Linear Algebra and its Applications 613 (2021) 1–23

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa

What fraction of an S_n-orbit can lie on a hyperplane? ☆

Jiahui Huang, David McKinnon *, Matthew Satriano

University of Waterloo, Department of Pure Mathematics, Waterloo, Ontario, Canada N2L 3G1

ARTICLE INFO

Article history:
Received 5 March 2020
Accepted 9 December 2020
Available online 14 December 2020
Submitted by J.M. Landsberg

MSC:
primary 20B30
secondary 05A05

Keywords:
Symmetric group
Hyperplane
Permutation

ABSTRACT

Consider the S_n-action on $\mathbb{R}^n$ given by permuting coordinates. This paper addresses the following problem: compute $\max_{v,H} |H \cap S_n v|$ as $H \subset \mathbb{R}^n$ ranges over all hyperplanes through the origin and $v \in \mathbb{R}^n$ ranges over all vectors with distinct coordinates that are not contained in the hyperplane $\sum x_i = 0$. We conjecture that for $n \geq 3$, the answer is $(n-1)!$ for odd n, and $n(n-2)!$ for even n. We prove that if p is the largest prime with $p \leq n$, then $\max_{v,H} |H \cap S_n v| \leq \frac{n!}{p}$. In particular, this proves the conjecture when n or $n-1$ is prime.

Contents

☆ The last two authors were partially supported by Discovery Grants from the Natural Sciences and Engineering Research Council.

* Corresponding author.
E-mail addresses: j346huang@edu.uwaterloo.ca (J. Huang), dmckinnon@uwaterloo.ca (D. McKinnon), msatrian@uwaterloo.ca (M. Satriano).

https://doi.org/10.1016/j.laa.2020.12.011

1. Introduction

Given a linear action of a Lie group $\mathcal{G}$ on a finite-dimensional vector space W, a question of central importance is to determine when the quotient $W/\mathcal{G}$ is smooth. This problem and variants of it have a long history in invariant theory with fundamental classification results having been obtained in [1,3,5–7,9,11]. In a recent preprint [4], *cf.* [8], Edidin and the third author considered the problem of giving an effective group theoretic characterization for when $W/\mathcal{G}$ is smooth, and related it to a variant of the following concrete question.

Question 1.1. *Let G be a finite group and V a finite-dimensional G-representation over a field k. Let $V = \bigoplus_i V_i$ be the decomposition into irreducible representations. What is*

$$\max_{v,H} |H \cap Gv|$$

as $H \subset V$ ranges over all hyperplanes through the origin, and $v \in V \smallsetminus \bigcup_i V_i$ ranges over all vectors whose orbit satisfies $|Gv| = |G|$?

In [4], the authors were primarily concerned with the case where $G = S_n$ and $k = \mathbb{R}$, and obtained bounds sufficient for their purposes, but the question of a general bound remained. We make the following conjecture:

Conjecture 1.2. *Let $n \geq 3$. As v ranges over all vectors in $\mathbb{R}^n$ with distinct coordinates not in the hyperplane $\sum_i x_i = 0$, and as $H \subset \mathbb{R}^n$ ranges over all hyperplanes through the origin, we have*

$$\max_{v,H} |H \cap S_n v| = \begin{cases} (n-1)!, & n \text{ is odd} \\ n(n-2)!, & n \text{ is even} \end{cases}$$

Let us motivate how these specific bounds arise.

Example 1.3. Given any $v = (c_1, \ldots, c_n)$ with distinct coordinates, consider the hyperplane H whose normal vector is $(c_n, c_n, \ldots, c_n, -\sum_{i=1}^{n-1} c_i)$. Then H contains $(c_{\sigma(1)}, \ldots, c_{\sigma(n-1)}, c_n)$ for all $\sigma \in S_{n-1}$, so $|H \cap S_n v| \geq (n-1)!$. ⋄

Example 1.4. Let $n \geq 3$. Consider the vector $v = (1, 2, \ldots, n)$ and the hyperplane H with normal vector $(-\sum_{i=2}^{n-1} i, -\sum_{i=2}^{n-1} i, 1+n, \ldots, 1+n)$. Then H contains every element of

$S_n v$ whose first two coordinates sum to $1+n$. When n is odd, there are $(n-1)!$ such elements of $S_n v$. When n is even, there are $n(n-2)!$ such elements. $\diamond$

By the above two examples, the bounds in Conjecture 1.2 are the smallest possible. The main result of this paper is:

Theorem 1.5. *Let $n \geq 2$ and let p be the largest prime with $p \leq n$. Then*

$$\max_{v,H} |H \cap S_n v| \leq \frac{n!}{p}.$$

In particular, if $n = p$ or $n = p+1$, then Conjecture 1.2 is true.

The proof of Theorem 1.5 involves tools from algebraic geometry, representation theory, combinatorics, and graph theory. The proof proceeds in several steps. Using techniques from algebraic geometry, we reduce the problem to one concerning intersections of hyperplanes with a specific curve C. We then divide the proof into two cases depending on whether or not the irreducible components C_i of C have dihedral stabilizers. We handle the non-dihedral case using techniques from combinatorics and representation theory. The dihedral case is the most involved. We construct a graph whose vertices are the irreducible components $C_i \subset H$. Assuming the existence of a hyperplane H that violates Theorem 1.5, we show the existence of a vertex C_0 whose neighbors have large degree relative to C_0. A careful analysis of the second order neighborhood of C_0 yields a contradiction.

Additionally, we prove the following two results. The first shows that the conjecture holds for generic v and the second gives an inductive statement, showing that the case of even n follows from that of odd n.

Proposition 1.6. *Let $n \geq 2$. There is a nonempty Zariski open subset U of $\mathbb{R}^n$ such that for any $v \in U \subset \mathbb{R}^n$, we have $\max_H |H \cap S_n v| = (n-1)!$ as H ranges over all hyperplanes in $\mathbb{R}^n$.*

Proposition 1.7. *Let $k \leq n$ be positive integers. If $\max_H |H \cap S_n v| \leq n!/k$, then for all $m \geq n$, we have $\max_H |H \cap S_m v| \leq m!/k$.*

In particular, if Conjecture 1.2 holds for an odd number n, then it is also holds for $n+1$.

Acknowledgments

It is our pleasure to thank Jason Bell, Ilya Bogdanov, Dan Edidin, Matt Kennedy, Heydar Radjavi, and Jerry Wang for helpful conversations. This paper is the outcome of an NSERC-USRA project; we thank NSERC for their support through the USRA program.

2. Proof of Propositions 1.6 and 1.7

We begin this section by analyzing the behaviour of $\max_H |H \cap S_n v|$ where v is a generic vector:

Proof of Proposition 1.6. By Example 1.3, we know that for every v with distinct coordinates, there exists a hyperplane H with $|H \cap S_n v| \geq (n-1)!$. So, it remains to show that for generic v we have $|H \cap S_n v| \leq (n-1)!$ for every hyperplane H. Let $v = (x_1, \ldots, x_n)$ where $x_1, \ldots x_n$ are indeterminates, and let Ω be the set of all subsets of $S_n v$ consisting of $(n-1)!+1$ elements. For every $\omega \in \Omega$, let M_ω be the matrix whose columns are the vectors in the set ω (with some ordering of the set ω whose choice will not affect the proof). Then we must show that for every $\omega \in \Omega$, the $n \times n$ minors of M_ω do not simultaneously vanish. Let $V_\omega \subseteq \mathbb{A}^n$ be the variety defined by the simultaneous vanishing of the $n \times n$ minors of M_ω. We need to show $\bigcup_{\omega \in \Omega} V_\omega \neq \mathbb{A}^n$, so it is enough to show that for each $\omega \in \Omega$, there exists some $v \in \mathbb{A}^n$ with $v \notin V_\omega$.

We prove this by induction. When $n = 2$, we must have $\omega = \{(x_1, x_2), (x_2, x_1)\}$, so $v = (1, 0)$ will suffice.

Now suppose $n > 2$. Consider the appearance of x_n in the rows of M_ω. If x_n shows up at least once in each row, let $v = (0, \ldots, 0, 1)$; then the column vectors in M_ω will contain the standard basis vectors, so the $n \times n$ minors will not vanish. If x_n does not appear in some row, then it only occurs in at most $n-1$ of the rows; hence, some row contains at least $\frac{(n-1)!+1}{n-1} > (n-2)!$ copies of x_n. By permuting rows of M_ω, we may assume there is a subset of $\omega' \subset \omega$ such that $|\omega'| = (n-2)!+1$ and every vector in ω' has x_n as its last entry.

