton or a commutative totally symmetric set of the same size. In the first case, it follows that ρ has cyclic image. Indeed, for $n \geq 5$ the normal closure of $\sigma_1\sigma_3^{-1}$ is B_n' and $B_n/B_n' \cong \mathbb{Z}$. Thus, we may henceforth assume that $\rho(Z_n)$ is a commutative totally symmetric set of cardinality $|Z_n|$.

By Proposition 4.3, $\rho(Z_n)$ is conjugate to a modification of Z_n, Z_n^*, $\widehat{Z}_n$, or $\widehat{Z}_n^*$. Let us consider the first case. Up to conjugating ρ, we may assume that $\rho(Z_n)$ is exactly a modification of Z_n, that is,

$$\rho(\sigma_i) = \sigma_i^\ell z^s$$

for all odd i.

For i even, we then have that $\rho(\sigma_i)$ is conjugate to $\sigma_i^\ell z^s$. So each such $\rho(\sigma_i)$ is equal to $H_{a_i}^\ell z^s$, where H_{a_i} is the half-twist about a curve a_i.

It is a fact that if H_a and H_b are the half-twists about curves a and b in D_n, and they satisfy the braid relation

$$H_a^\ell H_b^\ell H_a^\ell = H_b^\ell H_a^\ell H_b^\ell,$$

then $i(a,b) = 2$ and $\ell = \pm 1$; see [5, Lemma 4.9]. Up to the exceptional automorphism ϵ, we may assume that $\ell = 1$. It then further follows that $s = 0$, since an automorphism of B_n must preserve word length, that is, it respects the abelianization $B_n \to \mathbb{Z}$.

It also follows that the sequence of curves

$$c_1, a_2, c_3, a_4, \ldots$$

is a chain, meaning that consecutive curves intersect twice and all other pairs of curves have trivial geometric intersection. Up to automorphisms of B_n, we then have (by change of coordinates)

$$\rho(\sigma_i) = \sigma_i$$

for all i. In other words, up to modifying ρ by automorphisms, it is the identity. This completes the proof in the first case. Using similar reasoning, we rule out the other three possibilities for $\rho(Z_n)$, completing the proof.

From braid groups to mapping class groups. Chen–Mukherjea [11] use a similar approach to classify homomorphisms from the braid group B_n to the mapping class group $\mathrm{Mod}(S_g)$ when $g < n - 2$. As a corollary, they partially recover the result of Aramayona–Souto [1] classifying homomorphisms $\mathrm{Mod}(S_g) \to \mathrm{Mod}(S_h)$ for $h < 2g$.

5. Finite quotients of braid groups and mapping class groups

In 1947, Emil Artin [3] proved that for $n \geq 5$ every non-cyclic homomorphism $B_n \to \Sigma_n$ is standard. This means that up to conjugacy, the map sends σ_i to the transposition $(i\ \ i{+}1)$ for all i. His proof uses Bertrand's postulate, a deep fact from number theory which states that every interval $[n, 2n]$ contains a prime number. Artin wrote: "it would be preferable if a proof could be found that does not make use of this fact."

Kolay [17] found in 2021 a short, elementary proof of Artin's theorem, and in fact proved more. In the statement, we say that a quotient map is minimal if there is no quotient map whose codomain has smaller cardinality.

Theorem 5.1 (Kolay)**.** *Let $n \geq 3$. Up to conjugacy, there is a unique minimal non-cyclic quotient of B_n, namely, the standard map $B_n \to \Sigma_n$ for $n \neq 4$ or the standard map $B_4 \to \Sigma_3$.*

In 2019 the second author of this paper had asked: *What is the smallest non-cyclic quotient of B_n? Is it Σ_n?* Kolay's theorem answers this in the affirmative.

We give Kolay's stunningly simple proof below. The main ingredients are (1) a large collapsing set in B_n and (2) the orbit-stabilizer theorem. While Artin never defined collapsing sets, he certainly had all of the tools to prove Kolay's theorem. It is remarkable that 74 years passed in between the two works.

Before Kolay's work, partial answers to the second author's question were given by Chudnovsky–Kordek–Li–Partin [12], Caplinger–Kordek [8], and Scherich–Verberne [22].

5.1. **Braid groups.** In this section we explain Kolay's proof of Theorem 5.1, and in the next we explain how Kolay applied the same ideas to the case of the mapping class group.

