Università degli Studi di Padova

Dipartimento di Matematica "Tullio Levi–Civita"
Corso di Dottorato in Scienze Matematiche
XXXIII Ciclo

Invariable generation of finite groups

Daniele Garzoni

Sotto la supervisione di

Prof. Andrea Lucchini

A Mamma e Papà

Abstract

In this thesis we prove various results in the subject of invariable generation of finite groups. We also prove some upper bounds to the number of conjugacy classes of certain primitive permutation groups, which are then applied to a problem of invariable generation.

In Chapter 1, we present a first description of the results and of the objectives of the thesis, in order to orient and motivate the reader. In Chapter 2, we provide a general introduction to the subject of invariable generation of groups. The heart of the thesis consists of Chapters 3–7. In these chapters we prove all the main results, we provide context for each result, and we propose questions and conjectures.

Acknowledgements

I would like to thank my supervisor, Professor Andrea Lucchini, who has guided me in these three years.
Andrea, thank you for your suggestions, your ideas, and your support. You always let me pursue my interests and my research, at my own pace, with no pressure. This has made my experience very much stimulating and enjoyable.

I would like to thank Professor Nick Gill, who gave me a warm welcome during the month I spent at the University of South Wales.
Nick, thank you for working and discussing math with me. Your smile and your positive energy are contagious, in many ways. It has been a pleasure playing football with you!

Contents

Chapter 1

Main results of the thesis

In this chapter we present a first description of the main results of this thesis. We are not presenting here the statements of *all* the main results; we prefer, instead, to do this in the introduction of Chapters 3–7. In this way, we believe that the results can be contextualized in a more satisfactory way.

First of all, we define the main concept of this thesis.

Definition. *Let G be a group and let X be a subset of G. We say that X invariably generates G if $\langle x^{g(x)}, x \in X\rangle = G$ for every choice of $g(x) \in G$. We write in this case $\langle X\rangle_I = G$.*

In words, we are free to replace every element of X by an arbitrary conjugate, and we must always obtain a generating set for G. We note that, if G is abelian, then $\langle X\rangle_I = G$ if and only if $\langle X\rangle = G$.

This concept was introduced by Dixon [Dix92] with motivations from computational Galois theory. In Chapter 2 we will review these motivations, and we will present some known results in the area.

1.1 Minimal invariable generating sets

In Chapter 3 we will present the content of the paper [GL20], which was written in collaboration with Andrea Lucchini.

Given a finite group G and a subset X of G, we say that X is a *minimal generating set* of G if X generates G, but no proper subset of X generates G. We denote by $d(G)$ (resp. $m(G)$) the smallest (resp. largest) possible size of a minimal generating set of G.

We can extend these definitions to invariable generation in the obvious way. Namely, X is a *minimal invariable generating set* of G if X invariably generates G, but no proper subset of X invariably generates G. We denote by $d_I(G)$ (resp. $m_I(G)$) the smallest (resp. largest) possible size of a minimal invariable generating set of G.

In Chapter 3 we will prove various results regarding minimal invariable generating sets of finite groups, especially in relation to mimimal generating sets. For instance:

Theorem 1. *Let G be a finite soluble group. Then, $m_I(G) = m(G)$.*

Theorem 2. *Let G be a finite soluble group, and let t be an integer such that $d_I(G) \leqslant t \leqslant m_I(G)$. Then, G contains a minimal invariable generating set of size t.*

Theorems 1 and 2 appear respectively as Theorems 3.1.1 and 3.1.3 in Chapter 3. It is known that Theorem 2 is true in the case of usual generation, without the solubility assumption. More precisely, for every finite group G and for every $d(G) \leqslant t \leqslant m(G)$, G contains a minimal generating set of size t. This is a consequence of a result known as Tarski irredundant basis theorem [Tar75]. The question whether Theorem 2 is true without the solubility assumption remains open.

We will prove various other results of this flavour. We will also ask the following question, which appears as Question 3.1.2 in Chapter 3.

Question 3. *Let G be a finite group. Is it true that $m_I(G) \leqslant m(G)$?*

In Section 3.10 we will make some remarks which put the results of the chapter in comparison to the case of classical generation.

1.2 On the number of conjugacy classes of a primitive permutation group with nonabelian socle

The content of Chapter 4 is joint work with Nick Gill.

In Chapter 4, we will leave for a while invariable generation, in order to prove some bounds to the number of conjugacy classes of certain primitive permutation groups. We believe that these results are interesting in their own, and will be applied in Chapter 5 to a problem of invariable generation of S_n, which we will discuss in Section 1.3.

For a finite group G, let $k(G)$ denote the number of conjugacy classes of G. Our main result is as follows:

Theorem 4. *Let G be a finite primitive permutation group of degree n with nonabelian socle. Then of the following holds:*

(1) $k(G) < n/2$, and $k(G) = O(n^\delta)$ for some absolute $\delta < 1$.

(2) G belongs to explicit families of examples.

See Theorem 4.1.1 for a precise statement. In particular, the cases contemplated in item (2) are "related" either to two infinite families of almost simple

primitive permutation groups, or to further finitely many almost simple primitive permutation groups.

This theorem refines and complements some results in the area. Specifically, Maróti [Mar05] proved that, if G is a primitive permutation group of degree n, then $k(G) \leqslant p(n)$, where $p(n)$ denotes the number of partitions of n. This bound is attained by S_n in its action on n points. Moreover, he proved that if the socle of G is not a direct product of alternating groups, then $k(G) \leqslant n^6$.

See Chapter 4 for further background results and further comments and questions. For instance: In item (1) of Theorem 4, we have $k(G) = o(n)$. Can we prove that $k(G) = o(n)$ in further cases? See Question 4.1.4 and Conjecture 4.4.2.

1.3 Large minimal invariable generating sets of S_n

The content of Chapter 5 is joint work with Nick Gill.

For a finite group G, recall the invariant $m_I(G)$ introduced in Section 1.1. In Chapter 5 we will estimate $m_I(G)$ in case $G = S_n$. Our main result is as follows; this appears as Theorem 5.1.1 in Chapter 5.

Theorem 5. *Let $n \geqslant 5$ be an integer. Then*

$$\frac{n}{2} - \log n < m_I(S_n) < \frac{n}{2} + \Delta(n) + O\left(\frac{\log n}{\log\log n}\right),$$

where $\Delta(n)$ is the number of divisors of n.

It is well known that $\Delta(n) = n^{o(1)}$, therefore we deduce that $m_I(S_n)$ is asymptotic to $n/2$ as $n \to \infty$.

Theorem 1, stated in Section 1.1, implies that, for $n \leqslant 4$, $m_I(S_n) = m(S_n)$. We will prove the following as a consequence of Theorem 5.

Corollary 6. *Let $n \geqslant 5$ be an integer. Then, $m_I(S_n) < m(S_n)$.*

In particular, this gives a positive answer to Question 3 in case $G = S_n$.

Theorem 4, stated in the previous section, is an essential ingredient in the proof of the upper bound in Theorem 5.

We will ask some questions, mainly of combinatorial flavour, towards possible improvements of Theorem 5; see Questions 5.1.4 and 5.1.6.

1.4 Connected components in the invariably generating graph

In Chapter 6 we will present the content of the preprint [Gar20b].

Given a finite group G, the *invariably generating graph* $\Lambda(G)$ of G is the undirected graph whose vertices are the conjugacy classes of G different from

$\{1\}$, and two vertices x^G and y^G are adjacent if $\langle x, y\rangle_I = G$. This graph was introduced in [Gar20a]. Moreover, we define $\Lambda^+(G)$ as the graph obtained by removing the isolated vertices of $\Lambda(G)$.

We will prove the following result, which appears as Theorem 6.1.1 in Chapter 6.

Theorem 7. *For every positive integer n, there exists a finite group G such that $\Lambda^+(G)$ has more than n connected components.*

This result should be seen in comparison to the analogous graph for the case of usual generation. Given a finite group G, the *generating graph* $\Gamma(G)$ of G is the undirected graph whose vertices are the non-identity elements of G, and two vertices x and y are adjacent if $\langle x, y\rangle = G$. Moreover, we let $\Gamma^+(G)$ be the graph obtained by removing the isolated vertices of $\Gamma(G)$. No example of G is known for which $\Gamma^+(G)$ is disconnected, which represents a sharp difference with respect to Theorem 7.

See Chapter 6, and in particular Section 6.3, for further comments, related to clique number and chromatic number of the graph $\Lambda(G)$.

1.5 On the probability of generating invariably a finite simple group

In Chapter 7 we will present the content of the preprint [GM20], which was written in collaboration with Eilidh McKemmie.

Given a finite group G and a subset A of G, denote by $\mathbf{P}_{\mathrm{inv}}(G, A)$ the probability that, if $y \in G$ is chosen uniformly at random, there exists $x \in A$ such that $\langle x, y\rangle_I = G$. In case $A = \{x\}$, we will write $\mathbf{P}_{\mathrm{inv}}(G, x)$ instead of $\mathbf{P}_{\mathrm{inv}}(G, \{x\})$.

We will prove various results regarding $\mathbf{P}_{\mathrm{inv}}(G, A)$ in case G is a nonabelian finite simple group. For instance:

Theorem 8. *Let G be a nonabelian finite simple group. There exist an absolute constant $\epsilon > 0$ and an element $x \in G$ such that $\mathbf{P}_{\mathrm{inv}}(G, x) \geqslant \epsilon$.*

This appears as Theorem 7.1.1 in Chapter 7. In case G is a finite simple group of Lie type of bounded rank, we will prove a strengthening of Theorem 8 (see Theorem 7.1.2), which implies the following

Theorem 9. *Let G be a finite simple group of Lie type of bounded rank. Then, two randomly chosen elements of G invariably generate G with probability bounded away from zero.*

See Theorem 7.1.3 for a more precise statement. Theorem 9 nearly finishes a problem originally addressed by Dixon [Dix92]. See Section 2.2 in Chapter 2, and also Subsection 7.1.1 in Chapter 7, for more context and details in this direction.

We also wish to find "small" subsets A of G such that $\mathbf{P}_{\mathrm{inv}}(G, A)$ is close to 1 as the order of G grows.

Theorem 10. *Let G be a nonabelian finite simple group. With explicit exceptions (e.g., $G_2(q)$ with q a power of 3), there exists a subset A of G of size at most 6 such that $\mathbf{P}_{\mathrm{inv}}(G, A)$ tends to 1 as $|G| \to \infty$.*

The explicit exceptions are perhaps the most surprising part. For instance:

Theorem 11. *Let $G = G_2(q)$ with q a power of 3. Then $\mathbf{P}_{\mathrm{inv}}(G, G) = 1/2 + O(1/q)$.*

Theorems 10 and 11 are particular cases of the more general Theorem 7.1.4, which presents a strong dichotomy. These results admit nice interpretations in terms of the graph $\Lambda_e(G)$, whose vertices are the nontrivial elements of G, and two vertices are adjacent if $\langle x, y \rangle_I = G$. (This graph is obtained from the graph $\Lambda(G)$ — which we introduced in Section 1.4 — by replacing conjugacy classes by elements.)

See Subsection 7.1.2 for many comments in this direction, and many comparisons to the case of classical generation, specifically to celebrated results of Guralnick–Kantor [GK00].

Chapter 2

General introduction

In this chapter, we will review the motivations for the study of invariable generation, and we will present some known results in the area. Although our aim is to provide a general introduction to the subject, sometimes we will prioritize the results which, in a way or another, have some connections to the results presented in Chapter 1.

2.1 Motivation

The concept of invariable generation was introduced by Dixon [Dix92] with motivations from computational Galois theory. We now review these motivations.

Let $f \in \mathbf{Z}[X]$ be a monic separable polynomial of degree n. Let K be a splitting field of f over $\mathbf{Q}$, and set $\mathrm{Gal}(f) := \mathrm{Gal}(K/\mathbf{Q})$, the *Galois group* of f. Then $\mathrm{Gal}(f)$ acts faithfully on the n roots of f, and this gives an embedding $\mathrm{Gal}(f) \leqslant S_n$.

Question 2.1.1. *How can we recognize whether* $\mathrm{Gal}(f) = S_n$*?*

Dixon's idea was the following (in fact the general setup was known well before Dixon's paper; see for instance the beautiful survey of Stevenhagen–Lenstra [SL96] for a historical overview). Let p be a prime, and assume $p \nmid \mathrm{disc}(f)$. We reduce the coefficients of f modulo p, and we decompose the obtained polynomial into irreducible factors in $\mathbf{F}_p[X]$; say $\overline{f} = g_1 \cdots g_t$, with $g_i \in \mathbf{F}_p[X]$ irreducible of degree n_i. Since $p \nmid \mathrm{disc}(f)$, the irreducible factors are pairwise distinct. Note that $n_1 + \cdots + n_t = n$; in particular we can associate to the prime p a partition $\mathrm{cycle}(p) := (n_1, \ldots, n_t)$ of n.

Theorem 2.1.2. *(Frobenius)* $\mathrm{Gal}(f)$*, viewed as a subgroup of* S_n*, contains an element with cycle type* $\mathrm{cycle}(p)$*.*

Now pick primes $p_1, \ldots, p_N$ such that $p_i \nmid \mathrm{disc}(f)$ for every i. By Theorem 2.1.2, we know that $\mathrm{Gal}(f)$ contains elements $\sigma_1, \ldots, \sigma_N$ with cycle type $\mathrm{cycle}(p_1), \ldots, \mathrm{cycle}(p_N)$, respectively.

We deduce the following implication:

$$\langle \sigma_1, \dots, \sigma_N \rangle_I = S_n \implies \mathrm{Gal}(f) = S_n.$$

It is really necessary to consider invariable generation here, since we know only the cycle type of the elements (i.e., we know the elements only up to conjugation in S_n). This gives a method to answer Question 2.1.1.

However, what happens if N is large and $\{\sigma_1, \dots, \sigma_N\}$ does not invariably generate S_n? We cannot select primes forever. The answer is given by another theorem of Frobenius (proved around 1880), which was subsequently generalized by Chebotarev in his celebrated Density Theorem.

Theorem 2.1.3. *(Frobenius density theorem) Fix a partition $\mathfrak{p}$ of n. The natural density of primes p such that $p \nmid \mathrm{disc}(f)$ and $\mathrm{cycle}(p) = \mathfrak{p}$ is equal to the proportion of elements of $\mathrm{Gal}(f)$ whose cycle type is equal to $\mathfrak{p}$.*

Assume now we know that "few" random elements of S_n invariably generate S_n with good probability. Then pick "many" primes at random, and discard the primes dividing $\mathrm{disc}(f)$. Assume the corresponding elements of $\mathrm{Gal}(f)$, given by Theorem 2.1.2, do not invariably generate S_n. By Theorem 2.1.3, we can conclude that, *likely*, $\mathrm{Gal}(f) \neq S_n$.

Formally, we have constructed a Monte Carlo algorithm in order to answer Question 2.1.1, which depends on the properties of invariable generation of S_n.

2.2 Invariable generation of S_n

The efficiency of the algorithm constructed in the previous section depends on the following question: How small can we take t, so that t elements of S_n, chosen uniformly at random, invariably generate S_n with good probability?

The first result in this direction was obtained by Dixon [Dix92], who showed that $O((\log n)^{1/2})$ random elements of S_n invariably generate S_n with probability tending to 1 as $n \to \infty$. A fundamental contribution was given by Luczak–Pyber [LP93], who showed the existence of an absolute constant c such that c random elements of S_n invariably generate S_n with probability bounded away from zero. The exact value of c turned out to be four. This was proved by Pemantle–Peres–Rivin [PPR16] and Eberhard–Ford–Green [EFG17].

Theorem 2.2.1. *(1) [PPR16] The probability that four random elements of S_n invariably generate S_n is bounded away from zero.*

(2) [EFG17] The probability that three random elements of S_n invariably generate S_n tends to zero as $n \to \infty$.

We remark that the probability that a bounded number of random elements of S_n invariably generate S_n is bounded away from 1 — so Theorem 2.2.1(1) is best possible. The reason is as follows. A random permutation has a fixed point with probability bounded away from zero (specifically, with probability tending

to $1 - e^{-1}$). Therefore, if we pick a bounded number of random elements, they all will have a fixed point with probability bounded away from zero. We can conjugate these elements so that they fix the same point, and in particular we see that the elements do not invariably generate S_n.

2.3 General considerations

In this section, we present the basic theory of invariable generation of groups (that is, the theory that follows somewhat quickly from the definition). Near the end of the section we will also recall some more advanced results in the area.

2.3.1 Connection with permutation groups

There is an intimate and very important connection between invariable generation and the theory of permutation groups. For a group G and a subgroup H of G, set

$$\widetilde{H} = \bigcup_{g \in G} H^g.$$

(When we use the notation $\widetilde{H}$, the ambient group G will always be clear from the context.) We have the following basic lemma, which follows easily from the definitions. Recall that, if G is a group acting on a set Ω, an element of G is called a *derangement* if it fixes no point on Ω.

Lemma 2.3.1. *Let G be a group and let $X \subseteq G$. The following are equivalent:*

(i) $\langle X \rangle_I = G$.

(ii) For every proper subgroup H of G, $X \not\subseteq \widetilde{H}$.

(iii) For every transitive action of G on a set with at least two points, X contains a derangement.

Proof. The equivalence (i) $\iff$ (ii) follows from the definitions. The equivalence (ii) $\iff$ (iii) follows from the fact that, if H is the stabilizer of a point in a transitive action of G, then $\widetilde{H}$ is the set of elements of G fixing at least one point. □

Clearly, if G is a finite group, in item (ii) it is equivalent to consider only the maximal subgroups H of G.

The equivalence (i) $\iff$ (ii) says that the subsets $\widetilde{H}$ play the role that, in case of classical generation, is played by the subgroups H. Indeed, $\langle X \rangle = G$ if and only if $X \not\subseteq H$ for every proper subgroup H of G. This represents a complication — the subsets $\widetilde{H}$ are harder to handle than the subgroups H.

Nonetheless, the equivalence (i) $\iff$ (iii) establishes a strong connection between invariable generation and permutation groups. In particular, in order to study invariable generation we can invoke the quite advanced theory of permutation groups; see for instance Section 2.5.

It is clear from the definition that, if G is abelian, the concepts of generation and invariable generation coincide. In case of finite groups, one can extend this to nilpotent groups.

Lemma 2.3.2. *Assume G is a finite nilpotent group, and assume $X \subseteq G$. Then, $\langle X \rangle_I = G$ if and only if $\langle X \rangle = G$.*

Proof. In a finite nilpotent group G, every maximal subgroup H is normal, so $\widetilde{H} = H$. The statement follows immediately from this and Lemma 2.3.1. □

One cannot push this beyond nilpotent groups. Indeed, we shall see, mainly in Chapter 3, that for finite soluble groups the concepts of generation and invariable generation present strong differences. See also Detomi–Lucchini [DL16] for results in this direction.

2.3.2 Is it true that $\langle G \rangle_I = G$?

Is every group invariably generated by any of its subsets? This is not immediately clear from the definition. A classical (and elementary) theorem of Jordan from 1872 asserts that every *finite* transitive permutation group G on a set with at least two points contains a derangement. By Lemma 2.3.1, we deduce the following.

Lemma 2.3.3. *Assume G is finite. Then, $\langle G \rangle_I = G$.*

If G is infinite, however, this needs not be the case. One of the most natural examples is probably the following. Let $G = \mathrm{GL}_n(\mathbf{C})$, with $n \geqslant 2$, and let H be the subgroup of G consisting of the upper triangular matrices. Jordan normal form implies that $\widetilde{H} = G$, so by Lemma 2.3.1 we see that G is not invariably generated by any of its subsets (equivalently, it is not invariably generated by itself).

It is therefore natural to study under which conditions an infinite group G is such that $\langle G \rangle_I = G$. This problem was first addressed by Wiegold [Wie76, Wie77]. More recently, interesting results in this direction have been proved by Kantor–Lubotzky–Shalev [KLS15], Gelander [Gel15] and Gelander–Meiri [GM17].

Several interesting questions remain open in this area. For instance, is $\mathrm{SL}_n(\mathbf{Z})$ invariably generated by itself for $n \geqslant 3$? ([KLS15, Section 5]).

2.4 Invariable generation of finite groups

Let G be a finite group. Lemma 2.3.3 tells us that $\langle G \rangle_I = G$. Can we say something more? For a finite group G, let $d_I(G)$ denote the number of elements needed to invariably generate G (we introduced this invariant in Section 1.1, using a slightly different terminology).

Kantor–Lubotzky–Shalev [KLS11] proved several results in this direction. First, they showed that Lemma 2.3.3 can be improved as follows.

Theorem 2.4.1. *[KLS11] Let G be a finite group. Then, $d_I(G) \leqslant \log_2 |G|$.*

The proof of this theorem uses the Classification of Finite Simple Groups. Note that the bound is attained by elementary abelian 2-groups. For special classes of groups, one can obtain much better results. The following was proved by Kantor–Lubotzky–Shalev [KLS11] and Guralnick–Malle [GM12b] independently.

Theorem 2.4.2. *[KLS11, GM12b] Let G be a nonabelian finite simple group. Then, $d_I(G) = 2$.*

The proof of this theorem also uses CFSG. This is not surprising at all — we know that every finite simple group is 2-generated (not necessarily invariably) *only* as an application of CFSG.

We refer to Detomi–Lucchini [DL16] for interesting results regarding $d_I(G)$ in case G is a finite soluble group; and to Tracey [Tra19] for interesting results regarding $d_I(G)$ in case G is a finite linear group or a finite permutation group.

Another problem addressed in [KLS11] concerns the probability of generating invariably a *general* finite group. For a finite group G, the *Chebotarev invariant* $C(G)$ of G is the expected value of the random variable n that is minimal subject to the requirement that n randomly chosen elements of G invariably generate G. Kowalski–Zywina [KZ12] conjectured that $C(G) = O(|G|^{1/2})$ for every finite group G. A slightly weaker result was proved in [KLS11], while the conjecture was proved in full by Lucchini [Luc18].

Theorem 2.4.3. *[Luc18] There exists an absolute constant $c > 0$ such that $C(G) \leqslant c|G|^{1/2}$ for every finite group G.*

The constant c was estimated by Lucchini–Tracey [LT17].

2.5 Boston–Shalev conjecture

In this section we discuss a topic which is not immediately related to invariable generation, namely, the solution of the so-called *Boston–Shalev conjecture* by Luczak–Pyber and Fulman–Guralnick.

We prefer to include a discussion here, for two reasons. First, this topic is very important in the study of invariable generation — for instance, some of the results stated in this chapter make use of these ideas, as well as all the results which we will prove in Chapter 7. Second, the methods developed by Fulman and Guralnick are important in their own, and have several applications in various directions — incidentally, the paper [FG12] is an essential ingredient in the proof of the results of Chapter 4.

Since in this thesis we are mainly interested in invariable generation, we will motivate the interest in the Boston–Shalev conjecture (which we will shortly state) as follows.

For a finite group G, denote by $\mathbf{P}_I(G, t)$ the probability that t elements of G, chosen uniformly at random, invariably generate G. Let $\mathcal{M}(G)$ be a set of

representatives for the conjugacy classes of maximal subgroups of G. It follows immediately from Lemma 2.3.1 that

$$1 - \mathbf{P}_I(G,t) = \frac{|\bigcup_{M\in\mathcal{M}(G)}(\widetilde{M})^t|}{|G|^t},$$

where

$$(\widetilde{M})^t = \underbrace{\widetilde{M}\times\cdots\times\widetilde{M}}_{t}.$$

In particular, we see that, roughly speaking, $\mathbf{P}_I(G,t)$ is "large" if and only if the subsets $(\widetilde{M})^t$ of G^t are "small". We are now ready to state the following conjecture.

Boston–Shalev conjecture. *There exists an absolute constant $\epsilon > 0$ such that, if G is a finite simple group and if H is a proper subgroup of G, then*

$$\frac{|\widetilde{H}|}{|G|} \leqslant 1-\epsilon.$$

In other words, if G is simple, the proportion of derangements of G in every transitive action of G on a set with at least two points is at least ϵ.

Now we know that the conjecture is true. This was proved by Dixon [Dix92] and Luczak–Pyber [LP93] for alternating groups, and by Fulman–Guralnick [FG03, FG12, FG17, FG18] for groups of Lie type. (By Jordan's theorem, the statement is true for groups of bounded order, therefore one can ignore the sporadic groups.)

Theorem 2.5.1. *Boston–Shalev conjecture is true.*

This theorem is important in the proof of Theorem 2.4.3, and will be crucial in all the results of Chapter 7. We remark that, in many cases, Luczak–Pyber and Fulman–Guralnick proved much stronger results than the mere statement of the conjecture; we refer to the relevant papers for details. Moreover, the results of Luczak–Pyber have been subsequently improved by Eberhard–Ford–Green [EFG16] and Eberhard–Ford–Koukoulopoulos [EFK16].

As mentioned already, the work by Luczak–Pyber and Fulman–Guralnick has several applications, which go beyond derangements and invariable generation. We refer to the introduction of [FG12] for some of these.

Chapter 3

Minimal invariable generating sets

The content of this chapter consists of the paper [GL20], which was written in collaboration with Andrea Lucchini. The introduction has been changed slightly, and part of it has been moved to Section 3.10 (the last section of this chapter); but the mathematical content is unchanged.

3.1 Introduction

Let G be a finite group, and let X be a subset of G. We say that X is a *minimal generating set* of G if X generates G, but no proper subset of X generates G. We denote by $d(G)$ (resp. $m(G)$) the smallest (resp. largest) possible size of a minimal generating set of G.

Moreover, X is a *minimal invariable generating set* of G if X invariably generates G, but no proper subset of X invariably generates G. We denote by $d_I(G)$ (resp. $m_I(G)$) the smallest (resp. largest) possible size of a minimal invariable generating set of G.

Our aim is to study minimal invariable generating sets of finite groups, especially in relation to mimimal generating sets.

3.1.1 Estimating $m_I(G)$

There are three clear inequalities:

$$d(G) \leqslant m(G), \quad d_I(G) \leqslant m_I(G), \quad d(G) \leqslant d_I(G).$$

Instead, it is not immediately clear how to compare $m(G)$ and $m_I(G)$. In case of soluble groups, we give an answer, as follows.

Theorem 3.1.1. *Let G be a finite soluble group. Then, $m_I(G) = m(G)$.*

The problem is wide open in case of general finite groups. We ask the following

Question 3.1.2. *Let G be a finite group. Is it true that $m_I(G) \leqslant m(G)$?*

In case of positive answer, we would get the following somewhat surprising chain of inequalities:

$$d(G) \leqslant d_I(G) \leqslant m_I(G) \leqslant m(G).$$

Although we are not able to answer Question 3.1.2, In Section 3.4 we will prove that the difference $m(G) - m_I(G)$ can be arbitrarily large. This statement is somewhat opposite to the known fact (proved in Kantor–Lubotzky–Shalev [KLS11] and Detomi–Lucchini [DL16]) that $d_I(G) - d(G)$ can be arbitrarily large. In Chapter 5, we will also prove that Question 3.1.2 has a positive answer in case $G = S_n$ (see Corollary 5.1.2).

We will also prove that there exist infinitely many nonabelian finite simple groups G for which $m_I(G) = m(G)$ (see Proposition 3.5.1).

3.1.2 The Tarski irredundant basis theorem

The Tarski irredundant basis theorem is a theorem in universal algebra, proved by Tarski [Tar75], which has the following curious consequence.

Theorem. *[Tar75] Let G be a finite group, and let t be an integer such that $d(G) \leqslant t \leqslant m(G)$. Then, G contains a minimal generating set of size t.*

The proof of this result is essentially an elementary and nice counting argument. One may ask what happens for invariable generation. For soluble groups, we can give an answer. The proof, however, is different and requires some machinery — specifically, the theory of *crowns*, see Section 3.2.

Theorem 3.1.3. *Let G be a finite soluble group, and let t be an integer such that $d_I(G) \leqslant t \leqslant m_I(G)$. Then, G contains a minimal invariable generating set of size t.*

It is not clear to us what one should conjecture in general. The proof given for soluble groups is *really* about soluble groups, and does not give much insight on what should happen, say, for a finite simple group.

Question 3.1.4. *Does Theorem 3.1.3 hold for every finite group G?*

3.1.3 $\mathcal{B}_I$-groups

Apisa–Klopsch [AK14] defined the class of *$\mathcal{B}$-groups* as the class of finite groups for which $d(G) = m(G)$. We then define the *$\mathcal{B}_I$-groups* as the finite groups for which $d_I(G) = m_I(G)$.

In Proposition 3.8.2 we will classify the soluble $\mathcal{B}_I$-groups. On the other hand, the problem of investigating the finite unsoluble $\mathcal{B}_I$-groups remains entirely open.

Let us make some remarks. It was proved in [AK14] that a $\mathcal{B}$-group is soluble. What is more, the following implication holds (see Proposition 3.8.2).

$$G \text{ (soluble) } \mathcal{B}\text{-group} \implies G \text{ soluble } \mathcal{B}_I\text{-group}. \tag{3.1.1}$$

On the other hand, a $\mathcal{B}_I$-group needs not be soluble (we shall see that A_5 is a $\mathcal{B}_I$-group), and moreover the implication (3.1.1) does not admit a converse. Indeed, we will construct many examples of soluble $\mathcal{B}_I$-groups that are not $\mathcal{B}$-groups. Interestingly, these have connections with "secretive" p-groups, introduced in Kovács–Neubüser–Neumann [KNN71] with different purposes. See Section 3.8 for details.

3.1.4 Invariable basis property

Again, we extend a definition from [AK14], and we say that a finite group G has the *invariable basis property* if every subgroup of G is a $\mathcal{B}_I$-group.

In Proposition 3.9.3 we will classify the Frattini-free soluble groups with the invariable basis property. We will also prove that there are only four examples of nonsoluble groups with the invariable basis property; namely, the finite simple groups $\mathrm{PSL}_2(5)$, $\mathrm{PSL}_2(8)$, ${}^2B_2(8)$, ${}^2B_2(32)$ (see Corollary 3.9.6).

3.1.5 Context

The invariant $m_I(G)$ has been introduced in the paper [GL20], which is reproduced in this chapter. What else is known about the invariants $d(G)$, $d_I(G)$ and $m(G)$?

The invariant $d(G)$ has been intensively studied in the literature. For instance, a deep and important theorem states that $d(G) = 2$ for every nonabelian finite simple group G. This was proved by Steinberg [Ste62] for groups of Lie type, and by Aschbacher–Guralnick [AG84] for sporadic groups, the case of alternating group being folklore.

The invariable counterpart $d_I(G)$ has been studied more recently, and in Section 2.4 we recalled some results regarding this invariant. For instance, Theorem 2.4.2 states that $d_I(G) = 2$ for every nonabelian finite simple group G, which represents a strengthening of the equality $d(G) = 2$.

The invariant $m(G)$ has received some attention in connection to the study of the Product Replacement Algorithm. For instance, the product replacement graph $\Gamma_k(G)$ is connected for every $k \geqslant d(G) + m(G)$ (see Pak [Pak01, Proposition 2.2.2]), and a random walk on this graph reaches a uniform distribution in $|G|^{O(m(G))}k^2 \log k$ steps (see Diaconis and Saloff-Coste [DSC98, p. 254]). Results regarding $m(G)$ have been obtained by Apisa–Klopsch [AK14] and Lucchini [Luc13a, Luc13b].

We refer to Section 3.10 for further remarks, concerning mainly the relation between generation and invariable generation.

3.1.6 Organization of the chapter and notation

In Section 3.2 we introduce some background material, mainly about the theory of crowns in finite soluble groups. In Section 3.3 we prove Theorem 3.1.1, discuss Question 3.1.2 and prove related results. In Section 3.4 we construct examples of groups G for which $m(G) - m_I(G)$ is arbitrarily large. In Section 3.5 we show that there are infinitely many nonabelian finite simple groups G such that $m_I(G) = m(G)$. In Section 3.6 we prove Theorem 3.1.3 and make further considerations. In Section 3.7 we introduce the analogue of the Frattini subgroup of a group from the point of view of invariable generation. We use this concept in Sections 3.8 and 3.9, where we discuss $\mathcal{B}_I$-groups and groups with the invariable basis property, respectively. In Section 3.10 we conclude with some remarks.

As in Chapter 2, for a subgroup H of G, we set

$$\widetilde{H} = \bigcup_{g \in G} H^g.$$

3.2 Preliminaries

3.2.1 Some elementary considerations

We begin with an easy lemma, part of which is contained in Lemmas 2.3.1 and 2.3.2 in Chapter 2 (however, for convenience we restate it here in full).

Lemma 3.2.1. *Let G be a finite group, X be a subset of G and N be an abelian normal subgroup of G. Let $\pi : G \to G/N$ denote the natural projection.*

(i) $\langle X \rangle_I = G$ if and only if $X \not\subseteq \widetilde{M}$ for every maximal subgroup M of G.

(ii) If G is nilpotent, then $\langle X \rangle_I = G$ if and only if $\langle X \rangle = G$.

(iii) If $\pi(X)$ invariably generates G/N, and $Y \subseteq N$ generates N as a G-module, then $\langle X \cup Y \rangle_I = G$.

Proof. Items (i) and (ii) are contained in Lemmas 2.3.1 and 2.3.2. A proof of item (iii) can be found for instance in [KLS15, Lemma 2.10]. □

We state another lemma, whose proof follows immediately from the definitions.

Lemma 3.2.2. *Let G be a finite group, and let $X = \{x_1, \dots, x_t\}$ be a subset of G. For every i, let C_i be the conjugacy class of G containing x_i. Then, X is a minimal invariable generating set of G if and only if the following conditions are both satisfied:*

(a) There exists a set of maximal subgroups $J = \{M_1, \dots, M_t\}$ of G such that, for every $i \neq j$, $C_i \cap M_j \neq \varnothing$.

(b) *No proper subgroup of G has non-empty intersection with C_i for all $i = 1, \dots, t$.*

In order to bound $m_I(G)$ from above, it is convenient to introduce another invariant, which we denote by $\iota(G)$, as follows.

For every maximal subgroup M of G, denote by M^* the set of G-conjugacy classes having non-empty intersection with M. Let

$$\mathcal{C}(G) = \{M^* \mid M \text{ maximal subgroups of } G\}.$$

We say that a subset $\{X_1, \dots, X_t\}$ of $\mathcal{C}(G)$ is *independent* if, for every $1 \leqslant i \leqslant t$, the intersection $\cap_{j \neq i} X_j$ properly contains $\cap_j X_j$. We denote by $\iota(G)$ the largest cardinality of an independent subset of $\mathcal{C}(G)$.

Lemma 3.2.3. $m_I(G) \leqslant \iota(G)$.

Proof. Let $m = m_I(G)$ and let $\{x_1, \dots, x_m\}$ be a minimal invariable generating set of G. For $1 \leqslant i \leqslant m$ let C_i be the conjugacy class of G containing x_i. For every $1 \leqslant i \leqslant m$, there exists a maximal subgroup M_i of G such that $\{C_1, \dots, C_{i-1}, C_{i+1}, \dots, C_m\} \subseteq M_i^*$ but $C_i \notin M_i^*$. It follows that $\{M_1^*, \dots, M_m^*\}$ is an independent subset of $\mathcal{C}(G)$, and therefore $m \leqslant \iota(G)$. □

3.2.2 Crowns

In the remainder of this section we shall review the notion and the properties of *crowns* in finite soluble groups. (This notion can be given for arbitrary finite groups, see [DL03], but we are interested only in the soluble case.) For more details, see for instance [BBE06, Section 1.3].

Let G be a finite soluble group, and let $\mathcal{V}_G$ be a set of representatives for the irreducible G-groups that are G-isomorphic to a complemented chief factor of G. For $A \in \mathcal{V}_G$ let $\mathrm{R}_G(A)$ be the smallest normal subgroup contained in $\mathrm{C}_G(A)$ with the property that $\mathrm{C}_G(A)/\mathrm{R}_G(A)$ is G-isomorphic to a direct product of copies of A and it has a complement in $G/\mathrm{R}_G(A)$. The factor group $\mathrm{C}_G(A)/\mathrm{R}_G(A)$ is called the A-crown of G, and it is the socle of $G/\mathrm{R}_G(A)$. The positive integer $\delta_G(A)$ defined by $\mathrm{C}_G(A)/\mathrm{R}_G(A) \cong_G A^{\delta_G(A)}$ is called the A-rank of G and it coincides with the number of complemented factors in any chief series of G that are G-isomorphic to A. Moreover $\mathrm{C}_G(A)/\mathrm{R}_G(A)$ is complemented in $G/\mathrm{R}_G(A)$, so that $G/\mathrm{R}_G(A) \cong A^{\delta_G(A)} \rtimes H$ with $H \cong G/\mathrm{C}_G(A)$.

Lemma 3.2.4. *Let G be a finite soluble group with trivial Frattini subgroup. There exist $A \in \mathcal{V}_G$ and a nontrivial normal subgroup U of G such that $\mathrm{C}_G(A) = \mathrm{R}_G(A) \times U$. If G is nonabelian then A can be chosen with the extra property of being a nontrivial G-module.*

Proof. By [BBE06, Lemma 1.3.6], there exists $A \in \mathcal{V}_G$ and a nontrivial normal subgroup U of G such that $\mathrm{C}_G(A) = \mathrm{R}_G(A) \times U$. Assume that A is a trivial G-module. Then $G = \mathrm{C}_G(A) = \mathrm{R}_G(A) \times U$. Write $\mathrm{R}_G(A) = H$, which is nontrivial if G is nonabelian. In this case there exist a crown $\mathrm{C}_H(B)/\mathrm{R}_H(B)$

and a nontrivial normal subgroup W of H such that $\mathrm{C}_H(B) = \mathrm{R}_H(B) \times W$. We have $\mathrm{C}_G(B) = \mathrm{C}_H(B) \times U$ and $\mathrm{R}_G(B) = \mathrm{R}_H(B) \times U$, so $\mathrm{C}_G(B) = \mathrm{R}_G(B) \times W$. This means that we may consider B in place of A. It is possible that also B is a trivial G-module. In that case $G = \mathrm{C}_G(B) = \mathrm{R}_G(B) \times W = \mathrm{R}_H(B) \times U \times W$, and we can repeat the previous argument with H replaced by $\mathrm{R}_H(B)$. Continuing in this way, we obtain a nontrivial irreducible G-module satisfying our statement, except in the case when G is abelian. □

The following lemma will be applied several times. It says that essentially we need to care only of what happens modulo U and modulo $\mathrm{R}_G(A)$.

Lemma 3.2.5. *[Luc18, Lemma 4 and Lemma 12] Assume that G is a finite group with trivial Frattini subgroup and let $C = C_G(A), R = \mathrm{R}_G(A), U$ be as in the statement of Lemma 3.2.4. If $K \leqslant G$ is such that $KU = KR = G$, then $K = G$. In particular, for $g_1, \dots, g_t \in G$ if $\langle g_1 U, \dots, g_t U \rangle_I = G/U$ and $\langle g_1 R, \dots, g_t R \rangle_I = G/R$, then $\langle g_1, \dots, g_t \rangle_I = G$.*

The following represents the main result for dealing with invariable generation of finite soluble groups.

Proposition 3.2.6. *[DL15, Proposition 8] Let K be a finite soluble group and let A be a faithful nontrivial irreducible K-module. We may consider A as a vector space over the field $F = \mathrm{End}_K(A)$. Suppose that $\langle y_1, \dots, y_t \rangle_I = K$. Let δ be a positive integer and let $w_1, \dots, w_t \in A^\delta$ with $w_i = (w_{1,i}, \dots, w_{\delta,i})$. Consider the matrix W whose i-th column is w_i:*

$$W = \begin{pmatrix} w_{1,1} & \cdots & w_{1,t} \\ \vdots & & \vdots \\ w_{\delta,1} & \cdots & w_{\delta,t} \end{pmatrix}.$$

Then $y_1 w_1, \dots, y_t w_t$ invariably generate $A^\delta \rtimes K$ if and only if the rows of W (seen as vectors of A^t) are linearly independent modulo $B = \{(u_1, \dots, u_t) \in A^t \mid u_i \in [y_i, A],\ i = 1, \dots, t\}$.

In particular, there exist elements $w_1, \dots, w_t \in A^\delta$ such that $y_1 w_1, \dots, y_t w_t$ invariably generate $A^\delta \rtimes K$ if and only if

$$\delta \leqslant nt - \dim B = \sum_{i=1}^{t} \dim_F C_A(y_i).$$

We restate this proposition in a slightly different form that will suit better our exposition.

Corollary 3.2.7. *In the notation of the previous proposition, for $1 \leqslant i \leqslant t$ let $A_i = [y_i, A]$ and $B_i = A/A_i$. Again we consider A, A_i, B_i as vector spaces over the field $F = \mathrm{End}_K(A)$. The entries of the i-th column of W may be seen modulo A_i, that is, may be seen as elements of B_i. Let Z denote this new matrix:*

$$Z = \begin{pmatrix} w_{1,1} + A_1 & \cdots & w_{1,t} + A_t \\ \vdots & & \vdots \\ w_{\delta,1} + A_1 & \cdots & w_{\delta,t} + A_t \end{pmatrix} =: \begin{pmatrix} b_{1,1} & \cdots & b_{1,t} \\ \vdots & & \vdots \\ b_{\delta,1} & \cdots & b_{\delta,t} \end{pmatrix}.$$

Then $y_1w_1, \dots, y_tw_t$ invariably generate $A^\delta \rtimes K$ if and only if the rows of Z (seen as vectors of $B_1 \times \cdots \times B_t$) are linearly independent.

3.3 Estimating $m_I(G)$

3.3.1 Soluble groups

Proposition 3.3.1. *Let G be a finite soluble group. There exists a minimal invariable generating set of cardinality $m = m(G)$.*

Proof. We argue by induction on $|G|$. By [Luc13a, Theorem 2], $m(G)$ coincides with the number of non-Frattini factors in a chief series of G. Since $m(G) = m(G/\operatorname{Frat}(G))$, we may assume $\operatorname{Frat}(G) = 1$. Let N be a minimal normal subgroup of G and let H be a complement of N in G. By induction there exist $m(G) - 1$ elements $h_1, \dots, h_{m(G)-1}$ that form a minimal invariable generating set for H. If n is a nontrivial element of N, then by Lemma 3.2.1 $\{h_1, \dots, h_{m(G)-1}, n\}$ is a minimal invariable generating set of G. □

This shows that, in the soluble case, $m(G) \leqslant m_I(G)$. We now prove the other inequality, that is, we prove Theorem 3.1.1. Here we use all preliminaries on crowns introduced in Section 3.2.

Proof of Theorem 3.1.1. In view of Proposition 3.3.1, we only need to show that if $\{x_1, \dots, x_t\}$ is a minimal invariable generating set of G, then $t \leqslant m$. The statement is trivially true if G is nilpotent since, as observed in Lemma 3.2.1, in this case the notion of generation and invariable generation coincide. So we may assume that G is soluble but not nilpotent. We prove our statement by induction on $|G|$. We may assume $\operatorname{Frat}(G) = 1$. Choose a nontrivial G-module $A \in \mathcal{V}_G$ such that $R = \mathrm{R}_G(A), U, C = \mathrm{C}_G(A)$ satisfy the property described in Lemma 3.2.4.

There exists a positive integer δ such that $U \cong_G A^\delta$. By [Luc13a, Theorem 2], $m = m(G)$ coincides with the number of non-Frattini factors in a chief series of G, hence $m = m(G/U) + \delta$. Up to reordering the indices, we may assume that there exists $s \leqslant t$ such that $x_1, \dots, x_s$ is a minimal invariable generating set of G modulo U. By induction $s \leqslant m(G/U) = m - \delta$.

We work now in $\overline{G} = G/R$ and, for every $g \in G$, we set $\overline{g} = gR$. We have $C/R = UR/R \cong U \cong A^\delta$ and $G/R \cong C/R \rtimes H/R$ where $K := H/R$ acts in the same say on each of the δ factors of $C/R \cong A^\delta$ and this action is faithful and irreducible. We may identify $\overline{G}$ with the semidirect product $A^\delta \rtimes K$ and we can write $\bar{x}_i = w_i y_i$ with $w_i \in U = A^\delta$ and $y_i \in K$. Since $\langle x_1U, \dots, x_sU\rangle_I = G/U$ and $K \cong G/C$ is an epimorphic image of G/U, we deduce that $\langle y_1, \dots, y_s\rangle_I = K$.

We want to apply Proposition 3.2.6 and Corollary 3.2.7, and we employ the notations used there. Moreover, for $1 \leqslant k \leqslant t$ we denote by $Z_{\mathrm{rem}(k)}$ the matrix

obtained by Z removing the k-th column:

$$Z_{\mathrm{rem}(k)} = \begin{pmatrix} b_{1,1} & \cdots & b_{1,k-1} & b_{1,k+1} & \cdots & b_{1,t} \\ \vdots & & \vdots & \vdots & & \vdots \\ b_{\delta,1} & \cdots & b_{\delta,k-1} & b_{\delta,k+1} & \cdots & b_{\delta,t} \end{pmatrix} =: \begin{pmatrix} \rho_{1,k} \\ \vdots \\ \rho_{\delta,k} \end{pmatrix}$$

(here the $\rho_{i,k}$ are row vectors; same below with the $\sigma_{i,k}$), and with $Z_{\mathrm{kee}(k)}$ the matrix obtained by Z keeping only the first k columns:

$$Z_{\mathrm{kee}(k)} = \begin{pmatrix} b_{1,1} & \cdots & b_{1,k} \\ \vdots & & \vdots \\ b_{\delta,1} & \cdots & b_{\delta,k} \end{pmatrix} =: \begin{pmatrix} \sigma_{1,k} \\ \vdots \\ \sigma_{\delta,k} \end{pmatrix}.$$

Since $\langle w_1 y_1, \ldots, w_t y_t \rangle_I = \overline{G} \cong A^\delta \rtimes K$, we have that the rows of Z are linearly independent. On the other hand, since $\{x_1, \ldots, x_t\}$ is a minimal invariable generating set of G, if $s < k \leqslant t$ then $\langle x_1, \ldots, x_s, x_{s+1}, \ldots, x_{k-1}, x_{k+1}, \ldots, x_t \rangle_I \neq G$. Therefore, by Lemma 3.2.5 $\langle \bar{x}_1, \ldots, \bar{x}_s, \bar{x}_{s+1}, \ldots, \bar{x}_{j-1}, \bar{x}_{j+1}, \ldots, \bar{x}_t \rangle_I \neq \bar{G}$, and consequently the rows of $Z_{\mathrm{rem}(k)}$ are linearly dependent.