By induction, we can specialize the variables $x_1, \ldots, x_{n-1}$ to be distinct real numbers in such a way that the column vectors in ω' span a space of dimension at least $n-1$. Choose x_n so that $\sum_i x_i \neq 0$ and $x_n \neq x_i$ for $i = 1, \ldots, n-1$. Since the column vectors of $M_{\omega'}$ have the same last coordinate, they are all contained in the hyperplane H constructed in Example 1.3. We have therefore shown that if the $n \times n$ minors of M_ω vanish, then the span of the column vectors of M_ω is H. However, since $|\omega| > (n-1)!$ and the x_i are distinct real numbers, some column vector of M_ω must have last coordinate not equal to x_n; this vector is not in H and therefore the $n \times n$ minors of M_ω do not simultaneously vanish. □

We turn next to Proposition 1.7.

Proof of Proposition 1.7. We prove the result by induction on n. We assume there exists $k \leq n-1$ such that for all $w \in \mathbb{R}^{n-1}$ with distinct coordinates not summing to 0, and all hyperplanes $H' \subset \mathbb{R}^{n-1}$, we have $|S_{n-1} w \cap H'| \leq (n-1)!/k$. Now, let $v = (v_1, \ldots, v_n) \in \mathbb{R}^n$ with distinct coordinates not summing to 0. Suppose there exists $T \subseteq S_n v$ and a hyperplane $H \subset \mathbb{R}^n$ such that

$$|T| = |T \cap H| > \frac{n!}{k}.$$

Since S_n is the disjoint union of the cosets $(in)S_{n-1}$ for $1 \leq i \leq n$, there exists i with $|T \cap (in)S_{n-1}v \cap H| > (n-1)!/k$. Relabeling the coordinates of $\mathbb{R}^n$ if necessary, we can assume $i = n$. Let

$$U = T \cap S_{n-1}v$$

and $\pi\colon \mathbb{R}^n \to \mathbb{R}^{n-1}$ be the projection map $\pi(x_1, \ldots, x_n) = (x_1, \ldots, x_{n-1})$. Note that $\pi(v)$ has distinct coordinates and that $\pi(U) \subseteq S_{n-1}\pi(v)$. Moreover, $\pi|_U\colon U \to \pi(U)$ is a bijection, so $|\pi(U)| > (n-1)!/k$.

If $v_1 + \ldots + v_{n-1} \neq 0$, then by induction, $\pi(U)$ is not contained in a hyperplane, and must therefore span $\mathbb{R}^{n-1}$. As a result, $\mathrm{Span}(U)$ is either a hyperplane or $\mathbb{R}^n$. Notice that U is contained in the hyperplane H' given by $cx_n = v_n(x_1 + \ldots + x_{n-1})$ with $c = v_1 + \ldots + v_{n-1}$, i.e. the hyperplane constructed in Example 1.3. Thus, $\mathrm{Span}(U) = H'$ and hence $H' = H \supset T$. However, $|S_n v \cap H'| = (n-1)!$ which implies $(n-1)! \geq |T| > n!/k$, a contradiction.

If $v_1 + \ldots + v_{n-1} = 0$, then $\pi(U)$ is contained in the hyperplane $x_1 + \ldots + x_{n-1} = 0$. Let $w = \pi(v) + (1, \ldots, 1) \in \mathbb{R}^{n-1}$. Then the coordinates of w are distinct and do not sum to 0, so by induction, $\pi(U + (1, \ldots, 1))$ spans $\mathbb{R}^{n-1}$. In particular, $\pi(U)$ spans the hyperplane $x_1 + \ldots + x_{n-1} = 0$. This implies that U is not contained in any affine space of dimension less than $n-2$. Notice that U is contained in the affine space A given by $x_1 + \ldots + x_{n-1} = x_n - v_n = 0$, and that A has dimension exactly $n-2$. Note further that A is not a linear space since $v_n \neq 0$, and so U is not contained in any linear space of dimension $n-2$. Thus, U must span an $(n-1)$-dimensional space and since U is contained in the hyperplane H' given by $x_1 + \ldots + x_{n-1} = 0$, we must have $\mathrm{Span}(U) = H'$, and so $H = H'$. However, we see $(in)S_{n-1}v \cap H' = \varnothing$ for all $i \neq n$. Thus, we again find $(n-1)! \geq |H' \cap T| = |T| > n!/k$, a contradiction. □

As a result of Proposition 1.7, we only need to consider the case when $n = p$ for the proof of Theorem 1.5.

3. An analysis via algebraic geometry

Let $v = (c_1, \ldots, c_n) \in \mathbb{R}^n$ with distinct coordinates. Consider the elementary symmetric functions $e_k(x_1, \ldots, x_n)$ for $1 \leq k \leq n$, and the following system of equations:

$$e_1(x_1, \ldots, x_n) = e_1(c_1, \ldots, c_n)$$

$$\vdots$$

$$e_t(x_1, \ldots, x_n) = e_t(c_1, \ldots, c_n)$$

The set $S_n v$ is precisely the solution set of this system when $t = n$.

It will be convenient to work in projective rather than affine space in what follows, so we homogenize these equations by adding an additional variable z:

$$\begin{aligned} e_1(x_1, \ldots, x_n) &= e_1(c_1, \ldots, c_n)z \\ &\vdots \\ e_t(x_1, \ldots, x_n) &= e_t(c_1, \ldots, c_n)z^t \end{aligned}$$

Definition 3.1. Let $C \subset \mathbb{P}^n_{\mathbb{C}}$ be the algebraic variety cut out by the above system of equations where we take $t = n - 1$. We refer to C as the *elementary symmetric curve associated to* v.

The elementary symmetric curve plays a fundamental role in this paper. Throughout the rest of this section, we let C be the elementary symmetric curve associated to v and let

$$C = C_1 \cup \cdots \cup C_r$$

be the decomposition of C into its irreducible components.

If there is a hyperplane H which contains many conjugates of v, then it will have a large intersection with C. If H intersects C properly (that is, in a finite set of points), then H cannot intersect C in more than $(n-1)!$ points, and therefore H cannot contain more than $(n-1)!$ conjugates of v.

Write $C = C_1 \cup \ldots \cup C_r$ as a union of irreducible curves. (Note that C cannot have any components of dimension greater than one, because its intersection with the hypersurface $e_n(x_1, \ldots, x_n) = e_n(c_1, \ldots, c_n)$ is a finite set of points, namely $S_n v$.) If H contains more than $(n-1)!$ conjugates of v, then it must contain some irreducible component C_i of C.

Lemma 3.2. *We have the following properties:*

(1) *Each C_i has dimension 1.*
(2) *If $H \subset \mathbb{R}^n$ is a hyperplane that intersects C properly, i.e. in a finite set of points, then $|H \cap C| \leq (n-1)!$.*
(3) *If $H \subset \mathbb{R}^n$ is a hyperplane satisfying $|H \cap C| > (n-1)!$, then H contains some C_i.*

Proof. Notice that the intersection of C with the hypersurface $e_n(x_1, \ldots, x_n) = e_n(c_1, \ldots, c_n)$ is a finite set of points, namely $S_n v$. Since intersecting with a hypersurface decreases dimension by at most 1, we see each C_i has dimension at most 1. On the other hand, C is defined as the intersection of $n-1$ hypersurfaces, so each C_i has dimension at least 1. This proves (1).

Statements (2) and (3) follow immediately from Bézout's Theorem since C is intersection of hypersurfaces of degrees $1, 2, \ldots, n-1$, and hence has degree $(n-1)!$. □

Lemma 3.3. *The S_n-action on $\mathbb{R}^n$ induces a transitive action on the set of irreducible components $\{C_1, \ldots, C_r\}$. Moreover, for each i, we have $\deg(C_i) = \frac{(n-1)!}{r}$ and $|C_i \cap S_n v| = \frac{n!}{r}$.*

Proof. Since $Y \cap C = S_n v$ consists of $n! = \deg(Y)\deg(C)$ distinct points, Y intersects C properly and transversely. In particular, Y cannot intersect C at any point of intersection of two irreducible components of C. So, we find

$$\deg(Y)\deg(C) = |Y \cap C| = \sum_i |Y \cap C_i| \leq \sum_i \deg(Y)\deg(C_i) = \deg(Y)\deg(C)$$

from which we see

$$|C_i \cap S_n v| = |Y \cap C_i| = \deg(Y)\deg(C_i)$$

for every i. It follows that $C_i \cap S_n v \neq \varnothing$ and so the S_n-action on $\{C_1, \ldots, C_r\}$ is transitive. As a result, each C_i has the same degree and contains the same number of elements of $S_n v$, so we must have $\deg(C_i) = \frac{(n-1)!}{r}$ and $|C_i \cap S_n v| = \frac{n!}{r}$. □

Lemma 3.4. *Let $n = p$ be prime and $\mathrm{Stab}(C_i)$ be the stabilizer of C_i under the S_p-action on the set of irreducible components of C. Then $\mathrm{Stab}(C_i)$ contains a p-cycle.*

Proof. By Lemma 3.3, S_p acts transitively on the set of irreducible components of C. By Lemma 3.3, $|\mathrm{Stab}(C_i)| = \frac{p!}{r} = p\deg(C_i)$, so p divides $|\mathrm{Stab}(C_i)|$. It follows from Cauchy's Theorem that $\mathrm{Stab}(C_i)$ contains an element π whose order is p; since $\pi \in S_p$, it is necessarily a p-cycle. □