Collapsing sets. Let G be a group. We say that a subset $X \subseteq G$ is a collapsing set if for every group homomorphism $f : G \to H$ the restriction $f|X$ is either injective or constant. This notion is a generalization of totally symmetric sets. Indeed, Lemma 2.1 implies that every totally symmetric set is a collapsing set.

We also remark that under any homomorphism, a collapsing set maps to a collapsing set. Therefore, there is an analogous blueprint for collapsing sets, an idea that seems to be unexplored.

Strong collapsing sets. Let G be a group and let $X = \{x_1, \ldots, x_k\} \subseteq G$ be a subset. We say that X is a strong collapsing set if

$$G/\left\langle\!\left\langle x_i x_j^{-1} \right\rangle\!\right\rangle$$

is abelian for all pairs $\{i, j\}$. This is the same as saying that the normal closure of each $x_i x_j^{-1}$ contains the commutator subgroup $[G, G]$. If the x_i are all conjugate, the $x_i x_j^{-1}$ lie in $[G, G]$ and so each $G/\left\langle\!\left\langle x_i x_j^{-1} \right\rangle\!\right\rangle$ must exactly be the abelianization of G. It follows from this that a strong collapsing set of conjugate elements is a collapsing set, since conjugate elements in an abelian group are equal.

A large collapsing set. Similar to Step 1 of the totally symmetric set blueprint, we will make use of a large strong collapsing set in B_n for $n \geq 5$.

For each unordered pair $I = \{i, j\} \subseteq \{1, \ldots, n\}$ let σ_I denote the half-twist in B_n given by the counter-clockwise exchange of the ith and jth marked points in the upper half of D_n. If I and J are distinct ordered pairs then $\sigma_I \sigma_J^{-1}$ is conjugate in B_n to exactly one of the following: $\sigma_1 \sigma_2^{-1}$, $\sigma_1 \sigma_3^{-1}$, or $b = \sigma_{\{1,3\}} \sigma_{\{2,4\}}^{-1}$.

We claim that for $n \geq 5$ the normal closure of any of these three elements in B_n is the commutator subgroup B_n'. For the first two elements, this is a standard fact. Similarly, the commutator $[b, \sigma_4]$ both lies in the normal closure of b and is conjugate to $\sigma_1 \sigma_2^{-1}$. It follows that the normal closure of b is again B_n'. (An alternate, but equivalent, proof of the claim is given by the well-suited arc criterion of Lanier and the second author [10, Lemma 6.2].)

It follows from the claim that the set of all σ_I is a strong collapsing set for B_n. We refer to this as the standard strong collapsing set for B_n and denote it X_n.

Two basic group theory facts. In the proof of Theorem 5.1 we will use the following fact:

> *(1) If $f : \mathbb{Z} \times H \to G$ is a group homomorphism, and t denotes a generator of $\mathbb{Z}$, then $f(t) \notin f(H)$ if and only if*
>
> $$|f(\mathbb{Z} \times H)| \geq 2|f(H)|.$$

This fact is true because both conditions are equivalent to the statement that $f(t)f(H)$ is a nontrivial coset of $f(H)$ in $f(\mathbb{Z} \times H)$. We will also use the following:

> *(2) If G is a group, then $Z(G)$ is nontrivial if and only if*
>
> $$|G| \geq 2|G/Z(G)|.$$

This is true because if $Z(G)$ is nontrivial then each coset has at least two elements. While both statements make sense for infinite groups, we will only apply them when G is finite.

Base cases. We can prove the $n = 3$ and $n = 4$ cases of Theorem 5.1 by direct inspection. Because the abelianization of B_n is cyclic, a non-cyclic quotient of B_n is non-abelian. The only non-abelian group of order 6 or less is Σ_3. Thus, all other finite non-abelian quotients of B_3 and B_4 have order strictly greater than 6.

To see that the standard maps $B_3 \to \Sigma_3$ and $B_4 \to \Sigma_3$ are unique up to automorphisms of Σ_3, we simply check that (up to automorphisms of Σ_3) the only ordered pair of elements of Σ_3 satisfying the braid relation is $((1\ 2), (2\ 3))$.