We claim that, for every $s \leqslant k < t$, adding to $Z_{\mathrm{kee}(k)}$ the $(k+1)$-th column increases dimension of the row space, that is,

$$\dim_F \langle \sigma_{1,k}, \ldots, \sigma_{\delta,k} \rangle < \dim_F \langle \sigma_{1,k+1}, \ldots, \sigma_{\delta,k+1} \rangle. \tag{3.3.1}$$

Indeed, if the dimension stays the same then the $(k+1)$-th column is useless, i.e., $\dim_F \langle \sigma_{1,t}, \ldots, \sigma_{\delta,t} \rangle = \dim_F \langle \rho_{1,k+1} \ldots, \rho_{\delta,k+1} \rangle$. But the left-hand side is equal to δ, while the right-hand side is strictly smaller than δ: contradiction. Hence the claim is proved.

Since $\dim_F \langle \sigma_{1,t}, \ldots, \sigma_{\delta,t} \rangle = \delta$ we deduce from (3.3.1) that $t - s \leqslant \delta$, from which $t \leqslant s + \delta \leqslant (m - \delta) + \delta = m$. □

It is easy to find examples of (nonsoluble) groups for which Proposition 3.3.1 fails, namely, examples of groups G for which $m_I(G) < m(G)$. For instance, $m(A_5) = 3$ while $m_I(A_5) = 2$ (this is because any invariable generating set of A_5 must contain an element of order 5 and an element of order 3).

3.3.2 A strengthening of Question 3.1.2

Recall Question 3.1.2, which asks whether it is true that $m_I(G) \leqslant m(G)$ for every finite group G.

One could ask whether the following strengthening of Question 3.1.2 is true: if $\{x_1, \ldots, x_t\}$ is a minimal invariable generating set of G, then there exist $g_1, \ldots, g_t \in G$ such that $\{x_1^{g_1}, \ldots, x_m^{g_t}\}$ is a minimal generating set of G. Although we are not able to exhibit a soluble counterexample, the following shows that the statement is not true in general.

Lemma 3.3.2. *Let $G = A_{29}$ and consider the following three elements:*

$$\begin{aligned} a =&(2,3,4)(5,6,7)(8,9,\dots,18)(19,20,\dots,29),\\ b =&(1,2)(3,4)(5,6,\dots,29),\\ c =&(1,2)(3,4,\dots,8)(9,10,\dots,29). \end{aligned}$$

The set $\{a,b,c\}$ is a minimal invariable generating set of G, but for every $x,y,z \in G$, $\{a^x, b^y, c^z\}$ is not a minimal generating set.

Proof. It can be easily seen that no proper subgroup of A_{29} contains conjugates of a, b and c, so $\langle a,b,c\rangle_I = A_{29}$. On the other hand $\langle a,b\rangle$ stabilizes $\{1,2,3,4\}$, $\langle b,c\rangle$ stabilizes $\{1,2\}$ and $\langle a^{(2,8)}, c\rangle$ stabilizes $\{3,4,5,6,7,8\}$ so $\{a,b,c\}$ is a minimal invariable generating set of A_{29}. Now we want to show that, in any way we conjugate a,b,c, two elements are sufficient in order to generate A_{29}. Without loss of generality we may assume that one of this conjugates is a. Let $x,y \in A_{29}$. If $\langle a, b^x\rangle \neq A_{29}$, then $\langle a, b^x\rangle$ stabilizes either $\{1,2,3,4\}$ (in which case $\{1,2,3,4\}$ is mapped into itself by x) or $\{1,5,6,7\}$ (in which case $\{1,2,3,4\}$ is mapped to $\{1,5,6,7\}$ by x). Without loss of generality we may assume that a,b and x stabilize $\{1,2,3,4\}$. If $\langle a, c^y\rangle \neq A_{29}$, then it stabilizes $\{2,3,4,5,6,7\}$ (and the 6-cycle in the decomposition of c^y permutes the elements of this subset). But then $\langle b^x, c^y\rangle = A_{29}$. Indeed two conjugates of b and c either generate A_{29} or stabilize the same subset of cardinality 2. But this second possibility does not occur for b^x and c^y, indeed the support of the 2-cycle in the decomposition of b^x is contained in $\{1,2,3,4\}$ while the support of the 2-cycle in the decomposition of c^y must be disjoint from $\{2,3,4,5,6,7\}$. □

If G is a finite group, then $d_I(G) \geqslant d(G)$ and the difference $d_I(G) - d(G)$ can be arbitrarily large. [KLS11, Proposition 2.5] states that, for every $r \geqslant 1$, there is a finite group G such that $d(G) = 2$ but $d_I(G) \geqslant r$. We do not know whether the (somewhat opposite) inequality $m(G) \geqslant m_I(G)$ is true, but in any case we may exhibit examples in which the difference $m(G) - m_I(G)$ is arbitrarily large.

A first example is given by the symmetric group S_n, in which case $m(S_n) - m_I(S_n) \to \infty$ as $n \to \infty$. We refer to Chapter 5, specifically Subsection 5.1.3, for the proof of this fact (which follows quickly from Liebeck–Shalev [LS96]), and for other results regarding $m_I(S_n)$.

We note that [LS96] uses CFSG. In the next section we will give a more elementary example showing that $m(G) - m_I(G)$ can be arbitrarily large. With this purpose, we recall that in [Luc13b] it is noticed that $m(A \times B) = m(A) + m(B)$ for every pair of finite groups A and B.

Question 3.3.3. *Is it true that $m_I(A) + m_I(B) = m_I(A \times B)$ for every pair (A,B) of finite groups?*

It is easy to see that the inequality $m_I(A)+m_I(B) \leqslant m_I(A\times B)$ always holds. Indeed, if $\{a_1,\dots,a_r\}$ is a minimal invariable generating set of A and $\{b_1,\dots,b_s\}$ is a minimal invariable generating set of B, then $\{(a_1,1),\dots,(a_r,1),(1,b_1),\dots,(1,b_s)\}$ is a minimal invariable generating set of $A \times B$. Regarding the equality, we are only able to prove a very partial result.

Proposition 3.3.4. *Assume that A and B are finite groups without common composition factors. Then $m_I(A \times B) = m_I(A) + m_I(B)$.*

Proof. Assume that $g_1 = (a_1, b_1), \ldots, g_m = (a_m, b_m)$ is an invariable generating set of $G = A \times B$. There exists $I \subseteq \{1, \ldots, m\}$ such that $\{a_i \mid i \in I\}$ is a minimal invariable generating set for A and $J \subseteq \{1, \ldots, m\}$ such that $\{b_j \mid j \in J\}$ is a minimal invariable generating set for B. Then $\{(a_k, b_k) \mid k \in I \cup J\}$ is an invariable generating set for $A \times B$. So $m_I(A \times B) \leqslant m_I(A) + m_I(B)$. □

3.4 An example: $m_I(A_5^n)$

Since $m(A_5) = 3$ and $m(A \times B) = m(A) + m(B)$ for every pair of finite groups A and B, we have $m(A_5^n) = 3n$. We are going to show that $m(A_5^n) - m_I(A_5^n) \to \infty$ as $n \to \infty$. Indeed we shall prove:

Proposition 3.4.1. $m_I(A_5^n) = n \cdot m_I(A_5) = 2n$.

Notice first that, by what we said in the previous section, $n \cdot m_I(A_5) \leqslant m_I(A_5^n)$, so it suffices to show $m_I(A_5^n) \leqslant 2n$. In fact we shall bound $\iota(A_5^n)$, which is enough in view of Lemma 3.2.3. Recall the notation $\mathcal{C}(G)$ introduced before Lemma 3.2.3.

Proposition 3.4.2. $\iota(A_5^n) \leqslant 2n$.

Proof. We have 5 conjugacy classes C_1, C_2, C_3, C_4, C_5 in A_5 with representatives $1, (1,2)(3,4), (1,2,3), (1,2,3,4,5), (1,5,4,3,2)$. Notice that $C_1 = C_1^{-1}$, $C_2 = C_2^{-1}$, $C_3 = C_3^{-1}$, while $C_5 = C_4^{-1}$ and $C_4 = C_5^{-1}$. Moreover a maximal subgroup of A_5 is isomorphic to A_4, S_3 or D_{10} and $\mathcal{C}(A_5)$ contains only two elements: $Y_1 = \{C_1, C_2, C_3\}$ and $Y_2 = \{C_1, C_2, C_4, C_5\}$. Let $\Omega = \{C_1, C_2, C_3, C_4, C_5\}$, $\Omega^* = \{C_1, C_3, C_4, C_5\}$, $\Delta = \Omega^n$ and $\Delta^* = (\Omega^*)^n$. Notice that we are identifying the elements of Δ with the conjugacy classes of A_5^n.

Let $G = A_5^n$. A maximal subgroup M of G can be of two different kinds:

(1) there exist $1 \leqslant i \leqslant n$ and a maximal subgroup Y of A_5 such that $(x_1, \ldots, x_n) \in M$ if and only if $x_i \in Y$ (*product type*).

(2) there exist $1 \leqslant i < j \leqslant n$ and $\phi \in \mathrm{Aut}(A_5)$ such that $(x_1, \ldots, x_n) \in M$ if and only if $x_j = x_i^{\phi}$ (*diagonal type*).

As a consequence, the elements of $\mathcal{C}(G)$ are of the following kinds:

(1) $A_i = \{(\omega_1, \ldots, \omega_n) \in \Delta \mid \omega_i \in Y_1\}$,

(2) $B_i = \{(\omega_1, \ldots, \omega_n) \in \Delta \mid \omega_i \in Y_2\}$,

(3) $C_{i,j} = \{(\omega_1, \ldots, \omega_n) \in \Delta \mid \omega_i = \omega_j\}$,

(4) $D_{i,j} = \{(\omega_1, \ldots, \omega_n) \in \Delta \mid \omega_i = \omega_j^{-1}\}$.

We now assume that $\{X_1, \dots, X_t\}$ is an independent subset of $\mathcal{C}(G)$ and we set $\Delta_i = X_1 \cap \cdots \cap X_i$, $\Delta_i^* = \Delta_i \cap \Delta^*$. Moreover let Λ_i be the set of the $j \in \{1, \dots, n\}$ such that $\omega_j \notin \{C_4, C_5\}$ for every $(\omega_1, \dots, \omega_n) \in \Delta_i$.

We may assume that there exist a, b such that

- $\diamond$ If $i \leqslant a$ then there exists $I_i = (r_i, s_i)$ such that $X_i \in \{C_{r_i,s_i}, D_{r_i,s_i}\}$.
- $\diamond$ If $a < i \leqslant a + b$ then $X_i = A_r$ for some r.
- $\diamond$ If $a + b < i$ then $X_i = B_r$ for some r.

For $i \leqslant a$, let ρ_i be the smallest equivalence relation on $\{1, \dots, n\}$ containing all the pairs (r_j, s_j) with $j \leqslant i$. We may assume, up to reordering the indices, that there exists $a_1 \leqslant a$ such that for every $2 \leqslant i \leqslant a_1$ the relation ρ_{i-1} is finer than ρ_i, while $\rho_i = \rho_{a_1}$ if $i > a_1$. We can describe how Δ_{a_1} looks like. Assume that $B_1, \dots, B_l$ are the equivalence classes of the relation ρ_{a_1}. Then Δ_{a_1} is a product of l "diagonal subsets", each of cardinality 5: if $i_1, i_2 \in B_j$ for some $1 \leqslant j \leqslant l$, then there exists $\epsilon_{i_1,i_2} = \pm 1$ such that $\omega_{i_2} = \omega_{i_1}^{\epsilon_{i_1,i_2}}$ for every $(\omega_1, \dots, \omega_n) \in \Delta_{a_1}$. In particular, since $l \leqslant n - a_1$, we have

$$|\Delta_{a_1}| = 5^l \leqslant 5^{n-a_1} \text{ and } |\Delta_{a_1}^*| \leqslant 4^{n-a_1}.$$

Now assume $a_1 < i \leqslant a$. There exists an equivalence class B_j of ρ_{a_1} containing r_i and s_i and $\eta = \pm 1$, that $\omega_{s_i} = \omega_{r_i}^{\eta}$ for every $(\omega_1, \dots, \omega_n) \in X_i$. As we noticed above, there already exists $\epsilon = \epsilon_{r_i,s_i}$ such that $\omega_{s_i} = \omega_{r_i}^{\epsilon}$ for every $(\omega_1, \dots, \omega_n) \in \Delta_{a_1}$. We must have $\eta = -\epsilon$ (otherwise $\Delta_{a_1} \cap X_i = \Delta_{a_1}$), and consequently $\omega_{s_i} = \omega_{r_i} = \omega_{r_i}^{-1}$, (i.e. $\omega_{r_i} \notin \{C_4, C_5\}$) for every $(\omega_1, \dots, \omega_n) \in \Delta_i$. In particular

$$|\Delta_i^*| \leqslant \frac{|\Delta_{i-1}^*|}{2}.$$

Notice also that $|\Lambda_{a_1}| = 0$ and $|\Lambda_a| \geqslant a_2$, where we set $a_2 = a - a_1$.

Now assume $a < i \leqslant a+b$: again when we consider the intersection $\Delta_{i-1} \cap X_i$ we add the restriction that ω_i cannot belong to $\{C_4, C_5\}$, so $i \notin \Lambda_a$ (otherwise $\Delta_a \cap X_i = \Delta_{a_1}$) and

$$|\Delta_i^*| \leqslant \frac{|\Delta_{i-1}^*|}{2}.$$

Moreover $|\Lambda_{a+b}| \geqslant a_2 + b$.

Finally let $a + b < i$. We may assume that there exists c_1 such that $X_i = B_r$ with $r \in \Lambda_{a+b}$ if and only if $i \leqslant a + b + c_1$. If $a + b < i \leqslant a + b + c_1$, then

$$|\Delta_i^*| \leqslant \frac{|\Delta_{i-1}^*|}{2}.$$

We must have

$$1 \leqslant |\Delta_{a+b+c_1}^*| \leqslant \frac{4^n \cdot 3^{c_2}}{4^{a_1} \cdot 2^{a_2+b+c_1} \cdot 4^{c_2}}$$

and consequently $2a_1+a_2+b+c_1 \leqslant 2n$. Set $c_2 = c-c_1$. Notice that $a_2+b+c_2 \leqslant n$ (since $c_2 \leqslant |\{1, \dots, n\} \setminus \Lambda_{a+b}| \leqslant n - a_2 - b$) and $c_1 + c_2 \leqslant n$ (since there are at most n maximal subgroups of kind B_r), hence $2c_2 + a_2 + b + c_1 \leqslant 2n$. But then $2t = (2a_1 + a_2 + b + c_1) + (2c_2 + a_2 + b + c_1) \leqslant 4n$, from which $t \leqslant 2n$. $\square$

3.5 $m_I(G) = m(G)$ with G nonabelian simple

In this section we will exhibit infinitely many nonabelian finite simple groups G for which $m_I(G) = m(G)$ holds.

Proposition 3.5.1. *Assume p is a prime such that the following conditions are both satisfied: $p \equiv 1$ mod 40 and $p \equiv 2$ mod 3. Then $m_I(\mathrm{PSL}_2(p)) = m(\mathrm{PSL}_2(p)) = 3$.*

Notice that there exist infinitely many primes p satisfying the conditions in the statement. Indeed, every prime $p \equiv 41 \bmod 120$ satisfies them, and there exist infinitely many such primes by Dirichlet's theorem on arithmetic progressions. We remark that, with analogous proof, the statement holds also for $p \equiv -1 \bmod 40$ and $p \equiv 1 \bmod 3$. We proceed with the proof of the proposition.

Proof. In [Jam13] it was shown that, for $p > 31$, $m(\mathrm{PSL}_2(p)) = 3$, hence it remains to prove $m_I(\mathrm{PSL}_2(p)) = 3$. Let $G_p = \mathrm{PSL}_2(p)$. The subgroup structure of this group is well known, and we refer the reader to [Suz82, Chapter 3, Section 6] for detailed information. In particular, the condition $p \equiv 1 \bmod 40$ implies that the isomorphism classes of maximal subgroups of G_p are exactly the following: dihedral groups D_{p-1} and D_{p+1} of order $p-1$ and $p+1$, a Borel subgroup B of order $p(p-1)/2$, $H = A_5$ and $K = S_4$.

Consider $X = \{x, y, z\}$ where $|x| = 3$, $|y| = 4$ and $|z| = 5$. No proper subgroup of G_p contains elements of order 3, 4 and 5, hence X is an invariable generating set (the conditions on p imply that, while B and D_{p-1} contain elements of order 4 and 5, they do not contain elements of order 3). Moreover, in G_p every element of order coprime to p can be conjugate inside a fixed dihedral group, hence whenever $|a| = |b| \geqslant 3$ and $\gcd(p, |a|) = 1$, we have $a^{G_p} \cap \langle b \rangle \neq \varnothing$. Then, order considerations imply that any two elements of X can be conjugate inside a suitable maximal subgroup of G_p. This shows that $m_I(G_p) \geqslant 3$.

For the other inequality, we will show $\iota(G_p) \leqslant 3$, so that $m_I(G_p) \leqslant 3$ by Lemma 3.2.3. All subgroups isomorphic to B are conjugate, and all involutions are conjugate, hence $\mathcal{C}(G_p)$ consists of $D_{p-1}^*, D_{p+1}^*, B^*, H^*, K^*$. We have that $B^* \cap D_{p+1}^* = D_{p-1}^* \cap D_{p+1}^*$ is the conjugacy class of involutions, which belongs to every member of the list. Moreover, $D_{p-1}^* \subseteq B^*$. This easily implies $\iota(G_p) \leqslant 3$. □

3.6 The Tarski irredundant basis theorem

A nice result in universal algebra, due to Tarski and known with the name of Tarski irredundant basis theorem (see [Tar75], or [SB81, Theorem 4.4]), implies that, for every positive integer k with $d(G) \leqslant k \leqslant m(G)$, G contains a minimal generating set of cardinality k. A natural question is whether there exists a similar result for the invariable generation. Tarski's theorem relies on an elementary but clever counting argument which is quite flexible and can be

adapted to several different situations. However, as we shall see in this section, using this argument we are able to obtain only a weak and partial result. In order to see the problems in applying Tarski irredundant basis theorem to the invariable generation, we find it is interesting to sketch the proof of this partial result.

3.6.1 Trying to use Tarski's argument

Tarski's theorem is based on the notion of closure operator ([SB81, Definition 5.1]), which is a function C, from and to subsets of G, such that $X \subseteq C(X)$, $C(Y) \subseteq C(X)$ if $Y \subseteq X$, and $C(C(X)) = C(X)$. In case of generation, one defines $C(X) = \langle X \rangle$. For the argument, it is important that $C(X) = G$ if and only if X generates G (this is obviously true in the case when we define $C(X) = \langle X \rangle$). We should have this property also in the case of invariable generation. If $X = \{x_1, \ldots, x_t\}$, the first definition that comes to mind is then

$$C(X) = X \cup \left(\bigcap_{(g_1,\ldots,g_t)\in G^t} \langle x_1^{g_1}, \ldots, x_t^{g_t} \rangle \right).$$

Artificially, we have imposed $X \subseteq C(X)$, and monotonicity is immediate. What is not immediate from the definition, but straightforward to check, is that C is also idempotent. Moreover, it is not difficult to show that $C(X) = G$ if and only if $\langle X \rangle_I = G$. Therefore we have a closure operator, and we may be on the right track.

Now if we define, for $n, k \geqslant 1$,

$$C_n(X) = \bigcup_{Y\subseteq X, |Y|\leqslant n} C(Y), \quad C_n^1(X) = C_n(X), \quad C_n^{k+1}(X) = C_n(C_n^k(X)),$$

following [SB81] we may call a finite group G invariable n-ary if $C(X) = \bigcup_{i\in\mathbf{N}} C_n^i(X)$ for every subset X of G. Using this notion, it is possible to bound the "gap" that can occur between minimal invariable generating sets. More precisely, if we denote by $\mathrm{IrrB}_I(G)$ the set of the positive integers n such that G has a minimal invariable generating set of size n, we have the following

Theorem 3.6.1. *Let G be an invariable n-ary finite group, with $n \geqslant 2$. If $i < j$ with $i, j \in \mathrm{IrrB}_I(G)$ such that $\{i+1, \ldots, j-1\} \cap \mathrm{IrrB}_I(G) = \varnothing$, then $j - i \leqslant n-1$.*

Proof. Follows from the proof of [SB81, Theorem 4.4]. □

Corollary 3.6.2. *If G is an invariable 2-ary finite group then, for every $d_I(G) \leqslant k \leqslant m_I(G)$, there exists a minimal invariable generating set of size k.*

Notice that a finite group G is invariable 2-ary if the following holds: for every $X \subseteq G$, if $C_2(X) = X$ then $C(X) = X$.

We see some problems in this approach. The first is that, although Theorem 3.6.1 does give a bound, we are not able to give a structural interpretation of

the property of being invariable n-ary. Moreover, in case of nilpotent groups the closure operators defined for generation and for invariable generation need not coincide (remember that, instead, the notions of generation and invariable generation do coincide). Finally, the following result shows that Theorem 3.6.1 cannot give any absolute bound.

Lemma 3.6.3. *For every integer $n \geqslant 2$, there exists a finite group G which is not invariable n-ary.*

Proof. Let $n \geqslant 2$. Assume we prove that there exists a finite group G with the following property: $d_I(G) \geqslant n+1$ and there exists $g \in G$ that does not lie in any proper normal subgroup of G. Then, if we set $X = G \setminus \{g\}$, we have that $\langle X \rangle_I = G$ (since $|G \setminus \widetilde{M}| \geqslant |M|$ for every proper subgroup M of G). Hence $C(X) = G \not\subseteq X$. On the other hand, for every $x_1, \ldots, x_n \in G$, $N := \bigcap_{g_i \in G} \langle x_1^{g_1}, \ldots, x_n^{g_n} \rangle$ is a proper normal subgroup of G, hence $g \notin N$. This shows that $C_n(X) \subseteq X$, from which it follows that G is not invariable n-ary.

We are left to exhibit a group with the property described above. For a supersoluble example, consider $G = P \rtimes Q$, where $P \cong C_p^n$ and $Q \cong C_q$ for primes p and q, with q dividing $p-1$, and Q acts on each copy of C_p as multiplication in the field $\mathbf{F}_p$. It can be easily seen that $d_I(G) = n+1$. Moreover every proper normal subgroup of G is contained in P, so we can take in the role of g any element of $G \setminus P$. □

Summarizing, in order to apply Tarski irredundant basis theorem to the invariable generation we would need to define a closure operator C on the set of subsets of G with the following two properties:

(1) $C(X) = G$ if and only if $\langle X \rangle_I = G$,

(2) $C(X) = \bigcup_{i \in \mathbf{N}} C_2^i(X)$ for every subset X of G.

We are not able to find a closure operation satisfying (1) different from the one introduced above. However this fails property (2). So the question of extending Tarski's theorem to the invariable generation remains open (see Question 3.1.4).

3.6.2 Proof of Theorem 3.1.3

We now show that we can extend Tarski's theorem in case of soluble groups. In other words, we prove Theorem 3.1.3.

Proof of Theorem 3.1.3. We need to show that, if $\{x_1, \ldots, x_t\}$ is a minimal invariable generating set of G, with $t < m := m(G)$, then there exists a minimal invariable generating set of G of cardinality $t+1$.

The beginning of the proof is very similar to that of Theorem 3.1.1 given in Section 3.3. As we did there, we prove our statement by induction on $|G|$. We may assume that G is soluble but not nilpotent, the statement for nilpotent groups being easy to check (without applying Tarski's theorem). Again we may assume $\mathrm{Frat}(G) = 1$, and we choose a nontrivial G-module $A \in \mathcal{V}_G$ such that

$R = \mathrm{R}_G(A), U, C = \mathrm{C}_G(A)$ satisfy the property described in Lemma 3.2.4. We let δ be such that $U \cong_G A^\delta$. By [Luc13a, Theorem 2], $m = m(G/U) + \delta$. Up to reordering the indices, we may choose $s \leqslant t$ such that $x_1, \dots, x_s$ is a minimal invariable generating set of G modulo U.

In addition to what done in the proof of Theorem 3.1.1, we further choose, as we may, a complement H of U in G with $R \leqslant H$. For $1 \leqslant i \leqslant t$, we write $x_i = w_i h_i$ with $w_i \in U$ and $h_i \in H$.

We work in $\overline{G} = G/R$ and, for every $g \in G$, we set $\overline{g} = gR$. We may identify $\overline{G}$ with the semidirect product $A^\delta \rtimes K$ with $K = H/R$. Since $\langle x_1 U, \dots, x_s U \rangle_I = G/U$ and $K \cong G/C$ is an epimorphic image of G/U, we deduce that $\langle \overline{x}_1, \dots, \overline{x}_s \rangle_I = K$.

As in the proof of Theorem 3.1.1, we want to apply Proposition 3.2.6 and Corollary 3.2.7, and we employ the notations used there. A technical, but important, step here is that we pick complements for A_i in A. Namely, for $1 \leqslant i \leqslant t$ we decompose A as $A = A_i \oplus C_i$. Here the C_i play the role of the B_i in Theorem 3.1.1 and Corollary 3.2.7, with the advantage (not apparent yet) that they are subspaces of A. Then, if we denote by $c_{j,i}$ the projection of $w_{j,i} \in A$ onto C_i, the matrix Z of Corollary 3.2.7 is replaced by

$$Z = \begin{pmatrix} c_{1,1} & \cdots & c_{1,t} \\ \vdots & & \vdots \\ c_{\delta,1} & \cdots & c_{\delta,t} \end{pmatrix}$$

Moreover, for $1 \leqslant k \leqslant t$ the matrices $Z_{\mathrm{rem}(k)}$ and $Z_{\mathrm{kee}(k)}$ of Theorem 3.1.1 become here

$$Z_{\mathrm{rem}(k)} = \begin{pmatrix} c_{1,1} & \cdots & c_{1,k-1} & c_{1,k+1} & \cdots & c_{1,t} \\ \vdots & & \vdots & \vdots & & \vdots \\ c_{\delta,1} & \cdots & c_{\delta,k-1} & c_{\delta,k+1} & \cdots & c_{\delta,t} \end{pmatrix}$$

$$Z_{\mathrm{kee}(k)} = \begin{pmatrix} c_{1,1} & \cdots & c_{1,k} \\ \vdots & & \vdots \\ c_{\delta,1} & \cdots & c_{\delta,k} \end{pmatrix}$$

As in Theorem 3.1.1, the rows of Z (seen as vectors of $C_1 \times \cdots \times C_t \leqslant A^t$) are linearly independent, while, for every $s < k \leqslant t$, the rows of $Z_{\mathrm{rem}(k)}$ are linearly dependent.

One observation. The definition of C_i, hence of the $c_{j,i}$, depend upon the element h_i. Then, once the $c_{j,i}$ have been defined, the h_i have somewhat done their work (concerning generation modulo R), and we do not need to care about them anymore. Indeed, we only need to care of linear dependence of the rows of Z inside $C_1 \times \cdots \times C_t$ or, equivalently, inside A^t. Then also the C_i are not important anymore. This gives the possibility to suitably modify, to "clean" in some sense, the elements x_i without affecting their property of invariable generation.

As a first example, if we denote by $\widetilde{w}_i \in A^\delta \cong U$ the i-th column of Z, we may replace x_i with

$$\widetilde{x}_i = \begin{cases} \widetilde{w}_i h_i & \text{if } i \leqslant s, \\ \widetilde{w}_i & \text{otherwise.} \end{cases}$$

It is easy to see that $\{\widetilde{x}_1, \dots \widetilde{x}_t\}$ is a minimal invariable generating set of G. Indeed, by the choice of s the set invariably generates modulo U and it is not possible to remove one among the first s elements. On the other hand, all the considerations regarding the invariable generation modulo R are not affected, because they concern linear dependence of the rows of Z, and this matrix does not change in passing from x_i to $\widetilde{x}_i$. This implies that $\{\widetilde{x}_1, \dots \widetilde{x}_t\}$ invariably generates G minimally.

Now choose a subset $J = \{i_1, \dots, i_u\}$ of $\{1, \dots, s\}$ minimal with respect to the property that the δ vectors

$$(c_{j,i_1}, \dots, c_{j,i_u}, c_{j,s+1}, \dots, c_{j,t})$$

for $1 \leqslant j \leqslant \delta$ are linearly independent. The arguments applied in the previous paragraph imply that we still obtain a minimal invariable generating set if we replace $\widetilde{x}_j = \widetilde{w}_j h_j$ with h_j for every $j \in \{1, \dots, s\} \setminus J$. So from now on we will assume $\widetilde{w}_j = 0$ for every $j \in \{1, \dots, s\} \setminus J$.

If $J \neq \varnothing$, then we obtain a minimal invariable generating set of size $t+1$ by replacing $\widetilde{x}_{i_1} = \widetilde{w}_{i_1} h_{i_1}$ with the two elements $\widetilde{w}_{i_1}$ and h_{i_1}.

So we may assume $J = \varnothing$, from which it follows that the first s columns of Z are zero. For convenience, we may remove from the matrix Z such columns. The rank of the matrix clearly does not change. We call the matrix obtained in this way again Z.

$$Z = \begin{pmatrix} c_{1,s+1} & \cdots & c_{1,t} \\ \vdots & & \vdots \\ c_{\delta,s+1} & \cdots & c_{\delta,t} \end{pmatrix}$$

For $s < k \leqslant t$, we remove the first s columns in $Z_{\mathrm{rem}(k)}$ and $Z_{\mathrm{kee}(k)}$. Again, the rows of Z are linearly independent (i.e. $\operatorname{rank} Z = \delta$), while for $s < k \leqslant t$ the rows of $Z_{\mathrm{rem}(k)}$ are linearly dependent (i.e. $\operatorname{rank} Z_{\mathrm{rem}(k)} < \delta$).

Now the same argument as in the end of Theorem 3.1.1 shows that for $s \leqslant k < t$, $\operatorname{rank} Z_{\mathrm{kee}(k)} < \operatorname{rank} Z_{\mathrm{kee}(k+1)}$. Let

$$n_1 = \operatorname{rank} Z_{\mathrm{kee}(s+1)},\ n_2 = \operatorname{rank} Z_{\mathrm{kee}(s+2)} - \operatorname{rank} Z_{\mathrm{kee}(s+1)}, \dots,$$

$$n_{t-s} = \operatorname{rank} Z_{\mathrm{kee}(t)} - \operatorname{rank} Z_{\mathrm{kee}(t-1)}.$$

Notice that $n_1 + \cdots + n_{t-s} = \delta$, and $1 \leqslant n_i \leqslant n$ for every i. Let now $F = \mathrm{End}_H(A)$ and $n = \dim_F A$. Fixing a basis for A as an F-vector space, we may identify each element of A as a vector of F^n. Denote by e_i the vector of F^n all of whose entries are 0, expect the i-th which is 1, and consider the block

matrix

$$Y = \begin{pmatrix} e_1 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ e_{n_1} & 0 & \cdots & 0 \\ 0 & e_1 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & e_{n_2} & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & e_1 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & e_{t-s} \end{pmatrix}$$

Using the definition of the e_i, it is easy to check that we still obtain a minimal invariable generating set if we replace Z with Y. More precisely, if we consider the i-th column of Y as an element $\widetilde{y}_{s+i}$ of $U \cong A^\delta$, we get that $\{\widetilde{x}_1, \dots, \widetilde{x}_s, \widetilde{y}_{s+1}, \dots, \widetilde{y}_t\}$ is a minimal invariable generating set of G.

Assume first that there exists $i \in \{1, \dots, s-t\}$ with $n_i > 1$. Then $\widetilde{y}_{s+i} \in A^\delta$ has at least two nonzero entries, and it suffices to split $\widetilde{y}_{s+i}$ in two vectors: if we define

$$\widetilde{z}_1 = (0, \dots, 0, e_1, \dots, e_{n_i-1}, 0, 0, \dots 0),$$
$$\widetilde{z}_2 = (0, \dots, 0,\ 0\ , \dots,\ 0\ , e_{n_i}, 0, \dots 0),$$

then the set $\{\widetilde{x}_1, \dots, \widetilde{x}_s, \widetilde{y}_{s+1}, \dots, \widetilde{y}_t\} \cup \{\widetilde{z}_1, \widetilde{z}_2\} \setminus \{\widetilde{y}_{s+i}\}$ is a minimal invariable generating set of size $t+1$.

Assume finally $n_i = 1$ for every $i \in \{1, \dots, t-s\}$. In this case $t - s = \delta$. Since $t < m = m(H) + \delta$, we get $s < m(H)$. Then, by induction, there exists a minimal invariable generating set $\{\widetilde{k}_1, \dots, \widetilde{k}_{s+1}\}$ of H of cardinality $s+1$. It follows that $\{\widetilde{k}_1, \dots, \widetilde{k}_{s+1}, \widetilde{x}_{s+1}, \dots, \widetilde{x}_t\}$ is a minimal invariable generating set of G of cardinality $t+1$. □

3.7 The invariable Frattini

The Frattini subgroup $\mathrm{Frat}(G)$ of a finite group G is defined as the intersection of all maximal subgroups of G. An important feature of this subgroup is that it coincides with the elements of G that are useless in generating G. More precisely, $\mathrm{Frat}(G)$ coincides with the set of elements of G that can be dropped from every generating set of G (without compromising generation). This feature implies that the generation properties of G are essentially the same as those of $G/\mathrm{Frat}(G)$. Therefore, if we are interested in generation we can factor out $\mathrm{Frat}(G)$ with no harm. This considerably simplifies the situation, since the structure of Frattini-free groups is much more transparent than that of general groups (at least for soluble groups: think of how many times we applied Lemma 3.2.4 and Lemma 3.2.5).

Here we shall define the analogue of the Frattini subgroup from the point of view of the invariable generation. This will allow us to properly state the results of Section 3.8.

Consider the set $\Sigma = \Sigma(G)$ of all maximal members of the set of all $\widetilde{H}$, where H varies among the proper subgroups of G. Set $\mathrm{Frat}_I(G) = \bigcap_{\widetilde{M}\in\Sigma} \widetilde{M}$.

Lemma 3.7.1. $\mathrm{Frat}_I(G)$ *coincides with the set of elements of* G *that can be dropped from every invariable generating set.*

Proof. Assume $x \in \mathrm{Frat}_I(G)$ and assume $\{x\} \cup X$ invariably generates G for some set X. If X does not invariably generate G then $X \subseteq \widetilde{M}$ for some $\widetilde{M} \in \Sigma$, hence $\{x\}\cup X \subseteq \widetilde{M}$, against the assumption of invariable generation. Conversely, assume $x \notin \mathrm{Frat}_I(G)$: choose $\widetilde{M}$ such that $x \notin \widetilde{M}$. Then, by the maximality of $\widetilde{M}$ it follows that $\langle \{x\} \cup \widetilde{M}\rangle_I = G$, and clearly x cannot be omitted from this invariable generating set. □

By the previous lemma, $\mathrm{Frat}_I(G)$ plays, for the invariable generation, the same role played by the Frattini subgroup for the usual generation. Unfortunately $\mathrm{Frat}_I(G)$ needs not be a subgroup. For instance, if $G = A_5$ then $\mathrm{Frat}_I(G)$ is the set of all involutions of G — hence it generates G.

Notice that if $\widetilde{K} \in \Sigma(G)$, then K is a maximal subgroup of G and clearly if M is a maximal subgroup of G, then there exists a maximal subgroup K of G such that $\widetilde{K} \in \Sigma(G)$ and $M \subseteq \widetilde{M} \subseteq \widetilde{K}$, hence, by definition, $\mathrm{Frat}(G) \subseteq \mathrm{Frat}_I G$. This, if we want, is the reason why we can factor out $\mathrm{Frat}(G)$ also in the invariable setting.

Notice that $\mathrm{Frat}_I(G)$ is defined in a strange manner. Indeed, we do not intersect the $\widetilde{M}$'s for M running among all maximal subgroups of G; we take instead only the maximal sets among the $\widetilde{M}$'s. This is important for the proof of Lemma 3.7.1. However, we do not know whether this is really necessary, and we propose the following

Question 3.7.2. *For a finite group* G*, does* $\mathrm{Frat}_I(G)$ *coincide with the intersection of all* $\widetilde{M}$*, where* M *runs among all maximal subgroups of* G*?*

What we do know is that the two concepts are different a priori, meaning that there may exist maximal subgroups M_1 and M_2 such that $\widetilde{M_1}$ is properly contained in $\widetilde{M_2}$. For example in $G = A_6$ one can consider $M_1 \cong S_4$ and $M_2 \cong 3^2 : 4 \cong (S_3 \wr S_2) \cap A_6$. Then $\widetilde{M_2}$ is the set of the elements of G of order different from 5, while $\widetilde{M_1}$ does not contain elements of order 5 and moreover contains only one of the two conjugacy classes of elements of order 3. Hence $\widetilde{M_1} \subsetneq \widetilde{M_2}$. Nevertheless, once again this phenomenon cannot occur in the soluble world.

Lemma 3.7.3. *Assume that* G *is a finite soluble group and let* M_1, M_2 *be two maximal subgroups of* G*. If* $\widetilde{M_1} \subseteq \widetilde{M_2}$*, then* $\widetilde{M_1} = \widetilde{M_2}$*.*

Proof. We prove the statement by induction on the order of G. We may assume $\mathrm{Frat}(G) = 1$. Choose a nontrivial G-module $A \in \mathcal{V}_G$ such that $R =$

$\mathrm{R}_G(A), U, C = \mathrm{C}_G(A)$ satisfy the property described in Lemma 3.2.4. We further choose a complement H of U in G with $R \leqslant H$. We denote by $\mathcal{M}_1$ the set of the maximal subgroups of G containing U and by $\mathcal{M}_2$ the set of the maximal subgroups of G supplementing U. If $M \in \mathcal{M}_2$ then, by Lemma 3.2.5, $R \subseteq M$ and $M = WH^u$ with W a maximal H-submodule of U and $u \in U$. Assume now $\widetilde{M_1} \subseteq \widetilde{M_2}$. We consider the different cases:

(1) $M_1, M_2 \in \mathcal{M}_1$. In this case $\widetilde{M_1/U} \subseteq \widetilde{M_2/U}$, so by induction $\widetilde{M_1/U} = \widetilde{M_2/U}$, and consequently $\widetilde{M_1} = \widetilde{M_2}$.

(2) $M_1, M_2 \in \mathcal{M}_2$. We have $M_1 = W_1H^{u_1}$ and $M_2 = W_2H^{u_2}$. If $W_1 \neq W_2$, then $\langle M_1^{g_1}, M_2^{g_2} \rangle = G$ for every $g_1, g_2 \in G$, hence we cannot have neither the inclusion $\widetilde{M_1} \subseteq \widetilde{M_2}$ nor the inclusion $\widetilde{M_2} \subseteq \widetilde{M_1}$. If $W_1 = W_2$ then M_1 and M_2 are conjugates and $\widetilde{M_1} = \widetilde{M_2}$.

(3) $M_1 \in \mathcal{M}_1$ and $M_2 \in \mathcal{M}_2$. In this case $\langle M_1^{g_1}, M_2^{g_2} \rangle = G$ for every $g_1, g_2 \in G$ and, as above, we cannot have neither $\widetilde{M_1} \subseteq \widetilde{M_2}$ nor $\widetilde{M_2} \subseteq \widetilde{M_1}$. □

In particular, it follows from the previous lemma that Question 3.7.2 has an affirmative answer in case of finite soluble groups.

We make another little regression before going on with the next, more substantial, section. It is well known that if a prime p divides the order of a finite group G, then it divides also the order of $G/\operatorname{Frat}(G)$. In particular, $G \setminus \operatorname{Frat}(G)$ contains elements whose order is divisible by p. The analogue statement for invariable generation is false in general. For instance, if $G = A_5$ then $G \setminus \operatorname{Frat}_I(G)$ does not contain elements whose order is divisible by 2.

Notice that in the case of classical generation we can say a little more, namely, we can say that $G \setminus \operatorname{Frat}(G)$ contains elements of p-power order. This follows from the fact that it is always possible to lift an element without affecting the set of prime divisors of its order. For soluble groups, the corresponding "invariable" statement is true as well, although for the proof we invoke Hall's theorems.

Lemma 3.7.4. *Let G be a finite soluble group. If a prime p divides $|G|$, then the set $G \setminus \operatorname{Frat}_I(G)$ contains elements of p-power order.*

Proof. Consider a chief series of G, choose a nontrivial element from every complemented chief factor, and lift it to an element of G of prime power order. It is easy to check that these elements together form an invariable generating set; we may therefore extract a minimal invariable generating set X. If X did not contain any element of p-power order, then Hall's theorems would imply $X \subseteq \widetilde{K}$, where K is a Hall p-complement, contradicting the fact that $\langle X \rangle_I = G$. □

We apply this to prove a lemma that we will need in the following section. Unless otherwise stated, here and in the following sections modules are written multiplicatively, so that 1 denotes the identity element.

Lemma 3.7.5. *Let H be a finite soluble group, and let V be an H-module of finite p-power order. If $C_V(h) = 1$ for every $h \in H \setminus \mathrm{Frat}_I(H)$, then p does not divide $|H|$.*

Proof. Assume by contradiction that p divides $|H|$. Then, by Lemma 3.7.4 there exists $h \in H \setminus \mathrm{Frat}_I(H)$ of p-power order. Now we may construct $G = V \rtimes \langle h \rangle$. This is a finite p-group, hence $V \cap \mathrm{Z}(G) \neq 1$, from which $\mathrm{C}_V(h) \neq 1$, contradicting the hypothesis. □

3.8 $\mathcal{B}_I$-groups

A finite group G is called a $\mathcal{B}$-group if $d(G) = m(G)$. The letter $\mathcal{B}$ refers to the word "basis", since the property $d(G) = m(G)$ is a fundamental one for finite dimensional vector spaces. A classification of the Frattini-free $\mathcal{B}$-groups is given in [AK14, Theorem 1.4]: G is a Frattini-free $\mathcal{B}$-group if and only if one of the following holds:

(1) G is an elementary abelian p-group for some prime p;

(2) $P \rtimes Q$, where P is an elementary abelian p-group and Q is a nontrivial cyclic q-group, for distinct primes $p \neq q$, such that Q acts faithfully on P and the Q-module P is a direct sum of $m(G) - 1$ isomorphic copies of one simple module.

We may give a similar definition for the invariable generation: a finite group G is called a $\mathcal{B}_I$-group if $d_I(G) = m_I(G)$. It turns out that $\mathcal{B}$-groups are $\mathcal{B}_I$-groups (we include this statement in Proposition 3.8.2 below). Indeed, $\mathcal{B}$-groups are soluble. Moreover, $m(G) = d(G) \leqslant d_I(G) \leqslant m_I(G) = m(G)$, where the last equality follows from Theorem 3.1.1 (one can also check directly that the groups in (1) and (2) are $\mathcal{B}_I$-groups).

The converse implication is false. For example $d_I(A_5) = m_I(A_5) = 2$, so A_5 is a $\mathcal{B}_I$-group but not a $\mathcal{B}$-group. Another example is the following. Since $A_5 \cong \mathrm{SL}_2(4)$ we may consider $G = \mathrm{ASL}_2(4) \cong V \rtimes A_5$, where V is a 2-dimensional vector space over the field $\mathbf{F}_4$ with four elements. The elements of order 3 and 5 in A_5 act fixed-point-freely on V, so if $g \in G$ either $|g|$ divides 4 or g is conjugate to an element of order 3 or 5 in A_5. If X is an invariable generating set of G, then X contains necessarily an element of order 3, an element of order 5 and a 2-element with a nontrivial power in V; but three elements of this kind invariably generate G, so $d_I(G) = m_I(G) = 3$.

In this section we want to study the structure of soluble $\mathcal{B}_I$-groups. First notice that there exist soluble $\mathcal{B}_I$-groups that are not $\mathcal{B}$-groups. Indeed the quaternion group Q_8 is isomorphic to an irreducible subgroup of $\mathrm{GL}_2(3)$ and we may consider $G = V \rtimes Q_8$ where V is a 2-generated vector space over the field $\mathbf{F}_3$. The action of Q_8 on V is fixed-point-free, which implies that no element of G has order 6, so an invariable generating set of G must contain two elements of order 4 and one element of order 3, and consequently $d_I(G) = 3 = m_I(G)$.

It turns out that the soluble $\mathcal{B}_I$-groups which are not $\mathcal{B}$-groups are, in a sense, generalizations of the above example.

Lemma 3.8.1. *Assume that H is a finite soluble group and that N is a faithful irreducible H-module. Then $G = N \rtimes H$ is a $\mathcal{B}_I$-group if and only if the following conditions hold:*

(1) H is a $\mathcal{B}_I$-group.

(2) $C_N(h) = 1$ for every $h \in H \setminus \mathrm{Frat}_I(H)$.