Corollary 3.5. *Let $n = p$ be prime and $w = (\zeta, \ldots, \zeta^p)$ where $\zeta = e^{2\pi i/p}$. Then the complex linear span of any irreducible component C_0 of C contains σw for some $\sigma \in S_p$.*

Proof. Since $\mathrm{Stab}(C_0)$ contains a p-cycle π, the complex linear span of C_0 contains a subrepresentation of the permutation representation of $\langle \pi \rangle$. This subrepresentation is non-trivial since v has distinct coordinates. Thus, it contains a non-trivial complex irreducible $\langle \pi \rangle$-representation, which is necessarily spanned by σw for some $\sigma \in S_p$. □

We conclude this section with a key lemma used in the proof of Theorem 1.5. We know from Lemma 3.2 (3) that if $H \subset \mathbb{R}^n$ is a hyperplane with $|H \cap S_n v| > (n-1)!$, then H must contain an irreducible component C_i. In the proof of Theorem 1.5, we show that for each $C_i \subset H$, there are $n-1$ other irreducible components of C that are not contained in H. We then apply the following:

Lemma 3.6. *Let $H \subset \mathbb{R}^n$ be a hyperplane. Suppose that for each irreducible component C_i of C satisfying $C_i \subset H$, there are irreducible components $C_{i1}, \ldots, C_{i,n-1}$ with the following properties:*

(i) $H \cap C_{ij} \cap S_n v = \varnothing$ *and*
(ii) $C_{ik} = C_{j\ell}$ *if and only if* $i = j$ *and* $k = \ell$.

Then $|H \cap S_n v| \leq (n-1)!$.

Proof. Say H contains exactly m of the irreducible components of C. By Lemma 3.3 and Bézout's Theorem, we then have:

$$\begin{aligned} |H \cap S_n v| &\leq \sum_{C_i \subset H} \frac{n!}{r} + \sum_{\substack{C_k \not\subset H \\ C_k \cap H \cap S_n v \neq \varnothing}} \frac{(n-1)!}{r} \\ &\leq m\frac{n!}{r} + (r - mn)\frac{(n-1)!}{r} \\ &= (n-1)! \quad \square \end{aligned}$$

4. Lemmas concerning 2-cycles, 3-cycles, and 2-2-cycles

In this section, we collect several results concerning the structure of hyperplanes that simultaneously contain v and σv, where σ is a 2-cycle, a 3-cycle, or a 2-2-cycle. We also prove Theorem 1.5 for $p = 3, 5$.

Lemma 4.1. *Let* $n \geq 3$ *and* $v = (c_1, \ldots, c_n) \in \mathbb{R}^n$ *with distinct coordinates. Let* $H = (a_1, \ldots, a_n)^\perp$ *be a hyperplane containing* v. *If* $\tau = (ij)$ *is a transposition and* $\tau v \in H$, *then* $a_i = a_j$.

Let $\sigma = (ijk)$ *be a 3-cycle and* π *an* n*-cycle. If* H *contains* $\pi^m v$ *and* $\sigma\pi^m v$ *for all* m, *then* $a_i = a_j = a_k$.

Proof. By permuting coordinates, we can assume $\tau = (12)$. Then $\tau v - v = (c_2 - c_1, c_1 - c_2, 0, \ldots, 0)$ is contained in H. Since $c_2 \neq c_1$, we have $a_1 = a_2$.

For the second claim of the lemma, we first permute coordinates to assume $\sigma = (123)^{-1} = (132)$. Then for all i, we have[1]

$$d_i := \sigma\pi^{-i}v - \pi^{-i}v = (c_{\pi^i(2)} - c_{\pi^i(1)}, c_{\pi^i(3)} - c_{\pi^i(2)}, c_{\pi^i(1)} - c_{\pi^i(3)}, 0, \ldots, 0) \in H.$$

Let m be such that $c_m = \min_l c_l$. Choose i, j, k so that $\pi^i(1) = m$, $\pi^j(2) = m$, and $\pi^k(3) = m$. We claim that d_i, d_j, and d_k span a space of dimension at least 2. Indeed, the first three entries of d_i, d_j, and d_k have signs $(+, *_1, -), (-, +, *_2), (*_3, -, +)$ respectively, where $*_l$ is unknown. So, if d_i is a multiple of d_j, then $*_1$ must be negative and $*_2$ must be positive. This then shows that d_k is not a multiple of d_j.

[1] Recall that if $\epsilon \in S_n$, then the j-th coordinate of $\epsilon(v)$ is $c_{\epsilon^{-1}(j)}$.

Next, note that d_i, d_j, d_k are contained in the two dimensional space $W = \{(w_1, \ldots, w_n) : w_1 + w_2 + w_3 = w_4 = \cdots = w_n = 0\}$. So, d_i, d_j, d_k span W and so $W \subset H$. In particular, $(1, -1, 0, \ldots, 0), (1, 0, -1, \ldots, 0) \in H$ which implies $a_1 = a_2 = a_3$. □

As an application of Lemma 4.1, we prove Theorem 1.5 for hyperplanes whose normal vector has a distinct entry.

Corollary 4.2. *Let $p \geq 3$ be a prime. Let $H = (a_1, \ldots, a_p)^\perp$ be a hyperplane and assume there exists i such that for all $j \neq i$ we have $a_j \neq a_i$. Then $|H \cap S_p v| \leq (p-1)!$ for all $v \in \mathbb{R}^p$ with distinct coordinates.*

Proof. After permuting coordinates, we may assume $i = 1$. We will prove the Corollary by applying Lemma 3.6. Let C be the elementary symmetric curve associated to v and let $C_1, \ldots, C_r$ be its irreducible components. For each C_i in H, let $C_{ij} := (1j)C_i$ where $j \neq 1$. Since $\sigma(C_i) \cap S_p v = \sigma(C_i \cap S_p v)$ for all $\sigma \in S_p$, the first part of Lemma 4.1 shows $H \cap C_{ij} \cap S_p v = \varnothing$. If C_i and C_k are contained in H, and if $C_{ij} = C_{kl}$, then $(1l)(1j)C_i = C_k$. If $j \neq l$, then $(1jl)C_i = C_k$; this is not possible by the second claim in Lemma 4.1, where we take $\pi \in \mathrm{Stab}(C_i)$ to be the p-cycle constructed in Lemma 3.4. If $j = l$, then $C_i = C_k$ and so $i = k$. The result follows by Lemma 3.6. □

The rest of this section is concerned with the case where H contains v and σv for some 2-2-cycle σ. We start with the following preliminary result and as an application, prove Theorem 1.5 for special classes of hyperplanes.

Lemma 4.3. *Let $p \geq 5$ be prime, $v \in \mathbb{R}^p$ have distinct coordinates, and C be the elementary symmetric curve associated to v with some irreducible component C_0. Suppose $(ij)(kl)$ is a 2-2-cycle and $H = (a_1, \ldots, a_p)^\perp$ is a hyperplane containing C_0, and $(ij)(kl)C_0$. If $a_i = a_k = a$ and $a_j = a_l = b$, then $a = b$.*

Proof. By permuting coordinates, we can assume $v \in C_0$. From Corollary 3.5, we know $\mathrm{Span}_{\mathbb{C}} C_0$ contains σw for some $\sigma \in S_p$, where $w = (\zeta, \ldots, \zeta^p)$ and $\zeta = e^{2\pi i/p}$. Thus H contains both σw and $(ij)(kl)\sigma w$. Subtracting we find $w - (ij)(kl)\sigma w \in H = (a_1, \ldots, a_p)^\perp$. Set $\alpha' = \sigma^{-1}(\alpha)$, we have

$$(a-b)(\zeta^{i'} - \zeta^{j'} + \zeta^{k'} - \zeta^{l'}) = a(\zeta^{i'} - \zeta^{j'}) + b(\zeta^{j'} - \zeta^{i'}) + a(\zeta^{k'} - \zeta^{l'}) + b(\zeta^{l'} - \zeta^{k'}) = 0.$$

Since $p \geq 5$ and i, j, k, l are distinct, we must have $a = b$. □

Corollary 4.4. *Let $p \geq 5$ be a prime and $H = (a_1, \ldots, a_p)^\perp$ be a hyperplane. Suppose $i_1, \ldots, i_m$ are distinct, $j_1, \ldots, j_n$ are distinct, $a_{i_1} = \cdots = a_{i_m}$, $a_{j_1} = \cdots = a_{j_n}$, and $a_{i_1} \neq a_{j_1}$. If $nm \geq p - 1$, then $|H \cap S_p v| \leq (p-1)!$ for all $v \in \mathbb{R}^p$ with distinct coordinates.*

Proof. Let C be the elementary symmetric curve associated to v. We will prove the corollary by applying Lemma 3.6. For each irreducible component C_0 of C satisfying $C_0 \subset H$, consider the $nm \geq p-1$ curves

$$\{(i_k, j_l)C_0 : 1 \leq k \leq m, 1 \leq l \leq n\}.$$

By Lemma 4.1, we see $H \cap (i_k, j_l)C_0 \cap S_p v = \varnothing$.