Extension of the $n = 4$ case. We also will require the following statement:

> If $f : B_4 \to G$ is a quotient map that is injective on X_4 then $|G| \geq 4!$. Further, if $|G| = 4!$ then $G = \Sigma_4$ and f is standard.

Since this can be easily proved with a computer, we omit the proof (although it is a fun exercise to do it by hand!).

Kolay's proof. We prove the theorem by induction on n, with the base cases $n = 3$ and $n = 4$ (and the extension of the latter) handled as above.

Let $f : B_n \to G$ be a non-cyclic quotient map. Let X_n be the standard strong collapsing set in B_n. The group $G = f(B_n)$ acts by conjugation on the conjugacy class of $f(\sigma_1)$ in G. To prove that $|G| \geq n!$, we will apply the orbit-stabilizer theorem to this action. A further analysis will give the statement that f is conjugate to the standard map to Σ_n.

For the orbit-stabilizer argument there are, naturally, two steps. Specifically, if $\mathcal{O} \subseteq G$ and $\mathcal{S} \subseteq G$ are the orbit and stabilizer of $f(\sigma_1)$ then we will show that

$$|\mathcal{O}| \geq \binom{n}{2} \quad \text{and} \quad |\mathcal{S}| \geq 2 \cdot (n-2)!.$$

Orbit. Since X_n is a strong collapsing set, and since G is not cyclic, it follows that $|f(X_n)| = |X_n| = \binom{n}{2}$. In particular $|\mathcal{O}| \geq \binom{n}{2}$.

Stabilizer. We consider the action of B_n on itself by conjugation. The stabilizer of σ_1 in B_n contains a subgroup

$$\langle \sigma_1, \sigma_3, \sigma_4, \ldots, \sigma_{n-1} \rangle \cong \mathbb{Z} \times B_{n-2}.$$

The image $f(\mathbb{Z} \times B_{n-2})$ is a subgroup of $\mathcal{S}$. We would like to bound the cardinality of this image from below. We treat two cases, according to whether $f(\sigma_1)$ lies in $f(B_{n-2})$.

By induction we may assume that $|f(B_{n-2})| \geq (n-2)!$. Indeed, if $f(B_{n-2})$ were cyclic, then f would be cyclic, contrary to assumption. We will use this assumption in both cases.

Case 1: $f(\sigma_1) \notin f(B_{n-2})$. By the first basic group theory fact above, we have

$$|\mathcal{S}| \geq |f(\mathbb{Z} \times B_{n-2})| \geq 2|f(B_{n-2})| \geq 2(n-2)!$$

as desired.

Case 2: $f(\sigma_1) \in f(B_{n-2})$. In this case $f(\mathbb{Z} \times B_{n-2}) = f(B_{n-2})$. Since f is nontrivial, $f(\sigma_1)$ is nontrivial. Since σ_1 lies in the centralizer of B_{n-2} it must be that $f(\sigma_1)$ lies in the center of $f(B_{n-2})$. In particular, $Z(f(B_{n-2}))$ is nontrivial.

We claim that $f(B_{n-2})/Z(f(B_{n-2}))$ is not cyclic. Indeed, if it were cyclic then $f(B_{n-2})$ would be abelian (for any group G, if $G/Z(G)$ is abelian then G is). The abelianization of

B_{n-2} is cyclic, and so any abelian quotient of it is cyclic. In particular, $f(B_{n-2})$ is cyclic. It then follows that $f(B_n)$ is cyclic, contrary to assumption.

By the claim, the group $f(B_{n-2})/Z(f(B_{n-2}))$ is a non-cyclic quotient of B_{n-2}. By induction, its order is bounded below by $(n-2)!$. By the second basic group theory fact, we have

$$|\mathcal{S}| \geq |f(\mathbb{Z} \times B_{n-2})| = |f(B_{n-2})| \geq 2|f(B_{n-2})/Z(f(B_{n-2}))| \geq 2(n-2)!$$

We may now complete the proof of the first statement. By the orbit-stabilizer theorem, we have

$$|G| \geq f(B_n) \geq |\mathcal{O}||\mathcal{S}| \geq \binom{n}{2} 2 \cdot (n-2)! = \frac{n(n-1)}{2} \cdot 2 \cdot (n-2)! = n!$$

The first statement, $n = 6$ case. For $n = 6$ the argument is the same, except we must use the extension of the base case $n = 4$. Since we may assume that f is injective on X_6, it is injective on the copy of X_4 associated to $B_4 \leqslant B_6$. Hence the size of the stabilizer $\mathcal{S}$ is bounded below by $2 \cdot 4!$.