Proof. Notice that if $\langle h_1, \dots, h_t \rangle_I = H$ and $1 \neq n \in N$, then by Lemma 3.2.1 $\langle h_1, \dots, h_t, n \rangle_I = G$. Moreover, by Proposition 3.2.6, there exist $n_1, \dots, n_t \in N$ such that $\langle h_1 n_1, \dots, h_t n_t \rangle_I = G$ if and only if $\mathrm{C}_N(h_i) \neq 1$ for some $1 \leqslant i \leqslant t$. So G is a $\mathcal{B}_I$-group if and only if all the minimal invariable generating sets of H have the same cardinality (i.e. H is a $\mathcal{B}_I$-group) and whenever an element h of H appears in some minimal invariable generating set of H (i.e. whenever $h \notin \mathrm{Frat}_I(G)$), then $\mathrm{C}_N(h) = 1$. □

Proposition 3.8.2. *Let G be a finite soluble group. Then G is a $\mathcal{B}_I$-group if and only if one of the following occurs:*

(1) G is a $\mathcal{B}$-group;

(2) $G/\mathrm{Frat}(G) \cong N \rtimes H$ where H is a $\mathcal{B}_I$-group, N is a faithful irreducible H-module and $C_N(h) = 1$ for every $h \in H \setminus \mathrm{Frat}_I(H)$. In particular, by Lemma 3.7.5, H and N have coprime orders.

Proof. We may assume $\mathrm{Frat}(G) = 1$. Let $m = m_I(G) = m(G)$ and $F = \mathrm{Fit}\, G$. We have that F has a complement H in G and that $F = N_1 \times \cdots \times N_t$ where N_i is an irreducible H-module. First we claim that $N_i \cong_G N_j$ for every $i \neq j$. Indeed assume for example $N_1 \ncong_G N_2$. Choose $1 \neq x_1 \in N_1$ and $1 \neq x_2 \in N_2$ and let $x = x_1 x_2$. Take a set $\{y_1, \dots, y_{m-2}\}$ of invariable generators of G modulo $N_1 N_2$ and consider $X = \{y_1 \dots, y_{m-2}, x\}$. Assume that there exist $g_1, \dots, g_{m-2}, g \in G$ such that $Y := \langle y_1^{g_1}, \dots, y_{m-2}^{g_{m-2}}, x^g \rangle \neq G$. It follows that Y is a common complement of N_1 and N_2, but this implies $N_1 \cong_G N_2$, a contradiction. So $\langle X \rangle_I = G$ and $d_I(G) \leqslant m - 1 < m = m_I(G)$, against the assumption that G is a $\mathcal{B}_I$-group. So our claim has been proved and we may assume $F \cong_G N^t$ for a suitable irreducible G-module N. Let $K = \mathrm{End}_H N$ and $n = \dim_K N$. Recall that $F = \mathrm{C}_G(F)$ and G/F is isomorphic to a subgroup of $\mathrm{GL}_n(q)$, being $q = |K|$. There are three cases:

(a) N is central. In this case $G = F$ is an elementary abelian p-group.

(b) N is non central and $n = 1$: in this case G/F is cyclic: but then $m_I(G/F) = d_I(G/F) = 1$, hence G/F is a q-group for some prime q not dividing $|N|$. We conclude that $G = F \rtimes Q$, where F is an elementary abelian p-group and Q is a nontrivial cyclic q-group, for distinct primes $p \neq q$, such that Q acts faithfully on F and the Q-module F is a direct sum of $m(G) - 1$ isomorphic copies of one simple module.

(c) N is non central and $n \neq 1$. We claim that this implies $t = 1$. Indeed $G = N^t \rtimes H$ satisfies the hypothesis of Proposition 3.2.6. Suppose $t \neq 1$, let $\{y_1, \ldots, y_{m-t}\}$ be an invariable generating set of H and take $y_{m-t+1} = \cdots = y_{m-1} = 1$. Since

$$\sum_{1 \leqslant i \leqslant m-1} \dim_K \mathrm{C}_N(y_i) \geqslant \sum_{m-t < i \leqslant m-1} \dim_K \mathrm{C}_N(y_i) = n(t-1) \geqslant 2(t-1) \geqslant t,$$

there exist $w_1, \ldots, w_{m-1} \in N^t$ such that $G = \langle y_1 w_1, \ldots, y_{m-1} w_{m-1} \rangle_I$, but then $d_I(G) \leqslant m - 1 < m_I(G)$, a contradiction.

In cases (a) and (b) G is a $\mathcal{B}$-group, and it was already observed that $\mathcal{B}$-groups are $\mathcal{B}_I$-groups. In case (c) we may apply Lemma 3.8.1 to conclude that G is a $\mathcal{B}_I$-group if and only if the condition in (2) is satisfied. □

To construct $\mathcal{B}_I$-groups that are not $\mathcal{B}$-groups, we have to look for non-cyclic $\mathcal{B}$-groups H admitting a faithful irreducible H-module N with the property that $\mathrm{C}_N(h) = 1$ for every $h \in H \setminus \mathrm{Frat}_I(H)$, and construct then $G = N \rtimes H$.

For example, the dicyclic group H of order 12 is a $\mathcal{B}_I$-group and has an irreducible and fixed-point-free action on the 2-dimensional vector space V over the field with 13 elements, so $N \rtimes H$ is a $\mathcal{B}_I$-group of order $12 \cdot 13^2$ which is not a $\mathcal{B}$-group.

We can also take H to be a non-cyclic p-group. In this case, however, the only possibility to have an irreducible fixed-point-free action is when $p = 2$ and H is a generalized quaternion group [Rob12, 10.5.5]. If we want examples with p odd, we need finite p-groups H with an irreducible action on a module N which is not fixed-point-free, but such that $\mathrm{C}_N(y) = 1$ for every $y \notin \mathrm{Frat}\, H = \mathrm{Frat}_I\, H$.

Interestingly, the p-groups with this property have been studied with different purposes. They have been called "secretive" in [KNN71]. Wall [Wal75] proved that for each prime p and integer $d \geqslant 2$ there exists a finite secretive p-group P with $d(P) = d$. Therefore we have several examples of soluble $\mathcal{B}_I$-groups which are not $\mathcal{B}$-groups.

Outside the soluble case we know almost nothing. The problem of investigating the finite unsoluble $\mathcal{B}_I$-groups is entirely open.

3.9 Invariable basis property

A group G has the basis property if and only if $d(H) = m(H)$ for every $H \leqslant G$. The groups with this property are classified in [AK14, Corollary A.1]. In a similar way we can say that G has the invariable basis property if $d_I(H) = m_I(H)$ for every $H \leqslant G$. If G has the invariable basis property, then every cyclic subgroup of G has prime-power order. The groups all of whose elements have prime-power order are called CP-groups. They are studied in [Hei06].

Lemma 3.9.1. *Let G be a finite group and let N be a soluble normal subgroup of G. Denote by t the number of non-Frattini factors lying below N in a chief series passing through N. If $g_1N, \dots, g_dN$ is a minimal invariable generating set of G/N, then G admits a minimal invariable generating set of cardinality $d+t$.*

Proof. The proof is by induction on t, so it suffices to prove this statement in the particular case when N is a non-Frattini minimal normal subgroup of G. In this case there exists a complement H of N in G. For every i, we can write $g_i = h_i n_i$ with $h_i \in H$ and $n_i \in N$. For $1 \neq n \in N$, by Lemma 3.2.1 $\{h_1, \dots, h_d, n\}$ is a minimal invariable generating set of G. □

Lemma 3.9.2. *Let G be a finite group and let N be a soluble normal subgroup of G. If G has the invariable basis property, then G/N also has the invariable basis property.*

Proof. Follows immediately from Lemma 3.9.1. □

Proposition 3.9.3. Suppose that G is a finite soluble group with $\mathrm{Frat}(G) = 1$. Then G satisfies the invariable basis property if and only if one of the following occurs.

(1) G is an elementary abelian p-group.

(2) $G = P \rtimes Q$, where P is an elementary abelian p-group and Q is a nontrivial cyclic q-group, for distinct primes $p \neq q$, such that Q acts faithfully on P and the Q-module P is a direct sum of isomorphic copies of one simple module.

(3) $G = N \rtimes H$, where H is a generalized quaternion group, the action of H on N is irreducible and $|N| = p^2$ where p is a prime with $p \equiv 3 \mod 4$. In this case H coincides with the Sylow 2-subgroup of $\mathrm{SL}_2(p)$.

Proof. G is in particular a $\mathcal{B}_I$-group, so it satisfies one of the two possibilities described in Proposition 3.8.2. If G is a $\mathcal{B}$-group, then G satisfies (1) or (2). Otherwise $G = N \rtimes H$ where N is H-irreducible, $\dim_{\mathrm{End}_H(N)} N \neq 1$, $\mathrm{C}_N(h) = 1$ for every $h \in H \setminus \mathrm{Frat}_I(H)$ and $(|H|, |N|) = 1$. Moreover G is a soluble CP-group so, by [Hig57, Theorem 1], G has order divisible by at most 2 primes. Since $(|H|, |N|) = 1$, we conclude that H has prime power order. Since every element of G has prime power order, we also deduce that H acts fixed-point freely on N. By [Rob12, 10.5.5], H is cyclic or generalized quaternion. However we may exclude the first case, since it implies $\dim_{\mathrm{End}_H(N)} N = 1$.

Let us first consider the case $H = Q_8$, the quaternion group. Assume $|N|$ is a power of p, being p an odd prime. Let $\mathbf{F}_p$ be the field with p elements. We have, up to equivalence, a unique faithful irreducible $\mathbf{F}_p Q_8$-representation, say ϕ_p, and this representation has degree 2. Indeed choose a, b in $\mathbf{F}_p$ such that $a^2 + b^2 = -1$. Then $\phi_p : Q_8 \to \mathrm{GL}_2(\mathbf{F}_p)$ is defined by setting

$$\phi_p(i) = \begin{pmatrix} a & b \\ b & -a \end{pmatrix}, \quad \phi_p(j) = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad \phi_p(k) = \begin{pmatrix} b & -a \\ -a & -b \end{pmatrix}.$$

Since

$$\phi_p(-1) = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix},$$

Q_8 acts fixed-point-freely on $N = \mathbf{F}_p^2$ and $G = N \rtimes Q_8$ is a $\mathcal{B}_I$-group. Notice that $\phi_p(i), \phi_p(j)$ and $\phi_p(k)$ have minimal polynomial $x^2 + 1$.

If $p \equiv 1 \mod 4$, then there is $c \in \mathbf{F}_p$ such that $c^2 = -1$, hence we may choose $(a, b) = (c, 0)$ and $\phi_p(i)$ has eigenvalues c and $-c$. In this case consider $X = N \rtimes \langle i \rangle$, where $N = \langle w_1, w_2 \rangle$ with $w_1^i = cw_1$ and $w_2^i = -cw_2$. We have $X = \langle w_1 + w_2, i \rangle_I$, so $2 = d_I(X) < m_I(X) = 3$ and G does not satisfy the invariable basis property. It follows that $p \equiv 3 \mod 4$.

Consider now the general case $G = V \rtimes Q_{2^n}$, where V is an elementary abelian p-group and Q_{2^n} is the generalized quaternion group

$$Q_{2^n} = \langle x, y \mid x^{2^{n-1}} = 1, y^2 = x^{2^{n-2}}, y^{-1}xy = x^{-1} \rangle.$$

Suppose that this group has the invariable basis property. In particular $K = \langle x^{2^{n-3}}, y \rangle \cong Q_8$ is a subgroup of Q_{2^n} and $V \rtimes K$ is a subgroup of G. Since G has the invariable basis property, $V \rtimes K$ is a $\mathcal{B}_I$-group, hence by Proposition 3.8.2 V is a faithful irreducible K-module. It follows that $|V| \leqslant p^2$ and Q_{2^n} can be identified with a subgroup of $\mathrm{GL}_2(p)$. Since $V \rtimes K$ has the invariable basis property, we conclude again $p \equiv 3 \mod 4$. Moreover, y is an element of order 4 of $\mathrm{GL}_2(p)$, hence its characteristic polynomial is $t^2 + 1$ and consequently $\det y = 1$. Let α and β be the eigenvalues of x (in the algebraic closure of $\mathbf{F}_p$). Since x and x^{-1} are similar matrices we have $\{\alpha, \beta\} = \{\alpha^{-1}, \beta^{-1}\}$, from which $\beta = \alpha^{-1}$ and $\det x = 1$. Hence $Q_{2^n} \leqslant \mathrm{SL}_2(p)$. We deduce from [Gor07, Chap. 2, Theorem 8.3 (ii)] that Q_{2^n} coincides with a Sylow p-subgroup of $\mathrm{SL}_2(p)$.

Conversely, it is not difficult to see that if G satisfies (1), (2) or (3), then G has the invariable basis property. □

Corollary 3.9.4. *Let G be a finite soluble group with the invariable basis property. Then $G = P \rtimes Q$, where P and Q are Sylow subgroups of G and the action of Q on P is fixed-point-free. In particular Q is cyclic or generalized quaternion.*

Proof. Let $F = \mathrm{Frat}(G)$. By Proposition 3.9.3, $G/F = X \rtimes Y$ where X is a p-group, Y is a q-group and p and q are distinct primes. Since F is nilpotent and contains no element of order $p \cdot q$, we deduce that F is either a p-group or a q-group. Assume by contradiction that F is a nontrivial q-group and let P be a Sylow p-subgroup of G. Clearly FP is a normal subgroup of G, so by the Frattini argument, $G = FP\mathrm{N}_G(P) = F\mathrm{N}_G(P) = \mathrm{N}_G(P)$. But then both P and F are normal in G, so $[P, F] = 1$ and G contains an element of order $p \cdot q$, a contradiction. Therefore F is a p-group and the statement follows. □

While we did not study unsoluble $\mathcal{B}_I$-groups, the invariable basis property is restrictive enough to allow, with the help of the results in [Hei06], a characterization of all groups having this property. In particular, there are only four unsoluble groups sharing it.

Lemma 3.9.5. *Let G be a nonabelian finite simple group. Then G has the invariable basis property if and only if it is isomorphic to one of the following:*

(1) $\mathrm{PSL}_2(5)$, $\mathrm{PSL}_2(8)$.

(2) ${}^2B_2(8)$, ${}^2B_2(32)$.

Proof. G must be a CP-groups, so by [Hei06, Proposition 3] G is isomorphic to one of the following:

(1) $\mathrm{PSL}_2(q)$ for $q = 5, 7, 8, 9, 17$.

(2) $\mathrm{PSL}_3(4)$.

(3) ${}^2B_2(8)$, ${}^2B_2(32)$.

However, $\mathrm{PSL}_2(7), \mathrm{PSL}_2(9), \mathrm{PSL}_2(17)$ and $\mathrm{PSL}_3(4)$ have a subgroup isomorphic to S_4: since $2 = d_I(S_4) < m_I(S_4) = 3$, these groups do not have the invariable basis property. We analyze the remaining cases:

- $G = \mathrm{PSL}_2(5) \cong A_5$. We have already noticed that $d_I(G) = m_I(G) = 2$. It can be easily seen that if H is a proper subgroup of G then either H is a p-group or H is non cyclic with $m_I(H) = 2$. Hence G has the invariable basis property.

- $G = \mathrm{PSL}_2(8)$. An element of G can have order 1, 2, 3, 7, 9 and there are three conjugacy classes of maximal subgroups: F_{56}, D_{18}, D_{14}. The minimal invariable generating sets of G are precisely the sets consisting of two elements, one of order 7, the other of order 3 or 9, so $d_I(G) = m_I(G) = 2$. It can be easily seen that if H is a proper subgroup of G then either H is a p-group or H is non cyclic with $m_I(G) = 2$.

- $G = {}^2B_2(8)$. An element of G can have order 1, 2, 4, 5, 7, 13 and there are four conjugacy classes of maximal subgroups: $2^{3+3} : 7$ (the Frattini subgroup has order 8, and the factor group over the Frattini subgroup has a unique minimal normal subgroup, of order 8), $13 : 4$, $5 : 4$, D_{14}. The minimal invariable generating sets of G are precisely the sets consisting of two elements x, y such that $\{|x|, |y|\} = \{4, 7\}, \{5, 7\}, \{5, 13\}$ or $\{7, 13\}$. Again it can be easily seen that if H is a proper subgroup of G then either H is a p-group or H is non cyclic with $m_I(G) = 2$.

- $G = {}^2B_2(32)$. An element of G can have order 1, 2, 4, 5, 25, 31, 41 and there are four conjugacy classes of maximal subgroups: $2^{5+5} : 31$ (the Frattini subgroup has order 32, and the factor group over the Frattini subgroup has a unique minimal normal subgroup, of order 32), $41 : 4$, $25 : 4$ (the Frattini subgroup has order 5), D_{62}. The minimal invariable generating sets of G are precisely the sets consisting of two elements x, y such that $\{|x|, |y|\} = \{5, 31\}, \{25, 31\}, \{25, 41\}$ or $\{31, 45\}$. Again it can be easily seen that if H is a proper subgroup of G then either H is a p-group or H is non cyclic with $m_I(G) = 2$. □

Corollary 3.9.6. *Let G be a finite nonsoluble group. Then G has the invariable basis property if and only if $G \in \{\mathrm{PSL}_2(5), \mathrm{PSL}_2(8), {}^2B_2(8), {}^2B_2(32)\}$.*

Proof. We have to prove only the direct implication. G is a CP-group so by [Hei06, Proposition 2], there are normal subgroups $1 \leqslant N \leqslant M \leqslant G$ of G such that G/M is soluble, $M/N = S$ is a finite nonabelian simple group and N is a 2-group. By Lemmas 3.9.2 and 3.9.5, $M/N \in \{\mathrm{PSL}_2(5), \mathrm{PSL}_2(8), {}^2B_2(8), {}^2B_2(32)\}$; we want to show $M = G$ and $N = 1$.

It follows from Propositions 4 and 5 in [Hei06] that $M = G$. Notice that S contains a subgroup isomorphic to the dihedral group of order $2 \cdot p$, with $p = 5$ if $S = \mathrm{PSL}_2(5)$, $p = 7$ if $S \in \{\mathrm{PSL}_2(8), {}^2B_2(8)\}$, $p = 31$ if $S = {}^2B_2(32)$. So there exists a subgroup H of G containing N and with the property that $H/N \cong D_{2p}$. Since H satisfies the invariable basis property, we deduce from Corollary 3.9.4 that H has a normal Sylow p-subgroup, say P, and consequently $N \leqslant \mathrm{C}_G(P)$. Since G cannot contain elements of order $2 \cdot p$, we conclude $N = 1$. □

3.10 Further remarks

In this chapter, one of our motivations was to see to what extent, and in which way, results about minimal generating sets can be extended to minimal invariable generating sets.

There is a well developed theory of generation of finite groups. On the one side of the story there are soluble groups. Gaschütz [Gas59] found a nice formula which computes the minimal number of generators of a finite soluble group. On the other side of the story there are nonabelian finite simple groups. As recalled in Subsection 3.1.5, it follows from the CFSG that every finite simple group is generated by two elements.

In a sense, it is possible to combine the two stories (soluble groups and simple groups) in order to obtain a theory of generation for all finite groups (see Detomi–Lucchini [DL03]).

When one deals with invariable generation, many complications occur. For instance, we saw in Subsection 2.3.1 that, roughly speaking, it is necessary to replace the subgroups H of G by the subsets $\widetilde{H}$ of G, which are harder to handle.

We may even take the following viewpoint. An obvious feature of generation is represented by the following implication:

$$\langle x, y\rangle = G \implies \langle x, xy\rangle = G. \tag{3.10.1}$$

Despite obvious and apparently innocent, this property lurks behind more or less all results related to generation.

It is very easy to find examples (for instance in the symmetric group S_3) showing that this property fails for the invariable generation. This constitutes a serious obstacle for extending proofs from the classical to the invariable setting. For instance, Tarski irredundant basis theorem, which we stated in Subsection 3.1.2, is a nice application of (3.10.1). We managed to extend Tarski's theorem in case of finite soluble groups, but we had to change completely the argument.

We take the chance to conclude with the following remark. Most results of this chapter concern finite soluble groups. This depends on the fact that, for soluble groups, in some cases the theory of crowns allow to reduce to questions of vector spaces and linear algebra, which is the ideal environment for (invariable) generation. However, the reader should be alerted that even for this class of groups the situation is not easy, and many apparently approachable questions still do not find an answer.

Chapter 4

On the number of conjugacy classes of a primitive permutation group with nonabelian socle

The content of this chapter is joint work with Nick Gill. This is the only chapter of the thesis which is not concerned with invariable generation of groups.

4.1 Introduction

Throughout this chapter, $k(G)$ denotes the number of conjugacy classes of a finite group G. Maróti [Mar05] proved that if G is a primitive permutation group G of degree n, then $k(G) \leqslant p(n)$, where $p(n)$ denotes the number of partitions of n. This bound is attained by S_n in its action on n points. Moreover, he proved that if the socle of G is not a direct product of alternating groups, then $k(G) \leqslant n^6$.

In this chapter, we want to improve this bound under the assumption that G has nonabelian socle. In Subsection 4.1.2, we will give more context and review more results in this area.

There are two special types of primitive groups which we wish to single out.

(A) Let G be the symmetric group S_d or the alternating group A_d on $d \geqslant 5$ letters. For every $1 \leqslant k < d/2$, G acts primitively on the set of k-subsets of $\{1, \ldots, d\}$. These are in number $\binom{d}{k}$.

(B) Let G be an almost simple group with socle $\mathrm{PSL}_d(q)$, and assume $G \leqslant \mathrm{P\Gamma L}_d(q)$. Then G acts primitively on the set of 1-subspaces of $\mathbf{F}_q^d$. These are in number $(q^d - 1)/(q - 1)$.

G	n	$k(G)$
M_{11}	11,12	10
M_{12}	12,12	15
M_{22}	22	12
$M_{22}.2$	22	21
M_{23}	23	17
M_{24}	24	26
A_7	15,15	9
$S_8 \cong \mathrm{SL}_4(2).2$	35	22
$\mathrm{PSL}_2(11)$	11,11	8
$\mathrm{SO}_8^-(2)$	119	60
$\mathrm{Sp}_8(2)$	120,136	81
$\mathrm{SO}_8^+(2)$	120	67
$\mathrm{Sp}_6(2)$	28, 36	30
$\mathrm{PSp}_4(3) \cong \mathrm{SU}_4(2)$	27,36,40,40	20
$\mathrm{PSp}_4(3).2$	27,36,40,40,45	25
$\mathrm{PSU}_4(3).(2 \times 2)$	112	59
$\mathrm{P\Gamma U}_4(3)$	112	61
$\mathrm{SU}_3(3)$	28	14
$\mathrm{SU}_3(3).2$	28	16

Table 4.1: Almost simple primitive permutation groups G of degree n (up to equivalence) for which $k(G) \geqslant \frac{n}{2}$, and for which the action is not isomorphic to an action in (A) or (B).

Our main result says that, if G is a primitive group with nonabelian socle, then either G has very few conjugacy classes, or else the action of G is "related" to (A) or (B) or to further finitely many almost simple primitive permutation groups. The precise statement is as follows.

Theorem 4.1.1. *Let G be a primitive permutation group of degree n with nonabelian socle, so $\mathrm{Soc}(G) \cong S^r$, with S nonabelian simple and $r \geqslant 1$. Then one of the following holds.*

(1) $k(G) < n/2$, and $k(G) = O(n^\delta)$ for some absolute $\delta < 1$.

(2) $G \leqslant A \wr S_r$, A is an almost simple primitive permutation group of degree m with socle S, G acts in product action on $n = m^r$ points, and one of the following holds:

(i) The action of A on m points is equivalent to an action in Table 4.1, and $k(G) < n^{1.31}$.

(ii) The action of A on m points is isomorphic to an action described in (A) or (B). In the (B)-case, $k(G) < n^{1.9}$.

See footnote[1] for some terminology. We will first prove Theorem 4.1.1 in case G is almost simple, and then deduce the general case. For convenience, we state separately the almost simple case (where we also give an explicit estimate for δ).

Theorem 4.1.2. *Let G be an almost simple primitive permutation group of degree n. Then one the following holds.*

(1) $k(G) < n/2$, and $k(G) = O(n^{3/4})$.

(2) Either the action of G is equivalent to an action in Table 4.1, or the action of G is isomorphic to an action described in (A) or (B). In the (B)-case, $k(G) < 100n$.

In item (1), the exponent $3/4$ is sharp, although in most cases $k(G) = o(n^{3/4})$ as $n \to \infty$; see Remark 4.2.14 for a precise statement.

In the proof of Theorem 4.1.2, an essential ingredient is the work of Fulman–Guralnick [FG12], which give upper bounds to the number of conjugacy classes of almost simple groups of Lie type.

We immediately make some clarifications regarding the statement of Theorem 4.1.1.

Remark 4.1.3. (i) We are not asserting that every case appearing in Theorem 4.1.1(2) does not satisfy item (1). For instance, assume $m = n$, and consider $G = S_d$ acting on $n = \binom{d}{k}$ points as in (A), and assume $cd \leqslant k \leqslant \frac{d}{2}$ for some fixed constant c. Then it is well known that $n = \binom{d}{k}$ is exponential in d, while the number of conjugacy classes of $G = S_d$ is of the form $O(1)^{\sqrt{d}}$. In particular $k(G) = n^{o(1)}$ as $d \to \infty$.

(ii) In Theorem 4.1.1(2)(ii), we can be more precise about the adjective *isomorphic*, as follows. If A is A_d or S_d, then either the action of A is *equivalent* to the action on k-subsets; or else $(d, m) = (6, 6)$ or $(6, 15)$. Moreover, if A is almost simple with socle $\mathrm{PSL}_d(q)$ and $A \leqslant \mathrm{P\Gamma L}_d(q)$, then the action of A is equivalent to the action on the 1-subspaces or $(d-1)$-subspaces of $\mathbf{F}_q^d$. For this, see Lemmas 4.2.8 and 4.2.12.

(iii) Whenever G is almost simple with socle isomorphic to both A_d and $\mathrm{PSL}_f(q)$, we have excluded from Table 4.1 both the groups in (A) and (B). For instance, $G = S_6$ has 11 conjugacy classes, and contains a subgroup $S_3 \wr S_2$ of index 10 acting transitively on 6 points; but this does not appear in Table 4.1 in view of the isomorphism $S_6 \cong \mathrm{P\Sigma L}_2(9)$. The same reasoning applies to the isomorphisms $\mathrm{SL}_2(4) \cong \mathrm{PSL}_2(5)$ and $\mathrm{PSL}_2(7) \cong \mathrm{SL}_3(2)$.

[1] In order to avoid confusion, we recall the following standard definitions. Assume $\rho_i : G \to \mathrm{Sym}(\Omega_i)$ for $i = 1, 2$. The representations ρ_1 and ρ_2 are called *isomorphic* if there exist a bijection $\phi : \Omega_1 \to \Omega_2$ and $\Psi \in \mathrm{Aut}(G)$ such that $(g\rho_1) \cdot \phi = \phi \cdot (g^\Psi \rho_2)$ for every $g \in G$ (composition of mappings is left-to-right). If in this definition $\Psi = 1$, the representations ρ_1 and ρ_2 are called *equivalent*.

4.1.1 When is $k(G) = o(n)$?

Theorem 4.1.1 implies in particular that, if the socle of G is nonabelian, then either $k(G) = o(n)$, or G is "known". Can we prove that $k(G) = o(n)$ in further cases?

We are particularly interested in the cases contemplated in Theorem 4.1.1(2) (i), for which we show $k(G) < n^{1.31}$. We first note that there are examples in which $k(G) > n^{1.08}$ for arbitrarily large n, in contrast to item (1); see Lemma 4.4.1.

Still, it would be interesting to understand precisely when this happens (since there are only finitely many almost simple groups to handle).

Question 4.1.4. *Let A be an almost simple primitive group on m points appearing in Table 4.1. Determine whether every primitive subgroup G of $A \wr S_r$ on $n = m^r$ points is such that $k(G) = o(m^r)$ as $r \to \infty$.*

We refer to Section 4.4 for comments in this direction. See in particular Conjecture 4.4.2, which would provide an answer to Question 4.1.4.

4.1.2 Context

There are many results in the literature which give upper bounds to the number of conjugacy classes of a finite groups in terms of various parameters. We recall some of these, focusing on permutation groups.

Kovács–Robinson [KR93] proved that every permutation group of degree n has at most 5^{n-1} conjugacy classes. This estimate was subsequently improved by Liebeck–Pyber, Maróti, and Garonzi–Maróti, as follows:

$$k(G) \leqslant 2^{n-1} \qquad \text{([LP97])}$$
$$k(G) \leqslant 3^{(n-1)/2} \qquad \text{([Mar05])}$$
$$k(G) \leqslant 5^{(n-1)/3} \qquad \text{([GM15]).}$$

We should mention that, in [KR93] and [LP97], various other upper bounds to $k(G)$ are proved, where G is not necessarily a permutation group.

There are easy examples showing that these estimates are somewhat close to best possible, even for transitive groups. Indeed, the subgroup $S_4^{n/4} \leqslant S_n$ has $5^{n/4}$ conjugacy classes; and the transitive subgroup $G = S_4 \wr C_{n/4} \leqslant S_n$ has at least $5^{n/4-o(n)}$ conjugacy classes (see Lemma 4.2.1).

For primitive groups, the situation is very different. Improving results from [LP97], Maróti [Mar05] proved that every normal subgroup of a primitive permutation group G of degree n has at most $p(n)$ conjugacy classes; and if the socle of G is not a direct product of alternating groups, then $k(G) \leqslant n^6$. (Recall that $p(n) = O(1)^{\sqrt{n}}$, and in fact the asymptotic behaviour of $p(n)$ is famously known by work of Hardy–Ramanujan.) Theorem 4.1.1 can be regarded as an improvement of this statement, in case the socle of G is nonabelian.

4.1.3 Abelian socle

In this thesis we do not address the case in which the socle of G is abelian. In this case, we still have the bound $k(G) \leqslant n^6$ from [Mar05].

There is a deep problem, known as *non-coprime $k(GV)$-problem*, which was addressed by Guralnick–Tiep [GT05] and which asks (in particular) for a characterization of the affine primitive permutation groups of degree n for which $k(G) > n$. We refer to Guralnick–Tiep [GT05], Guralnick–Maróti [GM13] and the references therein for results in this direction, partly motivated by the celebrated Brauer's $k(B)$-conjecture.

4.1.4 Organization of the chapter

In Section 4.2 we prove Theorem 4.1.2, in Section 4.3 we prove Theorem 4.1.1, and in Section 4.4 we discuss Question 4.1.4 and make further comments.

4.2 Almost simple groups

In this section we prove Theorem 4.1.2. Regarding item (1), we prove the inequality $k(G) < n/2$ in Subsections 4.2.2–4.2.4, and then we prove the asymptotic inequality $k(G) = O(n^{3/4})$ in Subsection 4.2.5.

First, we gather some results that we will use throughout.

4.2.1 Some preliminary results and notation

For a finite group G, let $P(G)$ be the minimal degree of a faithful permutation representation of G. If G is almost simple with socle S, then $P(G)$ coincides with the minimal degree of a faithful transitive permutation representation of G, and moreover $P(S) \leqslant P(G)$. The values of $P(G)$ for G a finite simple group are known; they are listed for instance in [GMPS15, Table 4].

We recall an elementary lemma which appears in [Gal70]. We will often apply this lemma with no mention.

Lemma 4.2.1. *If G is a finite group and H is a subgroup of G, then*

$$k(H)/|G : H| \leqslant k(G) \leqslant |G : H| \cdot k(H).$$

If moreover H is normal in G, then

$$k(G) \leqslant k(H) \cdot k(G/H).$$

In one occasion, we will need the following variant (see [KR93, p. 447]).

Lemma 4.2.2. *Let G be a finite group and let N be a normal subgroup of G. Then*

$$k(G) \leqslant |G : N| \cdot \#\{G\text{-conjugacy classes of } N\}.$$

We are now ready to begin the proof of Theorem 4.1.2.

G	n	$k(G)$
M_{11}	11,12	10
M_{12}	12,12	15
M_{22}	22	12
$M_{22}.2$	22	21
M_{23}	23	17
M_{24}	24	26

Table 4.2: Faithful primitive permutation representations of degree n for sporadic almost simple groups G such that $k(G) \geqslant \frac{n}{2}$.

4.2.2 Sporadic groups

Lemma 4.2.3. *Let G be almost simple with socle S, a sporadic simple group. Let M be a core-free maximal subgroup of G, and write $n = |G : M|$. If $k(G) \geqslant \frac{n}{2}$, then G and n are listed in Table 4.2.*

Proof. We go through the ATLAS [CCN+85]. □

4.2.3 Alternating groups

We recall some results that we will use. The first is an inequality of Pribitkin [Pri09], as follows.

Lemma 4.2.4. *Let $p(d)$ be the number of partitions of the integer d. Then*

$$p(d) < \frac{e^{\pi\sqrt{2d/3}}}{d^{3/4}}.$$

We will also need the following pair of inequalities which are an easy consequence of work of Robbins on the Stirling approximations [Rob55].

Lemma 4.2.5. *Let $d \geqslant 2$ be an integer. Then*

$$\sqrt{2\pi}d^{d+1/2}e^{-d} \leqslant d! \leqslant e\,d^{d+1/2}e^{-d}.$$

Finally we will need the following result of Praeger and Saxl which makes use of CFSG [PS80].

Lemma 4.2.6. *Let $G \leqslant S_d$ and suppose that G is primitive and does not contain A_d. Then $|G| < 4^d$.*

We now prove a lemma, which is known (see [FG12, Corollary 2.7]).

Lemma 4.2.7. *For all d, $k(A_d) \leqslant k(S_d)$.*

The proof can be easily modified to obtain a strict inequality here.

Proof. For $d \leqslant 9$ we can check this directly. Thus assume that $d > 9$. Note that $k(A_d)$ is equal to the number of even partitions of d plus the number of partitions of d into distinct odd numbers; on the other hand $k(S_d)$ is equal to the number of even partitions of d plus the number of odd partitions of d. Thus we need to show that the number of partitions of d into distinct odd numbers is less than or equal to the number of odd partitions of d.

Given $\ell \geqslant 10$, a positive even integer, observe that there are precisely $\lfloor \frac{\ell}{4} \rfloor$ ways to partition ℓ into two distinct odd numbers. On the other hand, let Y_ℓ be the set of partitions of ℓ into an odd number of even numbers. It is easy to see that $|Y_\ell| > \lfloor \frac{\ell}{4} \rfloor$: we simply use the partitions

$$\ell,\ (\ell - 4, 2^2),\ (\ell - 8, 2^4), \dots \text{ and } (\ell - 8, 4^2).$$

Now let X be the set of all partitions of d into at least two distinct odd numbers; let X_ℓ be the subset of all partitions of X for which the sum of the two largest parts is equal to ℓ. Since $d \geqslant 10$, observe that $\ell \geqslant 10$. To each partition $\mathfrak{p} \in X_\ell$ we can associate a partition which has the same parts as $\mathfrak{p}$, apart from the largest two, which are replaced by a partition from Y_ℓ. The counting above implies that we can do this in such a way that the association yields an injective function from X to a proper subset of the odd partitions of X.

If d is even, then this yields that $k(A_d) < k(S_d)$. If d is odd, we must also associate an odd partition with the partition of d consisting of a single part. But since our injective function is not onto, this can be done. The result follows. □

Now we can prove the main result of this subsection.

Lemma 4.2.8. *Let G be almost simple with socle $S \cong A_d$. Let M be a core-free maximal subgroup of G, and write $n = |G : M|$. If $k(G) \geqslant \frac{n}{2}$, then one of the following holds.*

(1) G and n are listed in Table 4.3.

(2) M is intransitive in its action on d points, thus $n = \binom{d}{k}$ for some integer k such that $1 \leqslant k < \frac{1}{2}d$.

(3) $(d, n) = (6, 6)$ or $(6, 15)$, and the action of G on the cosets of M is isomorphic, but not equivalent, to the action on the coset of a maximal intransitive subgroup.

In item (3), if $(d, m) = (6, 6)$ we have $G = A_6$ or S_6, and $M = S_5 \cap G$, where S_5 is a subgroup of S_6 acting primitively on 6 points. If $(d, m) = (6, 15)$, again $G = A_6$ or S_6, and $M = (S_2 \wr S_3) \cap G$, where $S_2 \wr S_3$ acts transitively (and imprimitively) on 6 points. (Note that in the latter case, if $G = A_6$ then $k(G) < n/2$.)

Proof. For $d \leqslant 8$, we use the ATLAS [CCN+85] together with GAP [GAP19] to obtain the given list. For $9 \leqslant d \leqslant 20$ we use GAP to check that no examples occur. Assume, then, that $d > 20$.

G	n	$k(G)$
A_5	6	5
S_5	6	7
$A_6 = \mathrm{PSL}_2(9)$	10	7
$A_6.2 = \mathrm{PGL}_2(9)$	10	11
$A_6.2 = S_6$	10	11
$A_6.2 = M_{10}$	10	8
$A_6.(2 \times 2) = \mathrm{P\Gamma L}_2(9)$	10	13
A_7	15,15	9
A_8	15,15	14
S_8	35	22

Table 4.3: Faithful primitive permutation representations of degree n for almost simple groups G with socle A_d such that $k(G) \geqslant \frac{n}{2}$, and the action is not isomorphic to an action in (A).

Let us suppose, first, that M is primitive in its action on d points. Then Lemmas 4.2.4–4.2.7 imply that it is sufficient to prove the following

$$\frac{e^{\pi\sqrt{2d/3}}}{d^{3/4}} < \frac{\sqrt{2\pi}d^{d+1/2}}{4 \cdot e^d \cdot 4^d}.$$

If we assume that the other inequality holds, we get

$$\begin{aligned} e^{\pi\sqrt{2d/3}} &\geqslant \frac{\sqrt{2\pi} \cdot d^{d+5/4}}{4 \cdot e^d \cdot 4^d} \\ \implies 2 \cdot e^{\pi\sqrt{2d/3}} \cdot (4e)^d &\geqslant d^{d+5/4} \\ \implies 2 \cdot e^{2.6\sqrt{d}} \cdot (4e)^d &\geqslant d^{d+5/4} \\ \implies 2 \cdot (5e)^{d+\sqrt{d}} &\geqslant d^{d+5/4} \\ \implies d &\leqslant 20. \end{aligned}$$

Since $d > 20$, the result follows.

Let us suppose, next, that M is imprimitive in its action on d points. Then

$$n = \frac{d!}{(k!)^\ell \ell!},$$

where $d = k\ell$ and $k, \ell \geqslant 2$. Now Lemma 4.2.5 implies that

$$\begin{aligned} n &\geqslant \frac{\sqrt{2\pi} \cdot d^{d+1/2} \cdot e^{-d}}{(ek^{k+1/2}e^{-k})^{\ell} \cdot e \cdot \ell^{\ell+1/2} \cdot e^{-\ell}} = \frac{\sqrt{2\pi} \cdot d^{d+1/2}}{e \cdot k^{(k+1/2)\ell} \cdot \ell^{\ell+1/2}} \\ &= \frac{\sqrt{2\pi}}{e} \cdot \frac{\ell^{d-\ell}}{k^{(\ell-1)/2}} \\ &\geqslant \frac{\ell^{d-\ell}}{k^{\ell/2}} \\ &= \frac{\ell^{d-\ell}}{(d/\ell)^{\ell/2}}, \end{aligned}$$

which implies that

$$\begin{aligned} &(d-\ell)\log(\ell) - \frac{\ell}{2}\log(d) + \frac{\ell}{2}\log(\ell) \leqslant \log(n) \\ &\qquad \implies (d - \frac{\ell}{2})\log(\ell) - \frac{\ell}{2}\log(d) \leqslant \log(n). \end{aligned}$$

If we fix d and set $f(\ell)$ to be the function on the left hand side of the final inequality, with $\ell \in (0, d)$, then one computes that $f'(\ell)$ is a decreasing function. In particular, $f(\ell)$ takes its minimum value in the range $\ell \in (0, d)$ either when ℓ is as large as possible or as small as possible.

If ℓ is as small as possible, then $\ell = 2$ and we obtain that

$$f(2) \leqslant \log(n) \iff d - \log(d) - 1 \leqslant \log(n).$$

If ℓ is as large as possible, then $\ell = \frac{d}{2}$ and we obtain that

$$f\left(\frac{d}{2}\right) \leqslant \log(n) \iff \frac{d}{2}\log(d) - \frac{3d}{4} \leqslant \log(n).$$

Now it is easy to check that, for $d \geqslant 4$,

$$d - \log(d) - 1 \leqslant \frac{d}{2}\log(d) - \frac{3d}{4},$$

and we conclude that

$$d - \log(d) - 1 \leqslant \log(n). \tag{4.2.1}$$

On the other hand, if $k(G) \geqslant \frac{n}{2}$, then Lemmas 4.2.4 and 4.2.7 imply that

$$\frac{n}{2} < \frac{e^{\pi\sqrt{2d/3}}}{d^{3/4}}.$$

Taking logs and using (4.2.1), we get

$$d < \frac{1}{4}\log d + \pi\sqrt{\frac{2d}{3}}\log e + 2,$$

which, since $d > 20$, is false. This concludes the proof. □

4.2.4 Groups of Lie type

For groups of Lie type, we will use results of Fulman and Guralnick giving bounds on the number of conjugacy classes [FG12].

Theorem 4.2.9. *Let G be a connected simple algebraic group of rank r over a field of positive characteristic. Let F be a Steinberg–Lang endomorphism of G with G^F a finite Chevalley group over the field $\mathbf{F}_q$. Then*

$$k(G^F) \leqslant \min\{27.2q^r, q^r + 68q^{r-1}\}.$$

Theorem 4.2.10. *Let G be an almost simple group with socle S, a Chevalley group of rank r defined over $\mathbf{F}_q$. Then $k(G) \leqslant 100q^r$.*

First we deal with exceptional groups.

Lemma 4.2.11. *Let G be almost simple with socle S, a simple exceptional group of Lie type. Then $k(G) < \frac{1}{2}|G : M|$ for all core-free maximal subgroups M of G.*

Proof. We use the values for $P(S)$ given in [GMPS15], as well as the fact that $k(G) \leqslant 100q^\ell$ (from Theorem 4.2.10).

If S is not a Suzuki group, then the value of $P(S)$ given in [GMPS15] rules out all groups except $G_2(3)$, $G_2(4)$, $G_2(5)$, ${}^3D_4(2)$, ${}^2F_4(2)'$ and ${}^2G_2(3^3)$. For these groups, we can verify $k(G) < P(S)/2$ using the ATLAS [CCN+85].

If S is a Suzuki group, then [Suz62] tells us that $k(S) = q+3$, and Lemma 4.2.1 implies that $k(G)$ is at most $(q+3)f$, where $q = 2^f$. On the other hand, $P(S) = q^2 + 1$ by [GMPS15]. Then $(q+3)f \geqslant \frac{1}{2}(q^2+1)$ if and only if $q = 8$ (recall that f is odd and $f \geqslant 3$). But if $q = 8$, [CCN+85] tells us that $S.3$ has 17 conjugacy classes and this case, too, is excluded. □

Next we deal with the case in which G has socle $\mathrm{PSL}_d(q)$.

Lemma 4.2.12. *Let G be almost simple with socle $S \cong \mathrm{PSL}_d(q)$. Let M be a core-free maximal subgroup of G, and write $n = |G : M|$. If $k(G) \geqslant \frac{n}{2}$, then one of the following holds.*

(1) G and n are listed in Table 4.4.

(2) $G \leqslant \mathrm{P\Gamma L}_d(q)$ and M stabilizes a 1-dimensional or a $(d-1)$-dimensional subspace of $\mathbf{F}_q^d$, thus $n = \frac{q^d-1}{q-1}$.

Note that, in Table 4.4, $n = 6$ appears for $G = \mathrm{SL}_2(4)$, but not for $G = \mathrm{PSL}_2(5)$ (even if $\mathrm{SL}_2(4) \cong \mathrm{PSL}_2(5)$). Similarly, $n = 5$ appears for $\mathrm{PSL}_2(5)$, but not for $\mathrm{SL}_2(4)$. The same reasoning applies to the isomorphism $\mathrm{PSL}_2(7) \cong \mathrm{SL}_3(2)$.

Proof. In this proof we use [Kan79]. The main theorem of this paper, together with Theorem 4.2.10, implies that, if $k(G) \geqslant n/2$, then either $d \leqslant 4$, or

$$(d, q) \in \{(5,2), (5,4), (5,8), (6,2), (7,2)\}, \tag{4.2.2}$$

G	n	$k(G)$
$\mathrm{SL}_2(4) = A_5$	6,10	5
$\mathrm{SL}_2(4).2 = S_5$	6,10	7
$\mathrm{PSL}_2(5) = A_5$	5,10	5
$\mathrm{PGL}_2(5) = S_5$	5,10	7
$\mathrm{PSL}_2(7)$	7,7	6
$\mathrm{PSL}_2(9) = A_6$	6,6	7
$\mathrm{PSL}_2(9).2 = S_6$	6,6,15,15	11
$\mathrm{PSL}_2(11)$	11,11	8
$\mathrm{SL}_3(2)$	8	6
$\mathrm{SL}_3(2).2$	8	9
$\mathrm{SL}_4(2) = A_8$	8, 28	14
$\mathrm{SL}_4(2).2 = S_8$	8,28,35	22

Table 4.4: Faithful primitive permutation representations of degree n for almost simple groups G with socle $\mathrm{PSL}_d(q)$ such that $k(G) \geqslant \frac{n}{2}$, and the action is not isomorphic to an action in (B) (see the remark after the statement of Lemma 4.2.12).

or $H := M \cap \mathrm{P\Gamma L}_d(q)$ is reducible, or H normalizes $\mathrm{PSp}_d(q)$.

Assume first $d \geqslant 5$. This rules out the case in which H normalizes $\mathrm{PSp}_d(q)$. Let us now consider the case where H is reducible, stabilizing a subspace of dimension m.