Next, if $C_0 \subset H$ and $(i_k, j_l)C_0 = (i_{k'}, j_{l'})C_0$, then H contains both C_0 and $(i_{k'}, j_{l'})(i_k, j_l)C_0$. Similarly, if H contains distinct irreducible component C_0' and C_0, and if $(i_k, j_l)C_0 = (i_{k'}, j_{l'})C_0'$, then H contains both C_0 and $(i_{k'}, j_{l'})(i_k, j_l)C_0$. By Lemmas 4.1 and 4.3, this is not possible as $(i_{k'}, j_{l'})(i_k, j_l)$ is either a 2-2-cycle or a 3-cycle; in the case of a 3-cycle, we apply Lemma 4.1 by taking $\pi \in \mathrm{Stab}(C_0)$ to be the p-cycle constructed in Lemma 3.4. □

As a further application, we prove Theorem 1.5 for $p = 3, 5$.

Corollary 4.5. *Let $p \in \{3, 5\}$ and $H = (a_1, \ldots, a_p)^\perp$ be a hyperplane of $\mathbb{R}^p$. If $v \in \mathbb{R}^p$ has distinct coordinates not summing to 0, then $|H \cap S_p v| \leq (p-1)!$.*

Proof. Since the coordinates of v do not sum to 0, we know $H \neq (1, \ldots, 1)^\perp$. For $p = 3$, our desired result then follows directly from Corollary 4.2. When $p = 5$, Corollary 4.2 reduces us to the case $H = (a, a, b, b, b)^\perp$ for some distinct $a, b \in \mathbb{R}$. Our result then follows from Corollary 4.4. □

We end this section with some more refined results concerning the structure of hyperplanes that contain v and σv with σ a 2-2-cycle.

Lemma 4.6. *Let $p \geq 5$ be prime, $v \in \mathbb{R}^p$ have distinct coordinates, and $H = (a_1, \ldots, a_p)^\perp$ be a hyperplane. Suppose $\sigma = (ij)(kl)$ is a 2-2-cycle and π is an p-cycle. Let $G \subset S_p$ be a subgroup that contains π and assume $\dim \mathrm{Span}(Gv) > 3$. If H contains Gv and σGv, then $a_i = a_j$ and $a_k = a_l$.*

Proof. By permuting coordinates we can assume $\pi = (12 \ldots p)$. Note that the subspace $\mathrm{Span}(Gv) \subset \mathbb{R}^p$ is invariant under the action of $\langle \pi \rangle \simeq \mathbb{Z}/p$. Since $\dim \mathrm{Span}(Gv) > 3$, when viewed as a complex $\mathbb{Z}/p$-representation, it contains $w_1 = (1, \zeta, \ldots, \zeta^{p-1})$, $w_2 = (1, \zeta^{-1}, \ldots, \zeta^{-(p-1)})$, and $w_3 = (1, \zeta^m, \ldots, \zeta^{m(p-1)})$ for some primitive p-th root of unity ζ and some $m \neq 0, \pm 1 \bmod p$.

For $d \in \{1, 2, 3\}$, let t_d be such that the i-th coordinate of $u_d := \zeta^{t_d} w_d$ is 1, where i is as in the statement of the lemma. Since $u_d \in \langle \pi \rangle w_d \subset \mathrm{Span}(Gv)$, we see $\sigma u_d \in \mathrm{Span}(\sigma Gv)$, so in particular,

$$\sigma u_d - u_d \in H.$$

Let x be the j-th coordinate of u_1. Then the k-th and l-th coordinates of u_1 are, respectively, x^a and x^b for distinct $a, b \in \{2, \dots, p-1\}$. The j-th, k-th, and l-th coordinates of u_2 are then x^{-1}, x^{-a}, and x^{-b}, respectively. The j-th, k-th, and l-th coordinates of u_3 are x^m, x^{ma}, and x^{mb}, respectively. Letting $\alpha = 1 - x$, $\beta = x^a - x^b$, $\alpha' = 1 - x^{-1}$, and $\beta' = x^{-a} - x^{-b}$, we find

$$(\dots, \alpha, \dots, -\alpha, \dots, \beta, \dots, -\beta, \dots) = \sigma u_1 - u_1 \in H$$

and

$$(\dots, \alpha', \dots, -\alpha', \dots, \beta', \dots, -\beta', \dots) = \sigma u_2 - u_2 \in H$$

where the omitted entries are 0, and the four non-zero entries are in the j, i, l, k-th positions respectively.

If $(\alpha, \beta), (\alpha', \beta')$ are linearly independent, then have $(\dots, 1, \dots, -1, \dots, 0, \dots, 0, \dots) \in H$ which implies $a_i = a_j$, and consequently $a_k = a_l$.

Next suppose (α, β) and (α', β') are linearly dependent. Then

$$\frac{x^a - x^b}{1 - x} = \frac{\beta}{\alpha} = \frac{\beta'}{\alpha'} = \frac{x^{-a} - x^{-b}}{1 - x^{-1}} = \frac{x^{1-a} - x^{1-b}}{x - 1}$$

and hence

$$x^a - x^b = x^{1-b} - x^{1-a}. \tag{4.7}$$

Since $b \neq a$, this is a contradiction by the linear independence of roots of unity over $\mathbb{Q}$, unless $a + b = 1 \bmod p$.

So, we may suppose $a + b = 1 \bmod p$. Consider

$$(\dots, \alpha'', \dots, -\alpha'', \dots, \beta'', \dots, -\beta'', \dots) = \sigma u_3 - u_3 \in H$$

where $\alpha'' = 1 - x^m$ and $\beta'' = x^{ma} - x^{mb}$. If $(\alpha, \beta), (\alpha'', \beta'')$ are linearly independent, we again arrive at our desired conclusion that $a_i = a_j$ and $a_k = a_l$, so we may assume $(\alpha, \beta), (\alpha'', \beta'')$ are linearly dependent. Then

$$\frac{1 - x}{x^a - x^{1-a}} = \frac{\alpha}{\beta} = \frac{\alpha''}{\beta''} = \frac{1 - x^m}{x^{ma} - x^{m(1-a)}}$$

and so

$$x^{ma} - x^{m(1-a)} - x^{ma+1} + x^{m-ma+1} - x^a + x^{1-a} + x^{m+a} - x^{1-a+m} = 0. \tag{4.8}$$

Let $f(x)$ be the polynomial (4.8), where the exponents are taken to be numbers between 0 and p by reducing mod p, and we now view x as an indeterminate. Since $f(x)$ has

integer coefficients, has degree less than p, and has a primitive p-th root of unity as a root, it is a constant multiple of the p-th cyclotomic polynomial. Note the term x^{ma} cannot be cancelled by any other term since $m \neq 0, \pm 1 \bmod p$, so $f(x)$ is a non-zero polynomial with at most 8 terms. In particular, it is not a multiple of the cyclotomic polynomial for $p \geq 11$.

When $p = 5, 7$, since $f(x)$ is a non-zero constant multiple of the p-th cyclotomic polynomial, some of the terms in $f(x)$ must cancel to yield exactly p terms all with the same non-zero coefficient. This is impossible, however, as $f(x)$ has 4 terms with coefficient equal to 1 and 4 terms with coefficient equal to -1. □

Lemma 4.9. *Let $p \geq 7$ be a prime, $\zeta = e^{2\pi i/p}$, $w = (\zeta, \zeta^2, \dots, \zeta^p)$, and $\sigma \in S_p$. If $a_1, \dots, a_p \in \mathbb{R}$ and $H = (a_1, \dots, a_p)^\perp$ is a hyperplane that contains both σw and $(ij)(kl)\sigma w$ with i, j, k, l distinct, then either*

$$\sigma^{-1}(i) + \sigma^{-1}(j) = \sigma^{-1}(k) + \sigma^{-1}(l) \mod p$$

or $a_i = a_j$ and $a_k = a_l$.