The second statement. To prove the stronger statement that any quotient of B_n with order $n!$ is the standard one, it suffices to show that the $f(\sigma_i)$ have order 2 (because the quotient $B_n \to B_n / \langle\langle \sigma_1^2 \rangle\rangle \cong \Sigma_n$ is the standard quotient). But this is true because in order to realize the lower bound $|\mathcal{S}| \geq 2 \cdot (n-2)!$ it must be true by induction that $f(B_{n-2})$ is the standard quotient. □

5.2. **Mapping class groups.** We now turn our attention to the analogue of Theorem 5.1 for mapping class groups. The natural action of $\mathrm{Mod}(S_g)$ on $H_1(S_g; \mathbb{F}_2)$ gives rise to a representation

$$\mathrm{Mod}(S_g) \to \mathrm{Sp}_{2g}(\mathbb{F}_2).$$

The order of the latter group is

$$|\operatorname{Sp}_{2g}(\mathbb{F}_2)| = 2^{g^2} \prod_{i=1}^{g} (2^{2i} - 1).$$

Remarkably, we again have that the most natural small quotient is the smallest.

Theorem 5.2 (Kielak–Pierro)**.** *Let $g \geq 1$. Up to conjugacy, there is a unique minimal non-cyclic quotient of* $\mathrm{Mod}(S_g)$*, namely, the standard map* $\mathrm{Mod}(S_g) \to \mathrm{Sp}_{2g}(\mathbb{F}_2)$*.*

This theorem was conjectured by Zimmermann [23] in 2012. Kielak–Pierro [16] proved it in 2019, using the approach established by Baumeister–Kielak–Pierro [4] in their work on the analogous problem about outer automorphisms of free groups.

The Kielak–Pierro proof of Theorem 5.2 relies on the classification of finite simple groups and the representation theory of the mapping class group, as well as the deep work of Berrick–Gebhardt–Paris [6], which itself uses the Matsumoto presentation of the mapping class group. It is astonishing that Kolay's argument for the braid group applies with little modification to prove the same theorem.

Again, the keys to Kolay's proof of Theorem 5.2 are the construction of a large strong collapsing set in $\mathrm{Mod}(S_g)$, and an orbit-stabilizer argument.

The general outline is closely analogous to the argument for the braid group. Even the description of the large collapsing set constructed in Step 1 is similar. Several new tools are required. We introduce these in turn as we go.

Reduction to the open case. Let S_g^1 denote the surface with one boundary component obtained from S_g by removing the interior of an embedded disk. Since $H_1(S_g^1;\mathbb{F}_2)$ is naturally isomorphic to $H_1(S_g;\mathbb{F}_2)$ we also have a natural map

$$\mathrm{Mod}(S_g^1) \to \mathrm{Sp}_{2g}(\mathbb{F}_2).$$

By filling the disk back in, we also obtain a quotient map

$$\mathrm{Mod}(S_g^1) \to \mathrm{Mod}(S_g).$$

In general if (within some class of groups) G is the smallest quotient of a group M_1, and M is a quotient of M_1 that also has G as a quotient, then G is the smallest quotient of M (in that class of groups). Thus to prove Theorem 5.2 it suffices to prove the analogous statement for $\mathrm{Mod}(S_g^1)$.

The base case. The base case for the $\mathrm{Mod}(S_g^1)$-version of Theorem 5.2 is the case $g = 1$. In this case Dehn [13, p. 172] proved that $\mathrm{Mod}(S_g^1) \cong B_3$. By Theorem 5.1, the smallest non-cyclic quotient of this group is $\Sigma_3 \cong \mathrm{Sp}_2(\mathbb{F}_2)$, as desired. (One way to prove the last isomorphism is to use the formula for the cardinality of $\mathrm{Sp}_2(\mathbb{F}_2)$ and apply Theorem 5.1!).