Assume first $2 \leqslant m \leqslant d-2$. Then $n = |G : M| > q^{m(d-m)} \geqslant q^{2d-4}$. If $\frac{n}{2} \leqslant 100q^{\ell} = 100q^{d-1}$, then we get $q^{d-3} < 200$. We want to whittle down the possibilities, as follows. [FG12, Proposition 3.6] states that $k(\mathrm{PSL}_d(q)) \leqslant 2.5q^{d-1}$. This, together with the knowledge of $|\mathrm{Out}(S)|$ and Lemma 4.2.1, reduces easily to the cases $(d,q) = (5,2), (5,3), (5,4), (6,2)$. The same argument and [Kan79] rules out the cases $(d,q) = (5,8), (7,2)$ in (4.2.2). We can deal with the remaining cases with GAP [GAP19].

Assume now $m \in \{1, d-1\}$. The case in which $G \leqslant \mathrm{P\Gamma L}_d(q)$ appears in item (2) of the statement. If $G \not\leqslant \mathrm{P\Gamma L}_d(q)$, then M is a novelty and $|G : M| \geqslant q^{2d-3}$, and the GAP calculation from the previous paragraph rules out all possibilities.

Let us turn, then, to study what happens when $d \in \{2,3,4\}$. We make use of the counts given in [Mac81].

When $d = 2$, [Mac81] implies that

$$k(\mathrm{PSL}_2(q)) = \frac{1}{(q-1,2)}\big(q + 4(q-1,2) - 3\big) \text{ and } k(\mathrm{PGL}_2(q)) = q + (2, q-1).$$

We use this in combination with the explicit list of maximal subgroups in $\mathrm{PSL}_2(q)$ to conclude that either

(1) $q \leqslant 11$; or

(2) $q = 16$ and M is the normalizer of a torus, or a subfield subgroup such that $M \cap S \cong \mathrm{PGL}_2(\sqrt{q})$; or

(3) $q \in \{25, 49, 81, 64, 256\}$ and M is a subfield subgroup such that $M \cap S \cong \mathrm{PGL}_2(\sqrt{q})$.

Using [GAP19] we get the possibilities in Table 4.4.

Next assume that $d = 3$. If q is odd, then using [Mac81] we see that $k(\mathrm{PSL}_3(q)) \leqslant q^2 + q$, and this, along with [Kan79], allows us to conclude that $q \leqslant 9$. These possibilities can all be excluded using [CCN$^+$85]. If q is even, then [Mac81] implies that $k(G) \leq 2f(q^2 + q + 10)$ where $q = 2^f$. Using the list of subgroups in [BHRD13] this is enough to conclude that $q \leqslant 16$. Now [GAP19] excludes the remainder.

Finally, assume that $d = 4$. If q is odd, then [Mac81] implies that $k(G) \leq 2f(q^3+q^2+5q+21)$ where $q = p^f$. We use [Kan79] to conclude that $q = 3$. This final case is ruled out with [CCN$^+$85]. If q is even, then [Mac81] implies that $k(\mathrm{SL}_4(q)) = q^3 + q^2 + q$ and, again, we use the list of subgroups in [BHRD13] to conclude that $q \leqslant 16$. Now [GAP19, BHRD13, CCN$^+$85] rule out all except the listed exceptions for $q = 2$. □

Finally we deal with almost simple classical groups with socle not isomorphic to $\mathrm{PSL}_d(q)$ — we will only include the case $S \cong \mathrm{Sp}_4(2)' \cong \mathrm{PSL}_2(9)$.

Lemma 4.2.13. *Let G be almost simple with classical socle S, $S \ncong \mathrm{PSL}_d(q)$ or $S \cong \mathrm{Sp}_4(2)' \cong \mathrm{PSL}_2(9)$. Let M be a core-free maximal subgroup of G, and write $n = |G : M|$. If $k(G) \geqslant \frac{n}{2}$, then G and n are listed in Table 4.5.*

Proof. In order to exclude some potential examples, our basic strategy will be to use the bound $k(G) \leqslant |G : S|k(S)$ from Lemma 4.2.1, and try to show that this is smaller than $P(S)/2$. In order to bound $k(S)$, we will use the results in [FG12] for specific families, as follows.

Suppose that $S \cong \mathrm{PSU}_d(q)$. In this case [FG12, Proposition 3.10] implies that $k(S) \leqslant 8.26q^{d-1}$, and we use the values for $P(S)$ given in [GMPS15] to obtain that either S is in

$$\{\mathrm{PSU}_5(2), \mathrm{PSU}_6(2), \mathrm{PSU}_7(2), \mathrm{PSU}_5(3), \mathrm{PSU}_5(4)\}$$

or else $d \leqslant 4$. Groups with the five possible socles with $d > 4$ can be ruled out using [GAP19].

If $S = \mathrm{PSU}_4(q)$, then [Mac81] implies that

$$k(S) = \begin{cases} \frac{1}{4}(q^3 + q^2 + 7q + 23), & q \equiv 3 \pmod 4; \\ \frac{1}{2}(q^3 + q^2 + 7q + 9), & q \equiv 1 \pmod 4; \\ q^3 + q^2 + 3q + 2, & q \equiv 0 \pmod 2. \end{cases}$$

Thus, in any case, $k(G) \leqslant 2f(q^3 + q^2 + 7q + 23)$ where $q = p^f$. Since $P(S) = (q + 1)(q^3 + 1)$, by [GMPS15], we conclude that $q \in \{2, 3, 4, 5, 8, 16\}$. If $q \leqslant 5$,

G	n	$k(G)$
$\mathrm{SO}_8^-(2)$	119	60
$\mathrm{Sp}_8(2)$	120,136	81
$\mathrm{SO}_8^+(2)$	120	67
$\mathrm{Sp}_6(2)$	28, 36	30
$\mathrm{PSp}_4(3) = \mathrm{PSU}_4(2)$	27,36,40,40	20
$\mathrm{PSp}_4(3).2 = \mathrm{PSU}_4(2).2$	27,36,40,40,45	25
$\mathrm{PSU}_4(3).(2 \times 2)$	112	59
$\mathrm{P\Gamma U}_4(3)$	112	61
$\mathrm{Sp}_4(2)' = \mathrm{PSL}_2(9)$	6,6,10	7
$\mathrm{Sp}_4(2)'.2 = \mathrm{PGL}_2(9)$	10	11
$\mathrm{Sp}_4(2) = S_6$	6,6,10,15,15	11
$\mathrm{Sp}_4(2)'.2 = M_{10}$	10	8
$\mathrm{Sp}_4(2)'.(2 \times 2) = \mathrm{P\Gamma L}_2(9)$	10	13
$\mathrm{SU}_3(3)$	28	14
$\mathrm{SU}_3(3).2$	28	16

Table 4.5: Faithful primitive permutation representations of degree n for almost simple classical groups G with socle $S \ncong \mathrm{PSL}_d(q)$, or $S \cong \mathrm{PSL}_2(9)$, such that $k(G) \geqslant \frac{n}{2}$.

then [GAP19] yields the listed cases. If $q \in \{8, 16\}$, then $G = \mathrm{P\Gamma U}_4(q)$. The case $q = 8$ is eliminated by [BCP97]; we now consider the case $q = 16$. We want to apply Lemma 4.2.2 with $G = \mathrm{P\Gamma U}_4(16)$ and $N = \mathrm{SU}_4(16)$. Consider the split torus T of N of order 17^3, which intersects 284 nontrivial N-classes. By looking at eigenvalues, we see that none of these classes is fixed by the standard field automorphism σ of order 8 normalizing T. We deduce from Lemma 4.2.2 that $k(G) \leqslant 8(q^3 + q^2 + 3q + 2 - 284/2) = 34080$, which is enough to conclude that $k(G) < P(S)/2$.

If $G = \mathrm{PSU}_3(q)$, then [Mac81] implies that

$$k(S) = \begin{cases} q^2 + q + 2, & q \not\equiv 2 \pmod 3; \\ \frac{1}{3}(q^2 + q + 12), & q \equiv 2 \pmod 3. \end{cases}$$

Thus, in any case, $k(G) \leqslant 2f(q^2 + q + 12)$ where $q = p^f$. Since $P(S) = q^3 + 1$ if $q \neq 5$, by [GMPS15], we conclude that $q \leqslant 9$ or $G = \mathrm{P\Gamma U}_3(16)$. For $q \leqslant 9$, we obtain the listed examples using [GAP19] and [CCN$^+$85]. For $G = \mathrm{P\Gamma U}_3(16)$, the same argument used for the case $\mathrm{P\Gamma U}_4(16)$ works.

Suppose that $S \cong \mathrm{PSp}_d(q)$. If $d \geqslant 6$, then we use [FG12, Theorems 3.12 and 3.13] along with the values for $P(S)$ given in [GMPS15] to conclude that S is one of the following:

$$\{\mathrm{PSp}_6(3), \mathrm{Sp}_6(2), \mathrm{Sp}_8(2), \mathrm{Sp}_{10}(2)\}.$$

We use [GAP19] and [CCN$^+$85] to check these cases and obtain the listed ex-

amples.

If $S = \mathrm{PSp}_4(q)'$, then we use [Wal63] (for q odd) and [Eno72] (for q even) to establish that

$$k(\mathrm{Sp}_4(q)) = \begin{cases} q^2 + 5q + 10, & q \text{ odd;} \\ q^2 + 2q + 3, & q \text{ even.} \end{cases}.$$

This, combined with [GMPS15], implies that $q \leqslant 9$. Now [GAP19] and [CCN$^+$85] yield the listed examples.

Suppose that $S \cong \mathrm{P}\Omega_{2\ell+1}(q)$. Here we assume that $\ell \geqslant 3$ and that q is odd. Now [FG12, Theorem 3.19] along with the values for $P(S)$ given in [GMPS15] imply that $S = \mathrm{P}\Omega_7(3)$. This final case can be excluded using [CCN$^+$85].

Suppose that $S \cong \mathrm{P}\Omega^{\pm}_{2\ell}(q)$ with q odd. We make use of [FG12, Theorems 3.16 and 3.18] along with the values for $P(S)$ given in [GMPS15] to obtain that

$$S \in \{\mathrm{P}\Omega^{\pm}_{10}(3), \mathrm{P}\Omega^{\pm}_{8}(3), \mathrm{P}\Omega^{+}_{8}(5), \mathrm{P}\Omega^{+}_{8}(7)\}.$$

In $\mathrm{P}\Omega^{+}_{8}(5)$ and $\mathrm{P}\Omega^{+}_{8}(7)$, the outer automorphism group is S_4, and a subgroup of S_4 has at most 5 conjugacy classes, therefore by Lemma 4.2.1 we get $k(G) \leqslant 5k(S)$, which is enough to rule out these possibilities.

We use [BCP97] and [GAP19] to rule out the cases where $S = \mathrm{P}\Omega^{\pm}_{10}(3)$ or $\mathrm{P}\Omega^{\pm}_{8}(3)$.

Suppose that $S \cong \Omega^{\pm}_{2\ell}(q)$ with q even. We make use of [FG12, Theorem 3.22] along with the values for $P(S)$ given in [GMPS15] to obtain that

$$S \in \{\Omega^{\pm}_{10}(2), \Omega^{\pm}_{8}(2), \Omega^{+}_{8}(4)\}.$$

We use [CCN$^+$85] for the groups with $q = 2$, and we get the listed examples. We can rule out $\Omega^{+}_{8}(4)$ using [BCP97]. □

4.2.5 Proof of Theorem 4.1.2

Let G be an almost simple primitive permutation group of degree n. Putting together Lemmas 4.2.3, 4.2.8, 4.2.11, 4.2.12 and 4.2.13, we get that either $k(G) < n/2$, or we are in case (2) of Theorem 4.1.2 (regarding Table 4.1, recall Remark 4.1.3(iii)).

Note that, if the action of G is isomorphic to an action in (B), then $k(G) < 100n$ follows immediately from Theorem 4.2.10.

It remains to prove the asymptotic statement, that is, either $k(G) = O(n^{3/4})$, or the action of G is isomorphic to an action in (A) or (B). We assume that this latter condition does not hold, and we want to show $k(G) = O(n^{3/4})$.

We may assume that G is sufficiently large along the proof. Let M be the stabilizer of a point in the action of G on n points; in particular $|G : M| = n$. Write $S = \mathrm{Soc}(G)$.

Assume first $S \cong A_d$, and assume M is transitive on d points; we want to show $k(G) = n^{o(1)}$ as d tends to infinity. By Lemma 4.2.4 (or by Hardy–Ramanujan asymptotic formula), we have $k(G) = O(1)^{\sqrt{d}}$. On the other hand, by Lemmas 4.2.5 and 4.2.6, if M is primitive on d points then $n \geqslant (d/O(1))^d$;

and by (4.2.1) in the proof of Lemma 4.2.8, if M is imprimitive then $n \geqslant c^d$ for some constant c. Therefore $k(G) = n^{o(1)}$ if $S \cong A_d$.

Assume now $S \cong \mathrm{PSL}_d(q)$. We have $k(G) = O(q^{d-1})$ by Theorem 4.2.10. If $H := M \cap \mathrm{P\Gamma L}_d(q)$ is reducible in the action on $\mathbf{F}_q^d$, one possibility is that it stabilizes a k-space for some $2 \leqslant k \leqslant d-2$, and so $n > q^{2d-4}$. If $d \to \infty$, we see that $k(G) = o(n^{3/4})$; and if d is bounded, we see that $k(G) = O(n^{3/4})$ (we actually have $k(G) = o(n^{3/4})$ as $q \to \infty$ except for the case $(d,k) = (4,2)$). The remaining possibility is that $G \not\leqslant \mathrm{P\Gamma L}_d(q)$ and M is the stabilizer of a flag (pair of incident point-hyperplane) or antiflag (pair of complementary point-hyperplane). But in this case $n > q^{2d-3}$, and the previous computation is sufficient for $d \geqslant 4$; and for $d = 3$, $k(G) = O(n^{2/3})$.

If $d = 2m \geqslant 4$ and H normalizes $\mathrm{PSp}_{2m}(q)$, then $n \geqslant \frac{1}{(m,q-1)} q^{m^2-m}(q^3-1)(q^5-1)\cdots(q^{2m-1}-1)$ and we can easily check that $k(G) = o(n^{3/4})$.

If now H is irreducible and does not normalize $\mathrm{PSp}_d(q)$, we can apply the main theorem of [Kan79]. We see easily that $k(G) = o(n^{3/4})$ as $d \to \infty$. If d is bounded instead, then we can assume q is large and in particular [Kan79] implies that $n \geqslant q^{(d-1)(d-2)/2}$, which proves $k(G) = o(n^{3/4})$ in case $d \geqslant 5$. In case $d \leqslant 4$, we can use the list of maximal subgroups of $\mathrm{PSL}_d(q)$ given in [BHRD13] in order to prove $k(G) = o(n^{3/4})$ (if H is irreducible, $n \gg q^{3/2}$ for $d = 2$; $n \gg q^4$ for $d = 3$; and $n \gg q^5$ for $d = 4$).

Assume finally S is a group of Lie type and $S \ncong \mathrm{PSL}_d(q)$. In this case we want to show $k(G) = O(P(S)^{3/4})$, which implies the statement, since $P(S) \leqslant n$. This can be checked combining $k(G) = O(q^r)$ (where r is the untwisted rank of S) with the value of $P(S)$ given in [GMPS15]. In fact, we get $k(G) = o(P(S)^{3/4})$ unless $S \cong \mathrm{PSU}_4(q)$. (We remark that, in the latter case, $P(S)$ is equal to the number of totally singular 2-subspaces of $\mathbf{F}_{q^2}^4$; we also use [BHRD13] in order to see that $n \gg q^5$ for every other faithful primitive action of G.)

This concludes the proof of Theorem 4.1.2.

Remark 4.2.14. In Theorem 4.1.2(1), we actually showed that $k(G) = o(n^{3/4})$ as $n \to \infty$ unless $S \cong \mathrm{PSL}_4(q)$ and G acts on the set of 2-subspaces of $\mathbf{F}_q^4$, or $S \cong \mathrm{PSU}_4(q)$ and G acts on the set of totally singular 2-subspaces of $\mathbf{F}_{q^2}^4$. In these cases, we have $n \sim q^4$ and $k(S) \asymp q^3$, therefore $k(S) \asymp n^{3/4}$.

4.3 The general case

In this section we prove Theorem 4.1.1. We first prove a lemma.

Lemma 4.3.1. *Let G be a finite almost simple group with socle S. Then either $S \cong A_5, A_6, \mathrm{PSL}_2(7), \mathrm{PSL}_2(11)$, or $4 \cdot k(G)^2 < |S|$. Moreover, $k(G)^3 = O(|S|)$.*

We note that we actually have $k(G)^3 = o(|S|)$ as $|S| \to \infty$, except for the case $S \cong \mathrm{PSL}_2(q)$.

Proof. We first prove $4 \cdot k(G)^2 < |S|$, with the listed exceptions.

Assume first $S \cong A_d$. Then Lemmas 4.2.4, 4.2.5 and a straightforward computation imply that it is sufficient to show

$$3.2 \cdot e^{5.2\sqrt{d}+d} < d^{d+2},$$

which is true for $d \geqslant 10$. For $d \leqslant 9$, direct check gives the exceptions in the statement.

Assume now $S \cong \mathrm{PSL}_d(q)$. Using the bound $k(S) \leqslant 2.5q^{d-1}$ from [FG12], Lemma 4.2.1, and the fact that $|G : S| \leqslant 2f(d, q-1)$, with $q = p^f$, we see that it is sufficient to show

$$100f^2(d, q-1)^3 < q^{d(d-1)/2-2d+2}(q^2-1)\cdots(q^d-1).$$

If $d \geqslant 4$, we can easily verify that this is true. For $d = 3$, [Mac81] tells us that $k(S) \leqslant q^2 + q$. We compute that it is enough to show

$$16f^2(3, q-1)^3(q+1) < q(q-1)(q^3-1),$$

which can be verified unless $q = 2, 4$. The case $q = 2$ is in the statement (since $\mathrm{SL}_3(2) \cong \mathrm{PSL}_2(7)$), while the case $q = 4$ can be excluded with [GAP19].

If $d = 2$, we use the exact value of $k(S)$ (recalled in the proof of Lemma 4.2.12), in order to reduce to the cases $q \leqslant 16$ or $q = 25, 27, 32, 64, 81, 128, 256$. Then we use [GAP19] and we get the cases $q = 4, 5, 7, 9, 11$ in the statement.

Assume S is classical and $S \ncong \mathrm{PSL}_d(q)$. Here one can prove that $4k(G)^2 < |S|$ using the upper bounds for $k(S)$ given in [FG12]. One can also argue as follows (but this is not necessary). If G appears in Table 4.1, we can make a direct check. If G is not in Table 4.1, then Theorem 4.1.2 tells us that $k(G) < P_m(G)/2$, where $P_m(G)$ denotes the smallest index of a core-free maximal subgroup of G. Now it is known (see [KL90b, p. 178]) that $P(S) \leqslant |S|^{1/2}$. In particular, whenever $P_m(G) = P(S)$, we can immediately conclude $4k(G)^2 < |S|$. Certainly we have $P(S) \leqslant P_m(G)$. Using the value of $P(S)$ given in [GMPS15] (see also [Coo78], where an explicit M for which $|S : M| = P(S)$ is given), and consulting [KL90b], we deduce that $P_m(G) = P(S)$ unless $S \cong \mathrm{PSU}_3(5)$, $S \cong \mathrm{Sp}_4(q)$ with q even, $S \cong \mathrm{P}\Omega_8^+(q)$, or $S \cong \mathrm{P}\Omega_{2m}^+(3)$ with $m \geqslant 4$. (If $S \cong \mathrm{PSU}_3(5)$, $|S : M| = P(S)$ where M is isomorphic to A_7; if $S \cong \mathrm{P}\Omega_{2m}^+(3)$, M is the stabilizer of a nondegenerate 1-space.) We can exclude the unitary case with [CCN+85]; in the symplectic case we can use $k(\mathrm{Sp}_4(q)) = q^2 + 2q + 3$ (see the proof of Lemma 4.2.13); in the orthogonal cases we can use the bound $k(\mathrm{P}\Omega_{2m}^+(q)) \leqslant 14q^m$ given in [FG12].

Assume S is exceptional. In the proof of Lemma 4.2.11 we actually proved $k(G) < P(S)/2$, therefore we conclude by the argument of the previous paragraph.

Assume finally S is sporadic. We use [CCN+85] to conclude $4k(G)^2 < |S|$.

It remains to prove the asymptotic statement, that is, $k(G)^3 = O(|S|)$ (and indeed $k(G)^3 = o(|S|)$ if $S \ncong \mathrm{PSL}_2(q)$). We may assume that S is sufficiently large, and the statement is easy to check, using Lemma 4.2.4 and Theorem 4.2.10. □

We need other three technical lemmas.

Lemma 4.3.2. *Assume $S \cong A_d$, or S is the socle of some group appearing in Table 4.1. If $S \leqslant B \leqslant A \leqslant \mathrm{Aut}(S)$, then $k(B) \leqslant k(A)$, unless $A = \Omega_8^+(2).S_3$.*

Proof. If $S \cong A_d$, the statement follows from Lemma 4.2.7, and by direct check in case $d = 6$. If S is the socle of some group appearing in Table 4.1, we use [CCN+85]. □

Lemma 4.3.3. *Assume G is almost simple with socle $S \cong \mathrm{PSL}_d(q)$, with $d \geqslant 3$, and let m denote the number of flags (that is, pairs of incident point-hyperplane) in $\mathbf{F}_q^d$. Then, $k(G) < m/2$ and $k(G) = O(m^{2/3})$.*

We note that we actually have $k(G) = o(m^{2/3})$ as $m \to \infty$, except in case $d = 3$.

Proof. We begin with the inequality $k(G) < m/2$. If $G \not\leqslant \mathrm{P\Gamma L}_d(q)$, then G acts primitively on the set of flags, and the statement follows from Theorem 4.1.2. Assume now $G \leqslant \mathrm{P\Gamma L}_d(q)$. Then G acts primitively on the set of 2-subspaces of $\mathbf{F}_q^d$. It is easy to see that the number of 2-subspaces is smaller than the number of flags. Assume $d \geqslant 4$. Then, by Lemma 4.2.12, either $k(G) < m/2$, or G appears in Table 4.4. Examining Table 4.4, we see that $k(G) < m/2$ also in the latter case.

We are left with the case $d = 3$. We have $k(G) \leqslant 100q^2$ by Theorem 4.2.10, and moreover $m > q^3$. In particular, if $k(G) \geqslant m/2$ then $q < 200$. We whittle down a bit the possibilities. Write $a = (3, q-1)$. By [Mac81] and Lemma 4.2.1, we deduce $k(G) \leqslant |G : S| \cdot (q^2 + q + 5a - 5)/a < 2|G : S| \cdot q^2$. Therefore, if $q = p^f$, we have $q < 8af$. Using $q < 200$, we see that we are reduced to the cases $q \leqslant 27$ and $q = 32, 64$, which can be checked with [GAP19] (if $q \neq 2, 4, 8, 16$, it is enough to show that $4f(q^2+q+5a-5)$ is smaller than m, without computing the actual value of $k(G)$).

The asymptotic statement $k(G) = O(m^{2/3})$ can be checked easily using $k(G) = O(q^{d-1})$. □

Lemma 4.3.4. *Let A be an almost simple primitive group of degree m with socle S, and assume A is not in the possibilities of Theorem 4.1.2(2). Let $S \leqslant B \leqslant A$. Then, $k(B) < m/2$. Moreover, for every fixed $\alpha > 3/4$, if S is sufficiently large then $k(B) < m^\alpha$.*

Proof. We begin with the inequality $k(B) < m/2$. Write $m = |A : M|$ for some core-free maximal subgroup M of A. If $B = A$, the claim is true by Theorem 4.1.2. Assume by contradiction there exists B such that $k(B) \geqslant m/2 = |B : B \cap M|/2$. Let T be a core-free subgroup of B, maximal with respect to the property that $B \cap M \leqslant T$ and that T is core-free in B (that is, T does not contain S).

Note that the subgroups of B properly containing T must contain S. Then choose C such that $T < C \leqslant B$ and T is maximal in C. In particular, C acts primitively on the cosets of T, and moreover, by Lemma 4.2.1, $|B : C|k(C) \geqslant$

$k(B) \geqslant |B : C||C : T|/2$, whence $k(C) \geqslant |C : T|/2$. Therefore we can apply Theorem 4.1.2. The first possibility is that C appears in Table 4.1. By Lemma 4.3.2 and $k(B) \geqslant m/2$, we deduce $A = \Omega_8^+(2).S_3$. Then by [CCN+85] $m \geqslant 3600$, which contradicts $k(B) \geqslant m/2$. By Lemma 4.3.4, we also see that it cannot be $S \cong A_d$. By Theorem 4.1.2 and Lemma 4.2.12, the only remaining possibility is that $S \cong \mathrm{PSL}_d(q)$, $C \leqslant \mathrm{P\Gamma L}_d(q)$ and T is the stabilizer of a 1-space or $(d-1)$-space. In particular, $B \cap M$ stabilizes a 1-space or a $(d-1)$-space.

Assume first $A \leqslant \mathrm{P\Gamma L}_d(q)$. By assumption, M is not the stabilizer of a 1-space or $(d-1)$-space. Then, there is no other possibility for M (in such a way that $B \cap M$ fixes a 1-space or $(d-1)$-space), which is a contradiction. Assume finally $A \not\leqslant \mathrm{P\Gamma L}_d(q)$. Then the only possibility is that M is the stabilizer of a flag or antiflag. In particular, m is larger than the number of flags in $\mathbf{F}_q^d$, which contradicts Lemma 4.3.3. This final contradiction proves that $k(B) < m/2$ for every $S \leqslant B \leqslant A$.

Now we want to show that, for every fixed $\alpha > 3/4$, if S is sufficiently large then $k(B) < m^\alpha$ for every $S \leqslant B \leqslant A$.

By Theorem 4.1.2, we have $k(A) = O(m^{3/4})$. Assume $k(B) \geqslant m^\alpha$. We want to show that S has bounded order (in other words, we want to show that, if S is sufficiently large, we get a contradiction). By taking S large, we have $k(B) > k(A)$. Much of the argument of the first part of the proof carries unchanged, except that we have the inequality $|B : C|k(C) \geqslant k(B) \geqslant |B : C|^\alpha \cdot |C : T|^\alpha$, from which $k(C) \geqslant |C : T|^\alpha \cdot |B : C|^{\alpha-1}$. Note that $|B : C| \leqslant |\mathrm{Out}(S)|$ and $|C : T| \geqslant P(S)$. Using [GMPS15, Table 4], we easily see that $|\mathrm{Out}(S)| = P(S)^{o(1)}$ as $|S| \to \infty$ (the statement being obvious in case $S \cong A_d$), from which we get that, for every fixed $\beta < \alpha$, $k(C) \geqslant |C : T|^\alpha \cdot |B : C|^{\alpha-1} > |C : T|^\beta$ if S is sufficiently large. In particular we may take $\beta > 3/4$, and by Theorem 4.1.2, we deduce that C and $|C : T|$ must appear in item (2) of the theorem. Then, the argument that we used in the first part of the proof, together with Lemma 4.3.3, gives a contradiction. □

4.3.1 Proof of Theorem 4.1.1

We can now prove Theorem 4.1.1. We will apply many times Lemma 4.2.1, usually with no mention. Moreover, we will often use the following theorem from [LP97], which we recalled in the introduction of this chapter.

Theorem 4.3.5. *Let $r \geqslant 1$ and let $P \leqslant S_r$. Then, $k(P) \leqslant 2^{r-1}$.*

Let now G be a primitive permutation group of degree n with nonabelian socle $\mathrm{Soc}(G) \cong S^r$, with S simple.

In the following proof, a permutation group G of degree n in *product action* refers to a group $G \leqslant A \wr S_r$, where A is almost simple primitive on m points with socle S and G acts on $n = m^r$ points (so we do not include the actions that sometimes are called *holomorph compound* and *compound diagonal*).

Proof of Theorem 4.1.1. **Assume first** the action of G is not product action; we want to show $k(G) < n/2$ and $k(G) = O(n^\delta)$ for some absolute $\delta < 1$. We

begin with the first inequality.

We have $r \geqslant 2$ and either $n = |S|^r$, or $r = \ell t$ with $\ell \geqslant 2$, $t \geqslant 1$ and $n = |S|^{(\ell-1)t}$. In particular $n \geqslant |S|^{r/2}$. Furthermore, $G \leqslant \mathrm{Aut}(S) \wr S_r$. Then, by Lemma 4.2.1 and Theorem 4.3.5, $k(G) \leqslant k(G \cap \mathrm{Aut}(S)^r) \cdot 2^{r-1}$. Now $G \cap \mathrm{Aut}(S)^r$ admits a normal series of length r in which every factor is almost simple with socle S; therefore, by Theorem 4.3.5, $k(G \cap \mathrm{Aut}(S)^r) \leqslant f(S)^r$, where $f(S) = \max\{k(A) : S \leqslant A \leqslant \mathrm{Aut}(S)\}$. We deduce that it is enough to show that

$$2f(S) < |S|^{1/2}.$$

By Lemma 4.3.1, this is true unless $S \cong A_5, A_6, \mathrm{PSL}_2(7), \mathrm{PSL}_2(11)$. Assume then we are in one of these cases. If $n = |S|^r$ or $n = |S|^{(\ell-1)t}$ with $\ell \geqslant 3$, then $n \geqslant |S|^{2r/3}$, hence by the same argument as above we have $k(G) < n/2$ provided

$$2f(S) < |S|^{2/3}.$$

We can check that this is true. Therefore we are reduced to the case in which $S \in \{A_5, A_6, \mathrm{PSL}_2(7), \mathrm{PSL}_2(11)\}$, $r = 2t$ and $n = |S|^t$.

Assume first $t = 1$, and let $h(S)$ be the maximum number of conjugacy classes of a primitive group on $|S|$ points with socle S^2. We can use [GAP19] in order to compute that $h(S) < |S|/2$.

Next we deal with any $t \geqslant 1$. We have $G \leqslant D \wr S_t$, where D has socle S^2 and is primitive on $|S|$ points. Then $k(G) \leqslant k(G \cap D^t) \cdot 2^{t-1}$. Now $G \cap D^t$ admits a normal series of length t in which every factor has socle S^2 and is primitive on $|S|$ points; in particular $k(G \cap D^t) \leqslant h(S)^t < (|S|/2)^t$ and therefore $k(G) < |S|^t/2 = n/2$, as wanted.

We turn now to the asymptotic statement; namely, $k(G) = O(n^\delta)$ for an absolute $\delta < 1$. We assume that n is sufficiently large and we show $k(G) \leqslant n^\delta$ (which is equivalent up to enlarging δ). We will show in various places that $k(G) \leqslant n^{\delta'}$ for various δ'. In order to simplify notation, we will always use the same symbol δ — one should just take the maximum.

Assume first S is sufficiently large. By Lemma 4.3.1, we have $f(S) < |S|^{0.35}/2$. Using $n \geqslant |S|^{r/2}$, we deduce $k(G) < n^{0.7}$.

Assume now S has bounded order. If $S \ncong A_5, A_6, \mathrm{PSL}_2(7), \mathrm{PSL}_2(11)$, by Lemma 4.3.1 we have $2f(S) < |S|^{1/2}$, and in particular

$$k(G) < (2 \cdot f(S))^r < |S|^{r\delta/2} \leqslant n^\delta$$

for some $\delta < 1$ absolute (since $|S|$ is bounded).

Assume then $S \cong A_5, A_6, \mathrm{PSL}_2(7), \mathrm{PSL}_2(11)$. If $n = |S|^r$ or $n = |S|^{(\ell-1)t}$ with $\ell \geqslant 3$, then $n \geqslant |S|^{2r/3}$ and, as already observed, $f(S) < |S|^{2/3}/2$; therefore the same argument as above applies. The remaining case is $l = 2$ and $r = 2t$. We already observed that $2h(S) < |S|$, from which we get

$$k(G) < (2 \cdot h(S))^t < |S|^{t\delta} = n^\delta$$

for some $\delta < 1$ absolute.

Assume now the action of G is product action, and assume we are not in case (2) of the statement. We want to show $k(G) < n/2$ and $k(G) = O(n^\delta)$ for some $\delta < 1$ absolute. We begin with the first inequality. We have $G \leqslant A \wr S_r$, $n = m^r$, and A is an almost simple group with socle S admitting a primitive action on m points, which is not among the possibilities of Theorem 4.1.2(2).

Note that $k(G) \leqslant k(G \cap A^r) \cdot 2^{r-1}$, and $G \cap A^r$ admits a normal series of length r in which each factor is isomorphic to a subgroup $S \leqslant B \leqslant A$. By Lemma 4.3.4, $k(B) < m/2$ for every $S \leqslant B \leqslant A$, and therefore $k(G \cap A^r) < (m/2)^r$ and $k(G) < n/2$, as wanted.

The asymptotic statement $k(G) = O(n^\delta)$ for some $\delta < 1$ is proved as we did for the case in which $n = |S|^r$ or $n = |S|^{(\ell-1)t}$, dividing the cases $|S|$ sufficiently large and $|S|$ bounded. If S is sufficiently large, by Lemma 4.3.4 we have $k(B) < m^{0.8}/2$ for every $B \leqslant S \leqslant A$, and therefore $k(G) < n^{0.8}/2$. If S has bounded order, we only need to use $k(B) < m/2$ for every $S \leqslant B \leqslant A$, which holds again in view of Lemma 4.3.4.

Assume now we are in case (2)(i) of the statement; we want to show $k(G) < n^{1.31}$. We have $G \leqslant A \wr S_r$ and A is almost simple acting primitively on m points.

Let us consider first the case in which $A = M_{12}$ acting primitively on $m = 12$ points. If $r \geqslant 4$, [GM15] tells us that a subgroup of S_r has at most $5^{(r-1)/3} < 5^{r/3}$ conjugacy classes. In particular, using that $k(A) = 15$, we deduce that $k(G) < 15^r \cdot 5^{r/3}$, which we verify to be at most $n^{1.31}$. If $r \leqslant 3$, we use that a subgroup of S_r has at most r conjugacy classes, so $k(G) \leqslant 15^r \cdot r$, which is less than $n^{1.31}$ for $r \leqslant 3$.

Let us consider now all other cases. By Lemma 4.3.2 we have $k(B) \leqslant k(A)$ for every $S \leqslant B \leqslant A$. Then $k(G) \leqslant k(A)^r \cdot 2^{r-1} < (2k(A))^r$, so we only need to show that $2k(A) \leqslant m^{1.31}$. This can be checked easily going through all cases in Table 4.1 (but leaving out the case of M_{12} acting on 12 points).

Assume finally we are in case (2)(ii) of the statement, and the action of A is isomorphic to an action in (B); in particular $m = (q^d - 1)/(q - 1)$. We want to show $k(G) < n^{1.9}$.

If $r \geqslant 4$, by Theorem 4.2.10 we have $k(G) \leqslant (100q^{d-1})^r \cdot 5^{(r-1)/3} < (100 \cdot 5^{1/3})^r \cdot n$, hence we are done provided $100 \cdot 5^{1/3} \leqslant m^{0.9}$, that is, $m \geqslant 303$. If $r \leqslant 3$, we use $k(G) \leqslant (100q^{d-1})^r \cdot r$, and we see that $m \geqslant 303$ is enough also in these cases.

Therefore we assume $m < 303$; this leaves us with the cases $d = 6, 7, 8$ and $q = 2$; or $d = 5$ and $q \leqslant 3$; or $d = 4$ and $q \leqslant 5$, or $d = 3$ and $q \leqslant 16$; or $d = 2$ and $q < 302$.

We whittle down slightly the possibilities for $d = 2$. In the proof of Lemma 4.2.12, we recalled the exact value of $k(\mathrm{PSL}_2(q))$ and $k(\mathrm{PGL}_2(q))$. Using this and $q < 302$, it is easy to deduce that $k(A) \leqslant 8(q + 1) = 8m$. By the same computation as above, we are done provided $8 \cdot 5^{1/3} \leqslant m^{0.9}$, that is, $m \geqslant 19$. Therefore if $d = 2$ then we may assume $q \leqslant 17$.

Now we deal with all the remaining cases (for $d \leqslant 8$). We only need to show that $k(A) \cdot 5^{1/3} \leqslant m^{1.9}$, which can be checked with [GAP19]. □

4.4 Further comments

4.4.1 Theorem 4.1.1(2)(i)

In Theorem 4.1.1(2)(i), we proved $k(G) < n^{1.31}$. Can we get better bounds? Since we have finitely many possibilities for the almost simple primitive group A of degree m, we fix A and m, and we want to estimate $k(G)$ where $G \leqslant A \wr S_r$ is primitive, mainly when r is large.

First, we show that it is not always true that $k(G) = o(n)$ as $n \to \infty$ (and in fact it is not even true that $k(G) = O(n)$).

Lemma 4.4.1. *Consider $A = M_{12}$ acting primitively on 12 points, and consider $G = A \wr C_r$ acting on $n = 12^r$ points, where C_r is cyclic of order r. If r is large enough, then $k(G) > n^{1.08}$.*

Proof. By Lemma 4.2.1, we have

$$k(G) \geqslant \frac{k(A)^r}{r}.$$

Since $k(A) = 15$, this is easily seen to be larger than $n^{1.08}$ for r large enough. □

The same argument shows that $k(G) > n^\alpha$ for some absolute $\alpha > 1$ whenever A and m in Table 4.1 are such that $k(A) > m$ (but in the table, A and m are replaced by G and n). This happens rarely; specifically, when

$$(A, m) \in \{(M_{12}, 12), (M_{24}, 24), (\mathrm{Sp}_6(2), 28)\}.$$

Let us consider now the case in which $k(A) \leqslant m$ (by looking at Table 4.1, this is equivalent to $k(A) < m$). By Lemma 4.3.2, we have $k(B) \leqslant k(A)$ for every subgroup $S = \mathrm{Soc}(A) \leqslant B \leqslant A$. We assume $r \geqslant 4$, so that by [GM15] a subgroup of S_r has at most $5^{(r-1)/3} < 5^{r/3}$ conjugacy classes. Then, we have $k(G) < (k(A)5^{1/3})^r$, and whenever $k(A) \cdot 5^{1/3} < m$ we get $k(G) < n^\delta$ for some absolute $\delta < 1$. The condition $k(A) \cdot 5^{1/3} < m$ holds in some cases, but not quite in all.

Therefore one should try to change the argument. We make the following conjecture.

Conjecture 4.4.2. Let A be an almost simple primitive group on m points appearing in Table 4.1, and assume $k(A) < m$. Then, for every primitive subgroup $G \leqslant A \wr S_r$ on $n = m^r$ points, $k(G) = o(m^r)$ as $r \to \infty$.

In order to address Conjecture 4.4.2, it seems relevant to estimate the number of conjugacy classes in wreath products (although G needs not be a full wreath product, which is a complication).

4.4.2 Theorem 4.1.1(2)(ii)

Let us briefly comment the situation in Theorem 4.1.1(2)(ii). We have $G \leqslant A \wr S_r$ where A is almost simple primitive on m points. [Mar05] tells us that $k(G) \leqslant p(n)$, which can be attained if the action of A is isomorphic to an action in (A). Assume now the action of A is isomorphic to an action in (B); then we showed $k(G) < n^{1.9}$. We did not make any particular effort in order to sharpen this bound.

Question 4.4.3. *Let A be an almost simple primitive group isomorphic to a group in (B), and assume $G \leqslant A \wr S_r$ is primitive on $n = m^r$ points. Improve the bound $k(G) < n^{1.9}$.*

In case $G = A$, Theorem 4.1.2 tells us that $k(G) < 100n$ (this follows immediately from Theorem 4.2.10). Probably a much better bound holds also in the case $G = A$.

Recall that, if q is odd, and if $A = \mathrm{PGL}_2(q)$, then $k(A) = q + 2$. If we take for instance $q = 5$, by the same argument as in Lemma 4.4.1 we see that $k(A \wr C_r) > n^{1.08}$ for r sufficiently large (as a coincidence, we get the same bound as in Lemma 4.4.1). Therefore it is not true in general that $k(G) = O(n)$.

On the other hand, by the usual bound $k(G) < (100q^{d-1})^r \cdot 2^r$, we note that, for every fixed $\epsilon > 0$, $k(G) < n^{1+\epsilon}$ if $\mathrm{PSL}_d(q)$ is sufficiently large (and this holds independently of r).

Chapter 5

Large minimal invariable generating sets of S_n

The content of this chapter is joint work with Nick Gill.

5.1 Introduction

We want to estimate the invariant $m_I(G)$, introduced in Chapter 3, in case G is the symmetric group S_n. Our main theorem is as follows.

Theorem 5.1.1. *Let $n \geqslant 5$ be an integer. Then*

$$\frac{n}{2} - \log n < m_I(S_n) < \frac{n}{2} + \Delta(n) + O\left(\frac{\log n}{\log\log n}\right),$$

where $\Delta(n)$ is the number of divisors of n.

(Logarithms are in base 2.) It is well known that $\Delta(n) = n^{o(1)}$, therefore Theorem 5.1.1 implies that $m_I(S_n) \sim n/2$ as $n \to \infty$.

In the upper bound in Theorem 5.1.1, we will in fact prove a more explicit estimate, which will have the following consequence.

Corollary 5.1.2. *Let $n \geqslant 5$ be an integer. Then, $m_I(S_n) < m(S_n)$.*

Notice that, by Theorem 3.1.1, we have $m_I(S_n) = m(S_n)$ for $n \leqslant 4$. In particular, we get a positive answer to Question 3.1.2 in case $G = S_n$.

5.1.1 Method of proof: lower bound

In order to prove the lower bound in Theorem 5.1.1, we will exploit Lemma 3.2.2.

Each conjugacy class C of S_n corresponds to a particular partition, $\mathfrak{p}_C$, of the integer n. On the other hand, if M is an *intransitive* subgroup of S_n, then M is the stabilizer of some i-subset of $\{1, \ldots, n\}$.

We say that *the integer i is a partial sum of the partition $\mathfrak{p} = (a_1, \ldots, a_t)$* if we can write $i = a_{j_1} + a_{j_2} + \cdots + a_{j_\ell}$ for some $1 \leqslant j_1 < \cdots < j_\ell \leqslant t$. It is clear that the intersection $C \cap M$ is non-empty if and only if i is a partial sum of $\mathfrak{p}_C$. We will prove the following:

Proposition 5.1.3. *Let $n \geqslant 5$ be an integer. There is a set X of partitions of n with the following properties:*

(1) there is no integer $1 \leqslant i \leqslant \frac{1}{2}n$ which is a partial sum in x for every $x \in X$;

(2) for every $x \in X$, there exists an integer $1 \leqslant i \leqslant \frac{1}{2}n$ which is a partial sum in y for every $y \in X \setminus \{x\}$;

(3) $|X| > \frac{1}{2}n - \log n$.

By Lemma 3.2.2, wee see that Proposition 5.1.3 is almost enough to yield the lower bound in Theorem 5.1.1 straight away. To complete the proof of that lower bound, we must take care of Lemma 3.2.2(b) for proper transitive subgroups of S_n.

Our feeling is that the construction we give in our proof of Proposition 5.1.3 is pretty close to give as large a set X as is possible.

Question 5.1.4. *Is it true that the largest cardinality of a set X of partitions of n satisfying properties (1) and (2) of Proposition 5.1.3 is at most $\frac{n}{2} - \log n + O(1)$?*

Note that we certainly have $|X| \leqslant \frac{1}{2}n$.

5.1.2 Method of proof: upper bound

For a finite group G, let $k(G)$ denote the number of conjugacy classes of G.

In order to prove the upper bound in Theorem 5.1.1, we will exploit Lemma 3.2.3. In particular, we will give an upper bound to $\iota(G)$.

For every $i = 1, \ldots, t$, let M_i be a maximal subgroup of G. With notation as in Lemma 3.2.3, assume that $\{M_1^*, \ldots, M_t^*\}$ is independent. The following hold:

- For every $j = 1, \ldots, t$, M_j has non-empty intersection with at least $t - 1$ non-trivial G-conjugacy classes, and therefore $k(M_j) \geqslant t$.
- For every $i \neq j$, M_i and M_j are not G-conjugate.

We will prove the following proposition.

Proposition 5.1.5. *Suppose that $\{M_1, \ldots, M_t\}$ is a set of maximal subgroups of the symmetric group S_n such that (a) $k(M_i) \geqslant \frac{1}{2}n$ for every i; (b) if $i \neq j$, then M_i and M_j are not S_n-conjugate. Then*

$$t \leqslant \frac{n}{2} + \Delta(n) + O\left(\frac{\log n}{\log\log n}\right), \tag{5.1.1}$$

where $\Delta(n)$ is the number of divisors of n.

By the considerations preceding the proposition, we see that Proposition 5.1.5 gives an upper bound to $\iota(S_n)$. In particular, the upper bound in Theorem 5.1.1 follows immediately from Proposition 5.1.5.

The main point in the proof of Proposition 5.1.5 is to deal with the family of almost simple primitive subgroups of S_n; see Theorem 5.3.1. The key ingredient is Theorem 4.1.2, which we proved in Chapter 4.

We remark that, although Proposition 5.1.5 only states an upper bound for the number of maximal subgroups with at least $\frac{1}{2}n$ conjugacy classes, results in Section 5.3 outline specific families of maximal subgroups. In particular, the first two terms of (5.1.1) correspond to the intransitive and imprimitive subgroups of S_n, respectively.

This is important because our original aim in this chapter was to prove that $|m_I(G) - \frac{n}{2}| = O(\log n)$. We have managed this with the lower bound but not with the upper, precisely because $\Delta(n) - 2$, which is the number of imprimitive subgroups of S_n, is not $O(\log n)$. To achieve our original aim, it would be sufficient to establish that, in the following question, $t \leqslant n/2 + O(\log n)$. We state the question in terms of properties of S_n — it is easy enough to recast it as a number-theoretic question concerning partitions, similar to Question 5.1.4 above. We use notation as in Lemma 3.2.3.