Proof. For ease of notation, we let $\alpha' = \sigma^{-1}(\alpha)$ for $\alpha = 1, \dots, p$. First note that H contains the element

$$\sigma w - (ij)(kl)\sigma w = (\dots, \zeta^{j'} - \zeta^{i'}, \dots, \zeta^{i'} - \zeta^{j'}, \dots, \zeta^{l'} - \zeta^{k'}, \dots, \zeta^{k'} - \zeta^{l'}, \dots)$$

where the omitted entries are 0, and the four non-zero entries are in the j, i, l, k-th positions respectively. Since $(a_1, \dots, a_p)$ is a real vector, H also contains the complex conjugate vector

$$(\dots, \zeta^{-j'} - \zeta^{-i'}, \dots, \zeta^{-i'} - \zeta^{-j'}, \dots, \zeta^{-l'} - \zeta^{-k'}, \dots, \zeta^{-k'} - \zeta^{-l'}, \dots).$$

Now, if the two vectors $(\zeta^{j'} - \zeta^{i'}, \zeta^{l'} - \zeta^{k'}), (\zeta^{-j'} - \zeta^{-i'}, \zeta^{-l'} - \zeta^{-k'})$ are linearly independent, then $(\dots, 1, \dots, -1, \dots, 0, \dots, 0, \dots) \in H$, which means $a_i = a_j$, from which it follows that $a_k = a_l$. Otherwise,

$$\frac{\zeta^{j'} - \zeta^{i'}}{\zeta^{l'} - \zeta^{k'}} = \frac{\zeta^{-j'} - \zeta^{-i'}}{\zeta^{-l'} - \zeta^{-k'}}$$

and hence

$$-\zeta^{j'-k'} + \zeta^{j'-l'} + \zeta^{-k'+i'} - \zeta^{-l'+i'} + \zeta^{k'-j'} - \zeta^{l'-j'} - \zeta^{k'-i'} + \zeta^{l'-i'} = 0.$$

Consider the polynomial

$$f(z) := -z^{j'-k'} + z^{j'-l'} + z^{-k'+i'} - z^{-l'+i'} + z^{k'-j'} - z^{l'-j'} - z^{k'-i'} + z^{l'-i'}$$

where we view the exponents as numbers between 0 and p by reducing mod p. Since $\deg f(z) < p$ and since $f(z)$ is a polynomial with integer coefficients satisfying $f(\zeta) = 0$, it must be the case that $f(z)$ is a constant multiple of the p-th cyclotomic polynomial. For $p \geq 11$, since $f(z)$ has at most 8 terms, this forces $f(z) = 0$; in particular two terms of $f(z)$ must cancel. Similarly, for $p = 7$, we know $f(z)$ has at most 6 terms, and so two terms in the above expression must cancel. In all cases, when $p \geq 7$, we must have $z^{j'-k'} = z^{l'-i'}$ since i, j, k, l are distinct mod p. So, $j' - k' = l' - i'$ mod p, and hence $i' + j' = k' + l'$ mod p. □

5. Theorem 1.5 in the non-dihedral case

Given the algebro-geometric results in Section 3, the proof of Theorem 1.5 is divided into two cases, depending on whether or not the stabilizer of C_1 is the dihedral group D_{2p} with $2p$ elements. In this section, we prove the following result, which handles the non-dihedral case:

Theorem 5.1. *Let p be a prime and $v \in \mathbb{R}^p$ have distinct coordinates that do not sum to 0. If $\mathrm{Stab}(C_1) \not\cong D_{2p}$, then $\max_H |H \cap S_n v| = (p-1)!$.*

Given a subgroup G' of G, we let $N_G(G')$ denote the normalizer of G' in G. We recall the following two theorems, which we use to obtain a structure result for $\mathrm{Stab}(C_1)$.

Theorem 5.2 *(Burnside, [2]). For p prime, a transitive subgroup of S_p is either doubly transitive or contains a normal Sylow p-subgroup.*

Theorem 5.3 *([10, Exercise 2.6]). If G is a doubly transitive subgroup of S_n, then the permutation representation $\mathbb{R}^n$ is the direct sum of two irreducible G-representations: the trivial representation and the standard representation of S_n.*

Proposition 5.4. *Let p be prime, $v \in \mathbb{R}^p$ have distinct coordinates that do not sum to 0. Suppose $v \in C_0 \subset H$ where H is a hyperplane of $\mathbb{R}^p$ and C_0 is an irreducible component of the elementary symmetric curve associated to v. Then*

$$\mathrm{Stab}(C_0) = \langle \pi, \sigma \rangle \subset N_{S_p}(\langle \pi \rangle)$$

where π is a p-cycle, and σ is a power of some $(p-1)$-cycle.

Proof. To ease notation, let $G = \mathrm{Stab}(C_0)$. By Lemma 3.4, G contains a p-cycle π and hence is a transitive subgroup of S_p. Note that $\langle \pi \rangle$ is a Sylow p-subgroup of G.

Our first goal is to show $G \subset N_{S_p}(\langle \pi \rangle)$. If this is not the case, then $\langle \pi \rangle$ is not normal in G, and so G is doubly transitive by Theorem 5.2. Notice that

$$(1, \ldots, 1) = \frac{1}{\sum_i v_i} \sum_{i=0}^{p-1} \pi^i v \in \mathrm{Span}(\langle \pi \rangle v) \subset \mathrm{Span}(Gv)$$

so $\mathrm{Span}(Gv)$ contains the trivial representation. Since $v \in \mathrm{Span}(Gv)$ and v has distinct coordinates, we see $\mathrm{Span}(Gv)$ cannot equal the trivial representation. It follows then from Theorem 5.3 that $\mathrm{Span}(Gv) = \mathbb{R}^p$. On the other hand,

$$\mathbb{R}^p = \mathrm{Span}(Gv) \subset \mathrm{Span}(GC_0) = \mathrm{Span}(C_0) \subset H$$

which contradicts the fact that H is a hyperplane. We have therefore proven our claim that $G \subset N_{S_p}(\langle\pi\rangle)$.

Next, one readily checks that $N_{S_p}(\langle\pi\rangle) = \langle\pi, \tau\rangle$ where τ is a $(p-1)$-cycle such that $\tau^{-1}\pi\tau = \pi^k$ with k a generator for $(\mathbb{Z}/p)^*$. In particular, $N_{S_p}(\langle\pi\rangle) \simeq \langle\pi\rangle \rtimes (\mathbb{Z}/p)^*$ where $(\mathbb{Z}/p)^*$ acts on $\langle\pi\rangle \simeq \mathbb{Z}/p$ in the natural way. Since G is a subgroup of $N_{S_p}(\langle\pi\rangle)$ that contains π, we see $G = \langle\pi\rangle \rtimes Q$, where Q is a subgroup of $(\mathbb{Z}/p)^*$. It follows that $G = \langle\pi, \sigma\rangle$ where $\sigma = \tau^i$ for some i. □

Given the above structure result for $\mathrm{Stab}(C_0)$, we next understand how $\mathbb{C}^p$ decomposes as a $\mathrm{Stab}(C_0)$-representation.

Lemma 5.5. *Let p be prime, $\pi \in S_p$ be a p-cycle, and G be a subgroup of $N_{S_p}(\langle\pi\rangle)$. Then every non-trivial complex irreducible G-subrepresentation of the permutation representation $\mathbb{C}^p$ has dimension $|G/\langle\pi\rangle|$.*

Proof. As in the proof of Proposition 5.4, we know $G = \langle\pi, \sigma\rangle$ where $\sigma^{-1}\pi\sigma = \pi^k$. Fix a primitive p-th root of unity ζ. Decomposing $\mathbb{C}^p$ into irreducible subrepresentations of $\langle\pi\rangle \simeq \mathbb{Z}/p$, we have $\mathbb{C}^p = \bigoplus_{i\in\mathbb{Z}/p} V_i$ where $V_i = \mathrm{Span}(\omega_i)$ and $\pi\omega_i = \zeta^i\omega_i$. We find $\pi\sigma\omega_i = \sigma\pi^k\omega_i = \zeta^{ik}\sigma\omega_i$ and hence $\sigma V_i = V_{ik}$. So, the non-trivial irreducible G-subrepresentations of $\mathbb{C}^p$ are given by $\mathrm{Span}(GV_i) = \bigoplus_j V_{ik^j}$, where the sum runs over $0 \le j < \mathrm{ord}(k)$ and $\mathrm{ord}(k)$ is the order of k in $(\mathbb{Z}/p)^*$, i.e. the order of $G/\langle\pi\rangle$. □

Proof of Theorem 5.1. By Corollary 4.5, we may assume $p \ge 7$. Let C be elementary symmetric curve associated to v and let $C_1, \ldots, C_r$ be its irreducible components. By Lemma 3.2, we may assume that H contains an irreducible component of C; without loss of generality, $v \in C_1 \subset H$. Letting $G = \mathrm{Stab}(C_1)$, we know from Proposition 5.4 that $G = \langle\pi, \sigma\rangle$ where π is a p-cycle and $\sigma^{-1}\pi\sigma = \pi^k$. Since $G \ncong D_{2p}$, the order of σ, $\mathrm{ord}(\sigma)$, cannot equal 2.

Next, note that $\mathrm{Span}(\langle\pi\rangle v)$ contains the trivial representation, as $(1, \ldots, 1) = \frac{1}{\sum_i v_i}\sum_{i=0}^{p-1} \pi^i v$. On the other hand, $\mathrm{Span}(\langle\pi\rangle v)$ cannot equal the trivial representation since it contains v, which has distinct coordinates. So, $\mathrm{Span}(Gv)$ contains both the trivial and a non-trivial G-subrepresentation of $\mathbb{R}^p$.