A well-suited curve criterion. We now turn our attention to Step 1. In the braid group case, the construction of the large strong collapsing set used the well-known fact that for $i \neq j$ the quotient

$$B_n \,/\, \langle\!\langle \sigma_i\sigma_j^{-1} \rangle\!\rangle$$

is cyclic. To mimic this step, we will use the following fact:

> *Let $f \in \mathrm{Mod}(S_g^1)$ and suppose c is a curve with $i(c, f(c)) = 1$. Then the normal closure $\langle\!\langle f \rangle\!\rangle$ contains the commutator subgroup $\mathrm{Mod}(S_g^1)'$ and*
>
> $$\mathrm{Mod}(S_g^1)/\,\langle\!\langle f \rangle\!\rangle$$
>
> *is cyclic.*

This fact is an instance of the well-suited curve criterion of Lanier and the second author [20, Lemma 2.1]. While their argument is given explicitly for $\mathrm{Mod}(S_g)$, it applies verbatim for $\mathrm{Mod}(S_g^1)$. The key points are that—like $\mathrm{Mod}(S_g)$—the abelianization of $\mathrm{Mod}(S_g^1)$ is cyclic for $g \geq 1$ and—like $\mathrm{Mod}(S_g)$— the group $\mathrm{Mod}(S_g^1)$ has a generating set consisting of Dehn twists about curves that have pairwise intersection at most 1. (In the construction of the large strong collapsing set for B_n we noted that we could have used the well-suited arc criterion in the proof. Similarly, it is true here that we can give a proof that mimics the braid group case more closely. We leave it to the reader to decide which approach they prefer.)

The hyperelliptic involution and mod 2 homology. A hyperelliptic involution of S_g^1 is a homeomorphism ι of order two with $2g+1$ fixed points. The quotient $S_g^1/\langle\iota\rangle$ is D_{2g+1}, the disk with $2g+1$ marked points. These marked points are the images of the fixed points of ι. We denote the set of marked points by P.

Let D_{2g+1}° denote the disk with $2g+1$ punctures obtained by removing P. The homology group $H_1(D_{2g+1}^\circ) \cong (\mathbb{F}_2)^{2g+1}$ has a canonical generating set, namely, the classes represented by small loops around the punctures. This gives rise to a canonical homomorphism

$$H_1(D_{2g+1}^\circ) \to \mathbb{F}_2,$$

whereby each of these generators maps to 1. We denote the kernel by $H_1(D^\circ_{2g+1};\mathbb{F}_2)^{\text{even}}$. The elements of this kernel are exactly the ones represented by simple closed curves in D_{2g+1} surrounding an even number of marked points.

We will define a map

$$\Psi : H_1(S^1_g;\mathbb{F}_2) \to H_1(D^\circ_{2g+1};\mathbb{F}_2)$$

as follows. Given $v \in H_1(S^1_g;\mathbb{F}_2)$, we may represent v by a simple closed curve c that avoids the fixed points of ι. The image of c in D_{2g+1} represents an element of $H_1(D^\circ_{2g+1};\mathbb{F}_2)$.

Arnol'd [2] gave the following (easy-to-prove but) remarkable fact:

The map Ψ is an isomorphism

$$\Psi : H_1(S^1_g;\mathbb{F}_2) \xrightarrow{\cong} H_1(D^\circ_{2g+1};\mathbb{F}_2)^{even}$$

The map Ψ^{-1} can be described as follows. Given an element v of $H_1(D^\circ_{2g+1};\mathbb{F}_2)^{\text{even}}$ we represent it by a simple closed curve c, and $\Psi^{-1}(v)$ is the class represented by one component of the preimage of c.

Say that a subset of P is even if it has an even number of elements. By the above discussion we have natural bijections

$$H_1(D^\circ_{2g+1};\mathbb{F}_2)^{\text{even}} \leftrightarrow H_1(S^1_g;\mathbb{F}_2) \leftrightarrow \{\text{even subsets of } P\}$$

Further, the nonzero elements of $H_1(S^1_g;\mathbb{F}_2)$ correspond to the nonempty even subsets of P.

The large collapsing set. Let us represent D_{2g+1} as a disk with the points of P lying on a circle. For each nonempty subset $A \subseteq P$ there is, up to isotopy, a unique curve c_A in D_{2g+1} that bounds a convex disk containing exactly the points of P contained in A.

If A is even then the preimage of c_A in S^1_g has exactly two components. We choose one of these (arbitrarily) and call it $\tilde{c}_A$.