Question 5.1.6. *For every $i = 1, \ldots, t$, let M_i be a subgroup of S_n, which is either intransitive or imprimitive. Assume that $\{M_1^*, \ldots, M_t^*\}$ is independent. How large t can be?*

In truth, we believe that, at least for large enough n, a minimal invariable generating set of S_n of size $m_I(S_n)$ should concern only intransitive subgroups (in the sense that the set J from Lemma 3.2.2 should contain only intransitive subgroups). This would imply $m_I(S_n) \leqslant n/2$, and the problem of determining $m_I(S_n)$ would be reduced to the purely combinatorial problem addressed in Proposition 5.1.3 and Question 5.1.4.

5.1.3 Context

We first remark that, for large enough n, Corollary 5.1.2 follows quickly from results by Liebeck–Shalev [LS96]. The reason is as follows. In [LS96], it is proved that S_n has at most $\frac{n}{2} + o(n)$ conjugacy classes of maximal subgroups. Therefore, $m_I(S_n) \leqslant \iota(G) \leqslant \frac{n}{2} + o(n)$. On the other hand, it is easy to see that $m(S_n) \geqslant n - 1$ — just consider the transpositions $(1, 2), \ldots, (n-1, n)$. In particular, we have $m(S_n) - m_I(S_n) \to \infty$ as $n \to \infty$.

Using CFSG, Whiston [Whi00] proved that in fact $m(S_n) = n - 1$. But more is true. Cameron–Cara [CC02] showed that a minimal generating set of S_n of size $n-1$ is very restrictive: either it is made of $n-1$ transpositions, or it is made of a transposition, some 3-cycles, and some double transpositions (see [CC02, Theorem 2.1] for a precise statement).

One can hardly hope for a similar "elegant" result for $m_I(S_n)$, for the simple reason that a minimal invariable generating set of S_n of size t must contain t elements whose cycle types have no common partial sum. Still, it is true that

in the proof of the lower bound in Theorem 5.1.1, we feel somewhat restricted about the choice of the relevant partitions — but we are not able to make any precise statement in this direction.

5.1.4 Organization of the chapter and notation

In Section 5.2 we prove the lower bound on $m_I(G)$ given in Theorem 5.1.1. In Section 5.3 we prove the upper bound on $m_I(G)$ given in Theorem 5.1.1, along with Corollary 5.1.2.

We will use exponential notation for partitions, so the partition $(a_1^{n_1}, a_2^{n_2}, \ldots, a_t^{n_t})$ has n_1 parts of length a_1, n_2 parts of length $a_2, \ldots$, and n_t parts of length a_t. For a positive real number x, $\log(x)$ denotes a logarithm in base 2, and in one occasion we will write $\exp_2\{x\} := 2^x$. For a positive integer x, $\Delta(x)$ denotes the number of divisors of x.

5.2 The lower bound

In this section we prove the lower bound in Theorem 5.1.1. We first prove a lemma, then we prove Proposition 5.1.3, and finally we give a proof for the lower bound.

Lemma 5.2.1. *Let n and i be positive integers, with $i < n/3$. Then there exist a partition $\mathfrak{p}_{i,n}$ of n with the following properties:*

(1) If $n \neq 4i+2$ and $(n,i) \neq (8,1)$, then $\mathfrak{p}_{i,n}$ does not have i and $n-i$ as partial sums, and everything else is a partial sum.

(2) If $n = 4i+2$ or $(n,i) = (8,1)$, then $\mathfrak{p}_{i,n}$ does not have $i, n-i$ and $\frac{n}{2}$ as partial sums, and everything else is a partial sum.

Proof. Define

$$\mathfrak{p}_{i,n} = (1^{i-1}, i+1, (i+2)^j, (i+1)^k, c)$$

where $j \in \{0,1\}$, $k \geqslant 0$, and $i+1 \leqslant c \leqslant 2i+1$. To complete the definition we must specify j, k and c. To do this we consider a partial sum q, adding from left to right: we first sum the 1's and $(i+1)$ to obtain $q = 2i$. Now there are three cases:

(1) If $n - q \leqslant 2i+1$, then we set $c = n - q$.

(2) If $n - q = 2i+2$, then we set $k = 1$, $j = 0$ and $c = i+1$.

(3) If $n - q \geqslant 2i+3$, then we set $j = 1$, $k = 0$ and set $q = 2i + (i+2) = 3i+2$.

In the first and second cases, we are done; notice that the partition has the stated properties (in the first case we use the fact that $i < n/3$ to obtain $i+1 \leqslant c \leqslant 2i+1$ as required). If we are in the third case, then we proceed in a loop as follows:

(1) If $n - q \leqslant 2i + 1$, then we set $c = n - q$.

(2) If $n - q \geqslant 2i + 2$, then we set $k = k + 1$ and set $q = q + (i + 1)$.

It turns out that there is one situation — when $n = 4i+4$ and $(n, i) \neq (8, 1)$ — where our definition needs to be adjusted. In this case, we make the following definition:

$$\mathfrak{p}_{i,n} = (1^{i-1}, i + 1, i + 3, i + 1).$$

Now our definition is complete. We now let m be an integer such that $1 \leqslant m \leqslant n/2$ and we study when m is a partial sum of $\mathfrak{p}_{i,n}$ with a view to proving items (1) and (2) of the lemma.

Both items are clear for $m \leqslant 2i$, thus we may assume that $m \geqslant 2i + 1$. In particular this means that $n \geqslant 4i + 2$.

If $n = 4i + 2$, then we are in item (2) of the lemma, $\mathfrak{p}_{i,n} = (1^{i-1}, (i + 1)^3)$, and the statement holds.

If $n \geqslant 4i + 3$, note first that $j = 1$. Suppose, first, that $k = 0$. There are two possibilities: first, if $c \neq i + 2$, then $\mathfrak{p}_{i,n} = (1^{i-1}, i + 1, i + 2, c)$, and the statement holds. If instead $c = i + 2$, then $n = 4i + 4$. If $(n, i) = (8, 1)$, then $\mathfrak{p}_{1,8} = (2, 3, 3)$ and we are in item (2) of the lemma. Otherwise, we are in the exceptional case in our definition where $\mathfrak{p}_{i,n} = (1^{i-1}, i + 1, i + 3, i + 1)$, and the result holds.

We are left with the case in which $k \geqslant 1$, i.e. $\mathfrak{p}_{i,n}$ contains at least two $(i + 1)$'s (excluding c which may also equal $i + 1$).

We work here by induction on m: assuming that some $2i \leqslant m < n/2$ can be written *without using* c, we want to show that the same holds for $m + 1$. Since $m \geqslant 2i \geqslant i + 1$, in writing m without c we have certainly used at least one of $i + 1$ and $i + 2$. We now divide into three cases.

(1) In writing m we have not used all 1's. Then add a 1.

(2) In writing m we have not used $i + 2$. Then remove an $i + 1$ and add $i + 2$.

(3) In writing m we have used all 1's and $i + 2$. Suppose, first, that at least two $(i + 1)$'s have not been used; then remove all 1's, remove $i + 2$ and add two $(i + 1)$'s and we are done. On the other hand, suppose (for a contradiction) that in writing m as a partial sum all but one of the $i + 1$'s have been used. Then $c + (i + 1) > n/2$ and, since $c \leqslant 2i + 1$, we obtain that $n/2 < 3i + 2$. However the partial sum m has used all 1's, one $i + 1$ and one $i + 2$, so $m \geqslant 3i + 2$. Since $m < n/2$, we get $n/2 > 3i + 2$, which is a contradiction. □

5.2.1 Proof of Proposition 5.1.3

Now we prove Proposition 5.1.3. The proof we give below is constructive — we define an explicit set X with the given properties. We have decided not to define the set X outside of this proof, as the construction is built up in pieces as the proof proceeds.

In deducing the lower bound in Theorem 5.1.1, we will be interested in the properties of the partitions of X listed in the statement of Proposition 5.1.3, rather than their explicit construction. The paragraphs involving exceptions to this are labelled **(C1)**, **(C2)**, **(C3)** and **(C4)** in the following proof.

Proof of Proposition 5.1.3. Throughout the proof, we will use the notation $\mathfrak{p}_{t,n}$ to refer to the partitions in the statement of Lemma 5.2.1 (so we allow any partition having the properties of the statement).

If $5 \leqslant n \leqslant 10$, we have $n/2 - \log n < 2$, and the statement is easy to check. Therefore assume $n \geqslant 11$.

(C1) For $n = 11$, we set $x_1 = (2^2, 3, 4)$, $x_2 = (1, 3^2, 4)$, $x_3 = (1^2, 9)$. For $n = 12$, we set $x_1 = (2^2, 3, 5)$, $x_2 = (1, 3, 4^2)$, $x_3 = (1^2, 10)$. The statement holds by setting $X = \{x_1, x_2, x_3\}$.

From now on we assume that $n > 12$. This has the advantage that in the proof that follows, all partitions of the form $\mathfrak{p}_{i,\ell}$ that we consider will have $\ell > 8$ or $i > 1$, and so we need not worry about the case $(\ell, i) = (8, 1)$ mentioned in Lemma 5.2.1.

For $1 \leqslant t < n/3$, define

$$x_t = \mathfrak{p}_{t,n}.$$

(C2) We want to modify partition x_1. Namely, define

$$x_1 = \begin{cases} (3, 5, 2^k, 4^j) & \text{if } n = 14, 16 \\ (3, 2^k, 4^j) & \text{if } n = 13, 15, 17 \\ (3^\ell, 7, 2^k, 4^j) & \text{if } n \geqslant 18 \end{cases}$$

where $\ell = 1$ or 2 according to whether n is even or odd, and $j \in \{0, 1\}$ is defined by the condition that x_1 has an odd number of cycles of even length (and k is consequently uniquely defined). It is easy to see that, in every case, every $2 \leqslant i \leqslant n/2$ is partial sum in x_1.

Now, in order to go further, we will use a slightly different method. The partitions we are going to define will depend on a parameter j. We could define all of them at once, but to give an idea of the overall strategy, let us go through the first step explicitly.

We set $\alpha_1 = \lceil n/6 - 1 \rceil$ and, for every integer $\lceil n/3 \rceil \leqslant t \leqslant t_1 = 5n/12$ we define

$$x_t = (\mathfrak{p}_{\alpha_1, \alpha_1 + t}, c_t).$$

where $c_t = n - \alpha_1 - t$. Let us justify this definition:

(a) Observe first that $1 \leqslant \alpha_1 < (\alpha_1 + t)/3$ hence the partition $\mathfrak{p}_{\alpha_1, \alpha_1 + t}$ is well-defined.

(b) Next note that

$$\begin{aligned} c_t &= n - \alpha_1 - t \\ &> n - \left\lceil \frac{n}{6} - 1 \right\rceil - \frac{5n}{12} \\ &> \frac{5n}{12} \geqslant t > \alpha_1, \end{aligned}$$

and so α_1 and t are not partial sums of x_t.

(c) We can easily check that either $4\alpha_1+2 > \alpha_1+t$ or else $(\alpha_1, t, n) = (1, 5, 12)$. The second possibility is excluded by our assumption $n > 12$. The first possibility implies that Lemma 5.2.1(1) holds, and so all numbers up to $\alpha_1 + t$ are partial sums, apart from α_1 and t.

(d) Finally observe that

$$\begin{aligned} \alpha_1 + t &\geqslant \left\lceil \frac{n}{6} - 1 \right\rceil + \left\lceil \frac{n}{3} \right\rceil \qquad (5.2.1) \\ &\geqslant \left\lceil \frac{n}{2} - 1 \right\rceil. \end{aligned}$$

We conclude that all numbers up to n are partial sums in x_t apart from α_1, t and (possibly) $n/2$. In fact, checking (5.2.1) more carefully, it is clear that $\alpha_1 + t \geqslant \lfloor n/2 \rfloor$ unless $n \equiv 0 \pmod 6$ and $t = n/3$.

Conclusion 1: For $n/3 < t \leqslant 5n/12$, the partition x_t admits all partial sums up to n except α_1 and t.

Conclusion 2: For $t = n/3$, the partition x_t admits all partial sums up to n except α_1 and t and (if n is even) $n/2$.

Now our aim is to extend this definition to other parameters t_i that are larger than t_1. More precisely, for integer $1 \leqslant j \leqslant \log(n/6)$ we define

$$\alpha_j = \left\lceil \frac{n}{2^{j-1} \cdot 6} - 1 \right\rceil, \qquad t_j = \frac{(2^{j-1} \cdot 6 - 1)n}{2^j \cdot 6}.$$

For $j = 1$ this is consistent with the previous definition. Now for integers $2 \leqslant j \leqslant \log(n/6)$ and $\lfloor t_{j-1} \rfloor + 1 \leqslant t \leqslant \lfloor t_j \rfloor$ we define

$$x_t = (\mathfrak{p}_{\alpha_j, \alpha_j + t}, c_t)$$

where $c_t = n - \alpha_j - t$. Now, similarly to before, we must check four properties. Assume that $j \geqslant 2$.

(a) Observe that $1 \leqslant \alpha_j < (\alpha_j + \lfloor t_{j-1} \rfloor + 1)/3$ and so the partition $\mathfrak{p}_{\alpha_j, \alpha_j + t}$ is well-defined.

(b) Next note that $c_t = n - \alpha_j - t > t > \alpha_j$ and so α_j and t are not partial sums of x_t.

(c) Notice that the second case of Lemma 5.2.1 does not occur. Indeed, $4\alpha_j+2$ is strictly smaller than $\alpha_j + \lfloor t_{j-1} \rfloor + 1$. Moreover, it is easy to check that the case $(\ell, i) = (8, 1)$ cannot occur.

(d) Finally observe that $\alpha_j + \lfloor t_{j-1} \rfloor + 1 \geqslant \lfloor n/2 \rfloor$, and we conclude that all numbers up to n are partial sums in x_t apart from α_1 and t.

Conclusion 3: Set $m = \lfloor \log(n/6) \rfloor$. For $j = 2, \ldots, m$ and $\lfloor t_{j-1} \rfloor + 1 \leqslant t \leqslant \lfloor t_j \rfloor$, the partition x_t admits all partial sums up to n, except α_j and t.

We have now constructed $\lfloor t_m \rfloor$ partitions of n; set $X_0 = \{x_1, \ldots, x_{\lfloor t_m \rfloor}\}$. Notice that, by the choice of m, $2^{m+1} \cdot 6 > n$. Then

$$\begin{aligned} |X_0| = \lfloor t_m \rfloor &> \frac{(2^{m-1} \cdot 6 - 1)n}{2^m \cdot 6} - 1 \\ &= \frac{n}{2} - \frac{n}{2^m \cdot 6} - 1 \\ &> \frac{n}{2} - 3. \end{aligned} \tag{5.2.2}$$

Now we will remove some elements from X_0. First, observe that $\alpha_j < n/6$ for every j; we start by taking the subset X obtained by removing x_{α_j} for every $j \geqslant 1$.

(C3) Lemma 5.2.1, and the three conclusions listed above imply that, for each t satisfying $1 \leqslant t \leqslant t_m$, the partition x_t is the unique partition in X which does not admit t as a partial sum. Now we divide into two cases.

(1) There exists an integer belonging to the interval $(t_m, n/2]$ which is a partial sum for all $x_t \in X$. Then, the minimum such integer is $\lfloor t_m \rfloor + 1$; we add one further partition to X:

$$z = (1^{\lfloor t_m \rfloor}, n - \lfloor t_m \rfloor).$$

(2) No integer in $(t_m, n/2]$ is partial sum in all x_t's. Observing that $t_m \leqslant n/2-1$, Lemma 5.2.1 and the three conclusions above imply that $t_m = n/2 - 1$, i.e., $n = 2^m \cdot 6$. In this case we could leave X unchanged, and the statement would be proved. However, we prefer to immediately modify the set X. Notice that $\alpha_m = 1$ and $t_{m-1} = n/2 - 2$; it follows that $x_1 = \mathfrak{p}_{1,n} \notin X$, and X contains a unique partition, namely x_{t_m}, of the form $(\mathfrak{p}_{1,1+t}, c_t)$. Then, we remove such partition and we reintegrate the partition x_1 in X. Moreover, we add to X one further partition

$$z = (1^{n/2-2}, n/2 + 2).$$

(C4) Our construction is finished. Let us make one observation, before concluding the proof. In case (2) above, by construction $x_1 \in X$. We claim that the same holds in case (1). Indeed, one can easily check that $\alpha_m = 1$ if and only if $n = 2^m \cdot 6$, and otherwise $\alpha_j > 1$ for every j. Therefore, in case (1), in our procedure we did not remove x_1 from X_0, hence clearly $x_1 \in X$.

We are now ready to conclude the proof of the statement. The considerations above imply that items (1) and (2) of the statement hold. Regarding item (3),

$$\begin{aligned} |X| &\geqslant |X_0| + 1 - \log(n/6) \\ &\geqslant \frac{n}{2} - 2 - \log n + \log 6 \\ &> \frac{n}{2} - \log n. \end{aligned}$$

The proposition is now proved. □

We now deduce the lower bound of Theorem 5.1.1 from Proposition 5.1.3.

Proof of the lower bound of Theorem 5.1.1. For $5 \leqslant n \leqslant 10$, we have $n/2 - \log n < 2$. Since certainly $m_I(S_n) \geqslant 2$, the statement holds and we may assume $n \geqslant 11$.

Consider the set X of partitions constructed in the proof of Proposition 5.1.3. In this proof we will consider the elements of X as conjugacy classes of S_n. We want to show that X is a minimal invariable generating set for S_n.

It is easy to check the statement for $n = 11, 12$ (see the paragraph **(C1)** in the proof of Proposition 5.1.3). Assume now $n \geqslant 13$. By Proposition 5.1.3(1), the classes of X cannot have non-empty intersection with an intransitive subgroup of S_n. On the other hand, by Proposition 5.1.3(2), if we drop one class from X, then the remaining classes have non-empty intersection with some intransitive subgroup.

Now we deal with transitive groups. Note that X is not contained in A_n, since x_1 corresponds to an odd permutation (see the paragraphs **(C2)** and **(C4)**). Moreover, a power of x_1 corresponds to a cycle of prime length fixing at least 3 points, which belongs to no primitive group different from A_n and S_n by a classical theorem of Jordan. Assume now the classes of X have non-empty intersection with a maximal imprimitive subgroup, preserving a partition of $\{1, \ldots, n\}$ made of $r > 1$ blocks of size $k > 1$. Recall that X contains a partition $z = (1^{n-\ell}, \ell)$, with $\ell = n/2 + 2$ or $\ell = n - \lfloor t_m \rfloor$ (see the paragraph **(C3)**). By (5.2.2) in the proof of Proposition 5.1.3, we have $n/2 - 3 < \lfloor t_m \rfloor < n/2$, and in particular $n/2 < \ell < n/2+3$. We have that k must divide ℓ. Since k also divides n, we get $k < 6$. If $n \neq 15$, then x_1 cannot preserve blocks of size at most 5. If $n = 15$, we note that X contains $x_4 = \mathfrak{p}_{4,15}$, and we may take $\mathfrak{p}_{4,15} = (1^3, 5, 7)$, which does not preserve any nontrivial partition of $\{1, \ldots, 15\}$. The proof is now concluded. □

5.3 The upper bound

In this section we prove the upper bound in Theorem 5.1.1, and we prove Corollary 5.1.2. The main ingredient is the following result, which follows quickly from Theorem 4.1.2. Recall that $k(G)$ denotes the number of conjugacy classes of a finite group G.

n	G	$k(G)$
22	$M_{22}.2$	21
40	$\mathrm{PSU}_4(2).2$	25
45	$\mathrm{PSU}_4(2).2$	25

Table 5.1: Some maximal almost simple primitive subgroups, G of S_n, for which $k(G) \geqslant \frac{n}{2}$. In every case there is a single S_n-conjugacy class of primitive subgroups isomorphic to G.

Theorem 5.3.1. *Let G be a maximal almost simple primitive subgroup of S_n, and assume $k(G) \geqslant \frac{n}{2}$. Then one of the following occurs:*

(1) G is listed in Table 5.1;

(2) $G = A_n$, or $G = S_d$ and the action of G on n points is isomorphic to the action on the set of k-subsets of $\{1, \ldots, d\}$ for some $2 \leqslant k < d/2$.

(3) $G = \mathrm{P\Gamma L}_d(q)$ and the action of G on n points is isomorphic to the action on the set of 1-subspaces of $\mathbf{F}_q^d$.

Note that the subgroups mentioned at item (2) satisfy $n = \binom{d}{k}$ for some integer k with $1 \leqslant k < d/2$; and the subgroups mentioned at item (3) satisfy $n = (q^d - 1)/(q - 1)$.

Proof. The statement follows from Theorem 4.1.2, by checking with [GAP19] which of the entries in Table 4.1 correspond to maximal subgroups of S_n. □

As we observed in the introduction of this chapter, the upper bound in Theorem 5.1.1 follows immediately from Proposition 5.1.5, which we prove now.

Proof of Proposition 5.1.5. We make use of the families of maximal subgroups given in the Aschbacher–O'Nan–Scott theorem, in particular the description given in [LPS88].

(1) **Intransitive subgroups**: There are exactly $\lfloor \frac{n}{2} \rfloor$ conjugacy classes of these.

(2) **Imprimitive subgroups**: There are $\Delta(n) - 2$ of these, where Δ is the divisor function.

(3) **Affine subgroups**: There is at most 1 conjugacy class of these.

(4) **Almost simple subgroups**: If M is almost simple and $k(M) \geqslant \frac{n}{2}$, then it is among the possibilities listed by Theorem 5.3.1, as follows.

  (a) There are three possibilities for degrees $22, 40, 45$ listed in Table 5.1.

(b) There is at most one conjugacy class of maximal subgroups isomorphic to $\mathrm{P\Gamma L}_d(q)$ whenever $n = \frac{q^d-1}{q-1}$; we let a_n be the number of pairs (q, d) where q is a prime power, d is a positive integer, and $\frac{q^d-1}{q-1} = n$.

(c) There is at most one conjugacy class of maximal subgroups with socle A_d whenever $n = \binom{d}{k}$ for some k; we let b_n be the number of pairs (d, k) where d and k are positive integers with $k \leqslant d/2$ and $\binom{d}{k} = n$.

(5) **Diagonal subgroups**: Theorem 4.1.1 states that $k(M) < \frac{n}{2}$ in this case, so we can ignore these subgroups.

(6) **Product action subgroups**: In this case we have maximal subgroups isomorphic to $S_d \wr S_k$, where $n = d^k$ and $k > 1$. For fixed values of d and k, there is one conjugacy class, thus the number of conjugacy classes in S_n is equal to the number of pairs (d, k) where d and k are positive integers with $k > 1$ and $n = d^k$; we write this number as c_n.

(7) **Twisted wreath subgroups**: These are never maximal, as they are defined to be subgroups of groups with a product action [LPS88] and so can be ignored (and in any case, $k(M) < \frac{n}{2}$ by Theorem 4.1.1).

Observe that the number of conjugacy classes of maximal subgroup in S_n that are either imprimitive, affine, or given in Table 5.1 of Theorem 5.3.1 is at most $\Delta(n) - 1$. Therefore, if $\{M_1, \ldots, M_t\}$ is a set of maximal subgroups as in the statement, we have

$$t \leqslant \left\lfloor \frac{n}{2} \right\rfloor + \Delta(n) + a_n + b_n + c_n - 1. \tag{5.3.1}$$

(We will use this in the proof of Corollary 5.1.2.) In order to prove Proposition 5.1.5, it is clearly enough to show that

$$a_n + b_n + c_n = O\left(\frac{\log n}{\log\log n}\right).$$

To bound a_n, observe that if

$$\frac{q_1^{d_1} - 1}{q_1 - 1} = \frac{q_2^{d_2} - 1}{q_2 - 1},$$

then q_1 and q_2 must be coprime. We obtain that a_n must be bounded above by the number of distinct prime divisors of $n - 1$. In [Rob83] it is proved that this number is

$$O\left(\frac{\log(n-1)}{\log\log(n-1)}\right),$$

whence the same upper bound holds for a_n.

To bound b_n we refer to a result of Kane [Kan07], which asserts that[1]

$$b_n = O\left(\frac{\log n \log\log\log n}{(\log\log n)^3}\right).$$

To bound c_n, we first recall (see [Apo76, Theorem 13.12]) that, for a positive integer x,

$$\Delta(x) \leqslant \exp_2\left\{\frac{(1+o(1))\log x}{\log\log x}\right\}.$$

Now consider the prime factorization of n: $n = p_1^{a_1}\cdots p_t^{a_t}$. If $n = d^k$ then $p_1\cdots p_t$ divides d and k divides $a := \gcd\{a_1,\dots,a_t\}$. Therefore, the number of choices for k is at most the number of divisors of a different from 1. Now note that $a \leqslant \log n$, and therefore

$$c_n \leqslant \exp_2\left\{\frac{(1+o(1))\log\log n}{\log\log\log n}\right\}.$$

In particular we see that each of a_n, b_n, c_n is $O(\log n/\log\log n)$. This proves the proposition. □

We conclude with the proof of Corollary 5.1.2.

Proof of Corollary 5.1.2. Since $\{(1,2),(2,3),(3,4),\dots,(n-1,n)\}$ is a minimal generating set of size $n-1$, it is enough to show that $m_I(S_n) < n-1$. We will prove the stronger bound $\iota(S_n) < n-1$ (recall Lemma 3.2.3).

From (5.3.1) in the proof of Proposition 5.1.5, we deduce that

$$\iota(S_n) \leqslant \left\lfloor \frac{n}{2}\right\rfloor + \Delta(n) + a_n + b_n + c_n - 1,$$

therefore it is sufficient to show that $\Delta(n) + a_n + b_n + c_n < n/2$. Very weak estimates are enough here. First assume that $n \geqslant 71$.

As remarked in the proof of Proposition 5.1.5, a_n is bounded above by the number of distinct prime divisors of $n-1$, which is at most $\log n$. Moreover, c_n is bounded above by $\max\{\Delta(x) : x \leqslant \lfloor\log n\rfloor\}$, which is at most $\log n$. Let us consider b_n. Let $(d_1,k_1),\dots,(d_b,k_b)$ be pairs such that $\binom{d_i}{k_i} = n$ for all $i = 1,\dots,b$. Order so that $i<j$ implies that $k_i < k_j$ and observe that then $k_b \geqslant b$ and $d_b \geqslant 2b$. This implies that $n \geqslant \binom{2b}{b} > 2^b$. In particular $b < \log n$.

Finally we need to bound $\Delta(n)$. For every real number $a \in (0,n]$, we have $\Delta(n) < n/a + a$. By choosing $a = \sqrt{n}$, we deduce $\Delta(n) < 2\sqrt{n}$.

Therefore $\Delta(n) + a_n + b_n + c_n < 2\sqrt{n} + 3\log n$, and it is sufficient to show that $2\sqrt{n} + 3\log n \leqslant n/2$. Since $n \geqslant 71$, this is indeed the case.

[1] Singmaster's conjecture [Sin71] asserts that b_n is bounded above by an absolute constant; de Weger proposes that in fact this constant can be taken to be 4 [dW97], and evidence for the veracity of this conjecture is given in [BBDW17]; in particular this is known to be true if $n \leqslant 10^{60}$.

For $n \leqslant 70$ we use [GAP19] to find that, except when $n \in \{5, 6, 8, 12\}$, S_n has less than $n-1$ conjugacy classes of maximal subgroup and the result follows immediately.

For the remaining cases, say that two cycle types are *equivalent* if one is a power of the other one; e.g., $(2, 2)$ is equivalent to (4), $(2, 3, 3)$ is equivalent to $(2, 1^6)$, etc. We readily see that, if $\{M_1^*, \ldots, M_t^*\}$ is independent, then, for every i, M_i intersects non-trivially at least $t-1$ pairwise non-equivalent cycle types.

For $n \in \{5, 8, 12\}$, there are exactly $n-1$ conjugacy classes of maximal subgroups of S_n. However, in each case, there is one which does not intersect at least $n-2$ pairwise non-equivalent cycle types, and the result follows. (For $n = 5$ we may take $\mathrm{AGL}_1(5)$, for $n = 8$ we may take $\mathrm{PGL}_2(7)$, and for $n = 12$ we may take $\mathrm{PGL}_2(11)$.)

For $n = 6$, there are 6 conjugacy classes of maximal subgroups. One of these, with representative $\mathrm{PGL}_2(5)$, does not intersect 5 pairwise non-equivalent cycle types. In particular, we deduce that the only independent set of size 5 can possibly be

$$\Omega = \{A_6^*, (S_2 \wr S_3)^*, (S_3 \wr S_2)^*, S_5^*, (S_4 \times S_2)^*\}. \tag{5.3.2}$$

We easily verify that the intersection of the first four M^* in (5.3.2) consists of the conjugacy classes (1^6) and $(2^2, 1^2)$, both of which are contained in $(S_4 \times S_2)^*$. This shows that Ω is not independent, and therefore $\iota(S_6) < 5$, as wanted. $\square$

Chapter 6

Connected components in the invariably generating graph

The content of this chapter consists of the preprint [Gar20b].

6.1 Introduction

In this chapter, it is convenient to consider invariable generation by *conjugacy classes*, rather than invariable generation by *elements* (this is really only a matter of terminology). Specifically, given a finite group G and a set $X = \{C_1, \dots, C_t\}$ of conjugacy classes of G, we say that X invariably generates G if $\langle x_1, \dots, x_t \rangle = G$ for every $x_1 \in C_1, \dots, x_t \in C_t$, and we write in this case $\langle X \rangle_I = G$.

In [Gar20a], the following definition was given. For a finite group G, the *invariably generating graph* $\Lambda(G)$ of G is the undirected graph whose vertices are the conjugacy classes of G different from $\{1\}$, and two vertices C and D are adjacent if $\langle C, D \rangle_I = G$. If G is not invariably 2-generated, $\Lambda(G)$ is the empty graph. Even when G is invariably 2-generated, the graph $\Lambda(G)$ can have isolated vertices (e.g., when G is not cyclic, the classes contained in the Frattini subgroup); define $\Lambda^+(G)$ as the graph obtained by removing the isolated vertices of $\Lambda(G)$. In this chapter we prove the following result.

Theorem 6.1.1. *For every positive integer n, there exists a finite group G such that $\Lambda^+(G)$ has more than n connected components.*

Theorem 6.1.1 should be seen in comparison to the analogous graph for the case of usual generation; see Subsection 6.1.1.

In the proof of Theorem 6.1.1, G is a suitable direct power of a nonabelian finite simple group S. We use $S = \mathrm{PSL}_2(q)$, although there are other possible choices. A crucial ingredient is that $\Lambda^+(S)$ is bipartite, which follows from the fact that S admits a 2-covering (Lemma 6.2.4). See Section 6.3 for definitions

and further comments in this direction, related to clique number and chromatic number of $\Lambda^+(S)$.

We will also give a suitable lower bound to the number of connected components of $\Lambda^+(G)$ in our examples (see Theorem 6.2.8), which is not strictly necessary for the mere proof of Theorem 6.1.1. We will get the bound as a consequence of the following result.

Theorem 6.1.2. *Let $S = \mathrm{PSL}_2(q)$, and let C_1, C_2 be conjugacy classes of S chosen uniformly at random. Then*

$$\mathbf{P}(\langle C_1, C_2\rangle_I = S) = 1/2 + O(1/q).$$

The proof of Theorem 6.1.2 is straightforward, since subgroups and conjugacy classes of $\mathrm{PSL}_2(q)$ are known very explicitly. It is interesting that the asymptotic behaviour of $\mathbf{P}(\langle C_1, C_2\rangle_I = S)$ is equal to the asymptotic behaviour of $\mathbf{P}(\langle x_1, x_2\rangle_I = S)$, where $x_1, x_2 \in S$ are random *elements*. The latter statement will follow from the results of Chapter 7; see in particular Subsection 7.6.1.

6.1.1 Comparison to usual generation

For a finite group G, the *generating graph* $\Gamma(G)$ of G is the undirected graph whose vertices are the nonidentity elements of G, and two vertices x and y are adjacent if $\langle x, y\rangle = G$. This graph has been intensively studied in the last two decades; see Burness [Bur19] and Lucchini–Maróti [LM09] for many results in this context.

Again, the graph $\Gamma(G)$ can have isolated vertices, and we consider the graph $\Gamma^+(G)$ obtained by removing the isolated vertices of $\Gamma(G)$. It is known that $\Gamma^+(G)$ is connected in several cases (see Burness–Guralnick–Harper [BGH20], Crestani–Lucchini [CL13a, CL13b]). By contrast, no example of G is known for which $\Gamma^+(G)$ is disconnected, which determines a sharp difference with respect to Theorem 6.1.1.

We recall that this difference does not occur for nilpotent groups. Indeed, in a finite nilpotent group the concepts of generation and invariable generation coincide (see Lemma 2.3.2).

6.1.2 Organization of the chapter and notation

In Section 6.2 we prove Theorems 6.1.1 and 6.1.2, and in Section 6.3 we make further comments and propose some problems.

The asymptotic notation $f = O(g)$ means that $|f| \leqslant Cg$ for some constant C (so f might also be negative). For a real number x, we set $\exp_2\{x\} := 2^x$.

6.2 Proof of Theorems 6.1.1 and 6.1.2

6.2.1 Direct powers of finite simple groups

Throughout this subsection, S denotes a nonabelian finite simple group. We review some properties of invariable generation of direct powers of S, which reflect some interesting properties of the corresponding invariably generating graphs. The key tool is an elementary criterion due to Kantor and Lubotzky [KL90a], which we recall.

Denote by $\Psi_2(S)$ the set of all pairs (C_1, C_2), where C_i is a conjugacy class of S, and $\langle C_1, C_2\rangle_I = S$. Theorem 2.4.2 implies that $\Psi_2(S) \neq \varnothing$.

Let now t be a positive integer, and let C and D be conjugacy classes of S^t, with $C = C_1 \times \cdots \times C_t$ and $D = D_1 \times \cdots \times D_t$ (and C_i and D_i are conjugacy classes of S). Consider the matrix

$$A_{C,D} = \begin{pmatrix} C_1 & C_2 & \cdots & C_t \\ D_1 & D_2 & \cdots & D_t \end{pmatrix}.$$

Lemma 6.2.1. *We have that $\langle C, D\rangle_I = S^t$ if and only if the following conditions are both satisfied:*

(i) Each column of $A_{C,D}$ belongs to $\Psi_2(S)$, and

(ii) No two columns of $A_{C,D}$ lie in the same orbit for the diagonal action of $\mathrm{Aut}(S)$ on $\Psi_2(S)$.

Proof. See [KL90a, Proposition 6], and also [DL15, Lemma 20]. □

Let $\beta = \beta(S)$ be the largest integer for which S^β is invariably 2-generated. Lemma 6.2.1 implies that $\beta(S)$ is equal to the number of orbits for the diagonal action of $\mathrm{Aut}(S)$ on $\Psi_2(S)$. We note the following fact.

Lemma 6.2.2. *We have*

$$\frac{|\Psi_2(S)|}{|\mathrm{Out}(S)|} \leqslant \beta(S) \leqslant |\Psi_2(S)|.$$

Proof. The second equality is clear, and the first follows from the fact that $\mathrm{Inn}(S) \cong S$ acts trivially in the relevant action, hence each orbit has size at most $|\mathrm{Out}(S)|$. □

We expect $|\mathrm{Out}(S)|$ to be much smaller than $|\Psi_2(S)|$ for every sufficiently large nonabelian finite simple group S. Therefore, $|\Psi_2(S)|$ should be, in some sense, a good approximation for $\beta(S)$.

We make a one-paragraph digression in order to compare to the case of classical generation. Let $\delta = \delta(S)$ be the largest integer for which S^δ is 2-generated. Unlike for $\beta(S)$, there is an exact formula for $\delta(S)$, namely, $\delta(S) = \phi_2(S)/|\mathrm{Aut}(S)|$, where $\phi_2(S)$ denotes the number of ordered pairs $(x, y) \in S^2$ such that $\langle x, y\rangle = S$. (This goes back to Hall [Hal36] in 1930s and has been

widely used.) The difference is that the diagonal action of $\mathrm{Aut}(S)$ on the set of generating pairs of elements is semiregular (i.e., only the identity fixes a generating pair), while this needs not be the case for the action of $\mathrm{Aut}(S)$ on the set of invariable generating pairs of conjugacy classes.

Lemma 6.2.1 describes quite precisely the graph $\Lambda(S^\beta)$. Indeed, any arc in the graph is obtained as follows (and only in this way). Construct a $2 \times \beta$ matrix, in which the columns form a set of representatives for the $\mathrm{Aut}(S)$-orbits on $\Psi_2(S)$. Then the first row is adjacent to the second row in $\Lambda^+(S^\beta)$ (here we are identifying a conjugacy class $C_1 \times \cdots \times C_\beta$ of S^β with a row vector $(C_1, \ldots, C_\beta)$). Since $\mathrm{Aut}(S^\beta) \cong \mathrm{Aut}(S) \wr \mathrm{Sym}(\beta)$ acts by automorphisms on $\Lambda^+(S^\beta)$, we also see that $\Lambda^+(S^\beta)$ is arc-transitive.

6.2.2 The case $S = \mathrm{PSL}_2(q)$

In this subsection we choose $S = \mathrm{PSL}_2(q)$, with $q \geqslant 4$. For reader's convenience, we recall some well known facts. See [Suz82, Chapter 3.6] for the description of the maximal subgroups of S.

Lemma 6.2.3. *Let $S = \mathrm{PSL}_2(q)$, where $q \geqslant 4$ is a power of the prime p. Set $d = (2, q-1)$.*

(1) S contains a unique conjugacy class of involutions and, for p odd, two conjugacy classes of elements of order p.

(2) Assume $3 \leqslant \ell \mid (q \pm 1)/d$. There are $\phi(\ell)/2$ conjugacy classes of elements of order ℓ in S, where ϕ is Euler's totient function.

(3) The number of conjugacy classes of S is $(q + 4d - 3)/d$.

Note that item (3) was used also in the proof of Lemma 4.2.12. Here we will also need items (1) and (2).

Proof. We sketch a proof. (1) Assume first q is odd, and let us deal with involutions. Let $\varepsilon = 1$ if $q \equiv 1 \bmod 4$, and $\varepsilon = -1$ otherwise. By explicit matrix computation, we find that the number of involutions of S is $q(q+\varepsilon)/2$. This coincides with the number of conjugates of a dihedral subgroup of order $q - \varepsilon$, so we deduce that all involutions of S are conjugate. Next we deal with elements of order p (also in the case q even). The image in S of the subgroup of $\mathrm{SL}_2(q)$ consisting of the upper unitriangular matrices is a Sylow p-subgroup P of S. This is contained in a subgroup B, consisting of the image of the upper triangular matrices of $\mathrm{SL}_2(q)$. Let $1 \neq x \in P$. We verify that $x^S \cap P = x^B$, and we compute that conjugating x by elements of B can only multiply the upper-right entry by every nonzero square of $\mathbf{F}_q$. This proves (1).

(2) Assume $3 \leqslant \ell \mid (q \pm 1)/d$. All cyclic subgroups of order ℓ are conjugate in S. Assume $x \in S$ has order ℓ. We have that $\mathrm{N}_S(\langle x \rangle)$ is dihedral of order $2(q \pm 1)/d$, from which $x^S \cap \langle x \rangle = \{x, x^{-1}\}$. This proves (2).

(3) An element of S has either order p, or order dividing $(q \pm 1)/d$. Then (3) follows from (1), (2), and the formula

$$\sum_{\ell \mid (q \pm 1)/d} \frac{\phi(\ell)}{2} = \frac{q \pm 1}{2d}.$$

The statement is proved. □

The following lemma represents the main observation regarding $\Lambda^+(S)$.

Lemma 6.2.4. *The graph $\Lambda^+(S)$ is bipartite.*

Proof. It is well known that S admits a 2-covering, that is, a pair of proper subgroups (H, K) such that

$$S = \bigcup_{g \in S} H^g \cup \bigcup_{g \in S} K^g = \widetilde{H} \cup \widetilde{K}.$$

We take H a dihedral subgroup of order $2(q+1)/d$, and K a Borel subgroup of order $q(q-1)/d$, with $d = (2, q-1)$. A conjugacy class contained in $\widetilde{H} \cap \widetilde{K}$ is isolated in $\Lambda(S)$, and a class contained in $\widetilde{H} \setminus \widetilde{K}$ can be adjacent in $\Lambda^+(S)$ only to a class contained in $\widetilde{K} \setminus \widetilde{H}$. This gives a partition of $\Lambda^+(S)$ in two parts. □

Let $\mathscr{P}_1$ and $\mathscr{P}_2$ be the parts of $\Lambda^+(S)$ given in the proof of Lemma 6.2.4. We note that, for every conjugacy class C of S and for every $\sigma \in \mathrm{Aut}(S)$, $\{C, C^\sigma\}$ does not invariably generate S. (A way to see this is that the sets $\widetilde{H}$ and $\widetilde{K}$ from the proof of Lemma 6.2.4 are preserved by every automorphism of S.) In particular, for every $(C_1, C_2) \in \Psi_2(S)$, (C_1, C_2) and (C_2, C_1) belong to different $\mathrm{Aut}(S)$-orbits. We also note that the parts $\mathscr{P}_1$ and $\mathscr{P}_2$ are invariant under the action of $\mathrm{Aut}(S)$. We deduce the following

Lemma 6.2.5. *$\beta = \beta(S)$ is even, and for each vertex $C = C_1 \times \cdots \times C_\beta$ of $\Lambda^+(S^\beta)$, there exists a subset $\Omega = \Omega(C)$ of $\{1, \ldots, \beta\}$ of size $\beta/2$ such that for every $i \in \Omega$, $C_i \in \mathscr{P}_1$, and for every $i \notin \Omega$, $C_i \in \mathscr{P}_2$.*

We can finally prove the key result.

Theorem 6.2.6. *The graph $\Lambda^+(S^\beta)$ has at least $\frac{1}{2} \cdot \binom{\beta}{\beta/2}$ connected components.*

Proof. For a vertex $C = C_1 \times \cdots \times C_\beta$ of $\Lambda^+(S)$, let $\Omega(C)$ be the set from Lemma 6.2.5. Then, C can be adjacent only to vertices D such that $\Omega(D) = \{1, \ldots, \beta\} \setminus \Omega(C)$. In particular, the number of connected components of $\Lambda^+(S^\beta)$ is at least half the number of $\beta/2$-subsets of $\{1, \ldots, \beta\}$, which proves the statement. □

It is not difficult to establish that $\beta(S)$ tends to infinity as $|S| \to \infty$ (that is, $q \to \infty$), thereby proving Theorem 6.1.1. In the next subsection we will obtain a better estimate for $\beta(S)$.

6.2.3 Bounds

We want to estimate $\beta(S)$, where $S = \mathrm{PSL}_2(q)$. We will find the asymptotic behaviour of $|\Psi_2(S)|$, and then apply Lemma 6.2.2.

Theorem 6.2.7. *Let $S = \mathrm{PSL}_2(q)$ and $d = (2, q-1)$. We have*

$$|\Psi_2(S)| = \frac{q^2}{2d^2} + O(q).$$

(For q odd the first term of the expression is not an integer, but still the statement makes sense.)

Proof. In this proof, when we say that a conjugacy class C *intersects* a subgroup H, we mean $C \cap H \neq \varnothing$. We refer to [Suz82, Chapter 3.6] for the description of the maximal subgroups of S. We need to count the pairs of conjugacy classes (C_1, C_2) which invariably generate S. We ignore the pairs where either C_1 or C_2 is made of elements of order p, or of order at most 2. By Lemma 6.2.3, the number of these pairs is $O(q)$.

By this choice, up to swapping the indices, C_1 intersects a cyclic subgroup of order $(q-1)/d$, and C_2 intersects a cyclic subgroup of order $(q+1)/d$. Given C_1 and C_2 with this property, we have that C_1 and C_2 invariably generate S unless one of the following occurs:

(i) C_1 and C_2 intersect a certain maximal subgroup of order at most 60, and there are at most five possibilities for such subgroup.

(ii) C_1 and C_2 intersect a maximal subgroup conjugate to $\mathrm{PSL}_2(q^{1/r})$ where r is an odd prime (and q is an r-th power).

(In (ii), we are not considering subgroups $\mathrm{PGL}_2(q^{1/2})$. Indeed, any class of elements of $\mathrm{PGL}_2(q^{1/2})$ of order prime to q intersects a cyclic subgroup of S order $(q-1)/d$; and this cannot occur for C_2.) Clearly there are $O(1)$ possibilities for (C_1, C_2) satisfying (i). The number of conjugacy classes of $\mathrm{PSL}_2(q^{1/r})$ is $O(q^{1/r})$; therefore, for fixed r, the number of possibilities for the pair (C_1, C_2) satisfying (ii) is $O(q^{2/r})$. Summing through odd prime r, we get $O(q^{2/3})$ (note that there are at most $\log\log q$ possibilities for r).

Using Lemma 6.2.3, we get the following formula for $|\Psi_2(S)|$ (the factor 2 at the beginning comes from the fact that we may also have C_1 intersecting a cyclic subgroup of order $(q+1)/d$, and C_2 intersecting a cyclic subgroup of order $(q-1)/d$).

$$\begin{aligned} |\Psi_2(S)| &= 2 \cdot \sum_{\substack{\ell_1 \mid (q-1)/d \\ \ell_2 \mid (q+1)/d}} \frac{\phi(\ell_1)}{2} \frac{\phi(\ell_2)}{2} + O(q) \\ &= \frac{q^2-1}{2d^2} + O(q) = \frac{q^2}{2d^2} + O(q). \end{aligned}$$

(In the first equality we used the formula $\sum_{\ell \mid n} \phi(\ell) = n$, and the fact that $\phi(\ell_1)\phi(\ell_2) = \phi(\ell_1\ell_2)$ for coprime ℓ_1 and ℓ_2.) □

We can rephrase Theorem 6.2.7 in probabilistic language, that is, we can prove Theorem 6.1.2.