If $\sigma = 1$, then $G = \langle\pi\rangle$ and since $r = p!/|G| = (p-1)!$, we see from Lemma 3.3 that $\deg(C_1) = 1$, i.e. the curve C_1 is a line. Since the non-trivial irreducible $\langle\pi\rangle$-subrepresentations of $\mathbb{R}^p$ are all 2-dimensional, it follows that $\dim \mathrm{Span}(\langle\pi\rangle v) \ge 3$. In particular, the line C_1 cannot contain $\langle\pi\rangle v$.

So, we may assume $\mathrm{ord}(\sigma) \geq 3$. Let $H = (a_1, \ldots, a_p)^\perp$. Again by Lemma 4.1, $(ij)H = H$ if and only if $a_i = a_j$. Since the coordinates of v do not sum to 0, not all of the a_i are equal. From this, it is straightforward to check that there are at least $p-1$ distinct transpositions $\tau_1, \ldots, \tau_{p-1}$ such that $\tau_k H \neq H$. For each C_i contained in H, let $C_{ij} = \tau_j C_i$. We will check that the conditions in Lemma 3.6 are satisfied, and conclude $|H \cap S_p v| \leq (p-1)!$.

It follows directly from Lemma 4.1 that $H \cap C_{ij} \cap S_p v = \varnothing$. Next, suppose H contains C_i and C_j, and that we have $C_{ik} = C_{jl}$. If $k \neq l$, then H contains both C_j and $C_i = \tau_k \tau_l C_j$. Now, $\tau_k \tau_l$ cannot be a 3-cycle as this would contradict Lemma 4.1. So, $\tau_k \tau_l$ must be a 2-2-cycle, in which case we note that $\mathrm{Span}(Gv)$ contains the trivial and a non-trivial irreducible G-subrepresentation of $\mathbb{R}^p$. So $\dim \mathrm{Span}(Gv) \geq \mathrm{ord}(\sigma) + 1 > 3$ by Lemma 5.5, which gives a contradiction by Lemma 4.6. It follows that $k = l$, so $C_i = C_j$ and $i = j$. □

6. Completing the proof of Theorem 1.5

To finish the proof of Theorem 1.5, we must now handle the case not covered by Theorem 5.1 and Corollary 4.5, namely when $p \geq 7$ and the irreducible components of C have stabilizers isomorphic to D_{2p}, the dihedral group with $2p$ elements. Note that by Lemma 3.3, this implies the irreducible components of C have degree 2.

Let $H = (b_1, \ldots, b_p)^\perp$ be any hyperplane. We fix the following notation. Let $\{1, \ldots, p\} = \lambda_1 \cup \cdots \cup \lambda_K$ be the partition defined by the domains on which the function $j \mapsto b_j$ are constant. In other words, we have distinct $a_1, \ldots, a_K \in \mathbb{R}$ such that $b_j = a_J$ if and only if $j \in \lambda_J$. Let

$$m := \min_J |\lambda_J| \quad \text{and} \quad |\lambda_M| = m$$

for some fixed choice of M. By Corollaries 4.2 and 4.4, we can assume

$$2 \leq m < \sqrt{p-1}.$$

If $a_M \neq 0$, we may scale to assume $a_M = 1$.

We prove Theorem 1.5 by studying properties of a graph Γ which we now define. Throughout the rest of Section 6, we fix two distinct elements $i, k \in \lambda_M$ and let

$$T := \{(ij) : j \notin \lambda_M\} \cup \{(kj) : j \notin \lambda_M\}.$$

Let $\Gamma := \Gamma_{ik}$ be the graph whose vertices and edges are defined as follows. Let C be the elementary symmetric curve associated to v. The vertices of Γ are the irreducible components C_a of C for which $C_a \subset H$. Let $|\Gamma|$ denote the number of vertices in Γ. If $C_1, C_2 \in \Gamma$ are two vertices, we write $C_1 \sim C_2$ when C_1 and C_2 are connected by an edge. The edges of Γ are defined by

$$C_1 \sim C_2 \quad \iff \quad (ij)(kl)C_1 = C_2 \text{ for distinct } j, l \notin \lambda_M;$$

here i, k are the elements that we have fixed above.

We observe that for each $C_0 \in \Gamma$, if $\sigma, \tau \in T$ and $\sigma C_0 = \tau C_0$, then $\sigma^{-1}\tau \in \mathrm{Stab}(C_0) \simeq D_{2p}$. Since $\sigma^{-1}\tau$ is a product of two transpositions, it is not a p-cycle nor is it a product of $(p-1)/2$ disjoint 2-cycles, so $\sigma^{-1}\tau = 1$. Thus, we find

$$|\{\sigma C_0 : \sigma \in T\}| = |T| = 2(p - m).$$

For the rest of the section, we let

$$w = (\zeta, \ldots, \zeta^p),$$

where $\zeta = e^{2\pi i/p}$.

Lemma 6.1. *If $C_1, C_2 \in \Gamma$ and $C_1 \sim C_2$, then $(ij)(kl)C_1 = C_2$ for a unique pair (j, l).*

Proof. Suppose $(ij)(kl)C_1 = (ij')(kl')C_1 = C_2 \subset H$, where $j \neq j'$ or $l \neq l'$. Then

$$(ij)(kl)(ij')(kl') \in \mathrm{Stab}(C_1) \simeq D_{2p}.$$

By Corollary 3.5, $\mathrm{Span}_{\mathbb{C}} C_1$ contains σw for some $\sigma \in S_p$. Then H contains σw, $(ij)(kl)\sigma w$, and $(ij')(kl')\sigma w$, so by Lemma 4.9, the following two equations hold:

$$\begin{aligned} \sigma^{-1}(i) + \sigma^{-1}(j) &= \sigma^{-1}(k) + \sigma^{-1}(l) \mod p \\ \sigma^{-1}(i) + \sigma^{-1}(j') &= \sigma^{-1}(k) + \sigma^{-1}(l') \mod p. \end{aligned}$$

Subtracting the equations, we find

$$\sigma^{-1}(j') - \sigma^{-1}(j) = \sigma^{-1}(l') - \sigma^{-1}(l) \mod p.$$

This implies that $j = j'$ if and only if $l = l'$, and hence $j \neq j'$. In addition $j' = l'$ implies $j = l \mod p$ and $i = k \mod p$, which is also not true. Recall that neither i nor j' is equal to k or l mod p.

Putting these observations together we see that $(ij)(kl)(ij')(kl')$ is not the identity, as it sends j' to j. We see that $(ij)(kl)(ij')(kl')$ also does not permute $p \geq 7$ elements. So as an element of D_{2p}, it must be a product of $(p-1)/2$ disjoint transpositions, which is only possible when $p = 7$ and $(p-1)/2 = 3$. However, $(ij)(kl)(ij')(kl')$ is an even permutation so this is also not possible when $p = 7$. We have thus established our claim. □

Finally, we let

$$\mathcal{T} = \{\sigma C_0 : \sigma \in T, C_0 \in \Gamma\}.$$

Note that $\Gamma \cap \mathcal{T} = \varnothing$ by Lemma 4.1.

We divide the proof of Theorem 1.5 into two cases depending on the size of $\mathcal{T}$. The following result easily dispenses with the case where $\mathcal{T}$ is big.

Lemma 6.2. *If $|\mathcal{T}| \geq (p-1)|\Gamma|$, then $|H \cap S_p v| \leq (p-1)!$.*

Proof. If D is an irreducible component of C and $D \subset H$, then $|H \cap D \cap S_p v| = \frac{p!}{r}$. If $D \not\subset H$, then by Bézout's Theorem and Lemma 3.3, we have $|H \cap D \cap S_p v| \leq \frac{(p-1)!}{r}$. Furthermore, if $D \in \mathcal{T}$, then by Lemma 4.1, we have $H \cap D \cap S_p v = \varnothing$. Putting these bounds together, and making use of the fact that $\Gamma \cap \mathcal{T} = \varnothing$, we find

$$\begin{aligned} |H \cap S_p v| &\leq \sum_{D \subset H} \frac{p!}{r} + \sum_{\substack{D \not\subset H \\ D \notin \mathcal{T}}} \frac{(p-1)!}{r} + \sum_{\substack{D \not\subset H \\ D \in \mathcal{T}}} 0 \\ &\leq \frac{p!}{r}|\Gamma| + (r - |\Gamma| - (p-1)|\Gamma|)\frac{(p-1)!}{r} \\ &= (p-1)!. \quad \square \end{aligned}$$

The goal of the rest of Section 6 is to prove that $|\mathcal{T}| \geq (p-1)|\Gamma|$, and hence Theorem 1.5 holds in light of Lemma 6.2. To this end, we assume throughout the rest of Section 6 that

$$|\mathcal{T}| < (p-1)|\Gamma|$$

and aim to arrive at a contradiction.