We will show that the set of Dehn twists

$$X_g = \{T_{\tilde{c}_A} \mid A \subseteq P \text{ even}, \quad A \neq \emptyset\}$$

is a strong collapsing set in $\mathrm{Mod}(S^1_g)$.

In order to prove this, we take A and B to be distinct nonempty even subsets of P, and we assume that

$$f : \mathrm{Mod}(S^1_g) \to G$$

is a homomorphism with

$$f(T_{\tilde{c}_A}) = f(T_{\tilde{c}_B}).$$

For any choices of A and B, there exists an arc c in D_{2g+1} that connects two marked points, that intersects c_A in one point, and that is disjoint from c_B. To check this, we consider two cases, according to whether or not $A \cup B$ is a proper subset of P or not. In the first case, let $p \in P \setminus (A \cup B)$; we take c to connect p to a point $q \in A$. If $A \cup B = P$, then since A and B are even there is a point p in $A \cap B$, and c connects any such p to a point $q \in A \setminus B$.

The preimage of c in S^1_g is a simple closed curve $\tilde{c}$ with $i(\tilde{c}, \tilde{c}_A) = 1$ and $i(\tilde{c}, \tilde{c}_B) = 0$. It follows that

$$i(\tilde{c}, T_{\tilde{c}_A} T^{-1}_{\tilde{c}_B}(\tilde{c})) = i(\tilde{c}, T_{\tilde{c}_A}(\tilde{c})) = 1.$$

By the above well-suited curve criterion, we conclude that X_g is a strong collapsing set, as desired.

A derivative collapsing set. Say that two elements of the strong collapsing set X_g are dual if the corresponding curves have intersection number 1 (equivalently if they have algebraic intersection number 1). We define

$$X_g' = \{(x, y) \in X_g \times X_g \mid x \text{ is dual to } y\}.$$

The set X_g' is a collapsing set in the following sense: if $f : \mathrm{Mod}(S_g^1) \to G$ is a non-cyclic homomorphism, then each $f(x, y)$ is an ordered pair of distinct elements, and if the f-image of any two elements of X_g' coincide, they all coincide. Both of these statements follow from the fact that X_g is a collapsing set.

Since X_g is in bijection with the nonzero elements of $H_1(S_g^1; \mathbb{F}_2)$, it has $2^{2g} - 1$ elements. Given one nonzero element of $H_1(S_g^1; \mathbb{F}_2)$, there is a codimension-1 affine subspace of $H_1(S_g^1; \mathbb{F}_2)$ corresponding to dual elements in $H_1(S_g^1; \mathbb{F}_2)$. Thus we have

$$|X_g'| = (2^{2g} - 1) \cdot 2^{2g-1}$$

We are finally ready for the proof of the Kielak–Pierro theorem.

Proof of Theorem 5.2. We proceed by induction on g. We already checked the base case $g = 1$. Assume then that the theorem holds for $\mathrm{Mod}(S_g^1)$ with $g \geq 1$. We will show that it holds for $\mathrm{Mod}(S_{g+1}^1)$.

Let $f : \mathrm{Mod}(S_{g+1}^1) \to G$ be a non-cyclic quotient. We use essentially the same orbit-stabilizer argument that was used in the proof of Theorem 5.1. The group G acts by conjugation on the set of ordered pairs of elements of G. Because X_{g+1}' is a collapsing set in the sense described above, the orbit is bounded below by

$$|f(X_{g+1}')| = |X_{g+1}'| = (2^{2(g+1)} - 1) \cdot 2^{2(g+1)-1} = (2^{2(g+1)} - 1) \cdot 2^{2g+1}$$

The stabilizer contains a copy of the image of $\mathrm{Mod}(S_g^1)$. By induction this gives a lower bound of

$$|f(\mathrm{Mod}(S_g^1)| \geq 2^{g^2} \prod_{i=1}^{g} (2^{2i} - 1).$$

Multiplying these together gives the desired bound

$$2^{(g+1)^2} \prod_{i=1}^{g+1} (2^{2i} - 1).$$

For this bound to be realized, the elements of X_g must map to elements of order 2. From there it follows that the quotient is $\mathrm{Sp}_{2(g+1)}(\mathbb{F}_2)$. This completes the proof.