Proof of Theorem 6.1.2. The statement follows from Lemma 6.2.3(3) and Theorem 6.2.7. □

At this point we can estimate $\beta(S)$ and get a lower bound to the number of connected components of $\Lambda^+(S^\beta)$, thereby proving Theorem 6.1.1. The symbol $o(1)$ is understood with respect to the limit $q \to \infty$.

Theorem 6.2.8. *Let $S = \mathrm{PSL}_2(q)$. We have*

$$q^{2-o(1)} \leqslant \beta(S) \leqslant \frac{q^2}{2} + O(q).$$

Let $N(S)$ denote the number of connected components of $\Lambda^+(S^\beta)$. Then

$$N(S) \geqslant \exp_2\left\{q^{2-o(1)}\right\}.$$

Proof. The order of $\mathrm{Out}(S)$ is at most $2\log q$. By Lemma 6.2.2 and Theorem 6.2.7 we get

$$q^{2-o(1)} \leqslant \beta \leqslant \frac{q^2}{2} + O(q), \tag{6.2.1}$$

which proves the first part of the statement. By Theorem 6.2.6, Stirling's approximation and (6.2.1), we get

$$\begin{aligned} N(S) \geqslant \frac{1}{2} \cdot \binom{\beta}{\beta/2} &= (1+o(1)) \cdot \frac{2^\beta}{(2\pi\beta)^{1/2}} \\ &\geqslant \exp_2\left\{q^{2-o(1)}\right\}, \end{aligned}$$

which proves the last part of the statement. □

6.3 Further comments

Recall that $\Gamma^+(G)$ is the graph obtained by removing the isolated vertices from the generating graph $\Gamma(G)$ of G. Crestani–Lucchini [CL13b] showed that, if G is a 2-generated direct power of a nonabelian finite simple group, then $\Gamma^+(G)$ is connected.

In particular, Theorem 6.2.8 says that the result of [CL13b] does not hold for invariable generation. Nevertheless, a combinatorial proof along the lines of [CL13b, Theorem 3.1] might be feasible in order to show the following: If a finite simple group S is such that $\Lambda^+(S)$ is connected and not bipartite, then $\Lambda^+(S^t)$ is connected for every $t \leqslant \beta(S)$. Unfortunately, the fact that $\Lambda^+(S)$ is connected is known essentially only for alternating groups [Gar20a], therefore the result at this stage would not be of wide use.

We also remark that we are currently unable to construct examples of soluble groups G for which $\Lambda^+(G)$ is disconnected.

Question 6.3.1. *Let G be a finite soluble group which is invariably 2-generated. Is the graph $\Lambda^+(G)$ connected?*

Crestani–Lucchini [CL13a] showed that this is true for the graph $\Gamma^+(G)$ (and in particular Question 6.3.1 has a positive answer for nilpotent groups).

6.3.1 $\Lambda^+(S)$ bipartite

For the proof of Theorem 6.2.8, the only important fact about $S = \mathrm{PSL}_2(q)$ is that the graph $\Lambda^+(S)$ is bipartite, which follows from the fact that S admits a 2-covering (see the proof of Lemma 6.2.4). Recall that, given a finite group G, a 2-*covering* of G is a pair (H, K) of proper subgroups such that

$$G = \bigcup_{g \in G} H^g \cup \bigcup_{g \in G} K^g.$$

The 2-coverings of the finite simple groups have been well studied; see Bubboloni [Bub10], Bubboloni–Lucido [BL02], Bubboloni–Lucido–Weigel [BLW06, BLW11], Pellegrini [Pel13]. In particular, all finite simple groups admitting a 2-covering are known, except for some classical groups in small dimension.

We have the following clear implications:

$$\begin{aligned} S \text{ admits a 2-covering} &\implies \Lambda^+(S) \text{ is bipartite} \\ &\implies \Lambda^+(S) \text{ has no triangles} \end{aligned} \tag{6.3.1}$$

(These implications are a particular case of the inequalities in (6.3.2) below.) The reverse of the first implication in (6.3.1) does not necessarily hold. For instance, A_9 does not admit a 2-covering (it was proved in [Bub10] that A_n admits a 2-covering if and only if $4 \leqslant n \leqslant 8$). On the other hand, it is not difficult to show that $\Lambda^+(A_9)$ is bipartite. This might be one of only finitely many exceptions.

Problem 6.3.2. *Determine the finite simple groups S for which $\Lambda^+(S)$ is bipartite (resp., contains no triangles). Up to finitely many cases, do the reverse implications in* (6.3.1) *hold?*

These considerations can be viewed more generally as follows. For a noncyclic finite group G, let $\kappa(G)$ be the clique number of $\Lambda^+(G)$, that is, the largest order of a complete subgraph of $\Lambda^+(G)$. Let $\tau(G)$ be the chromatic number of $\Lambda^+(G)$, that is, the least number of colours needed to colour the vertices of $\Lambda^+(G)$ in such a way that adjacent vertices get different colours. Let $\gamma(G)$ be the normal covering number of G, that is, the least number of proper subgroups of G such that each element of G lies in some conjugate of one of these subgroups. The following inequalities hold:

$$\kappa(G) \leqslant \tau(G) \leqslant \gamma(G). \tag{6.3.2}$$

(These are "invariable" versions of inequalities studied for instance in [LM09].) The implications in (6.3.1) can be stated as follows for a general noncyclic finite

group G: $\gamma(G) \leqslant 2 \implies \tau(G) \leqslant 2 \implies \kappa(G) \leqslant 2$. (We note that, for every finite group G, $\gamma(G) \geqslant 2$ and, by Theorem 2.4.2, if S is nonabelian simple then $\kappa(S) \geqslant 2$.) Problem 6.3.2 asks whether, up to finitely many exceptions, $\gamma(S) = 2 \iff \tau(S) = 2 \iff \kappa(S) = 2$.

The invariants $\kappa(G)$ and $\gamma(G)$ have been studied; see for instance Britnell–Maróti [BM13], Bubboloni–Praeger–Spiga [BPS13] and Garonzi–Lucchini [GL15].

As a final remark, the fact that $\Lambda(G)$ can have no triangles is somewhat strange, in comparison to classical generation. Indeed, for every 2-generated finite group G of order at least 3, the generating graph $\Gamma(G)$ contains a triangle, and indeed "many" triangles. This follows from the crucial fact that if $\langle x, y\rangle = G$ then $\langle x, xy\rangle = \langle xy, y\rangle = G$. As we already observed in Section 3.10, the fact that this property fails for invariable generation represents an annoying obstacle in order to extend results from the classical to the invariable setting.

Chapter 7

On the probability of generating invariably a finite simple group

The content of this chapter consists of the preprint [GM20], which was written in collaboration with Eilidh McKemmie.

7.1 Introduction

For a finite group G and a subset A of G, denote by $\mathbf{P}_{\mathrm{inv}}(G, A)$ the probability that, if $y \in G$ is chosen uniformly at random, there exists $x \in A$ such that $\langle x, y\rangle_I = G$. In case $A = \{x\}$, we will write $\mathbf{P}_{\mathrm{inv}}(G, x)$ instead of $\mathbf{P}_{\mathrm{inv}}(G, \{x\})$.

We consider the case in which G is a nonabelian finite simple group. Our general aim is to find "small" subsets A of G such that $\mathbf{P}_{\mathrm{inv}}(G, A)$ is "large".

We state our first result, which in most cases will be asymptotically superseded by subsequent theorems. In Subsection 7.1.2 we will provide more context for these theorems, also in relation to "classical" generation.

Theorem 7.1.1. *Let G be a nonabelian finite simple group. There exist an absolute constant $\epsilon > 0$ and an element $x \in G$ such that $\mathbf{P}_{\mathrm{inv}}(G, x) \geqslant \epsilon$.*

As recalled in Theorem 2.4.2 of Chapter 2, every finite simple group is invariably generated by two elements. Therefore, we will only need to prove Theorem 7.1.1 for sufficiently large finite simple groups.

In light of Theorem 7.1.1, one would like to find many elements $y \in G$ with the property that $\mathbf{P}_{\mathrm{inv}}(G, y)$ is bounded away from zero uniformly. For groups of Lie type of bounded rank different from $G_2(3^a)$, one can take almost all elements.

Theorem 7.1.2. *Let G be a finite simple group of Lie type of untwisted rank r defined over the field with q elements,*[1] *and assume $G \ncong G_2(q)$ when $3 \mid q$. Then, $\mathbf{P}_{\mathrm{inv}}(G, y) \geqslant c/r + O(r^r/q)$ for an absolute constant $c > 0$ and for a proportion of elements $y \in G$ of the form $1 - O(r^r/q^{1/2})$.*

(In fact, if $G \ncong \mathrm{PSL}_2(q)$, in the last error term one can replace $q^{1/2}$ by q.) In Theorem 7.6.1 we will give an explicit value for c. The group $G = G_2(3^a)$ does not satisfy the statement; we will see in Theorem 7.1.4 that $\mathbf{P}_{\mathrm{inv}}(G, y) = 0$ for roughly half of the elements $y \in G$. However, we will show that $\mathbf{P}_{\mathrm{inv}}(G, y) \geqslant 1/6 + O(1/q)$ for the remaining half of the elements (Theorem 7.4.1). With a bit of care for the error term in case $G \cong \mathrm{PSL}_2(q)$, we get the following immediate consequence.

Theorem 7.1.3. *Let G be a finite simple group of Lie type of untwisted rank r defined over the field with q elements. Let $x_1, x_2 \in G$ be chosen uniformly at random. Then,*

$$\mathbf{P}(\langle x_1, x_2 \rangle_I = G) \geqslant c/r + O(r^r/q).$$

Of course, in the previous two theorems we are thinking of r fixed, and $q \to \infty$. In order to avoid confusion, we point out that with $u = O(z)$ we mean that $|u| \leqslant Cz$ for some constant C (so there is no assertion on the sign of u).

We will review the history of Theorem 7.1.3 in Subsection 7.1.1. An interesting purpose is to obtain sharp bounds in Theorem 7.1.3. Although in this thesis we do not pursue this goal, in Theorem 7.6.2 we will obtain a formula of the type $\mathbf{P}(\langle x_1, x_2 \rangle_I = G) = f(r) + d(r)/q$. The main term $f(r)$ is very explicit, depends only on the Weyl group, and can be computed essentially in an algorithmic way; it should be possible to compute it precisely for all exceptional groups.

Eberhard–Ford–Green [EFG17] and McKemmie [McK19] showed that, for alternating groups and for groups of Lie type of large rank over large fields, $\mathbf{P}(\langle x_1, x_2, x_3 \rangle_I = G)$ tends to zero. (Conjecturally, in groups of Lie type there should be no restriction on the field size.) Therefore, it is not possible to extend Theorem 7.1.2 to the other families of finite simple groups — not even for a proportion of elements bounded away from zero. In this sense, we can say that the element x of Theorem 7.1.1 is "special", unless G is of Lie type of bounded rank.

In Theorems 7.1.1 and 7.1.2 we have bounded $\mathbf{P}_{\mathrm{inv}}(G, x)$ away from zero. Next, we would like to get probabilities approaching 1. With this purpose, we consider more elements simultaneously. In most cases, only few elements are needed; in the remaining cases, the whole group is not enough.

Theorem 7.1.4. *Let G be a nonabelian finite simple group.*

- ◇ *Assume G is of Lie type and $G \ncong G_2(q)$ when $3 \mid q$. There exists a subset A_b of G such that $|A_b|$ and $\mathbf{P}_{\mathrm{inv}}(G, A_b)$ are as in Table 7.1.*

[1] This is not exactly precise for Suzuki and Ree groups, and in fact, not even for unitary groups. We will recall the exact definition of "q" in Section 7.3. Concretely, q is the parameter appearing in Table 7.1.

- ◇ *Assume G is alternating or classical. There exists a subset A_ℓ of G such that $|A_\ell|$ and $\mathbf{P}_{\mathrm{inv}}(G, A_\ell)$ are as in Table 7.1.*[2]

- ◇ *Assume $G = G_2(q)$ with $3 \mid q$, or $G = \mathrm{PSp}_{2m}(q)$ with q even and m sufficiently large, or $G = \mathrm{P}\Omega_{2m+1}(q)$ with q odd and m sufficiently large. Then $\mathbf{P}_{\mathrm{inv}}(G, G)$ is as in Table 7.1.*

Theorem 7.1.4 presents a strong dichotomy; it is worth stating this separately.

Corollary 7.1.5. *Let G be a nonabelian finite simple group.*

(1) Assume G is not as in Table 7.2. Then, there exists $A \subseteq G$ of size at most 6 such that $\mathbf{P}_{\mathrm{inv}}(G, A)$ tends to 1 as $|G| \to \infty$.

(2) Assume G is as in Table 7.2. Then $\mathbf{P}_{\mathrm{inv}}(G, G)$ is bounded away from 1. (Equivalently, $\mathbf{P}_{\mathrm{inv}}(G, y) = 0$ for a proportion of elements $y \in G$ bounded away from zero.)

Corollary 7.1.5 follows from Theorem 7.1.4 by setting $A = A_\ell$ for alternating groups, $A = A_b$ for exceptional groups, and $A = A_b \cup A_\ell$ for classical groups. Of course, every case in which the size of A_ℓ is equal to 1 represents a strengthening of Theorem 7.1.1. Combining this with Theorem 7.1.2, we see that in most cases asymptotically we can do better than Theorem 7.1.1. The improvement is complementary: while Theorem 7.1.2 does not hold in large rank, here we cannot have $|A_b| = 1$ in bounded rank (see Lemma 7.7.2). We note finally that the size of $|A_\ell|$ is sharp in every case (Lemmas 7.8.6 and 7.8.10), and that there are cases in which we need $|A_b| \geqslant 4$ (see Lemma 7.7.1; note that $|A_b| \leqslant 4$ unless $G = F_4(2^a)$). Here, with "sharp" we mean that if choose a set Y of smaller size, then $\mathbf{P}_{\mathrm{inv}}(G, Y)$ remains bounded away from 1 as the relevant parameters grow.

7.1.1 Context: Theorem 7.1.3

For convenience, we recall some results from Section 2.2. The first result of the flavour of Theorem 7.1.3 was obtained by Dixon [Dix92]. He showed that $O((\log n)^{1/2})$ random elements of S_n invariably generate S_n with probability tending to 1 as $n \to \infty$. Luczak and Pyber [LP93] showed that $O(1)$ random elements of S_n invariably generate with probability bounded away from zero. The exact value of $O(1)$ turned out to be four: Pemantle–Peres–Rivin [PPR16] proved that four elements are enough, while Eberhard–Ford–Green [EFG17] showed that three are not. The same results hold for alternating groups. McKemmie [McK19] extended these results to classical groups of large rank, leaving open the case of classical groups over small fields.

Theorem 7.1.3 addresses the case of groups of Lie type of bounded rank, which therefore nearly finishes the problem of invariable generation of finite

[2] In A_b, "b" stands for "bounded", and in A_ℓ, "ℓ" stands for "large". This is referred to the rank of the groups; the reason of this choice should be clarified by looking at the bounds in Table 7.1.

G	Conditions	Size of A_b or A_ℓ	Bounds
		$\lvert A_b\rvert$	
${}^2B_2(q^2)$		2	$\mathbf{P}_{\mathrm{inv}}(G, A_b) \geqslant 1 - O(r^r/q)$
${}^2G_2(q^2)$		2	
$G_2(q)$,	$3 \nmid q$	2	
${}^3D_4(q)$		2	
${}^2F_4(q^2)$		2	
$F_4(q)$	q odd	2	
$F_4(q)$	q even	6	
$E_6(q)$		2	
${}^2E_6(q)$		2	
$E_7(q)$		2	
$E_8(q)$		2	
$\mathrm{PSL}_n(q)$		2	
$\mathrm{PSU}_n(q)$		2	
$\mathrm{PSp}_{2m}(q)$	m even, q odd	2	
$\mathrm{PSp}_{2m}(q)$	m odd, q odd	3	
$\mathrm{PSp}_{2m}(q)$	q even	4	
$\mathrm{P}\Omega_{2m+1}(q)$	q odd	2	
$\mathrm{P}\Omega^-_{2m}(q)$		2	
$\mathrm{P}\Omega^+_{2m}(q)$	m odd	2	
$\mathrm{P}\Omega^+_{2m}(q)$	m even	4	
		$\lvert A_\ell\rvert$	
A_n		1	$\mathbf{P}_{\mathrm{inv}}(G, A_\ell) \geqslant 1 - O(n^{-0.08})$
$\mathrm{PSL}_n(q)$		1	$\mathbf{P}_{\mathrm{inv}}(G, A_\ell) \geqslant 1 - O(r^{-0.005})$
$\mathrm{PSU}_n(q)$		1	
$\mathrm{PSp}_{2m}(q)$	q odd	1	
$\mathrm{P}\Omega^-_{2m}(q)$		1	
$\mathrm{P}\Omega^+_{2m}(q)$		1	
$\mathrm{PSp}_{2m}(q)$	q even	2	$\mathbf{P}_{\mathrm{inv}}(G, A_\ell) \geqslant 1 - 6/q + O(r^{-0.005})$
$\mathrm{P}\Omega_{2m+1}(q)$	q odd	2	
$G_2(q)$	$3 \mid q$		$\mathbf{P}_{\mathrm{inv}}(G, G) = 1/2 + O(1/q)$
$\mathrm{PSp}_{2m}(q)$	q even, m large		$\mathbf{P}_{\mathrm{inv}}(G, G) \leqslant 1 - 1/4q^3$
$\mathrm{P}\Omega_{2m+1}(q)$	q odd, m large		$\mathbf{P}_{\mathrm{inv}}(G, G) \leqslant 1 - 1/6q$

Table 7.1: For groups of Lie type, r denotes the untwisted rank. It is understood that every nonempty entry in the fourth column applies to all subsequent empty lines.

G	Conditions
$G_2(q)$	$3 \mid q$
$\mathrm{PSp}_{2m}(q)$	q even fixed
$\mathrm{P}\Omega_{2m+1}(q)$	q odd fixed

Table 7.2:

simple groups by randomly chosen elements. (We note that, using a bounded number of random elements, the probability cannot approach 1; this follows from results by Fulman–Guralnick [FG03], and it is summarized for instance in [KLS11, Corollary 5.7].)

7.1.2 Context: Theorems 7.1.1, 7.1.2 and 7.1.4

It is convenient to visualize things as follows. Recall the definition of the generating graph $\Gamma(G)$, and of the invariably generating graph $\Lambda(G)$, which we gave in Chapter 6. Here we need to work with a variant of the graph $\Lambda(G)$, which we denote by $\Lambda_e(G)$, and which is obtained by $\Lambda(G)$ by replacing conjugacy classes by elements. More precisely, the vertices of $\Lambda_e(G)$ are the nontrivial elements of G, and two vertices x and y are adjacent if $\langle x, y\rangle_I = G$. (This graph was defined in [Gar20a].)[3]

In this language, Theorem 7.1.1 says that $\Lambda_e(G)$ contains large stars when G is simple.

For $x \in G$, let $\mathbf{P}(G, x)$ denote the probability that, if $y \in G$ is random, then $\langle x, y\rangle = G$: this is the "classical" version of our $\mathbf{P}_{\mathrm{inv}}(G, x)$. Set then $P^-(G) = \min\{\mathbf{P}(G, x) : 1 \neq x \in G\}$. Guralnick and Kantor [GK00] showed that $P^-(G) > 0$ for every finite simple group G, i.e., $\Gamma(G)$ has no isolated vertices. Moreover, in [GLSS99], following results from [GKS94] and [LS99], the behaviour of $P^-(G)$, where G is simple and $|G| \to \infty$, was determined.

It is easy to see that the Guralnick–Kantor result fails for invariable generation: $\Lambda_e(G)$ can have isolated vertices (for instance, a 3-cycle in A_n with n even). It seems to us that Theorems 7.1.1, 7.1.2 and 7.1.4 are essentially the best one can hope for in the invariable setting. Moreover, our results are among the first probabilistic statements concerning invariable generation of finite simple groups by *two* elements. We are aware only of [Sha98, Theorem 4.2], which is Theorem 7.1.4 in case $G = \mathrm{PSL}_n(q)$ and n large.

Corollary 7.1.5 can be clearly stated in terms of $\Lambda_e(G)$ as follows.

Corollary 7.1.6. *Let G be a finite simple group.*

(1) Assume G is not as in Table 7.2. Then, the proportion of isolated vertices of $\Lambda_e(G)$ tends to zero as $|G| \to \infty$. More precisely, if we remove a negligible proportion of vertices from $\Lambda_e(G)$, and if we remove some further edges, we obtain a graph which is the union of at most 6 stars.

(2) Assume G is as in Table 7.2. Then the proportion of isolated vertices of $\Lambda_e(G)$ is bounded away from zero.

Item (2) determines a sharp contrast with respect to the case of classical generation. Theorem 7.1.4 and Corollary 7.1.5 can be seen also as a sort of

[3] The graphs $\Lambda(G)$ and $\Lambda_e(G)$ share many properties; for instance, the number of connected components, and the value of the diameter in each connected component. In particular, the results of Chapter 6 hold for both graphs. However, the properties concerning proportions of vertices or edges can be very different. Since in this chapter we are working with random elements, we need to work with the graph $\Lambda_e(G)$.

"invariable" version of a concept introduced recently by Burness and Harper. In [BH19], the *total domination number* of a finite simple group G is defined as the total domination number of $\Gamma(G)$, i.e., the minimal size of a subset A of G such that, if $1 \neq y \in G$, there exists $x \in A$ such that $\langle x, y \rangle = G$. Corollary 7.1.5 can be thought of as an analogue for invariable generation — although again it is necessary to ignore a small proportion of elements.

7.1.3 Methods

We restate the main theorems in terms of certain subsets of the group.

Let G be a finite group, and let $x \in G$. We define $\mathscr{M}(x)$ as the union of all conjugates of maximal subgroups of G containing x. Equivalently, $\mathscr{M}(x)$ coincides with the union of all conjugacy classes of elements intersecting some maximal overgroup of x.

Lemma 7.1.7. *Let A be a subset of G. Then,*

$$1 - \mathbf{P}_{\mathrm{inv}}(G, A) = \frac{|\bigcap_{x \in A} \mathscr{M}(x)|}{|G|}.$$

Proof. Given $y \in G$, $\{x, y\}$ invariably generates G if and only if $y \notin \mathscr{M}(x)$. The statement follows. □

Therefore our business is to find elements x such that $\mathscr{M}(x)$ is small. This ultimately depends on two facts:

(i) existence of elements lying in few maximal subgroups, and

(ii) existence of maximal subgroups M of G such that $\widetilde{M} = \cup_{g \in G} M^g$ is small.

We must note that, in fact, (i) and (ii) perform only part of the job. Indeed, taking intersections of sets $\mathscr{M}(x)$ is rather more delicate, and will require much more work. Moreover, of course the proof of the upper bound to $\mathbf{P}_{\mathrm{inv}}(G, G)$ in Theorem 7.1.4 goes in the opposite direction.

Item (i) is a well studied topic (cf. [GK00, Wei92]). Often, for us applying results from these papers is convenient, rather than essential. Indeed, we are interested only in overgroups up to conjugation, which simplifies the situation. Moreover, in our probabilistic approach we can ignore the overgroups which are small (e.g. certain almost simple subgroups in groups of Lie type).

For what concerns item (ii), our main tool will be the positive solution of the Boston–Shalev conjecture by Luczak–Pyber and Fulman–Guralnick, which we discussed in Section 2.5.

Another key tool is the intimate connection between the properties of invariable generation of a group of Lie type and the structure of its Weyl group. We will make this precise in Section 7.3. In bounded rank, this will allow us to translate the main theorems in terms of maximal tori (see e.g. Theorem 7.3.10). We will exploit the connection also in large rank, where the asymptotic properties of the Weyl groups will be relevant.

We remark again that it is enough to prove Theorem 7.1.1 for sufficiently large finite simple groups, since Theorem 2.4.2 states that every finite simple group is invariably generated by two elements. What is more, for groups of Lie type we can divide the proof of Theorem 7.1.1 in two steps: first produce an element x_1 for groups of sufficiently large rank, and then produce an element x_2 for groups of bounded rank and sufficiently large fields.

Finally, we note that we are free to define the subsets A_b and A_ℓ from Theorem 7.1.4 only for sufficiently large finite simple groups. In fact, in groups of Lie type, for the set A_ℓ we may assume that r is sufficiently large, and for the set A_b we may assume $q \gg r^r$. Clearly, also the proof of Theorem 7.1.4 splits naturally into bounded rank and large rank.

7.1.4 Organization of the chapter and notation

In Section 7.2 we prove the main theorems for alternating groups. In Sections 7.3–7.7 we prove the main theorems for groups of Lie type of bounded rank. In Section 7.8 we deal with groups of Lie type of large rank.

The asymptotic notation $f = O(g)$ means that $|f| \leqslant Cg$ for some constant C (so f might also be negative). As in Chapter 2, for a subgroup H of G, we set

$$\widetilde{H} = \bigcup_{g \in G} H^g.$$

We will only use this notation in Section 7.3, where G will always be $X_\sigma = X(q)$ (see Section 7.3 for definitions).

7.2 Alternating groups

In this section we prove Theorem 7.1.4 for alternating groups. Since $|A_\ell| = 1$, this implies Theorem 7.1.1. Conceptually, the proof follows from [LP93]. We will make use of [EFG16] and [EFK16] in order to obtain better bounds.

If n is odd, choose $x \in G = A_n$ to be an n-cycle. If n is even, choose $x \in G$ to have cycle type $(n/2, n/2)$. Then set $A_\ell = \{x\}$. In the first case, the overgroups of x are transitive subgroups, while in the second case the overgroups of x are either transitive, or fix a set of size $n/2$.

By [EFK16, Theorem 1.1], the proportion of elements of A_n lying in proper transitive subgroups of A_n is $O(n^{-0.08})$. By [EFG16, Theorem 1.1], the same bound holds for the proportion of elements fixing a set of size $n/2$. Therefore

$$\frac{|\mathscr{M}(x)|}{|G|} = O(n^{-0.08}).$$

Then $\mathbf{P}_{\mathrm{inv}}(G, A_\ell) = 1 - O(n^{-0.08})$ by Lemma 7.1.7. (We mention that, in [EFK16] and [EFG16], much more precise estimates are proved than those used here.)

7.3 Groups of Lie type of bounded rank: preliminaries

In this section we introduce all the machinery that will lead us to the proof of the main theorems for groups of Lie type of bounded rank. We will prove Theorem 7.1.4 in Sections 7.4 and 7.5, and deduce Theorems 7.1.2 and 7.1.3 in Section 7.6.

We single out a special case.

Theorem 7.3.1. *Theorems 7.1.2, 7.1.3, 7.1.4 hold in case $G \cong \mathrm{PSL}_2(q)$.*

The subgroup structure of $\mathrm{PSL}_2(q)$ is very well known and it is easy to prove Theorem 7.3.1. We will do this at the beginning of Section 7.5.

The reason why we separate out this case is minor. Indeed, we will give an argument which works in general, but which gives error terms in q of type $O(1/q^{1/2})$ if $G \cong \mathrm{PSL}_2(q)$, and of type $O(1/q)$ otherwise (see Remark 7.3.5). We then prefer to consider $\mathrm{PSL}_2(q)$ separately, and deal with the other cases uniformly.

Let X be a connected simple linear algebraic group over an algebraic closure k of a finite field of characteristic p (for all this theory, our reference is [MT11]). Let σ be an endomorphism of X such that the set X_σ of fixed points of σ is a finite group, and such that the derived subgroup $[X_\sigma, X_\sigma] = X'_\sigma$ is a perfect group. Let T be a σ-stable maximal torus of X. Then σ acts naturally on the character group $\mathrm{Hom}(T, \mathrm{GL}_1)$. It turns out that the eigenvalues of σ on $\mathrm{Hom}(T, \mathrm{GL}_1) \otimes_{\mathbf{Z}} \mathbf{C}$ have all the same absolute value, which we denote by q, and which is a fractional power of p (cf. [MT11, Lemma 22.1 and Proposition 22.2]). We will write $X_\sigma = X(q)$, except in case X_σ is a Suzuki or a Ree group, in which case q is not an integer but q^2 is an integer, and we will write $X_\sigma = X(q^2)$. However, when we present general arguments which apply to all groups, we will write $X(q)$, without highlighting the difference for Suzuki and Ree groups. This notation conflicts slightly with Table 7.1, where for exceptional groups $X(q)$ denotes a finite simple group (while here $X(q)$ needs not be perfect). This will cause no confusion.

In Sections 7.3–7.7, we fix X, and we let σ vary — concretely, for classical groups we are fixing the rank and we are letting q go to infinity, and moreover we are dealing with the exceptional groups.

7.3.1 Subgroups of maximal rank

We begin by recalling a well known fact.

Theorem 7.3.2. *Assume $H \leqslant X$ is closed, connected and σ-stable, and assume $s \in H_\sigma$ is semisimple. Then, s belongs to a σ-stable maximal torus of H.*

Proof. We have that s is contained in a maximal torus S of H. Then $s \in S \leqslant \mathrm{C}_H(s)^\circ$. Since s is central in $\mathrm{C}_H(s)^\circ$, s is contained in every maximal torus of $\mathrm{C}_H(s)^\circ$ (this is also a maximal torus of H). Now $\mathrm{C}_H(s)^\circ$ is σ-stable and

connected, hence by Lang–Steinberg it contains a σ-stable maximal torus S' (cf. [MT11, Theorem 21.11]). We have $s \in S'$ and we are done. $\square$

A proper closed subgroup K of X is called of maximal rank if it contains a maximal torus of X. A subgroup of X_σ is called of maximal rank if it is of the form K_σ, where K is a σ-stable subgroup of maximal rank (by Lang–Steinberg, K contains a σ-stable maximal torus).

Now we list some notation that we will keep throughout Sections 7.3–7.7. We advise the reader to consult this list whenever he or she finds an unknown symbol, rather than to read now all the items. We prefer to amass here this notation, since we will use it several times in several different places.

- $\diamond$ p denotes the characteristic of the field k.
- $\diamond$ r denotes the rank of X (i.e, the dimension of a maximal torus). By Theorem 7.3.1, we may assume $r \geqslant 2$, but we will make the requirement explicit.
- $\diamond$ $\mathcal{M} = \mathcal{M}(X_\sigma)$ denotes the set of maximal subgroups of X_σ of the form K_σ, where K is a maximal σ-stable subgroup of X of maximal rank.
- $\diamond$ $\mathcal{M}_{\rm con} = \mathcal{M}_{\rm con}(X_\sigma)$ denotes the set of subgroups of X_σ of the form K_σ°, where K is a maximal σ-stable subgroup of X of maximal rank and K° denotes its connected component.[4]
- $\diamond$ For $x \in X_\sigma$, $\mathcal{M}(x)$ denotes the set of all conjugates of overgroups of x belonging to $\mathcal{M}$.
- $\diamond$ For $x \in X_\sigma$, $\mathcal{T}(x)$ denotes the set of maximal tori of X_σ contained in some K_σ°, where $K_\sigma \in \mathcal{M}(x)$ (note that K_σ° needs not contain x).
- $\diamond$ For a subset $A \subseteq X_\sigma$, $\mathbf{P}^*_{\rm inv}(X_\sigma, A)$ denotes the probability that, if $y \in X_\sigma$ is chosen uniformly at random, there exists $x \in A$ such that for every $g_1, g_2 \in X_\sigma$, $\langle x^{g_1}, y^{g_2} \rangle$ is not contained in any maximal subgroup of X_σ not containing X'_σ. (In particular $\mathbf{P}^*_{\rm inv}(X_\sigma, A) = \mathbf{P}_{\rm inv}(X_\sigma, A)$ if X_σ is perfect.)
- $\diamond$ $\Delta = \Delta(X_\sigma)$ denotes the set of elements y of X_σ which are regular semisimple, and such that if y belongs to a maximal subgroup M of X_σ, then either $X'_\sigma \leqslant M$, or $M = K_\sigma \in \mathcal{M}$ and $y \in K_\sigma^\circ$.
- $\diamond$ For a maximal torus S of X_σ, Δ_S denotes the set of elements of X_σ lying in a conjugate of S and in Δ.

We recall a theorem which is essential for our purposes. Note that the proportion in the statement is independent of X, hence the result can be applied to groups of growing Lie rank (indeed we will use it in Section 7.8).

[4] For a closed subgroup K of X, whenever we write K_σ° we mean $(K^\circ)_\sigma$.

Theorem 7.3.3. *[GL01, Theorem 1.1] The proportion of regular semisimple elements of $X(q)$ is $1 - O(1/q)$.*

For Suzuki and Ree groups the proportion is in fact $1 - O(1/q^2)$, but we will not use this. The proof of the following theorem is essentially contained in [FG03].

Theorem 7.3.4. *Assume $r \geqslant 2$. We have*

$$\frac{|\Delta|}{|X(q)|} = 1 - \frac{O(r^r)}{q}.$$

Proof. Clearly we can assume $r \leqslant q$, otherwise the statement is empty. Let A_1 be the set of elements which are not regular semisimple. Let A_2 be the set of elements which belong to maximal subgroups of $X(q)$ which do not contain any maximal torus and which do not contain $X(q)'$. Let A_3 be the set of elements which belong to $K_\sigma \setminus K^\circ$ for some maximal σ-stable subgroup K of X of maximal rank. We need to prove $|A_1 \cup A_2 \cup A_3|/|X(q)| = O(r^r/q)$. We have $|A_1|/|X(q)| = O(1/q)$ by Theorem 7.3.3.

We deal with A_2. Let Ω be the set of maximal subgroups of $X(q)$ which do not contain $X(q)'$, which do not contain any maximal torus of $X(q)$, and which are not subfield subgroups (cf. [FG03, Section 3]). If $M \in \Omega$ then M has $O(q^{r-1})$ conjugacy classes (see [FG03] and the proof of [FG12, Theorem 7.3]).

Assume now $x \in X(q)$ is regular semisimple. Then the $X(q)$-class of x has size $O(|X(q)|)/(q-1)^r$. By Theorem 7.3.3, we see that if $M \in \Omega$ then

$$\frac{|\widetilde{M}|}{|X(q)|} = \frac{O(q^{r-1})}{(q-1)^r} + \frac{O(1)}{q} = \frac{O(1)}{q},$$

where in the last equality we used $r \ll q$. It it known (cf. [LMS05, Theorem 1.3]) that the number of conjugacy classes of subgroups in Ω is bounded by a function of r. (Note that $X(q)$ surjects, with central kernel, onto an almost simple group generated by inner-diagonal automorphisms.) In case X is classical, by [GLT12, Theorem 1.2] we can take this function to be $O(r^6)$.

Now we deal with subfield subgroups. The argument given in [FG03, Lemma 3.7] shows that the proportion of elements lying in subfield subgroups is $O(r/q^{r/2})$ $+ O(1/q)$, which is $O(r/q)$ since $r \geqslant 2$. Therefore $|A_2|/|X(q)| = O(r^6/q)$.

Finally we deal with A_3. Let K be a maximal σ-stable closed subgroup of X of maximal rank. We claim that

$$\frac{|\bigcup_{g \in X(q)} (K_\sigma \setminus K^\circ)^g|}{|X(q)|} = \frac{O(|K_\sigma : K_\sigma^\circ|)}{q}.$$

This is essentially contained in the proof of [FG03, Proposition 4.2] (we use the same arguments and the same computations, except that we bound the size of a regular semisimple class by $O(X(q))/(q-1)^r$, and moreover we use $r \ll q$, so that $((q+1)/(q-1))^{r-1}$ is bounded).

At this point we can deduce that $|A_3|/|X(q)| = O(r^r/q)$. Indeed, it is known that $|K_\sigma : K_\sigma^\circ| \ll (r+1)!$, and moreover the number of $X(q)$-conjugacy classes

of maximal subgroups of maximal rank is linear in r, from which $|A_3|/|X(q)| = O(r^r/q)$. (These facts are known in a very precise way, cf. [LSS92] and [LS98]. We will recall them in Sections 7.4 and 7.5. The term $(r+1)!$ can occur for stabilizer of decompositions in classical groups.)

Putting together the bounds given for A_1, A_2 and A_3, we get the result. □

Remark 7.3.5. By the same argument as in the previous proof, the proportion of elements of $\mathrm{SL}_2(q)$ belonging to subfield subgroups is $O(1/q^{1/2})$. If q is a square, the proportion of elements inside a conjugate of $\mathrm{SL}_2(q^{1/2})$ is indeed of this form. This is the only reason for which we have considered separately this case.

We also note that, in bounded rank, an essential part of our method is to focus on regular semisimple elements. By work of Guralnick–Lübeck [GL01] and Fulman–Neumann–Praeger [FNP05], it is known that, except for Suzuki and Ree groups, the proportion of elements of $X(q)$ which are not regular semisimple is comparable to $1/q$ (up to constants). Therefore, with our method we cannot get error terms in q which are better than $O(1/q)$.

7.3.2 Maximal tori and Weyl group

There is a well known connection between the maximal tori of X_σ and the Weyl group of X, which now we recall (see [MT11, Section 25] for the general theory). Together with Theorems 7.3.3 and 7.3.4, this will enable us to translate the main theorems in terms of maximal tori. For the following discussion, see also [FG03].

Throughout this subsection, we fix a σ-stable maximal torus T, and let $W = \mathrm{N}_X(T)/T$ be the Weyl group of X with respect to T. Then σ acts on W. There is a bijection between X_σ-conjugacy classes of σ-stable maximal tori of X and W-conjugacy classes contained in the coset σW of the group $\langle\sigma\rangle \ltimes W$ (if σ acts trivially on W, these can be identified with the conjugacy classes of W). If $w \in W$, we denote by T_w any representative of the conjugacy class of maximal tori corresponding to the W-class of σw. We have $T_w = T^g$, where $g \in X$ is such that $g^\sigma g^{-1}$ maps to $w \in W$. Moreover $\mathrm{N}_{X_\sigma}(T_w)/(T_w)_\sigma \cong \mathrm{C}_W(\sigma w)$.

Let $\Psi \subseteq W$ be such that $\{\sigma w, w \in \Psi\}$ is a set of representatives for the W-classes in the coset σW. Let Ω be a subset of Ψ. Denote by $\mathbf{P}(W, \sigma, \Omega)$ the probability that a random element of σW is W-conjugate to σw for some $w \in \Omega$. In case $\Omega = \{w\}$, we will write $\mathbf{P}(W, \sigma, w)$ instead of $\mathbf{P}(W, \sigma, \{w\})$. Using that $\mathrm{N}_{X_\sigma}(T_w) \leqslant \mathrm{N}_{X_\sigma}((T_w)_\sigma)$, by a trivial union bound we get

$$\frac{|\bigcup_{w\in\Omega} \widetilde{(T_w)_\sigma}|}{|X_\sigma|} \leqslant \sum_{w\in\Omega} \frac{1}{|\mathrm{C}_W(\sigma w)|} = \mathbf{P}(W, \sigma, \Omega). \tag{7.3.1}$$

Despite being trivial, for q large this bound is accurate.

Theorem 7.3.6.

$$\frac{|\bigcup_{w\in\Omega} \widetilde{(T_w)_\sigma}|}{|X(q)|} = \mathbf{P}(W, \sigma, \Omega) + \frac{O(1)}{q}. \tag{7.3.2}$$

Proof. By Theorems 7.3.2 and 7.3.3 we have

$$\beta := \frac{|\bigcup_{w\in\Omega} \widetilde{(T_w)_\sigma}|}{|X(q)|} + \frac{|\bigcup_{w\in\Psi\setminus\Omega} \widetilde{(T_w)_\sigma}|}{|X(q)|} = 1 - \frac{O(1)}{q}.$$

(Note that $\beta \leqslant 1$ by (7.3.1).) Moreover

$$\begin{aligned}\frac{O(1)}{q} = 1 - \beta = &\left(\mathbf{P}(W,\sigma,\Omega) - \frac{|\bigcup_{w\in\Omega} \widetilde{(T_w)_\sigma}|}{|X(q)|}\right)\\ &+ \left(\mathbf{P}(W,\sigma,\Psi\setminus\Omega) - \frac{|\bigcup_{w\in\Psi\setminus\Omega} \widetilde{(T_w)_\sigma}|}{|X(q)|}\right),\end{aligned}$$

where by (7.3.1) both summands are nonnegative. In particular they are both $O(1/q)$, and the proof is concluded. □

This theorem was used in [FG03]. One of the main observations in this section is that X is fixed. Then W is fixed, and if Ω is nonempty the term at the right-hand side of (7.3.2) is always bounded away from zero: it is at least $1/|W| + O(1/q)$.

For a subset A of $X(q)$, let $T_1, \dots, T_\ell$ be a set of representatives of the $X(q)$-conjugacy classes of members of $\cap_{x\in A}\mathcal{T}(x)$ (possibly $\ell = 0$). Write $T_i = (T_{w_i})_\sigma$, where T_{w_i} is a σ-stable maximal torus of X and $w_i \in W = \mathrm{N}_X(T)/T$. Set $\Omega = \{w_1, \dots, w_\ell\}$.

Theorem 7.3.7. *Assume $r \geqslant 2$. We have*

$$1 - \mathbf{P}^*_{\mathrm{inv}}(X(q), A) = \mathbf{P}(W,\sigma,\Omega) + \frac{O(r^r)}{q}. \tag{7.3.3}$$

Proof. Reasoning as in Lemma 7.1.7, and using Theorem 7.3.4, we have

$$1 - \mathbf{P}^*_{\mathrm{inv}}(X(q), A) = \frac{|\bigcap_{x\in A}\bigcup_{M\in\mathcal{M}(x)} M|}{|X(q)|} + \frac{O(r^r)}{q}.$$

Now we look at the right-hand side of the above equation. Assume $y \in \Delta$. Then y is regular semisimple; let $S = \mathrm{C}_X(y)^\circ$ be its maximal torus in X. Assume $y \in K_\sigma$ for some $K_\sigma \in \mathcal{M}(x)$ and some $x \in A$ (and K is σ-stable of maximal rank.) By definition of Δ we have $y \in K^\circ$. By Theorem 7.3.2, y lies in some σ-stable maximal torus of K°, which is also a maximal torus of X, hence must coincide with S. In particular, if y lies in some member of $\mathcal{M}(x)$ for every $x \in A$, then S belongs to $\mathcal{T}(x)$ for every $x \in A$. Using Theorem 7.3.4, this shows that

$$1 - \mathbf{P}^*_{\mathrm{inv}}(X(q), A) = \frac{|\bigcup_{i=1}^{\ell} \widetilde{(T_{w_i})_\sigma}|}{|X(q)|} + \frac{O(r^r)}{q}. \tag{7.3.4}$$

Finally, the right-hand side of (7.3.4) is equal to the right-hand side of (7.3.3) by Theorem 7.3.6. □

We record a consequence of the previous proof.

Theorem 7.3.8. *Assume $\cap_{x\in A}\mathcal{T}(x) = \varnothing$. Then the set Δ contributes to $\mathbf{P}^*_{\mathrm{inv}}(X(q), A)$. In other words, for every $y \in \Delta$, there exists $x \in A$ such that, for every $g_1, g_2 \in X(q)$, $\langle x^{g_1}, y^{g_2}\rangle$ is not contained in any maximal subgroup of $X(q)$ not containing $X(q)'$.*

Proof. Follows from the previous proof. □

7.3.3 From X_σ to X'_σ

All the discussion above is about X_σ, which needs not be perfect. However in the end we want to prove our main theorems, which are about finite simple groups. We now establish the connection, showing also that the isogeny type of X is not relevant. Let X_{sc} be the group of simply connected type, and let $\pi : X_{\mathrm{sc}} \to X$ be the natural isogeny. Then σ lifts to a morphism $X_{\mathrm{sc}} \to X_{\mathrm{sc}}$ (cf. [MT11, Proposition 22.7]), which for convenience we still denote by σ. Write as usual $X_\sigma = X(q)$ and $(X_{\mathrm{sc}})_\sigma = X_{\mathrm{sc}}(q)$.

Lemma 7.3.9. *Let A be a subset of $X_{\mathrm{sc}}(q)$. Then*

$$\mathbf{P}_{\mathrm{inv}}(X(q)', A^\pi) = \mathbf{P}_{\mathrm{inv}}(X_{\mathrm{sc}}(q), A) = \mathbf{P}^*_{\mathrm{inv}}(X(q), A^\pi) + \frac{O(r^r)}{q}.$$

Proof. Let Z be the kernel of π. Then $(X_{\mathrm{sc}})_\sigma/Z_\sigma \cong ((X_{\mathrm{sc}})_\sigma)^\pi = X'_\sigma$ ([MT11, Proposition 24.21]). Moreover Z_σ is contained in every maximal subgroup of $(X_{\mathrm{sc}})_\sigma$, hence the first equality of the statement holds.

Now note that $\mathbf{P}_{\mathrm{inv}}(X_{\mathrm{sc}}(q), A) = \mathbf{P}^*_{\mathrm{inv}}(X_{\mathrm{sc}}(q), A)$, since $X_{\mathrm{sc}}(q)$ is perfect. Since Z is contained in every maximal torus of X_{sc}, π induces a bijection between σ-stable subgroups of maximal rank of X_{sc} and of X, which maps overgroups of $y \in A$ to overgroups of $y^\pi \in A^\pi$. Then the second equality follows from Theorem 7.3.7. □

In order to prove Theorem 7.1.4 for groups of Lie type of bounded rank, in view of Theorem 7.3.7 and Lemma 7.3.9 it is sufficient to choose X of some isogeny type, and prove the following statement.