Lemma 6.3. *If $D \in \Gamma$ is a vertex, let $d(D)$ be its degree in Γ. Then there exists $C_0 \in \Gamma$ such that*

$$\sum_{D \sim C_0} d(D) \geq \kappa + 2d(C_0),$$

where $\kappa = (p-2m)^2 - 1$.

Proof. We begin by counting the number of elements in $\mathcal{T}$. Note that if $C_1, C_2 \in \Gamma$ are distinct, then we cannot have $(ij)C_1 = (il)C_2$ since this would imply $j \neq l$ and $C_2 = (ijl)C_1$, contradicting Lemma 4.1. Next notice that if $(ij)C_1 = (kl)C_2$, then $(kl)(ij)C_1 = C_2$ and so Lemma 4.1 shows we must have $j \neq l$, i.e. $C_1 \sim C_2$. Conversely, if $C_1 \sim C_2$, then we have already established that there is a unique pair (j,l) for which $(ij)(kl)C_1 = C_2$; it follows that $(ij)C_1 = (kl)C_2$ and $(kl)C_1 = (ij)C_2$. Putting these observations together, we see that if e is the number of edges of Γ, then

$$|\mathcal{T}| = |T||\Gamma| - 2e = 2(p-m)|\Gamma| - 2e.$$

Since $|\mathcal{T}| < (p-1)|\Gamma|$, we have

$$e > \frac{(p-2m+1)|\Gamma|}{2}.$$

Suppose that $\sum_{D\sim C} d(D) < \kappa + 2d(C)$ for all vertices $C \in \Gamma$. Then we see

$$\sum_{C\in\Gamma}\sum_{D\sim C} d(D) < \sum_{C\in\Gamma}(\kappa + 2d(C)) = \kappa|\Gamma| + 2\sum_{C\in\Gamma} d(C) = \kappa|\Gamma| + 4e.$$

One readily checks that

$$\sum_{C\in\Gamma}\sum_{D\sim C} d(D) = \sum_{C\in\Gamma} d(C)^2.$$

By the Cauchy–Schwartz inequality, we see

$$\sum_{C\in\Gamma} d(C)^2 \geq |\Gamma|\left(\frac{1}{|\Gamma|}\sum_{C\in\Gamma} d(C)\right)^2 = \frac{4e^2}{|\Gamma|}.$$

Thus, $\kappa|\Gamma| + 4e > 4e^2/|\Gamma|$ and so

$$\begin{aligned}\kappa|\Gamma|^2 > 4e(e-|\Gamma|) &> 4|\Gamma|\frac{p-2m+1}{2}\left(|\Gamma|\frac{(p-2m+1)}{2} - |\Gamma|\right)\\ &= |\Gamma|^2(p-2m+1)(p-2m-1) = \kappa|\Gamma|^2,\end{aligned}$$

a contradiction. □

Throughout the rest of this section, we fix C_0, κ, and $d := d(C_0)$ as in Lemma 6.3. We prove

Proposition 6.4. $\sum_{D\sim C_0} d(D) \leq p - m + 2d$.

Assuming Proposition 6.4 for the moment, let us complete the proof of Theorem 1.5. By Lemma 6.3 and Proposition 6.4, we have $\kappa \leq \sum_{D\sim C_0} d(D) - 2d \leq p - m$. Now, if $p = 7$, then $2 \leq m < \sqrt{p-1}$ implies $m = 2$ and hence $\kappa = (7-4)^2 - 1 = 8 > 7 - 2 = p - m$, a contradiction. If $p > 7$, then

$$\kappa = (p-2m)^2 - 1 > (p - 2\sqrt{p})^2 - 1 > p - 2 \geq p - m,$$

again a contradiction.

The rest of Section 6 is devoted to the proof of Proposition 6.4. The proof is based on an analysis of the edges in the second-order neighborhood of C_0. By definition of C_0, it has d neighbors $C_1, \ldots, C_d$ such that the sum of the degrees of these neighbors is at least $\kappa + 2d$. We have $j_1, \ldots, j_d, l_1, \ldots, l_d \notin \lambda_M$ with $j_a \neq l_a$ such that

$$C_a := (ij_a)(kl_a)C_0.$$

For notational convenience, let $j_0 = k$ and $l_0 = i$ so that $C_0 = (ij_0)(kl_0)C_0$.

Lemma 6.5. *We have the following:*

(1) *For* $0 \le a \le d$,

$$\sigma^{-1}(i) + \sigma^{-1}(j_a) = \sigma^{-1}(k) + \sigma^{-1}(l_a) \mod p.$$

In particular, j_a *determines* l_a*, and* l_a *determines* j_a*.*

(2) $j_0, \dots, j_d$ *are distinct and* $l_0, \dots, l_d$ *are distinct.*

Proof. From Corollary 3.5, there exists $\sigma \in S_p$ such that $\sigma w = (\zeta^{\sigma^{-1}(1)}, \dots, \zeta^{\sigma^{-1}(p)}) \in \operatorname{Span} C_0$. It follows that the linear span of C_a contains $(ij_a)(kl_a)\sigma w$. Since C_0 and C_a are contained in H, Lemma 4.9 then tells us that $\sigma^{-1}(i) + \sigma^{-1}(j_a) = \sigma^{-1}(k) + \sigma^{-1}(l_a) \mod p$, proving (1).

To prove (2), first let $a, b \in \{1, \dots, d\}$ and assume $j_a = j_b$. From (1), we know $l_a = l_b$, and so $C_a = (ij_a)(kl_a)C_0 = (ij_b)(kl_b)C_0 = C_b$, so $a = b$. As for j_0, recall that $j_1, \dots, j_d \notin \lambda_M$ and $j_0 = k \in \lambda_M$, so they are necessarily distinct. □

We next define a set of pairs

$$R \subset \{(j, D) : j \notin \lambda_M, D \in \{C_0, \dots, C_d\}\}$$

that will be used to parameterize a subset of edges emanating from the C_a. Let $1 \le a \le d$. Then we define $(j_a, C_0) \in R$. We also define $(j, C_a) \in R$ if there exists l for which $C_a \sim (il)(kj)C_a$ and $\{j, l\} \cap \{j_a, l_a\} = \varnothing$. Consider the map

$$e \colon R \hookrightarrow \bigcup_{a=1}^{d} \{\text{edges out of } C_a\}$$

defined as follows: $e(j_a, C_0)$ is the edge between C_a and C_0; otherwise $e(j, C_a)$ is the edge between C_a and $(il)(kj)C_a$ where $l \notin \lambda_M$ is uniquely determined by Lemma 6.5 (1). Note that the map e is injective by Lemma 6.1.

Lemma 6.6. $\sum_{D \sim C_0} d(D) \le |R| + 2d$.

Proof. To prove the lemma, we fix $a \in \{1, \dots, d\}$ and consider every edge out of C_a. We show that there are at most 2 edges out of C_a which are not in the image of the map e. Hence, $\sum_{D \sim C_0} d(D)$, which is the total number of edges out of $C_1, \dots, C_d$, is at most $|R| + 2d$.

Consider an edge that is not in the image of e. Then it is of the form $C_a \sim (il)(kj)C_a$ with $\{j,l\} \cap \{j_a, l_a\} \neq \varnothing$. This breaks up into several cases:
Case 1: $j = j_a$. If $C_a \sim (il)(kj_a)C_a$, then l is uniquely determined by Lemma 4.9. Thus there is at most one edge, out of C_a, with $j = j_a$, that is not in the image of e.
Case 2: $l = l_a$. This is similar to Case 1.
Case 3: $j = l_a$ or $l = j_a$. Then since $C_0 = (ij_a)(kl_a)C_a \sim C_a$, and since j, l uniquely determine each other, we must have both $j = l_a$ and $l = j_a$. Thus $(il)(kj)C_a = C_0$ and this edge is equal to $e(j_a, C_0)$, so it is in the image of e.