A final lament. Kolay's proof of Theorem 5.2 goes through the derived collapsing set X_g'. An analogous argument can be used to give an unnecessarily complicated proof of Theorem 5.1 (about braid groups). The situation suggests to the authors that there should be a proof of Theorem 5.2 that uses X_g directly, and decreases genus in two inductive steps. We were not able to find such a proof. We implore the reader to find one.

6. Speculations and representations

As suggested to us by Kordek, there is a strong analogy between the collision-implies-collapse property and Schur's lemma from representation theory. We can give weight to this analogy as follows.

A homomorphism $f : G \to H$ induces a linear map $f_* : \mathbb{C}[G] \to \mathbb{C}[H]$. The vector spaces $\mathbb{C}[G]$ and $\mathbb{C}[H]$ come equipped with a G-action and an H-action, respectively, where both groups act by conjugation on the basis elements. If $X \subseteq G$ is a totally symmetric set with $|X| = k$, and G_X is the stabilizer of X in G, then $\mathbb{C}[X]$ is a representation of G_X. By total symmetry G_X surjects onto Σ_k, and so $\mathbb{C}[X]$ is a representation of G_X; in fact this representation factors through the permutation representation of Σ_k on $\mathbb{C}[X]$. On the other hand, the vectors space $\mathbb{C}[f(X)]$ is a representation of $f(G_X) \subseteq H_{f(X)}$. Lemma 2.1 implies that the latter representation has either the same dimension as $\mathbb{C}[X]$ or it has dimension 1. This statement can be derived from Schur's lemma using the following three facts: (1) the first representation factors through Σ_k, (2) f_* intertwines the two representations, and (3) the permutation representation of Σ_k is the direct product of two irreducible representations of Σ_k, namely, the standard representation and the trivial one.

Other groups. Because total symmetry can be understood within representation theory as above, we are led to speculate on which aspects of representation theory can be brought to bear in the theory of totally symmetric sets. To begin, we know that, while the representation theory of the symmetric group is rich in and of itself, there is a broad landscape of representations of various groups.

Question 1. *To what extent, and to what end, can the theory of totally symmetric sets be generalized to arbitrary groups besides Σ_k?*

As an example of what we have in mind, we note that there is no lift of the totally symmetric set $\{(1\ i)\} \subseteq \Sigma_n$ to a totally symmetric set in B_n. On the other hand, there is a lift to a cyclically symmetric set, that is, a set with an action of $\mathbb{Z}/(n-1)$. Similarly, there are large sets of Dehn twists in $\mathrm{Mod}(S_g)$ that carry an action by the dihedral group D_{2g}. How can these sets be used in the classification of homomorphisms between braid groups and mapping class groups?

Extending the analogy to representation theory. Because of the connection between totally symmetric sets and representation theory described above, it is natural to ask which notions from representation theory have analogues for totally symmetric sets.

Question 2. *Which of the concepts in representation theory—direct sum, direct product, tensor product, etc.—have analogues in the theory of totally symmetric sets?*

Already in their work, Salter and the first author give versions of sub-representations and induced representations for totally symmetric sets [9].

Multiple totally symmetric sets. The arguments presented in this paper are carried out by analyzing the action of a homomorphism on a single totally symmetric set or collapsing set. But many groups, such as the braid group, contain totally symmetric sets that are compatible in some sense (for instance, elements either commute or braid).

Question 3. *How can multiple totally symmetric sets in a group be used to give stronger constraints on homomorphisms than can be obtained with a single totally symmetric sets?*

One step in this direction is taken in the work of Scherich–Verberne, where they study homomorphisms of virtual, welded, and classical braid groups by considering multiple totally symmetric sets at once [22].

Bounds on representations. As we have seen, the fact that Σ_n has a totally symmetric set of cardinality $n-1$ can be used, along with the work of the first author and Salter, to give a lower bound on the dimension of a non-cyclic linear representation of Σ_n. We are curious to what extent this line of reasoning holds for other groups.

As a first test case, we might consider the monster group M. The smallest nontrivial representation of M has dimension $47 \cdot 59 \cdot 71 = 196,883$. Based on the case of the symmetric group, one might hope that this is because M contains a totally symmetric set of cardinality 196,883. But we already showed in Section 2.2 that M cannot contain a totally symmetric set whose cardinality is greater than 43.