Theorem 7.3.10. *Assume $r \geqslant 2$ and $r^r \ll q$, and assume $X(q) \ncong G_2(q)$ when $3 \mid q$. Then, there exists $A_b \subseteq X(q)'$ of size as in Table 7.1 such that $\cap_{x\in A_b}\mathcal{T}(x) = \varnothing$. If $3 \mid q$, then $\mathbf{P}_{\mathrm{inv}}(G_2(q), G_2(q)) = 1/2 + O(1/q)$.*

We will prove Theorem 7.3.10 in Sections 7.4 and 7.5. By the Borel–Tits theorem (see [BT71, Corollaire 3.9]), a maximal σ-stable subgroup of X of maximal rank is either parabolic, or its connected component is reductive. We make some general considerations regarding the second case.

7.3.4 Reductive subgroups of maximal rank

We fix a pair (T, B), where T is a σ-stable maximal torus of X, and B is a σ-stable Borel subgroup of X containing T. We let Φ be the root system with respect to T, and we denote by U_α, $\alpha \in \Phi$, the root subgroups with respect to T. We let $W = \mathrm{N}_X(T)/T$ be the Weyl group of X with respect to T. Throughout this subsection, we make the following

Assumption 7.3.11. *σ acts trivially on Φ.*

This is not essential, but the results are easier to state, and we will apply them only under this assumption. The following discussion is taken from [LSS92]. Let K be a closed connected reductive subgroup of X containing T. Then, $K = \langle T, U_\alpha, \alpha \in \Psi\rangle$ for a p-closed subset Ψ of Φ (see [MT11, Section 13] for this notion). Let $W(\Psi)$ be the Weyl group of K, i.e., the subgroup of W generated by the reflections in roots of Ψ. We have $\mathrm{N}_X(K)/K \cong \mathrm{N}_W(W(\Psi))/W(\Psi) =: W_\Psi$. Note that K is σ-stable, since σ acts trivially on Φ.

Assume now H is a σ-stable conjugate of K. In particular, there exists $g \in X$ such that $H = K^g$ and such that T^g is σ-stable. Then $g^\sigma g^{-1} \in \mathrm{N}_X(T) \cap \mathrm{N}_X(K)$, which maps to an element of W_Ψ that we denote by $\rho(K^g)$.

Lemma 7.3.12. *The map ρ defined above induces a well-defined bijection between $\{X_\sigma$-orbits on the σ-stable conjugates of $K\}$ and $\{$conjugacy classes in $W_\Psi\}$.*

Proof. This is [Car78, Propositions 1 and 2] in case σ acts trivially on Φ. □

In [Car78, Propositions 1 and 2] the general case (i.e., σ does not necessarily act trivially on Φ) is considered. This is more technical to state.

Next, given a σ-stable maximal torus S, we want to determine its closed connected reductive overgroups in X. Fix g such that $S = T^g$ and let w be the image of $g^\sigma g^{-1}$ in W.

Lemma 7.3.13. *The closed connected σ-stable reductive overgroups of S are in bijection with p-closed subset of Φ which are w-stable.*

Proof. If H is a closed connected σ-stable reductive overgroup of S, set $K := H^{g^{-1}}$. Then $K = \langle T, U_\alpha, \alpha \in \Psi\rangle$ for some (unique) p-closed subset Ψ of Φ. Moreover, $g^\sigma g^{-1}$ normalizes K and T, hence w fixes Ψ. Conversely, assume Ψ is w-stable and p-closed; then $\langle T, U_\alpha, \alpha \in \Psi\rangle^g$ is a closed connected reductive overgroup of S (see [MT11, Theorem 13.6]). For $\alpha \in \Psi$ we have $(U_\alpha^g)^\sigma = U_\alpha^{\sigma g^\sigma} = U_\alpha^{g^\sigma} = (U_{\alpha w})^g$. Since Ψ is w-stable, we deduce that $\langle T, U_\alpha, \alpha \in \Psi\rangle^g$ is σ-stable. □

See [Wei92, Theorem 5] for a more general statement, considering the case in which σ does not necessarily act trivially on Φ. This does not present serious changes: one replaces w by σw in the statement above.

At this point we divide the discussion between between exceptional and classical groups.

G	$\lvert x_1\rvert$	$\lvert x_2\rvert$
${}^2B_2(q^2)$	Φ'_8	$\Phi'_8(-q)$
${}^2G_2(q^2)$	Φ'_{12}	$\Phi'_{12}(-q)$
$G_2(q)$, $3 \nmid q$	Φ_3	$\Phi_3(-q)$
${}^3D_4(q)$	Φ_{12}	$(q^3+1)(q-1)/(2,q-1)$
${}^2F_4(q^2)$	Φ'_{24}	$\Phi'_{24}(-q)$
$F_4(q)$, q odd	Φ_{12}	Φ_8
$E_6(q)$	$\Phi_3\Phi_{12}/(3,q-1)$	$\Phi_1\Phi_2\Phi_8/\delta$
${}^2E_6(q)$	$\Phi_6\Phi_{12}/(3,q+1)$	$\Phi_1\Phi_2\Phi_8/\delta$
$E_7(q)$	$\Phi_2\Phi_{18}/(2,q-1)$	$\Phi_1\Phi_9/(2,q-1)$
$E_8(q)$	Φ_{30}	$\Phi_{30}(-q)$

Table 7.3: $A_b = \{x_1, x_2\}$ in Theorem 7.1.4 for exceptional groups different from $F_4(2^a)$.

7.4 Exceptional groups

In this section we will prove Theorem 7.3.10 (hence Theorem 7.1.4) for simple exceptional groups. We choose X of adjoint type. The subgroups of maximal rank of X_σ have been classified by Liebeck, Saxl and Seitz [LSS92].

We can assume that q is sufficiently large in the proof. This implies that every maximal torus S_σ of $X(q)$ contains regular semisimple elements. (In fact, more is true. By Theorem 7.3.3, almost all elements of $X(q)$ are regular semisimple. By Theorem 7.3.6, the proportion of elements in a conjugate of S_σ is bounded away from zero; therefore, almost all elements in a conjugate of S_σ are regular semisimple.) In particular, it follows that whenever $S_\sigma \leqslant M_\sigma^\circ$, with $M_\sigma \in \mathcal{M}$, then $S \leqslant M^\circ$.

We define A_b as the set of elements appearing in Table 7.3. There is a slight ambiguity in the notation. Namely, so far we have denoted $X(q) = X_\sigma$; recall however that the elements in Table 7.3 belong to the derived subgroup $X(q)'$. This should cause no confusion.

In the table, the case $F_4(q)$ with q even is missing. In this case the set A_b has size 6, hence for aesthetic reasons we have not included it. We will treat this case in detail in Subsection 7.4.6. Each element in Table 7.3 is regular semisimple. The existence of these elements follows from the general theory of the structure of maximal tori, cf. [MT11, Section 25]; in cases $E_6(q)$, ${}^2E_6(q)$ and $E_7(q)$ the order is adjusted so that indeed the elements belong to the derived subgroup. When necessary, we will provide more details regarding the elements along the proof. We write $\Phi_n = \Phi_n(q)$ for the n-th cyclotomic polynomial evaluated at q. Moreover $\Phi'_8 = \Phi'_8(q) = q^2 + \sqrt{2}q + 1$, $\Phi'_{12} = \Phi'_{12}(q) = q^2 + \sqrt{3}q + 1$, $\Phi'_{24} = \Phi'_{24}(q) = q^4 + \sqrt{2}q^3 + q^2 + \sqrt{2}q + 1$ (this notation is taken from [GM12a]). We will refer to [GM12a, Table 6] and [GM12b, Table 1] for the overgroups of many elements in Table 7.3, although we remark that these tables rely mostly on [Wei92].

7.4.1 Some twisted groups, and $E_8(q)$

In many cases we can exploit a very convenient situation. Indeed, consider the groups ${}^2B_2(q^2)$,${}^2G_2(q^2)$,${}^3D_4(q)$,${}^2F_4(q^2)$ and $E_8(q)$. Then by [GM12a, Table 6] we see that the element x_1 lies only in one maximal subgroup, namely $\mathrm{N}_{X(q)}(S_\sigma)$, where S_σ is the maximal torus of $X(q)$ containing x_1. Since q is large, S_σ contains regular semisimple elements, and in particular $\mathrm{N}_{X(q)}(S_\sigma) = \mathrm{N}_{X(q)}(S)$.[5] The connected component of $\mathrm{N}_X(S)$ is S, since S has finite index in its normalizer. By definition, we deduce that $\mathcal{T}(x)$ contains only the conjugates of S_σ.

Then, in order to prove Theorem 7.3.10, we just need to show that the element x_2 does not belong to any conjugate of $N_{X(q)}(S_\sigma)$. This is easily done by order considerations.

7.4.2 $E_6(q)$ and ${}^2E_6(q)$

We write $E_6(q) = {}^+E_6(q)$ and ${}^2E_6(q) = {}^-E_6(q)$. Consider ${}^\varepsilon E_6(q)$ with $\varepsilon \in \{+,-\}$. By [GM12b, Table 1], x_1 is contained only in $({}^3D_4(q) \times (q^2+\varepsilon q+1))$.3 (among the maximal subgroups of X_σ). The order of x_2 is $(q^4+1)/(q^2-1)$, divided by a small number δ. Set $h = (4, q-\varepsilon 1)$. For $\varepsilon = -$, x_2 is contained in a maximal subgroup $M = h.(\mathrm{P}\Omega^\varepsilon_{10}(q) \times (q-\varepsilon)/h).h$; and for $\varepsilon = +$, x_2 is contained in a parabolic subgroup with Levi complement of type D_5 (cf. [LSS92, Table 5.1]). By order considerations we see that if $\varepsilon = +$ then $\mathcal{M}(x_2)$ contains only parabolics of type D_5, while if $\varepsilon = -$ then $\mathcal{M}(x_2)$ contains the conjugates of M, and parabolics with Levi complement of type ${}^2D_4$. Using the knowledge of maximal tori of ${}^\varepsilon E_6(q)$ (see [DF91]), we deduce by order considerations that $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$, which proves Theorem 7.3.10.

In Subsections 7.4.3–7.4.6, we employ the notation of Subsection 7.3.4.

7.4.3 $G_2(q)$ with $3 \nmid q$

We immediately recall some facts regarding maximal tori of $G_2(q)$ that we will use also in Subsection 7.4.4. We have $W = W(G_2) \cong D_{12}$, hence by the general theory (cf. [MT11, Section 25]) there are six $G_2(q)$-classes of maximal tori, with representatives $T_1, \ldots, T_6$, and with orders $q^2-1, q^2-1, (q-1)^2, (q+1)^2, q^2+q+1, q^2-q+1$, respectively. We assume $T_i = (T_{w_i})_\sigma$, where w_1 is a reflection in a short root, w_2 is a reflection in a long root, $w_3 = 1$, $w_4 = -1$, $|w_5| = 3$, $|w_6| = 6$. For $i = 1, \ldots, 6$, we will write Δ_i instead of Δ_{T_i}.

By Theorem 7.3.6, the proportion of elements lying in a conjugate of T_1 or T_2 is equal to $O(1/q)$ plus the proportion of noncentral involutions of D_{12}, which is $1/2$. Consequently, the proportion of elements lying in a conjugate of T_i for some $i = 3, \ldots, 6$ is $1/2 + O(1/q)$. By Theorem 7.3.4, the same estimates hold for the proportion of the Δ_i's.

[5] We note that, in fact, $\mathrm{N}_{X(q)}(S_\sigma) = \mathrm{N}_{X(q)}(S)$ holds under the weaker hypothesis that S_σ is *nondegenerate*; see [Car93, Section 3.6] for this notion.

We begin the proof in case $3 \nmid q$. By [LSS92, Table 5.1] and by order considerations, $\mathcal{M}(x_1)$ contains only the conjugates of $\mathrm{SL}_3(q).2$, and $\mathcal{M}(x_2)$ contains only the conjugates of $\mathrm{SU}_3(q).2$. We need to show these two subgroups do not contain a common maximal torus (up to conjugacy). By order considerations, if there exists a common torus of $\mathrm{SL}_3(q).2$ and $\mathrm{SU}_3(q).2$, then it must be $T_1 = (T_{w_1})_\sigma$ or $T_2 = (T_{w_2})_\sigma$. For $i = 1, 2$, fix $g_i \in X = G_2$ such that $T_{w_i} = T^{g_i}$ (and $g_i^\sigma g_i^{-1}$ maps to $w_i \in W$). By Lemma 7.3.13, the closed connected reductive subgroups of G_2 containing T_{w_i} are precisely the subgroups $K(\Psi)^{g_i}$, where $K(\Psi) = \langle T, U_\alpha, \alpha \in \Psi \rangle$ and Ψ is p-closed and w_i-stable. Since $3 \nmid q$, by [MT11, Theorem 13.14] we deduce that every p-closed subset of Φ is closed; in particular there is only one p-closed subset Ψ of type A_2: the set of all long roots. Note that $W_\Psi = \mathrm{N}_W(W(\Psi))/W(\Psi) \cong C_2$, hence by Lemma 7.3.12 there are two corresponding $G_2(q)$-classes. Now $w_2 \in W(\Psi)$, while $w_1 \notin W(\Psi)$; then by Lemma 7.3.12 $K(\Psi)^{g_2}$ is $G_2(q)$-conjugate to $K(\Psi)$, and $K(\Psi)_\sigma \cong \mathrm{SL}_3(q)$, while $K(\Psi)_\sigma^{g_1} \cong \mathrm{SU}_3(q)$. Theorem 7.3.10 follows.

7.4.4 $G_2(q)$ with $3 \mid q$

We keep the notation from the beginning of Subsection 7.4.3. Let $G = G_2(q)$. We want to prove $\mathbf{P}_{\mathrm{inv}}(G, G) = 1/2 + O(1/q)$. We will prove the following more precise statement, which we will use in Section 7.6.

Theorem 7.4.1. *(1)* $\mathbf{P}_{\mathrm{inv}}(G, y) \geqslant 1/6 + O(1/q)$ *for a proportion of elements* $y \in G$ *of the form* $1/2 + O(1/q)$.

(2) $\mathbf{P}_{\mathrm{inv}}(G, y) = 0$ *for a proportion of elements* $y \in G$ *of the form* $1/2 + O(1/q)$.

Note that $y \in G$ contributes to $\mathbf{P}_{\mathrm{inv}}(G, G)$ if and only if $\mathbf{P}_{\mathrm{inv}}(G, y) > 0$.

Now we prove Theorem 7.4.1. There is an automorphism γ of $G_2(q)$ which induces a graph automorphism of order two on the Dynkin diagram, exchanging long and short roots. The set Ψ' of short roots is 3-closed (cf. [MT11, Proposition 13.15]). There are two conjugacy classes of subgroups $\mathrm{SL}_3(q).2$, with representatives H_1 and H_2, and two conjugacy classes of subgroups $\mathrm{SU}_3(q).2$, with representatives K_1 and K_2. We have $H_1^\gamma = H_2$ and $K_1^\gamma = K_2$. Moreover γ exchanges the classes of T_1 and T_2. Up to changing indices, T_i is contained in a conjugate of H_i and K_i. What is more, the only overgroups of T_5 (resp. T_6) are conjugates of H_1 and H_2 (resp. conjugates of K_1 and K_2).

Therefore, by definition of Δ, every element of Δ_3 (resp. Δ_4, resp. Δ_5) invariably generates with every element of Δ_6 (resp. Δ_5, resp. Δ_6). We observed that the proportion of elements belonging to Δ_i for some $i = 3, \ldots, 6$ is $1/2 + O(1/q)$. Moreover, for $i \in \{5, 6\}$, we have $|\Delta_i|/|G| = 1/6 + O(1/q)$. Therefore (1) is proved.

We move to (2). We want to show that

($\star$) for every $x \in G$, x belongs to a maximal subgroup containing a conjugate of T_1, and to a maximal subgroup containing a conjugate of T_2.

This implies that all elements y lying in a conjugate of T_1 or T_2 are such that $\mathbf{P}_{\mathrm{inv}}(G, y) = 0$. We observed that these elements have proportion $1/2 + O(1/q)$, hence in order to prove (2) we only need to prove ($\star$).

It is sufficient to focus on $i = 1$, since the two tori are exchanged by an automorphism of G. Representatives of the conjugacy classes of maximal subgroups containing T_1 are the following:

$$\{P, H_1, K_1, C\}.$$

Here P is a parabolic subgroup with respect to the short root of a base, and $C \cong (\mathrm{SL}_2(q) \circ \mathrm{SL}_2(q)).2$ is the centralizer of an involution in $G_2(q)$ (cf. [Kle88, Theorem A]).

Let $x \in G$. If x is unipotent, then x is contained in both conjugacy classes of parabolic subgroups. Assume then $x = su$, with $1 \neq s$ semisimple, u unipotent, and $[s, u] = 1$. Then $x \in \mathrm{C}_G(s) < G$. If $\mathrm{C}_G(s)$ is a maximal torus of even order, it is contained in a conjugate of C. The remaining classes of maximal tori have representatives T_5 (contained in H_1) and T_6 (contained in K_1). Examining [Kle88, Table II, p. 41], we see that all other possibilities for $\mathrm{C}_G(s)$ contain a central involution, hence are contained in a conjugate of C. Then ($\star$) is proved and we are done.

7.4.5 $E_7(q)$

By [GM12a, Table 6] and [GM12b, Table 1] we see that x_1 is contained only in a maximal subgroup ${}^2E_6(q)_{\mathrm{sc}}.D_{q+1}$ of X_σ, and x_2 is contained in two (conjugate) parabolics P and P' of type E_6, and in the normalizer of a common Levi complement L. Our aim is to show that $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$.

Claim 7.4.2. *Assume $g \in X = E_7$ and assume $g^\sigma g^{-1} \in \mathrm{N}_X(T)$ maps to $w \in W$. Assume Ψ and Ψ' are two p-closed subsets of Φ of type E_6. If $w \in \mathrm{N}_W(W(\Psi)) \cap \mathrm{N}_W(W(\Psi'))$, then either $w \in W(\Psi) \cap W(\Psi')$ or $w \notin W(\Psi) \cup W(\Psi')$.*

We first observe that Claim 7.4.2 implies $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$. Consider a maximal torus S of X_σ; assume $S = (T_w)_\sigma$, where $T_w = T^g$ and $g^\sigma g^{-1}$ maps to $w \in W$. By Lemma 7.3.13, the closed connected reductive subgroups of E_7 containing T_w are precisely the subgroups $K(\Psi)^g$, where $K(\Psi) = \langle T, U_\alpha, \alpha \in \Psi \rangle$ and Ψ is p-closed and w-stable (i.e., $w \in \mathrm{N}_W(W(\Psi))$). By Lemma 7.3.12 we see that Claim 7.4.2 implies that $S = (T_w)_\sigma$ cannot be contained in a maximal subgroup of type ${}^2E_6(q)_{\mathrm{sc}}D_{q+1}$, and at the same time in a Levi complement of type E_6, so that $S \notin \mathcal{T}(x_1) \cap \mathcal{T}(x_2)$ and (since S was arbitrary) $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$.

In order to prove Claim 7.4.2, we recall that $W = \langle x \rangle \times W^+$, where $|x| = 2$ and $W^+ \cong \mathrm{Sp}_6(2)$ is the "rotation subgroup", consisting of the element of W with determinant 1 in the action on $\mathbf{R}^7$. We will view the elements of W as pairs, according to this decomposition. If Ψ is a subset of type E_6, then $W(\Psi) \cong \mathrm{SO}_6^-(2)$. Clearly we cannot have $W(\Psi) \leqslant W^+$, since $W(\Psi)$ contains reflections.

Let K be the unique subgroup of $W(\Psi)$ of index 2, isomorphic to $\Omega_6^-(2)$. Then $K \leqslant W^+$. Let $H \cong \mathrm{SO}_6^-(2)$ be the normalizer of K in W^+; we have $H = K \rtimes \langle r \rangle$, where r is a reflection in a nonsingular vector (for the orthogonal geometry on $\mathbf{F}_2^6$). We have $W(\Psi) = \langle (x, r), K \rangle$ and $\mathrm{N}_W(W(\Psi)) = \langle x \rangle \times H$. Now if Ψ' is another subset of type E_6, we have $\Psi' = \Psi^g$ with $g \in W^+$, and consequently $W(\Psi') = \langle (x, r^g), K^g \rangle$ and $\mathrm{N}_W(W(\Psi')) = \langle x \rangle \times H^g$. We see that Claim 7.4.2 is equivalent to the following condition:

($\star$) Fix Ψ as above. Then, for every $g \in W^+$, $W(\Psi) \cap \mathrm{N}_W(W(\Psi^g)) \leqslant W(\Psi^g)$.

It is easy to see that ($\star$) is equivalent to

($\star\star$) Fix $K \leqslant W^+$ and $H \leqslant W^+$ as above. Then, for every $g \in W^+$, $K \cap H^g \leqslant K^g$.

($\star\star$) holds in general, in the following sense. Assume q is even, and recall that $\mathrm{Sp}_{2m}(q) \cong \mathrm{SO}_{2m+1}(q)$ (see Subsection 7.8.3 for some words about this isomorphism). Denote by V the $(2m+1)$-dimensional orthogonal module. Then ($\star\star$) is a particular case of the following lemma.

Lemma 7.4.3. *Assume W and W' are nondegenerate hyperplanes of V (not necessarily of the same sign). Then $\Omega(W) \cap \mathrm{SO}(W') \leqslant \Omega(W')$.*

Proof. Recall that $\Omega(W)$ can be characterized as the subset of $\mathrm{SO}(W)$ consisting of the elements g such that $\dim \mathrm{C}_W(g)$ is even (cf. [Wil09, p. 77]). We have $V = W \perp V^\perp$, and g acts trivially on $V^\perp$, therefore $\dim \mathrm{C}_W(g) = \dim \mathrm{C}_V(g) - 1$, which is independent of W. This proves the lemma. □

Claim 7.4.2 is proved and we are done.

7.4.6 $F_4(q)$

We first fix some notation taken from [Law99]. Let $\mathbf{R}^4$ be equipped with the usual orthonormal basis $e_1, \ldots, e_4$. We may take $\Phi \subseteq \mathbf{R}^4$ with set of positive roots

$$\Phi^+ = \{e_i \pm e_j, 1 \leqslant i < j \leqslant 4\} \cup \{e_i, 1 \leqslant i \leqslant 4\} \cup \{(e_1 \pm e_2 \pm e_3 \pm e_4)/2\}$$

and base

$$\Sigma = \{e_2 - e_3, e_3 - e_4, e_4, (e_1 - e_2 - e_3 - e_4)/2\}.$$

As in [Law99], we will write 1 in place of e_1, $1-2$ in place of $e_1 - e_2$, $+---$ in place of $(e_1 - e_2 - e_3 - e_4)/2$, etc. The corresponding reflections will be denoted by w_1, w_{1-2}, w_{+---}, etc. In [Law99] the complete list of maximal tori of $F_4(q)$ is given. In particular, for each $(\delta, \delta') \in \{+, -\}^2$, there are two conjugacy classes of maximal tori of order $(q^3 + \delta 1)(q + \delta' 1)$; we let $T^i_{\delta,\delta'}$, $i = 1, 2$, be representatives for the two classes (so for instance $T^1_{+,-}$ is a representative of a class of tori of

order $(q^3+1)(q-1)$). Assume $T^i_{\delta,\delta'} = (T_w)_\sigma$ with $w = w^i_{\delta,\delta'}$. With notation as in [Law99, pp. 93–96], we may choose

$$w^1_{+,+} = w^{(13)} = w_3w_{2-3}w_{1-2}w_4 \qquad w^2_{+,+} = w_{(20)} = w_1w_2w_4w_{+-+-}$$
$$w^1_{+,-} = w^{(15)} = w_4w_{3-4}w_{2-3} \qquad w^2_{+,-} = w_{(13)} = w_4w_{3-4}w_{+--+}$$
$$w^1_{-,+} = w^{(14)} = w_1w_{3-4}w_{2-3} \qquad w^2_{-,+} = w_{(15)} = w_{1-2}w_4w_{++--}$$
$$w^1_{-,-} = w^{(12)} = w_{3-4}w_{2-3} \qquad w^2_{-,-} = w_{(7)} = w_4w_{+---}$$

Here composition is right-to-left; this however makes no difference, because in a Weyl group every element is conjugate to its inverse, cf. [Car93, Corollary p. 45]. We are now ready to begin the proof. We divide the cases q even and q odd.

(a) Assume q is odd. By [GM12a, Table 6] we have that x_1 is contained only in a subgroup ${}^3D_4(q).3$; and by [GM12b, Table 1] x_2 is contained only in a subgroup $2.\Omega_9(q)$. We need to show that $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$. By order inspection, the only possibilities for $T(x_1) \cap \mathcal{T}(x_2)$ are the eight tori $T^i_{\delta,\delta'}$. By our choice (see [Law99, pp. 94–95]) the maximal tori of type 1 (i.e., the tori $T^1_{\delta,\delta'}$) are contained in $2.\Omega_9(q)$. This subgroup is obtained as the fixed points of a connected reductive subgroup of F_4 of type B_4. What we need to show is that none of the tori $T^1_{\delta,\delta'}$ belongs to a conjugate of ${}^3D_4(q).3$. Fix (δ,δ'), and fix $g \in X$ such that $T_{w^1_{\delta,\delta'}} = T^g$. There is a unique p-closed subset Ψ of Φ of type D_4, namely the set of all long roots (the set of all short roots is only 2-closed). Of course Ψ is fixed by every element of W. Correspondingly, by Lemma 7.3.13, $T_{w^1_{\delta,\delta'}}$ has a unique connected reductive overgroup of type D_4, namely $\langle T, U_\alpha, \alpha \in \Psi\rangle^g$. The fixed points of such a subgroup is of type $D_4(q)$ or ${}^2D_4(q)$. Indeed, this is true for every maximal torus of B_4. It follows that $T^1_{+,+}$ is contained in $D_4(q)$ or ${}^2D_4(q)$, but not in ${}^3D_4(q)$. This concludes the proof in case q is odd.

(b) Assume q is even. There is an automorphism γ of $F_4(q)$ which induces a graph automorphism of order two on the Dynkin diagram, sending $2-3$ to $+--$ and $3-4$ to 4. In this case, there are two conjugacy classes of maximal subgroups isomorphic to $\Omega_9(q)$: we pick representatives $B_4(1)$ and $B_4(2)$ for them. Similarly, there are two conjugacy classes of maximal subgroups isomorphic to ${}^3D_4(q).3$, $\mathrm{P}\Omega^+_8(q).S_3$, P, P': we pick representatives ${}^3D_4(i)$, $D_4(i)$, $P(i)$, $P'(i)$, $i = 1, 2$, in the respective cases. Here $P(1)$ (resp. $P(2)$) denotes a parabolic subgroup of type B_3 (resp. C_3); $P'(1)$ (resp. $P'(2)$) denotes a parabolic subgroup of type $A_2 \times \tilde{A}_1$ (resp. $\tilde{A}_2 \times A_1$), where $\tilde{A}_i$ denotes a subset consisting of short roots. (The reason why there are two classes of reductive subgroups as above is that there are subsets of Φ of type C_4 and D_4 which are 2-closed; see [MT11, Proposition 3.15].) We see that all the pairs of classes above are fused by γ, i.e., $B_4(1)^\gamma = B_4(2)$, and similarly for the others. There are also maximal subgroups $e.(\mathrm{PSL}^\varepsilon_3(q) \times \mathrm{PSL}^\varepsilon_3(q)).e.2$, with $\varepsilon \in \{+,-\}$ (one class for each sign; here $\mathrm{PSL}^+_3(q) = \mathrm{PSL}_3(q)$, $\mathrm{PSL}^-_3(q) = \mathrm{PSU}_3(q)$ and $e = (3, q - \varepsilon 1)$). We let R_ε be representatives of these classes.

We observe that for every $(\delta,\delta')\in\{+,-\}^2$, γ exchanges the classes of $T^1_{\delta,\delta'}$ and $T^2_{\delta,\delta'}$. This can be seen as follows. In [Gut72, Table 1], the action of γ on Φ is computed (see also (2.15) of the same paper). In particular we have

$$\begin{aligned}(w^1_{+,+})^\gamma &= w_{3+4}w_{+---}w_2w_{3-4}\\ (w^1_{+,-})^\gamma &= w_{3-4}w_4w_{+---}\\ (w^1_{-,+})^\gamma &= w_{1+2}w_4w_{+---}\\ (w^1_{-,-})^\gamma &= w_4w_{+---}\end{aligned}$$

At this point one computes that in each case $(w^1_{\delta,\delta'})^\gamma$ is W-conjugate to $w^2_{\delta,\delta'}$. This is immediate if $(\delta,\delta')=(-,-)$. In general, it is sufficient to prove that $(w^1_{\delta,\delta'})^\gamma$ is not conjugate to $w^1_{\delta,\delta'}$. This can be done for instance by showing that $(w^1_{\delta,\delta'})^\gamma$ and $w^1_{\delta,\delta'}$ have different root lengths inside the $(-\delta'1)$-eigenspace relative to the action on $\mathbf{R}^4$. It follows that $(T^1_{\delta,\delta'})^\gamma$ is $F_4(q)$-conjugate to $T^2_{\delta,\delta'}$.

Now note that for every (δ,δ'), each subgroup $B_4(i)$, ${}^3D_4(i)$, $D_4(i)$, $P(i)$ and $P'(i)$ contains members of at most one class of tori of type $T_{\delta,\delta'}$. By the same argument as in item (a), we may choose notation such that for every (δ,δ'), $T^1_{\delta,\delta'}$ belongs to a conjugate of $B_4(1)$, ${}^3D_4(1)$ and, possibly, $D_4(1)$ (but not $D_4(2)$). We now want to show that the maximal tori $T^1_{\delta,\delta'}$ can possibly belong only to conjugates of $P(1)$ and $P'(1)$, but not to conjugates of $P(2)$ and $P'(2)$. If we prove this, it will automatically follow that the tori $T^2_{\delta,\delta'}$ can belong only to type 2 subgroups. Note that by the previous considerations, if $T^1_{\delta,\delta'}$ belongs to a conjugate of $P(1)$ (resp. $P'(1)$), then it does not belong to a conjugate of $P(2)$ (resp. $P'(2)$).

By order considerations, $P(1)$ and $P(2)$ can contain the tori $T^i_{-,-}$ and $T^i_{+,-}$ for $i=1,2$. Moreover $P'(1)$ and $P'(2)$ can contain the tori $T^i_{-,+}$ and $T^i_{-,-}$ for $i=1,2$. No other embedding of the tori $T^i_{\delta,\delta'}$ in parabolic subgroups occurs. We see that $w^1_{-,-}=w^{(12)}$ and $w^1_{+,-}=w^{(15)}$ belong to the Weyl subgroup of type B_3 corresponding to removing $+--$ from the base Σ. Therefore we deduce that $T^1_{-,-}$ and $T^1_{+,-}$ belong to Levi complements of $F_4(q)$-conjugates of $P(1)$. Therefore the case P is done. We move to case P', which is similar. We have $w^1_{-,+}=w^{(14)}$. We see that both $w^{(12)}$ and $w^{(14)}$ belong to Weyl subgroups corresponding to a subset of type $A_2\times\tilde{A}_1$. Indeed $w^{(12)}$ lies in the natural one corresponding to removing 4 from the base Σ; and $w^{(14)}$ lies in the subset in which a base of A_2 is $\{2-3,3-4\}$ and a base of $\tilde{A}_1$ is $\{1\}$. Therefore, we obtain that $T^1_{-,-}$ and $T^1_{-,+}$ belong to Levi complements of $F_4(q)$-conjugates of $P'(1)$.

With all the information we have gathered, it is not difficult to deduce the proof of Theorem 7.3.10. With this aim, we choose a generator $x^i_{+,-}$ of the cyclic torus $T^i_{+,-}$, $i=1,2$, and a generator $x^i_{-,+}$ of the cyclic torus $T^i_{-,+}$, $i=1,2$. Moreover we choose elements x_1 and x_2 as in item (a); in particular x_1 belongs only to ${}^3D_4(1)$ and ${}^3D_4(2)$, and x_2 belongs only to $B_4(1)$ and $B_4(2)$. Our set of elements is therefore

$$A_b=\{x_1,x_2,x^1_{+,-},x^2_{+,-},x^1_{-,+},x^2_{-,+}\}.$$

We now want to show that there is not a maximal torus of $F_4(q)$ belonging to an overgroup of all these elements. By order considerations, if a torus belongs to $\mathcal{T}(x_1)$ and $\mathcal{T}(x_2)$, then it must be one of the eight tori $T^i_{\delta,\delta'}$. Assume $i = 1$: the argument is entirely symmetric and the case $i = 2$ proved in the same way. By our choice of notation, $x^2_{+,-}$ and $x^2_{-,+}$ belong to type 2 subgroups. However $T^1_{\delta,\delta'}$ does not belong to any of these. The only other ovegroups of maximal rank of $x^2_{+,-}$ and $x^2_{-,+}$ are, respectively, R_- and R_+. However, our torus $T^1_{\delta,\delta'}$ belongs to exactly one of these (depending on the value of δ). Therefore we have shown that the overgroups of our six elements cannot contain a common maximal torus, and the proof is concluded.

7.5 Classical groups of bounded rank

As promised at the beginning of Section 7.3, we immediately deal with $G \cong \mathrm{PSL}_2(q)$.

Proof of Theorem 7.3.1. Let $G = \mathrm{PSL}_2(q)$. The subgroup structure of G is well known, cf. [Suz82, Chapter 3.6]. Let $d = (2, q-1)$. Let $S_\pm$ be the set of elements of G with order strictly larger than 5 and dividing $(q \pm 1)/d$, and let $S = S_+ \cup S_-$. We have $|S_\pm|/|G| = 1/2 + O(1/q)$. Let F be the set of elements of G lying in subfield subgroups; we have $|F|/|G| = O(1/q^{1/2})$ (we observed this in Remark 7.3.5).

Let $A_b = \{x_1, x_2\}$, where x_1 has order $(q-1)/d$ and x_2 has order $(q+1)/d$. Every element of S_+ (resp. S_-) invariably generates with x_1 (resp. x_2). This proves Theorem 7.1.4. Every element of $S_+ \setminus F$ (resp. $S_- \setminus F$) invariably generates with every element of S_- (resp. S_+). This proves Theorem 7.1.2. Now consider $B = (S_+ \times S_-) \cup (S_- \times S_+) \subseteq G^2$; we have $|B|/|G|^2 = 1/2 + O(1/q)$. Moreover $B \setminus F^2$ consists of invariable generating pairs, which proves Theorem 7.1.3. (If q is a square, the elements of $\mathrm{PGL}_2(q^{1/2})$ belong all to S_-. Therefore, in the above proof we could worry only about elements belonging to $\mathrm{PSL}_2(q^{1/r})$ with r odd, and in Theorem 7.1.2 we could get an error term of type $O(1/q^{2/3})$; but we do not insist on this.) □

In the remainder of the section we prove Theorem 7.3.10 for the other classical groups (hence Theorem 7.1.4 for classical groups of bounded rank). We first make our choice for the type of X. Let X be one of the algebraic groups $\mathrm{SL}_n(k), \mathrm{Sp}_n(k)$, $\mathrm{SO}_n(k)$. We require that if $X = \mathrm{SO}_n(k)$ and $p = 2$ then n is even.

Denote by $V = k^n$ the natural module of X. Here $\mathrm{Sp}_n(k)$ is the group of isometries of a nondegenerate bilinear alternating form on V (n is even), while $\mathrm{SO}_n(k)$ is the connected component of the isometry group $\mathrm{GO}_n(k)$ of a quadratic form on V, with associated nondegenerate bilinear form. We have $|\mathrm{GO}_n(k) : \mathrm{SO}_n(k)| = 2$. These groups are well defined up to conjugation in $\mathrm{GL}_n(k)$, since all such forms are equivalent (for all this, see for instance [MT11, Section 1.2 and Definition 1.15]).

Let $\sigma : X \to X$ be a Steinberg morphism as in [LS98, p. 434], such that X_σ is one of the following finite groups:

$$X_\sigma = \mathrm{SL}_n(q), \mathrm{SU}_n(q), \mathrm{Sp}_n(q), \mathrm{SO}_n^\pm(q)\,(q\,\text{odd}), \Omega_n^\pm(q)\,(q\,\text{even}).$$

Specifically, $\sigma = \bar{\sigma}\tau$, where $\bar{\sigma}$ is a Frobenius morphism corresponding to the field automorphism $\alpha \mapsto \alpha^q$ of k (q power of p), and $\tau = 1$, or $X = \mathrm{SO}_n(k)$ and τ is conjugation by a reflection in a nonsingular vector, or $X = \mathrm{SL}_n(k)$ and τ is the inverse-transpose map (all this with respect to certain fixed bases). We specify that we are also contemplating the case $\mathrm{SO}_n(q)$ with n odd (in this case by our choice q is odd).

Except for $\mathrm{SO}_n^\pm(q)$, which has a derived subgroup $\Omega_n^\pm(q)$ of index 2, and except for other finitely many cases, the group X_σ is perfect. See [KL90b, Chapter 2] for the definition of $\Omega_n^\pm(q)$.

7.5.1 Subgroups of maximal rank

We need to understand the subgroups of maximal rank in X_σ. The proof of the following theorem is essentially taken from [LS98], with additional claims from [MT11, Section 13]. We prefer to sketch a proof since the result does not rely on the most difficult parts of [LS98].

Theorem 7.5.1. *Let X and σ be as above. Let M be a σ-stable closed subgroup of X of maximal rank. Then, M is contained in a σ-stable subgroup of the following types.*

(1) Stabilizer in X of a nonzero proper subspace of V. If $X = \mathrm{Sp}_n(k)$ or $\mathrm{SO}_n(k)$, the space is totally singular or nondegenerate. If it is nondegenerate, it can be chosen of even dimension.

(2) In case $X = \mathrm{SL}_n(k)$, stabilizer in X of a pair of proper subspaces U and W such that $\dim U + \dim W = \dim V$, $\dim U \neq \dim W$ and either $U \leqslant W$ or $U \cap W = 0$.

(3) Stabilizer in X of a decomposition $V = V_1 \oplus \cdots \oplus V_t$ with $t \geqslant 2$. If $X = \mathrm{Sp}_n(k)$ or $\mathrm{SO}_n(k)$, the spaces V_i are isometric, and either pairwise orthogonal and nondegenerate of even dimension, or $t = 2$ and the spaces are totally singular.

(4) $p = 2$, $X = \mathrm{Sp}_n(k)$ and $M \leqslant \mathrm{N}_X(\mathrm{SO}_n(k)) = \mathrm{GO}_n(k)$.

Proof. We assume for the first part of the proof that we are not in case $X = \mathrm{SL}_n(k)$ with σ involving the inverse-transpose map. In particular σ can be regarded as a semilinear map of V.

Assume first M fixes a proper nonzero σ-stable subspace W, and choose it to be of minimal dimension. If there is a form, then M fixes also $W^\perp$, which is σ-stable (cf. [LS98, Proposition 2.5]). Then M fixes $W \cap W^\perp$, which is σ-stable, hence by minimality W is either nondegenerate or totally isotropic. If the space

is nondegenerate of odd dimension ℓ, the only possibility is $X = \mathrm{SO}_n(k)$ and $M^\circ \leqslant \mathrm{SO}_\ell(k) \times \mathrm{SO}_{n-\ell}(k)$. If n is odd, then M stabilizes also a nondegenerate space of even dimension $n-\ell$. If n is even, instead, the stabilizer of W has rank $n/2-1$, hence it is not of maximal rank. Assume now W is totally isotropic. Then M fixes also the set of singular vectors of W, which is σ-stable, hence by minimality either W is totally singular, or $X = \mathrm{SO}_n(k)$, $p = 2$, n is even and W is a nonsingular 1-space. In the latter case, however, the stabilizer is isomorphic to $\mathrm{Sp}_{n-2}(k)$, which is not of maximal rank. In particular, if M fixes a proper nonzero σ-stable subspace, we are in case (1) of the statement.

Then we assume that M does not fix any proper nonzero σ-stable subspace of V. Let $H := M^\circ$ be the connected component of M. The proof of [LS98, Lemma 3.2] shows that either both M and H act homogeneously, or we are in case (3) of the statement. The parity requirement in (3) comes from the following reason: if $X = \mathrm{SO}_n(k)$ and the decomposition $V = V_1 \perp \cdots \perp V_t$ is isometric, then the connected component of the stabilizer of the decomposition is $\mathrm{SO}(V_1) \times \cdots \times \mathrm{SO}(V_t)$, which has maximal rank only if V_i has even dimension.

Therefore assume M and H acts homogeneously. By assumption M, hence H, contains a maximal torus S. Using the explicit description of maximal tori in classical groups, we note that as a kS-module V is the sum of 1-dimensional pairwise nonisomorphic modules. In particular, it follows that both M and H act irreducibly.

Since H is connected and it acts faithfully and irreducibly on V, it follows that H is reductive (cf. [MT11, Proposition 15.1]). Then $H = [H,H]\mathrm{Z}(H)^\circ$. It follows by Schur's lemma that $\mathrm{Z}(H) \leqslant \mathrm{Z}(X)$, which is a finite group, hence $\mathrm{Z}(H)^\circ \leqslant \mathrm{Z}(X)^\circ = 1$. In particular, $H = [H,H]$ is semisimple. In [MT11, Chapter 13], H is called a subsystem subgroup of X. One checks easily that the examples in [MT11, Theorem 13.12] give (well recognizable) reducible subgroups. In [MT11, Theorem 13.15], item (1) corresponds to item (4) in this theorem, and item (2) does not arise by assumption (if $X = \mathrm{SO}_n(k)$ and $p = 2$ then n is even). By [MT11, Theorem 13.14], there are no other possibilities for H. This concludes the proof, except $X = \mathrm{SL}_n(k)$ and σ involves the inverse-transpose map (the previous proof works also if σ is the identity).

Let us consider the remaining case. As in the proof of [LS98, Lemma 3.7], we view $\mathrm{SL}_n(k)$ as a subgroup of $\mathrm{SO}_{2n}(k)$: we may decompose the orthogonal module as $E \oplus F$, in such a way that the embedding of $\mathrm{SL}_n(k)$ is given by $g \mapsto \mathrm{diag}(g, g^{-T})$. It follows from the proof of [LS98, Lemma 3.7] that either we are in cases (1), (2) or (3) of the statement, or M and M° act homogeneously on E. Then we may proceed exactly as in the first part of the proof (once M° was shown to be homogeneous, the morphism σ was not used anymore). □

We can now descend to finite groups. Aschbacher [Asc84] classified the maximal subgroups of the finite classical groups, dividing them into nine classes, denoted by $\mathcal{C}_1, \ldots, \mathcal{C}_8, \mathcal{S}$. We refer the reader to [KL90b] for the detailed description of the first eight classes (although in the present chapter this does not make any difference, we remark that classes in [Asc84] and classes in [KL90b] differ slightly; we take [KL90b] as a reference).

We just recall that maximal subgroups from class $\mathcal{C}_1$ are subspace stabilizers; maximal subgroups from class $\mathcal{C}_2$ stabilize suitable direct sum decompositions of the natural module, so they are subgroups of $\mathrm{GL}_\ell(q) \wr S_{n/\ell} < \mathrm{GL}_n(q)$ for some $\ell < n$; and maximal subgroups from class $\mathcal{C}_3$ preserve an extension field structure on the natural module, so they are subgroups of $\mathrm{GL}_{n/b}(q^b) \rtimes \mathrm{Gal}(\mathbf{F}_{q^b}/\mathbf{F}_q) < \mathrm{GL}_n(q)$ for some prime b. In the following statement, we set

$$\mathrm{Cl}_n(q) = \mathrm{GL}_n(q), \mathrm{GU}_n(q), \mathrm{Sp}_n(q), \mathrm{SO}_n^\pm(q)\,(q\,\mathrm{odd}), \Omega_n^\pm(q)\,(q\,\mathrm{even})$$

in the various cases (this is unusual notation; however it will not be used elsewhere and it should not cause any confusion). When we write in brackets "class $\mathcal{C}_i$", we mean that the subgroup M_σ under consideration is contained in a maximal subgroup of X_σ of class $\mathcal{C}_i$.

Theorem 7.5.2. *Let X, σ and M be as in Theorem 7.5.1. Then, M_σ is contained in a maximal subgroup of X_σ from classes $\mathcal{C}_1$, $\mathcal{C}_2$ or $\mathcal{C}_3$, except $p = 2$, $X_\sigma = \mathrm{Sp}_n(q)$ and $M_\sigma \leqslant \mathrm{SO}_n^\pm(q) = N_{X_\sigma}(\Omega_n^\pm(q))$. Assume now M is contained in a subgroup as in Theorem 7.5.1(3).*

(1) If $X_\sigma = \mathrm{SL}_n(q)$ or the decomposition is isometric, then either $M_\sigma^\circ \leqslant \mathrm{Cl}_\ell(q)^t \cap X_\sigma$ (class $\mathcal{C}_2$), or $M_\sigma^\circ \leqslant \mathrm{Cl}_{n/b}(q^b) \cap X_\sigma$ (class $\mathcal{C}_3$), or $X_\sigma = \mathrm{SU}_n(q)$ with n even and $M_\sigma^\circ \leqslant \mathrm{GL}_{n/2}(q^2) \cap X_\sigma$ (class $\mathcal{C}_2$). In case X_σ is orthogonal, ℓ and n/b must be even.

(2) If the decomposition is totally singular, then $X = \mathrm{Sp}_n(k)$ or $\mathrm{SO}_n(k)$, n is even, and either $M_\sigma^\circ \leqslant \mathrm{GL}_{n/2}(q)$ (class $\mathcal{C}_2$), or $M_\sigma^\circ \leqslant \mathrm{GU}_{n/2}(q)$ (class $\mathcal{C}_3$).