We have therefore shown that for fixed $1 \leq a \leq d$, there are at most 2 edges not in the image of the map e, corresponding to Cases 1 and 2. □

To complete the proof of Proposition 6.4, we need only show $|R| \leq p - m$. This follows from:

Proposition 6.7. *The projection map*

$$R \longrightarrow \{1, 2, \ldots, p\} \setminus \lambda_M$$

defined by $(j, C_a) \mapsto j$ *is injective.*

We prove this after a preliminary lemma. For ease of notation, throughout the rest of this section, we let

$$j' := \sigma^{-1}(j)$$

for $j \in \{1, \ldots, p\}$. Consider the function $f : \{j + p\mathbb{Z} : j \neq 2k' - i' \mod p\} \to \mathbb{C}$ defined by

$$f(j) = \frac{\zeta^{i'} - \zeta^{j}}{\zeta^{k'} - \zeta^{i'+j-k'}}.$$

Lemma 6.8. f *is injective.*

Proof. Note that

$$f(j') = \frac{\zeta^{i'} - \zeta^{j'}}{\zeta^{k'} - \zeta^{i'+j'-k'}} = \frac{\zeta^{i'}}{\zeta^{k'}} \cdot \frac{1 - \zeta^{j'-i'}}{1 - \zeta^{i'+j'-2k'}}$$

Note further that since $i \neq k$, we have $i' \neq k'$ and so $-i' \neq i' - 2k' \mod p$. Thus, it suffices to show more generally that if $0 \leq a, b < p$ with $a \neq b$, then the function

$$g(x) = \frac{1 - \zeta^{a+x}}{1 - \zeta^{b+x}}$$

is injective for $x \in \{0, 1, \dots, p-1\} \smallsetminus \{p-b\}$. Now, if $g(x) = g(y)$ for some $x, y \in \{0, 1, \dots, p-1\} \smallsetminus \{p-b\}$, then

$$\frac{1-\zeta^{a+x}}{1-\zeta^{b+x}} = \frac{1-\zeta^{a+y}}{1-\zeta^{b+y}}$$

and hence

$$\zeta^{a+y} - \zeta^{a+x} + \zeta^{b+x} - \zeta^{b+y} = 0.$$

As a result, if we take the exponents of the polynomial $z^{a+y} - z^{a+x} + z^{b+x} - z^{a+x}$ to be integers between 0 and p by reducing mod p, then it must be the zero polynomial; indeed, it is divisible by the p-th cyclotomic polynomial but has degree less than p. In particular, the z^{a+y} term must cancel with z^{a+x} or z^{b+y}, and hence

$$a + y = a + x \text{ or } a + y = b + y \bmod p.$$

Since $a \neq b$, we see $x = y$, and so g is injective. □

Proof of Proposition 6.7. Let the a_J and b_j be as in the first few paragraphs of Section 6. Consider the binary operation $\odot : \{a_J : J \neq M\}^{\times 2} \to \mathbb{R}$ defined by

$$a_J \odot a_L = -\frac{a_L - a_M}{a_J - a_M}.$$

We will show that if $(j, C_a) \in R$, then

$$b_j \odot b_l = f(j' + i' - l'_a), \tag{6.9}$$

where l is the unique element satisfying $l' = i' + j' - k'$. Assuming this for the moment, we see j determines l, which then determines $b_j \odot b_l = f(j' + i' - l'_a)$. Since f is injective by Lemma 6.8, we see j determines l'_a. Since $l_0, \dots, l_d$ are distinct, by Lemma 6.5 (2), we find that there is at most one value $0 \leq a \leq d$ for which $(j, C_a) \in R$, thereby proving the proposition.

It remains to prove (6.9). We first consider elements of form $(j_a, C_0) \in R$. In this case, $j = j_a$, $l = l_a$, and $b_i = b_k = a_M \notin \{b_j, b_l\}$. Since H contains both C_0 and $C_a = (ij_a)(kl_a)C_0$, Lemma 4.9 shows that $i' + j' = k' + l' \bmod p$. Since $\sigma w \in \mathrm{Span}(C_0) \subset H$ and $(ij)(kl)\sigma w \in \mathrm{Span}((ij)(kl)C_0) \subset H$, we find

$$(\dots, \zeta^{j'} - \zeta^{i'}, \dots, \zeta^{i'} - \zeta^{j'}, \dots, \zeta^{l'} - \zeta^{k'}, \dots, \zeta^{k'} - \zeta^{l'}, \dots) = \sigma w - (ij)(kl)\sigma w \in H$$

where the omitted entries are 0, and the non-zero entries are in the j, i, l, k-th positions, respectively. As $H = (b_1, \dots, b_p)^{\perp}$, we have

$$b_j(\zeta^{j'} - \zeta^{i'}) + b_i(\zeta^{i'} - \zeta^{j'}) + b_l(\zeta^{l'} - \zeta^{k'}) + b_k(\zeta^{k'} - \zeta^{l'}) = 0$$

and so

$$b_l = \frac{(\zeta^{j'} - \zeta^{i'} + \zeta^{l'} - \zeta^{k'})a_M - b_j(\zeta^{j'} - \zeta^{i'})}{\zeta^{l'} - \zeta^{k'}}.$$

Since $l' = i' + j' - k' \bmod p$, we have $\zeta^{l'} = \zeta^{i'+j'-k'}$. Note that $j' \neq 2k' - i' \bmod p$ since otherwise we would have $k' = l' \bmod p$, which is not possible as $k \neq l$. As a result, $f(j')$ is well-defined and

$$b_l = a_M + (a_M - b_j)f(j').$$

As a result, we have our desired equality

$$b_j \odot b_l = -\frac{b_l - a_M}{b_j - a_M} = f(j') = f(j' + i' - l'_0).$$

We next consider an element of the form $(j, C_a) \in R$ for some $1 \leq a \leq d$. Then, by definition, we have $C_a \sim (il)(kj)C_a$ for some $\{j, l\} \cap \{j_a, l_a\} = \varnothing$. Let $w_a := (ij_a)(kl_a)\sigma w \in C_a \subset H$ and note $(il)(kj)w_a \in (ij)(kl)C_a \subset H$. So,

$$(\ldots, \zeta^{j'_a} - \zeta^{l'}, \ldots, \zeta^{l'_a} - \zeta^{j'}, \ldots., \zeta^{j'} - \zeta^{l'_a}, \ldots, \zeta^{l'} - \zeta^{j'_a}, \ldots) = w_a - (il)(kj)w_a \in H$$

where the omitted entries are 0, and the non-zero entries are in the i, k, j, l-th position, respectively. As a result,

$$b_i(\zeta^{j'_a} - \zeta^{l'}) + b_k(\zeta^{l'_a} - \zeta^{j'}) + b_j(\zeta^{j'} - \zeta^{l'_a}) + b_l(\zeta^{l'} - \zeta^{j'_a}) = 0.$$

It follows that

$$b_l = \frac{(\zeta^{l'} - \zeta^{j'_a} + \zeta^{j'} - \zeta^{l'_a})a_M - b_j(\zeta^{j'} - \zeta^{l'_a})}{\zeta^{l'} - \zeta^{j'_a}}$$

and hence

$$b_j \odot b_l = \frac{\zeta^{j'} - \zeta^{l'_a}}{\zeta^{l'} - \zeta^{j'_a}}.$$

It remains to prove this expression equals $f(j' + i' - l'_a)$.

Since $i' + j'_a = k' + l'_a \bmod p$ and $i' + j' = l' + k' \bmod p$, we have

$$l' - l'_a + i' = i' + (j' + i' - l'_a) - k' \mod p.$$

As a result,

$$f(j' + i' - l'_a) = \frac{\zeta^{i'} - \zeta^{j'-l'_a+i'}}{\zeta^{k'} - \zeta^{i'+(j'+i'-l'_a)-k'}} = \frac{\zeta^{i'} - \zeta^{j'-l'_a+i'}}{\zeta^{j'_a-l'_a+i'} - \zeta^{l'-l'_a+i'}} = \frac{\zeta^{l'_a} - \zeta^{j'}}{\zeta^{j'_a} - \zeta^{l'}}$$

thereby finishing the proof. □

Declaration of competing interest

None declared.

References

[1] O.M. Adamovich, E.O. Golovina, Simple linear Lie groups having a free algebra of invariants, Sel. Math. Sov. 3 (2) (1983/84) 183–220, selected reprints.
[2] W. Burnside, On some properties of groups of odd order, Proc. Lond. Math. Soc. 33 (1901) 162–185.
[3] Claude Chevalley, Invariants of finite groups generated by reflections, Am. J. Math. 77 (1955) 778–782.
[4] Edidin Dan, Matthew Satriano, An intrinsic characterization of cofree representations of reductive groups, https://arxiv.org/pdf/1905.04845.pdf, 2019.
[5] Victor G. Kac, Vladimir L. Popov, Ernest B. Vinberg, Sur les groupes linéaires algébriques dont l'algèbre des invariants est libre, C. R. Acad. Sci. Paris Sér. A-B 283 (12) (1976) A875–A878.
[6] Peter Littelmann, Koreguläre und äquidimensionale Darstellungen, J. Algebra 123 (1) (1989) 193–222.
[7] Gerald W. Schwarz, Representations of simple Lie groups with regular rings of invariants, Invent. Math. 49 (2) (1978) 167–191.
[8] Gerald W. Schwarz, Differential operators on quotients of simple groups, J. Algebra 169 (1) (1994) 248–273.
[9] Gerald W. Schwarz, Representations of simple Lie groups with a free module of covariants, Invent. Math. 50 (1) (1978/79) 1–12.
[10] Jean-Pierre Serre, Linear Representations of Finite Groups, Springer-Verlag, New York-Heidelberg, 1977, translated from the second French edition by Leonard L. Scott, Graduate Texts in Mathematics, vol. 42.
[11] G.C. Shephard, J.A. Todd, Finite unitary reflection groups, Can. J. Math. 6 (1954) 274–304.