Question 4. *What are the largest totally symmetric sets of the monster group? Can they give insight into the 196883-dimensional representation?*

It is tantalizing that totally symmetric sets might give new insights into the notoriously mysterious monster group. And similarly for other groups, discovered and not.

References

[1] Javier Aramayona and Juan Souto. Homomorphisms between mapping class groups. *Geom. Topol.*, 16(4):2285–2341, 2012.
[2] V. I. Arnol'd. A remark on the branching of hyperelliptic integrals as functions of the parameters. *Funkcional. Anal. i Priložen.*, 2(3):1–3, 1968.
[3] Emil Artin. Braids and permutations. *Ann. of Math. (2)*, 48:643–649, 1947.
[4] Barbara Baumeister, Dawid Kielak, and Emilio Pierro. On the smallest non-abelian quotient of $\mathrm{Aut}(F_n)$. *Proc. Lond. Math. Soc. (3)*, 118(6):1547–1591, 2019.
[5] Robert W. Bell and Dan Margalit. Braid groups and the co-Hopfian property. *J. Algebra*, 303:275–294, 2006.
[6] A. J. Berrick, V. Gebhardt, and L. Paris. Finite index subgroups of mapping class groups. *Proc. Lond. Math. Soc. (3)*, 108(3):575–599, 2014.
[7] Noah Caplinger. Large totally symmetric sets. *New York J. Math.*, 29:931–938, 2023.
[8] Noah Caplinger and Kevin Kordek. Small Quotients of Braid Groups. *arXiv e-prints*, page arXiv:2009.10139, September 2020.
[9] Noah Caplinger and Nick Salter. Totally symmetric sets in the general linear group. *arXiv preprint arXiv:2202.10216*, 2022.
[10] Lei Chen, Kevin Kordek, and Dan Margalit. Homomorphisms between braid groups. arXiv:1910.00712, Oct 2019.
[11] Lei Chen and Aru Mukherjea. From braid groups to mapping class groups. *Math. Z.*, 303(2):Paper No. 27, 26, 2023.
[12] Alice Chudnovsky, Kevin Kordek, Qiao Li, and Caleb Partin. Finite quotients of braid groups. *Geom. Dedicata*, 207:409–416, 2020.
[13] M. Dehn. Die Gruppe der Abbildungsklassen. *Acta Math.*, 69(1):135–206, 1938. Das arithmetische Feld auf Flächen.
[14] Joan L. Dyer and Edna K. Grossman. The automorphism groups of the braid groups. *American Journal of Mathematics*, 103(6):1151–1169, 1981.
[15] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups (PMS-49)*. Princeton University Press, 2012.
[16] Dawid Kielak and Emilio Pierro. On the smallest non-trivial quotients of mapping class groups. *Groups Geom. Dyn.*, 14(2):489–512, 2020.
[17] Sudipta Kolay. Smallest noncyclic quotients of braid and mapping class groups. *Geom. Topol.*, 27(6):2479–2496, 2023.
[18] Kevin Kordek, Lily Qiao Li, and Caleb Partin. Upper bounds for totally symmetric sets. *Involve*, 14(5):853–870, 2021.
[19] Kevin Kordek and Dan Margalit. Homomorphisms of commutator subgroups of braid groups. Bulletin of the London Mathematical Society, 54 (1) 95-111,, 2019.
[20] Justin Lanier and Dan Margalit. Normal generators for mapping class groups are abundant. *Comment. Math. Helv.*, 97(1):1–59, 2022.
[21] Michael Lönne. Fundamental groups of projective discriminant complements. *Duke Mathematical Journal*, 150(2):357 – 405, 2009.
[22] Nancy Scherich and Yvon Verberne. Finite image homomorphisms of the braid group and its generalizations. *arXiv e-prints*, page arXiv:2012.01378, December 2020.

[23] Bruno P. Zimmermann. On minimal finite quotients of mapping class groups. *Rocky Mountain J. Math.*, 42(4):1411–1420, 2012.

Dan Margalit, Department of Mathematics, Vanderbilt University, 1326 Stevenson Center Ln, Nashville, TN 37240, dan.margalit@vanderbilt.edu

Noah Caplinger, Department of Mathematics, University of Chicago, 5734 University Ave., Chicago, IL 60637