Proof. Note that if M stabilizes $V = V_1 \oplus \cdots \oplus V_t$ as in Theorem 7.5.1(3), then M° fixes each V_i (indeed the subgroup stabilizing each V_i is a closed subgroup of finite index of M, hence contains the connected component). What is more, if $X = \mathrm{SO}_n(k)$ then $M^\circ \leqslant \mathrm{SO}(V_1) \times \cdots \times \mathrm{SO}(V_t)$. Keeping in mind this observation, the proof of the theorem follows from the arguments in [LS98, Section 4], together with Theorem 7.5.1. □

We are ready to define the set A_b from Theorem 7.3.10 for classical groups.

7.5.2 Definition of the set A_b

Assume our finite classical group has natural module V of dimension $n \geqslant 2$, defined over the field with q^u elements, with $u = 2$ in case of unitary groups, and $u = 1$ otherwise. In view of Theorem 7.3.1, we can also assume $n \geqslant 3$ (although logically this is not relevant). In light of various isomorphisms between groups of small rank, we make the requirement $n \geqslant 3$ for unitary groups, $n \geqslant 4$ for symplectic groups and $n \geqslant 7$ for orthogonal groups (see [KL90b, Section 2.9]). Recall also that for orthogonal groups, if n is odd we have q odd (this is justified by the isomorphism $\mathrm{GO}_n(q) \cong \mathrm{Sp}_{n-1}(q)$ when q is even).

We define A_b as the set of elements appearing in Table 7.4. However the notation in Table 7.4 is ambiguous, and we need to explain it. In order to

do this, we recall that the conjugacy classes of maximal tori in finite classical groups have an interpretation in terms of (signed) partitions; see for instance [FG17, Section 5]. We quickly recall some facts.

In $\mathrm{SL}_n(q)$, a maximal torus T_w corresponding to a partition $w = (a_1, \dots, a_t)$ of n fixes a decomposition $V = V_1 \oplus \cdots \oplus V_t$, acting irreducibly on the a_i-th dimensional space V_i for every i. In $\mathrm{SU}_n(q)$, a maximal torus T_w corresponding to a partition $w = (a_1, \dots, a_t)$ of n fixes a decomposition $V = V_1 \perp \cdots \perp V_t$, where V_i is nondegenerate and of dimension a_i. If a_i is odd then T_w acts irreducibly on V_i; if a_i is even then T_w fixes $V_i = A_i \oplus B_i$, where A_i and B_i are totally singular (of dimension $a_i/2$), and T_w acts irreducibly on both. In $\mathrm{Sp}_{2m}(q)$, a maximal torus T_w corresponding to a signed partition $w = (a_1^{\varepsilon_1}, \dots, a_t^{\varepsilon_t})$ of m, with $\varepsilon_i \in \{+, -\}$, fixes a decomposition $V = V_1 \perp \cdots \perp V_t$, where V_i is nondegenerate and of dimension $2a_i$. If $\varepsilon_i = -$ then T_w acts irreducibly on V_i, while if $\varepsilon_i = +$ it acts irreducibly on two complementary totally singular subspaces. In orthogonal groups of even dimension $2m$, the same holds; just recall that in orthogonal groups of plus (resp. minus) type, the product of the signs of the cycles of w must be $+$ (resp. $-$). In orthogonal groups of odd dimension $2m+1$, a maximal torus T_w corresponding to a signed partition $w = (a_1^{\varepsilon_1}, \dots, a_t^{\varepsilon_t})$ of m fixes $V = U \perp W$, centralizing the nondegenerate 1-space U, and acting on W as explained for orthogonal groups in even dimension (the type of U is determined by the sign of W, i.e., by the product of the signs of w).

Now we can explain the notation in Table 7.4. Each element x is semisimple, and the corresponding entry in the table denotes the (conjugacy class of a) maximal torus containing x. In each partition, we have removed the external brackets for aesthetic reasons. Of course, the torus alone does not determine uniquely the element (not even its order).

We require that the element x has the following order on each irreducible fixed space. For convenience we will identify the spaces with the corresponding parts of the partition, as explained above.

We deal separately with linear and unitary groups, as we need to take care of the determinant. In $\mathrm{SL}_n(q)$, x_1 has order $(q^n - 1)/(q - 1)$, and x_2 has order $q^{n-1} - 1$. In $\mathrm{SU}_n(q)$, x_1 has order $(q^n + 1)/(q + 1)$ for n odd, and order $(q^n - 1)/(q + 1)$ for n even; while x_2 has order $q^{n-1} - 1$ for n odd, and order $q^{n-1} + 1$ for n even.

If $x \in \mathrm{Sp}_{2m}(q)$, then x has order $q^a + 1$ on each (a^-), and order $q^a - 1$ on each (a^+). If $x \in \Omega^{\pm}_{2m}(q)$, then x has order $(q^a + 1)/(2, q - 1)$ on each (a^-), and order $(q^a - 1)/(2, q - 1)$ on each (a^+). If $x \in \Omega_{2m+1}(q)$ the same holds; recall that x centralizes also a 1-space.

We further make two requirements. The element x_3 in $\mathrm{Sp}_4(q)$ with q even acts as $(1^-, 1^-)$. Accordingly, we require $x_3 = (g, g^2)$ with g of order $q + 1$. Similarly, the element x_3 in $\Omega^+_8(q)$ acts as $(2^-, 2^-)$, and we have $x_3 = (g, g^2)$ with g of order $(q^2 + 1)/(2, q - 1)$.

With these choices, it is not difficult to check that, if q is sufficiently large ($q \geqslant 10$ say), then every $x \in A_b$ is separable, i.e., it has distinct eigenvalues on the natural module. In fact we need to prove a little more.

In the following lemma, the constant 10 could be slightly decreased. This

lemma will be used to prove Lemma 7.5.4, in which the assumption $q \geqslant 10n^4$ could be much relaxed. However, in Theorem 7.3.10 we require $q \gg r^r$, hence there is no reason to insist for a sharp assumption here (indeed we could prove Lemma 7.5.3 with 10 replaced by an unknown constant, and assume $q \gg r^r$ in Lemma 7.5.4).

Lemma 7.5.3. *Let f be a positive integer. If $q \geqslant 10f$ and $x \in A_b$, then x^f is separable.*

Proof. The proof is straightforward. For $x \in A_b$, we need to prove two things:

(i) if $W \leqslant V$ is irreducible for T (the torus of x), then W is irreducible for x^f, and

(ii) if W_1 and W_2 are distinct irreducible modules for T, then x^f has distinct eigenvalues on W_1 and W_2.

Items (i) and (ii) imply that x^f has the same fixed spaces as the torus T of x, and therefore x^f is separable. The argument is essentially the same in all cases.

For item (i), we show that if $q \geqslant 10f$ then the order of x^f is large enough, so that x^f acts irreducibly on W. For instance, assume the torus T acts irreducibly on a nondegenerate module W of dimension $2a$ (and assume we are not in the unitary case). Then by our choices the order of x on W is larger than $(q^a+1)/2$. Then the order of x^f is larger than $(q^a+1)/2f$, which is strictly larger than $q^{a-1}+1$; in particular x does not fix any proper nondegenerate submodule of W. What is more, if x^f fixes a proper totally singular submodule U, then this has dimension $\ell \leqslant a$. If $\ell < a$, the same argument as above gives a contradiction, and if $\ell = a$, then the order of x^f would divide both $q^a - 1$ and $q^a + 1$, hence it would divide 2, which is false. In unitary groups, the argument is the same; and in case W is totally singular, the argument is similar (we note also that the dimension of any $\langle x^f \rangle$-submodule of W divides the dimension of W).

For item (ii), we first observe that if W_1 and W_2 have different dimension, the claim is obvious. There are cases that can be checked separately, namely x_2 in $\mathrm{SL}_2(q)$, x_2 in $\mathrm{SU}_3(q)$, x_3 in $\mathrm{Sp}_4(q)$ with q even, and x_3 in $\Omega_8^+(q)$. In all other cases, if W_1 and W_2 have equal dimension, say a, then they are totally singular and T acts irreducibly on both, with $W_1 \oplus W_2$ nondegenerate. Let us as assume we are not in the unitary case. If x has eigenvalues $\{\lambda, \lambda^q, \ldots, \lambda^{q^{a-1}}\}$ on W_1, then it has eigenvalues $\{\lambda^{-1}, \lambda^{-q}, \ldots, \lambda^{-q^{a-1}}\}$ on W_2. By our choices, $|\lambda| \geqslant (q^a-1)/2$. Assume by contradiction that the eigenvalues of x^f on W_1 and W_2 coincide. In particular $\lambda^f = \lambda^{-fqi}$ for some $1 \leqslant i \leqslant a-1$. Then $|\lambda|$ divides $f(q^i+1) \leqslant f(q^{a-1}+1)$, which contradicts $q \geqslant 10f$. The unitary case is similar (in this case the eigenvalues of x on W_2 are the q-th powers of the inverses of the eigenvalues on W_1). □

G	x_1	x_2	x_3	x_4
$\mathrm{SL}_n(q)$, $n \geqslant 2$	n	$n-1,1$		
$\mathrm{SU}_n(q)$, $n \geqslant 3$	n	$n-1,1$		
$\Omega_{2m+1}(q)$, $m \geqslant 3$, q odd	m^-	m^+		
$\Omega^-_{2m}(q)$, $m \geqslant 4$	m^-	$(m-1)^-, 1^+$		
$\mathrm{Sp}_{2m}(q)$, $m \geqslant 2$ even, q odd	m^-	$(m-1)^-, 1^+$		
$\mathrm{Sp}_{2m}(q)$, $m \geqslant 3$ odd, q odd	m^-	$(m-1)^-, 1^-$	m^+	
$\mathrm{Sp}_{2m}(q)$, $m \geqslant 2$, q even	m^-	$(m-1)^-, 1^+$	$(m-1)^-, 1^-$	m^+
$\Omega^+_{2m}(q)$, $m \geqslant 5$ odd	m^+	$(m-1)^-, 1^-$		
$\Omega^+_{2m}(q)$, $m \geqslant 4$ even	m^+	$(m-1)^-, 1^-$	$(m-2)^-, 2^-$	$(m-2)^-, 1^-, 1^+$

Table 7.4: $A_b = \{x_1, x_2, x_3, x_4\}$ in Theorem 7.1.4 for classical groups.

7.5.3 Overgroups of the elements of A_b

By Theorem 7.5.2, $\mathcal{M}$ consists of members from Aschbacher's classes $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$, plus the subgroups $\mathrm{SO}^{\pm}_n(q) < \mathrm{Sp}_n(q)$ for even q. Now we show that we can focus on overgroups of $x \in A_b$ from $\mathcal{M}_{\mathrm{con}}$.

Lemma 7.5.4. *Assume $q \geqslant 10n^4$. Assume $x \in A_b$ and let T be its maximal torus in X_σ. Let $M_\sigma \in \mathcal{M}(x)$. Then $x \in T \leqslant M_\sigma^\circ$. In particular, $\mathcal{T}(x)$ consists of the maximal tori contained in some overgroup of x belonging to $\mathcal{M}_{con}$.*

Proof. By Theorems 7.5.1 and 7.5.2, and by the fact that x is separable, we see that the statement is easy, unless M_σ preserves a direct sum decomposition, or is an extension field subgroup.

Note that x fixes at most 4 irreducible spaces on the natural module. In particular, if $x \in M_\sigma \leqslant \mathrm{GL}_k(q^u) \wr S_t$, then x induces a permutation of S_t having at most 4 cycles, and therefore having order at most n^4. We see that $x^f \in M_\sigma^\circ$ for some $f \leqslant n^4$. If $x \in M_\sigma$ and M_σ preserves an extension field subgroup, then the index of M_σ° in M_σ is at most n. In particular, $x^f \in M_\sigma^\circ$ for some $f \leqslant n$.

Since $q \geqslant 10n^4 \geqslant 10f$, by Lemma 7.5.3 x^f is separable, and in particular regular semisimple. Therefore, the maximal torus of x^f, which is T, is contained in M_σ°. □

We have the useful consequence that, for each element of A_b, we just need to determine the overgroups of its maximal torus from classes $\mathcal{C}_1$ and $\mathcal{C}_3$; in class $\mathcal{C}_3$ we only have to consider the linear subgroup $\mathrm{GL}_{n/b}(q^{ub}) < \mathrm{GL}_n(q^u)$. (As usual there is also the subgroup $\mathrm{SO}^{\pm}_{2m}(q) < \mathrm{Sp}_{2m}(q)$ for even q; here the fixed points of the connected component is $\Omega^{\pm}_{2m}(q)$.) The overgroups from class $\mathcal{C}_1$ are easily determined by looking at the action on the space. We now prove a lemma concerning extension field subgroups.

Lemma 7.5.5. *Assume $x \in \mathrm{GL}_{n/b}(q^b) < \mathrm{GL}_n(q)$ for some prime b dividing n. Assume x is separable over $\mathbf{F}_q$. Then each (irreducible) nonzero space fixed by x has dimension divided by b over $\mathbf{F}_q$.*

Proof. Assume U is an irreducible space for x over $\mathbf{F}_{q^b}$ of dimension a. Then U has dimension ab over $\mathbf{F}_q$. Moreover, x acts homogeneously on U over $\mathbf{F}_q$. Since x is separable, it follows that x acts irreducibly on U over $\mathbf{F}_q$. The lemma follows. □

Lemma 7.5.6. *Let b be prime.*

(1) The subgroup $\mathrm{GL}_{n/b}(q^b) \cap \mathrm{SL}_n(q) < \mathrm{SL}_n(q)$ or $\mathrm{GU}_{n/b}(q^b) \cap \mathrm{SU}_n(q) < \mathrm{SU}_n(q)$ contains a representative of all maximal tori corresponding to partitions of n in which each part has length divisible by b (recall that for unitary groups b must be odd, cf. [KL90b, Section 4.3]).

(2) Assume $2m/b$ is even. The subgroup $\mathrm{Sp}_{2m/b}(q^b) < \mathrm{Sp}_{2m}(q)$ or $\mathrm{SO}^{\pm}_{2m/b}(q^b) < \mathrm{SO}^{\pm}_{2m}(q)$ (q odd) or $\Omega^{\pm}_{2m/b}(q^b) < \Omega^{\pm}_{2m}(q)$ (q even) contains a representative of all maximal tori corresponding to signed partitions of m in which all parts have length divisible by b.

(3) The subgroup $\mathrm{GU}_m(q) < \mathrm{Sp}_{2m}(q)$ or $\mathrm{GU}_m(q) < \mathrm{SO}^{\pm}_{2m}(q)$ (q odd) or $\mathrm{GU}_m(q) < \Omega^{\pm}_{2m}(q)$ (q even) contains a representative of all maximal tori corresponding to signed partitions of m in which odd parts have minus sign and even parts have plus sign (recall that for m odd $\mathrm{GU}_m(q)$ embeds in minus type orthogonal groups, while for m even in plus type orthogonal groups).

What is more, in each case we have listed all conjugacy classes of maximal tori contained in the corresponding subgroup.

Proof. Let H be an extension field subgroup as in the statement. Since H is of maximal rank, each maximal torus of the ambient group contained in H is a maximal torus of H, hence it has a description in terms of its Weyl group. It is then not difficult to establish the lemma with the help of Lemma 7.5.5 (see [KL90b, Section 4.3] for a detailed description of the various embeddings). □

Note that in the previous lemma we have not included extension field subgroups of type $\mathrm{O}_{n/b}(q^b)$ with n/b odd. Indeed, we already observed in Theorem 7.5.2 that these are not of maximal rank.

The last case to consider is $M_\sigma = \mathrm{SO}^{\pm}_{2m}(q) < \mathrm{Sp}_{2m}(q)$ with q even (note that $M_\sigma^\circ = \Omega^{\pm}_{2m}(q)$). Recall the isomorphism $\mathrm{Sp}_{2m}(q) \cong \mathrm{SO}_{2m+1}(q)$. In the following lemma, we allow also the case of q odd, as the argument is the same.

Lemma 7.5.7. *Let $\varepsilon \in \{+,-\}$. The subgroup $\mathrm{SO}^{\varepsilon}_{2m}(q) < \mathrm{SO}_{2m+1}(q)$ contains a representative of all conjugacy classes of maximal tori corresponding to signed partitions in which the product of the signs is $\varepsilon 1$, and contains no other maximal torus.*

Proof. See [KL90b, Proposition 2.5.11]. □

7.5.4 Proof of Theorem 7.1.4

Proving Theorem 7.3.10 (hence Theorem 7.1.4 for classical groups of bounded rank) is now just a matter of checking. Indeed, thanks to Lemmas 7.5.4, 7.5.6 and 7.5.7, we know all the members of $\mathcal{T}(x)$ for every element $x \in A_b$ (the case of class $\mathcal{C}_1$ is easily understood, since x is separable). Let us prove the statement for unitary groups, for symplectic groups, and for $\Omega^+_{2m}(q)$ with m even.

For $x \in X_\sigma$, we denote by $\mathcal{M}_{\mathrm{con}}(x)$ the members of $\mathcal{M}_{\mathrm{con}}$ containing x (this notation is used only here). Moreover, we denote by $N^\pm_\ell$ the (fixed points of the connected component of) the stabilizer of a nondegenerate subspace of sign $\pm$ and of dimension ℓ, and by P_ℓ the stabilizer of a totally singular ℓ-space. Finally, in the following discussion we identify each torus with the corresponding (signed) partition; hence a partition can be contained in a subgroup.

(i) $\mathrm{SU}_n(q)$. If n is odd, $\mathcal{M}_{\mathrm{con}}(x_1)$ consists only of (unitary) extension field subgroups of type $\mathrm{GU}_{n/b}(q^b)$ (b odd). On the other hand $\mathcal{M}_{\mathrm{con}}(x_2)$ consists of $P_{(n-1)/2}$ and N_1. Therefore $\mathcal{T}(x_2)$ contains only the partitions with a 1-cycle. By Lemma 7.5.6, none of these belongs to $\mathcal{T}(x_1)$, hence $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$. Assume now n is even. Again $\mathcal{T}(x_2)$ contains the partitions with a 1-cycle. Moreover $\mathcal{M}_{\mathrm{con}}(x_1)$ consists of unitary extension field subgroups, and $P_{n/2}$. Then $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$ (note that $P_{n/2}$ contains the partitions with all even parts).

(ii) $\mathrm{Sp}_{2m}(q)$. If m is even and q is odd, then $\mathcal{M}_{\mathrm{con}}(x_1)$ consists of every $\mathrm{Sp}_{2m/b}(q^b)$, but not of $\mathrm{GU}_m(q)$ (by Lemma 7.5.6). On the other hand $\mathcal{M}_{\mathrm{con}}(x_2)$ consists only of P_1 and N_2. Hence $\mathcal{T}(x_2)$ contains the signed partitions with a 1-cycle. By Lemma 7.5.6, none of these partitions belongs to $\mathcal{T}(x_1)$, hence $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) = \varnothing$.

If m is odd and q is odd, the difference is that $\mathcal{M}_{\mathrm{con}}(x_1)$ contains $\mathrm{GU}_m(q)$, and as before every $\mathrm{Sp}_{2m/b}(q^b)$ (here b is odd). On the other hand $\mathcal{M}_{\mathrm{con}}(x_3)$ contains P_m, every $\mathrm{Sp}_{2m/b}(q^b)$, and does not contain $\mathrm{GU}_m(q)$. Moreover $\mathcal{M}_{\mathrm{con}}(x_2)$ contains only N_2. Then $\mathcal{T}(x_2)$ contains signed partition with a 1-cycle. But $(1^-, \ldots)$ is not contained in $\mathcal{T}(x_3)$, and $(1^+, \ldots)$ is not contained in $\mathcal{T}(x_1)$ (note that a partition contained in P_m has all positive cycles). Hence $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) \cap \mathcal{T}(x_3) = \varnothing$.

Assume now q is even. Note that $\mathrm{GU}_m(q)$ is contained in $\Omega^+_{2m}(q)$ or in $\Omega^-_{2m}(q)$ according to whether m is even or odd (cf. [KL90b, Section 4.3]). Moreover, by Lemma 7.5.7 we deduce that a partition contained in P_m is contained in $\Omega^+_{2m}(q)$. As a consequence we can ignore $\mathrm{GU}_m(q)$ and P_m. The result follows with arguments as above. Indeed, $\mathcal{T}(x_2) \cap \mathcal{T}(x_3)$ consists of the signed partitions w with a 1-cycle. If the product of the signs of w is 1 (resp. -1), then w does not belong to $\mathcal{T}(x_1)$ (resp. $\mathcal{T}(x_4)$). Therefore $\mathcal{T}(x_1) \cap \mathcal{T}(x_2) \cap \mathcal{T}(x_3) \cap \mathcal{T}(x_4) = \varnothing$.

(iii) $\Omega^+_{2m}(q)$ with m even. We see that $\mathcal{T}(x_4)$ contains $(1^+, \ldots)$, $(1^-, \ldots)$, $(2^-, \ldots)$. Now $\mathcal{M}_{\mathrm{con}}(x_3)$ consists of N_4^- and subgroups of type $\Omega^+_m(q^2)$. Of the three subpartitions listed above, only $(2^-, \ldots)$ can belong to $\Omega^+_m(q^2)$. Therefore, $\mathcal{T}(x_3) \cap \mathcal{T}(x_4)$ contains $(1^+, 1^-, \ldots)$ and $(2^-, \ldots)$. At this point we note that $\mathcal{M}_{\mathrm{con}}(x_1)$ contains P_m, $\mathrm{GU}_m(q)$ and every $\Omega^+_{2m/b}(q^b)$. Then $\mathcal{T}(x_1) \cap \mathcal{T}(x_3) \cap \mathcal{T}(x_4)$ consists of the partitions contained in $\Omega^+_m(q^2)$ and of type $(2^-, \ldots)$. We

observe that none of these belongs to $\mathcal{T}(x_2)$, and the proof is concluded.

The other cases are dealt with similarly. In fact we have left out precisely the cases which are easiest to check.

7.6 Proof of Theorems 7.1.2 and 7.1.3

This section is devoted to prove Theorems 7.1.2 and 7.1.3, which imply Theorem 7.1.1 for groups of Lie type of bounded rank. By Theorem 7.3.1, we may assume $r \geqslant 2$.

For the group $G = G_2(3^a)$, Theorem 7.1.3 follows from Theorem 7.4.1, which we proved in Subsection 7.4.4 (see also Subsection 7.6.1). We will now prove Theorem 7.1.2.

Let X and σ be chosen as in Sections 7.4 and 7.5. As we did for Theorem 7.1.4, it is convenient for us to prove a slightly different statement. In the following proof, for a subset $Y = \{x_1, \dots, x_t\}$ of $X(q)$, we will say that Y invariably generates at least $X(q)'$ if for every $g_1, \dots, g_t \in X(q)$, $\langle x_1^{g_1}, \dots, x_t^{g_t} \rangle$ is not contained in maximal subgroups not containing $X(q)'$.

Theorem 7.6.1. *Assume $r \geqslant 2$ and $X(q) \ncong G_2(q)$ when $3 \mid q$. Set $\alpha = 1/|W(E_8)|$ if X is exceptional, and $\alpha = 1/4r$ if X is classical. Then $\mathbf{P}^*_{\mathrm{inv}}(X(q), y) \geqslant \alpha + O(r^r/q)$ for a proportion of elements $y \in X(q)$ of the form $1 - O(r^r/q)$.*

Proof. Let A_b be the set of elements of $X(q)'$ of the statement of Theorem 7.3.10. By Theorems 7.3.8 and 7.3.10, we deduce that, for every $y \in \Delta$, y invariably generates at least $X(q)'$ with some $x \in A_b$. By our choice of the set A_b in the various cases, x is regular semisimple: let $S = (T_w)_\sigma \leqslant X(q)$ be its maximal torus, with $w \in W$.

We claim that $\{y, z\}$ invariably generates at least $X(q)'$ for every $z \in \Delta_S$. Assume not; then, by the definition of Δ we must have $\langle y^{g_1}, z^{g_2} \rangle \leqslant M$ for some $g_1, g_2 \in X(q)$ and some $M \in \mathcal{M}_{\mathrm{con}}$. Now z^{g_2} is regular semisimple; in particular its maximal torus in $X(q)$ is contained in M. But this torus is $X(q)$-conjugate to S, which contradicts the fact that $\{y, x\}$ invariably generates at least $X(q)'$.

In order to conclude the proof, we only need to lower bound the proportion of Δ and Δ_S in $X(q)$. We know that $|\Delta|/|X(q)| = 1 - O(r^r/q)$ by Theorem 7.3.4. Moreover $|\Delta_S|/|X(q)| = \mathbf{P}(W, \sigma, w) + O(r^r/q)$ by Theorems 7.3.4 and 7.3.6. Clearly $\mathbf{P}(W, \sigma, w) \geqslant 1/|W|$, which is at least $1/|W(E_8)|$ for exceptional groups. For classical groups, we can check easily that $\mathbf{P}(W, \sigma, w) \geqslant 1/4r$. The bound is attained for $(2^-, 2^-)$ in $W(D_4)$. Another case which gets close to the bound is $((m-1)^\pm, 1^\pm)$ in $W(C_m)$ for $m \geqslant 3$; the corresponding class has proportion $1/4(m-1)$ in $W(C_m)$. The proof is concluded. (Note that we may draw a random element of $W(C_m)$ or $W(D_m)$ in two steps: first draw a random permutation π of S_m, and then assign a sign to each cycle of π at random, with the obvious restriction in $W(D_m)$.) □

We note that, by the same proof, with more care one can improve the value of α in case X is exceptional.

Now we show that Theorem 7.1.2 follows. With notation as in the previous proof, we have that for every $y \in \Delta' := \Delta \cap X(q)'$, y invariably generates $X(q)'$ with every element of $\Delta'_S := \Delta_S \cap X(q)'$. Therefore we only need to lower bound the proportion of Δ' and Δ'_S in $X(q)'$. By our choice of the type of X, the index of $X(q)'$ in $X(q)$ is bounded (it is at most 3), hence $|\Delta'|/|X(q)'| = 1 - O(r^r/q)$. Moreover, with a similar reasoning as in Lemma 7.3.9, and using Theorem 7.3.6, we see that $|\cup_{g \in X(q)} (S \cap X(q)')^g|/|X(q)'| \geqslant \mathbf{P}(W, \sigma, w) + O(1/q)$, from which $|\Delta'_S|/|X(q)'| \geqslant \alpha + O(r^r/q)$.

7.6.1 An elaboration on Theorem 7.1.3

We obtain here a more precise estimate for the probability in Theorem 7.1.3. For simplicity, we take X of simply connected type (so $X_\sigma = X(q)$ is quasisimple). Let $\{T_1, \ldots, T_\ell\}$ be a set of representatives for the $X(q)$-conjugacy classes of maximal tori of $X(q)$. Write $T_i = (T_{w_i})_\sigma$ with $w_i \in W$, as we did in Subsection 7.3.2. Define a relation $\sim$ on $\{1, \ldots, \ell\}$ as follows. If $1 \leqslant i, j \leqslant \ell$, then $i \sim j$ if there are no conjugates of T_i and T_j with a common overgroup in $\mathcal{M}_{\mathrm{con}}$.

Theorem 7.6.2. *Assume $r \geqslant 2$. Let $x_1, x_2 \in X(q)$ be chosen uniformly at random. Then,*

$$\mathbf{P}(\langle x_1, x_2 \rangle_I = X(q)) = \sum_{\substack{(i,j) \\ i \sim j}} \mathbf{P}(W, \sigma, w_i)\mathbf{P}(W, \sigma, w_j) + \frac{d(r)}{q}$$

for some function $d(r)$.

Proof. Note that $\Delta = \cup_{i=1}^{\ell} \Delta_{T_i}$, a disjoint union. By the definition of Δ and by Theorem 7.3.2, whenever $i \sim j$ every element of Δ_{T_i} invariably generates $X(q)$ with every element of Δ_{T_j}. Clearly this is not true if $i \not\sim j$. By Theorems 7.3.4 and 7.3.6, we have $|\Delta_{T_i}|/|X(q)| = \mathbf{P}(W, \sigma, w_i) + O(r^r/q)$. The statement follows. □

Note that we have a nice and rather explicit expression for the main term of $\mathbf{P}(\langle x_1, x_2 \rangle_I = X(q))$, which one should be able to estimate with accuracy (for all exceptional groups it should be possible to compute the exact value). For instance, with notation as in Subsections 7.4.3 and 7.4.4, in case $G = G_2(3^a)$ we have $3 \sim 6$, $4 \sim 5$ and $5 \sim 6$. One deduces easily that $\mathbf{P}(\langle x_1, x_2 \rangle_I = G) = 1/9 + O(1/q)$. A very easy case is $\mathrm{SL}_2(q)$, where the probability is $1/2 + O(1/q)$ (for the error term, see the the proof of Theorem 7.3.1 at the beginning of Section 7.5).

We remark also that this quite an unusual way to address a problem of random generation. Indeed, in these problems one usually proves that there is a small chance to be trapped in a maximal subgroup — and, as a consequence, there is a large probability to generate. Here, on the other hand, we are directly *exhibiting* many pairs of elements which (invariably) generate, which is a sort of opposite approach. In the language of Subsection 7.1.2, we are exhibiting large complete bipartite subgraphs of the graph $\Lambda_e(X(q))$.

7.7 Lower bound to $|A_b|$

We show that there are cases in which we need $|A_b| \geqslant 4$ in Theorem 7.1.4 (note that only in $F_4(2^a)$ we used a set of size at least 5).

Lemma 7.7.1. *Assume q is even and $m \geqslant 2$. Let $G = \mathrm{PSp}_{2m}(q) = \mathrm{Sp}_{2m}(q)$, and let Y be a subset of G of size 3. Then, $\mathbf{P}_{\mathrm{inv}}(G, Y) \leqslant 1 - 1/2^m m! + O(1/q)$.*

Proof. Note that $2^m m! = |W(C_m)|$. Let y_1, y_2, y_3 be elements of G; we claim that $\Omega := \cap \mathcal{T}(y_i) \neq \varnothing$. Assume we prove the claim, and assume $S \in \Omega$. Then, the elements lying in a conjugate of S contribute to $1 - \mathbf{P}_{\mathrm{inv}}(G, Y)$. By Theorem 7.3.6, the proportion of elements lying in a conjugate of S is at least $1/|W(C_m)| + O(1/q)$. Therefore it is sufficient to prove the claim.

If the y_i act all reducibly, the maximal torus corresponding to $w = (1^+, \ldots, 1^+)$ belongs to Ω (since it is contained in every maximal subgroup from class $\mathcal{C}_1$). Then we may assume that y_1 acts irreducibly (and y_2 and y_3 do not). If none of y_2 and y_3 act as (m^+) (i.e., irreducibly on two complementary totally singular subspaces), then the torus corresponding to $w = (1^+, \ldots, 1^+, 1^-)$ lies in Ω: it lies in $\mathcal{T}(y_1)$ because it is contained to $\mathrm{SO}^-_{2m}(q)$; and it lies in $\mathcal{T}(y_2)$ and $\mathcal{T}(y_3)$ since it is contained in every subspace stabilizer except the stabilizer of a totally singular m-space. Hence we may assume y_2 acts as (m^+). Note that both y_1 and y_2 belong to $\mathrm{Sp}_{2m/b}(q^b)$ for every prime divisor b of m. Assume now y_3 lies in $\mathrm{SO}^\varepsilon_{2m}(q)$, with $\varepsilon \in \{+, -\}$ (it is well known that every element belongs to such a subgroup; cf. [Dye79]). Now observe that $\mathrm{Sp}_{2m/b}(q^b)$ and $\mathrm{SO}^\varepsilon_{2m}(q)$ contain a common maximal torus: that corresponding to $w = (m^\varepsilon)$, for instance. This concludes the proof. □

Next, we show that for groups of bounded rank we cannot have $|A_b| = 1$ in Theorem 7.1.4.

Lemma 7.7.2. *Let X be a connected simple linear algebraic group, σ a Steinberg morphism, and $x \in X_\sigma$. Then, x is contained in a subgroup of $X_\sigma = X(q)$ of maximal rank. In particular, $\mathbf{P}^*_{\mathrm{inv}}(X(q), x) \leqslant 1 - 1/|W| + O(1/q)$.*

Proof. Write $x = us$ for the Jordan decomposition into the unipotent part u and the semisimple part s. Every parabolic subgroup of X_σ contains a Sylow p-subgroup of X_σ; hence u belongs to a conjugate of every parabolic subgroup. Moreover $\mathrm{Z}(X_\sigma)$ is contained in every parabolic of X_σ. Therefore if $s \in \mathrm{Z}(X_\sigma)$ we have that x is contained in a parabolic of X_σ. Assume then $s \notin \mathrm{Z}(X_\sigma)$. Then $x \in \mathrm{C}_X(s) < X$, which is σ-stable. By Theorem 7.3.2, s is contained in a σ-stable maximal torus T of X, hence $T \leqslant \mathrm{C}_X(s)$ and $\mathrm{C}_X(s)$ is of maximal rank. The first part of the statement is proved. The last part follows from Theorem 7.3.6 and the fact that, if $S \in \mathcal{T}(x)$, then the elements lying in a conjugate of S contribute to $1 - \mathbf{P}^*_{\mathrm{inv}}(X(q), x)$. □

7.8 Groups of Lie type of large rank

In this section we prove Theorem 7.1.1 and Theorem 7.1.4 for groups of Lie type of large rank. We work with quasisimple groups G, rather that with the simple quotients $G/\mathrm{Z}(G)$ (this makes no difference since $\mathrm{Z}(G)$ is contained in every maximal subgroup of G). Throughout, G will be one of the groups $\mathrm{SL}_n(q)$, $\mathrm{SU}_n(q)$, $\mathrm{Sp}_n(q)$, $\Omega_n^{\pm}(q)$. We will denote by V the natural n-dimensional module for G. We can assume that n is sufficiently large along the proof.

Recall that in the bounded rank case, for classical groups we could focus on Aschbacher's classes $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$, with the exception $\mathrm{SO}_n^{\pm}(q) < \mathrm{Sp}_n(q)$ (Theorem 7.5.2). In the large rank case, the same happens. If G is a finite (quasi)simple classical group, denote by $\mathcal{M}' = \mathcal{M}'(G)$ the set of all maximal subgroups of G of classes $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3$, plus $\mathrm{SO}_n^{\pm}(q) < \mathrm{Sp}_n(q)$ with q even.

Theorem 7.8.1. *[FG12, Theorem 7.7] Let G be a finite (quasi)simple classical group of untwisted Lie rank r defined over the field with q elements. For r sufficiently large, the proportion of elements of G which lie in subgroups not belonging to $\mathcal{M}'$ is $O(q^{-r/3})$.*

7.8.1 General case

For large rank groups, in most cases we do not need much work (thanks to known results). Specifically, assume G is not symplectic in even characteristic, and not orthogonal in odd dimension. In these cases, we are going to prove Theorem 7.1.4 with $|A_\ell| = 1$, which of course implies Theorem 7.1.1.

Let $x \in G$ be the element defined in [GK00, Table II], and set $A_\ell = \{x\}$. By the proof of [GK00, Proposition 4.1] it follows that x is contained in no irreducible maximal subgroup of G. Let now Ω be the set of integers which occur as the dimension of a proper nonzero subspace of V fixed by x. By [GK00, Table II], we see that Ω has very small size (bounded absolutely from above), and Ω contains only integers ℓ such that both ℓ and $n-\ell$ are comparable to n (up to constants). By [FG18, Theorems 2.2, 2.3, 2.4, 2.5],

$$\frac{|\mathscr{M}(x)|}{|G|} = O(n^{-0.005}).$$

We deduce by Lemma 7.1.7 that $\mathbf{P}_{\mathrm{inv}}(G, x) = 1 - O(n^{-0.005})$, which concludes the proof.

Now we need to deal with the remaining cases. We devote one subsection to each. The difference in these cases is that every element belongs to maximal subgroups whose union of conjugates is large.

7.8.2 Orthogonal groups in odd dimension

Here we assume $G = \Omega_{2m+1}(q)$ with q odd. Let R^+ and R^- denote the union of the stabilizers of hyperplanes of plus and minus sign, respectively, and let P_1 denote the union of the stabilizers of singular 1-spaces. It is well known and easy

that every element of G has eigenvalue 1 on the natural module; in particular, $G = R^+ \cup R^- \cup P_1$. We can prove Theorem 7.1.1.

Proof of Theorem 7.1.1. Let $x \in G$ be as in [GK00, Table II]. By the proof of [GK00, Proposition 4.1] it follows that the only maximal overgroup of x is the stabilizer of a nondegenerate hyperplane of minus type. It follows from [FG17, Theorem 9.26] that the proportion of elements of G lying in R^- is bounded away from 1 absolutely (it is at most 0.93 for n sufficiently large). We conclude by Lemma 7.1.7. □

We will see in Lemma 7.8.6 that $\mathbf{P}_{\mathrm{inv}}(G, x)$ remains bounded away from 1 for every $x \in G$. Now we want to prove Theorem 7.1.4. Here we have a strong dichotomy between the cases q fixed and $q \to \infty$.

Recall that, in orthogonal groups, regular semisimple elements might have eigenvalues of multiplicity greater than 1 (i.e., they need not be separable). We prove a simple known lemma which specifies how this can happen.

Lemma 7.8.2. *Assume $g \in G$ is regular semisimple. Then g centralizes a nondegenerate 1-space, and $\dim \mathrm{C}_V(g) = 1$. Moreover, either g fixes no other nondegenerate 1-space, or it acts as -1 on a nondegenerate 2-space, and fixes no other nondegenerate 1-space.*

Proof. We observe first that any semisimple element $g \in G$ centralizes a nondegenerate 1-space. Indeed, V decomposes as a perpendicular direct sum of nondegenerate irreducible $\langle g \rangle$-submodules, which have either dimension 1, or have even dimension (and g has determinant 1 on any space of even dimension). Since V has odd dimension, it follows that g fixes a nondegenerate subspace U of V of dimension 1. In case g has eigenvalue -1 on U, it follows that on the orthogonal complement $U^\perp$ g has both the eigenvalues 1 and -1; so g centralizes a nondegenerate 1-space.

Recall that regular semisimple elements can be characterized as those elements which do not commute in G with any unipotent element. We now prove that if g is regular semisimple, then g does not centralize any nondegenerate space W of dimension 2. Indeed, assume it does. Applying the same argument of the previous paragraph to the orthogonal complement of W, we get that g centralizes on V a nondegenerate space of dimension 3. It follows that g commutes with a subgroup of G isomorphic to $\Omega_3(q)$ (cf. [KL90b, Lemma 4.1.1(ii)]). This is a contradiction since $\Omega_3(q) \cong \mathrm{PSL}_2(q)$ contains elements of order $\mathrm{char}(\mathbf{F}_q)$.

It is now easy to conclude that g cannot centralize any 2-space. Moreover, g cannot have a -1-eigenspace of dimension 3, otherwise again it would commute with a subgroup $\Omega_3(q)$. Putting together all the information, the statement is proved. □

We prove that, if q is large, with high probability an element is separable.

Theorem 7.8.3. *Assume $m \geqslant 3$. The proportion of elements of $G = \Omega_{2m+1}(q)$ which are separable is larger than $1 - 6/q$. These elements fix only one nondegenerate hyperplane.*

Proof. By Lemma 7.8.2, a separable element fixes only one nondegenerate 1-space, hence only one hyperplane. Therefore we only need to prove the first part of the statement.

By [GL01, Theorem 2.3], the proportion of regular semisimple elements in G is at least $1-2/(q-1)-2/(q-1)^2$. Since $2/(q-1)+2/(q-1)^2+1/(q-1) \leqslant 6/q$, by Lemma 7.8.2 we just need to prove that the proportion of elements which act as -1 on a nondegenerate 2-space W is at most $1/(q-1)$ (note that a regular semisimple element cannot have equivalent modules of dimension at least 2). Let E be the set of such elements.

For fixed W, there are at most $|\Omega_{2m-1}(q)|$ choices for the element (W determines whether on the orthogonal complement the element must belong to Ω or to $\mathrm{SO} \setminus \Omega$). Now we have to sum through all possible W's. Since $2(q-1) = |\mathrm{GO}_2^+(q)| < |\mathrm{GO}_2^-(q)| = 2(q+1)$, there are at most

$$\frac{2|\mathrm{GO}_{2m+1}(q)|}{|\mathrm{GO}_2^+(q)| \cdot |\mathrm{GO}_{2m-1}(q)|}$$

choices. In particular, we deduce that

$$|E| \leqslant \frac{2|\Omega_{2m-1}(q)| \cdot |\mathrm{GO}_{2m+1}(q)|}{|\mathrm{GO}_2^+(q)| \cdot |\mathrm{GO}_{2m-1}(q)|} = \frac{2|\Omega_{2m+1}(q)|}{|\mathrm{GO}_2^+(q)|} = \frac{|\Omega_{2m+1}(q)|}{q-1}.$$

The proof is finished. □

Now we can prove the lower bound to $\mathbf{P}_{\mathrm{inv}}(G, A_\ell)$ in Theorem 7.1.4.

Proof of the lower bound to $\mathbf{P}_{\mathrm{inv}}(G, A_\ell)$. Let $x = x_1$ be as in [GK00, Table II], and let x_2 act on the space as $(m \oplus m) \perp 1$ and having order $(q^m-1)/2$. Set $A_\ell = \{x_1, x_2\}$.

We claim that x_2 does not lie in maximal subgroups from classes $\mathcal{C}_2$ and $\mathcal{C}_3$. If this is true, then by Theorem 7.8.1 and [FG18, Theorem 2.5], the proportion of elements of G lying in conjugates of overgroups of both x_1 and x_2 is $|R^+ \cap R^-|/|G| + O(n^{-0.005})$. By Theorem 7.8.3, $|R^+ \cap R^-|/|G| \leqslant 6/q$, therefore $\mathbf{P}_{\mathrm{inv}}(G, A_\ell) \geqslant 1 - 6/q + O(n^{-0.005})$ by Lemma 7.1.7.

Therefore it suffices to prove the claim. Assume first $x_2 \in \mathrm{GL}_{n/b}(q^b) \rtimes C_b$ for some prime b dividing n. If m is large then $(q^m-1)/2b > q^{m/2}-1$. In particular, $x_2^b \in \mathrm{GL}_{n/b}(q^b)$ fixes only m-spaces, and only one 1-space, which is impossible.

Assume now $x_2 \in \mathrm{GL}_k(q) \wr S_t$ with $n = kt$ and $t > 1$. The element x_2 induces a permutation π of S_t having at most 3 cycles; hence the order of π, say ℓ, is at most n^3. Then $x_2^\ell \in \mathrm{GL}_k(q)^t$. Again, for m large x_2^ℓ fixes m-spaces and a 1-space. Provided $n > 3$, this is impossible for an element of $\mathrm{GL}_k(q)^t$, and the proof is finished. □

Now we want to bound $\mathbf{P}_{\mathrm{inv}}(G, G)$ from above. The key fact is that, for q fixed, the proportion of regular semisimple elements acting as -1 on a nondegenerate 2-space is bounded away from zero. We recall an important result.

Theorem 7.8.4. *The $m \to \infty$ proportion of separable elements of $\Omega_{2m+1}(q)$ is at least 0.348, and it is equal to the corresponding limiting proportion in the nontrivial coset of $\Omega_{2m+1}(q)$ in $\mathrm{SO}_{2m+1}(q)$.*

Proof. See [FG17, Theorems 7.19 and 7.24]. □

Theorem 7.8.5. *If m is sufficiently large, the proportion of elements of G which are regular semisimple, and which act as -1 on a nondegenerate 2-space of plus type, is at least $1/6q$. These elements fix hyperplanes of both signs, and fix a singular 1-space.*

Proof. The last statement is clear, since a 2-space of plus type contains a singular 1-space, and contains nondegenerate 1-spaces of square and non-square discriminant. Therefore we only need to prove the first part of the statement. The proof is similar (although opposite in spirit) to Theorem 7.8.3. Let E be the set of regular semisimple elements which act as -1 on a nondegenerate 2-space of plus type W. For fixed W, by Theorem 7.8.4 there at least $|\Omega_{2m-1}(q)|/3$ choices for the element on $W^\perp$. Then we have to sum through all W's. We have

$$|E| \geqslant \frac{|\Omega_{2m-1}(q)|}{3} \frac{|\mathrm{GO}_{2m+1}(q)|}{|\mathrm{GO}_2^+(q)| \cdot |\mathrm{GO}_{2m-1}(q)|} = \frac{|\Omega_{2m+1}(q)|}{3|\mathrm{GO}_2^+(q)|} \geqslant \frac{|\Omega_{2m+1}(q)|}{6q},$$

which concludes the proof. (Conceptually, there is nothing special here in considering a 2-space: the same argument applies to elements acting as -1 on a space of bounded dimension.) □

At this point it is easy to deduce the upper bound to $\mathbf{P}_{\mathrm{inv}}(G, G)$ in Theorem 7.1.4. Indeed, we already observed that $G = R^+ \cup R^- \cup P_1$. By Theorem 7.8.5, we have $|R^+ \cap R^- \cap P_1|/|G| \geqslant 1/6q$ for sufficiently large m, hence $\mathbf{P}_{\mathrm{inv}}(G, G) \leqslant 1 - 1/6q$ by Lemma 7.1.7.

We finally observe that we cannot have $|A_\ell| = 1$ in Theorem 7.1.4 (not even for $q \to \infty$).

Lemma 7.8.6. *R^+, R^- and P_1 have proportion in G bounded away from zero absolutely. In particular, for every $x \in G$, $\mathbf{P}_{\mathrm{inv}}(G, x)$ is bounded away from 1 absolutely.*

Proof. If we prove the first part of the statement, the last part will follow from Lemma 7.1.7 and the fact that $G = R^+ \cup R^- \cup P_1$. Therefore we only need to prove the the first part. For q fixed, we proved a stronger statement in Theorem 7.8.5. Now we deal with large q. By Theorem 7.8.3, the proportion of separable elements in G is $1 - O(1/q)$. If a separable element g fixes a nondegenerate hyperplane W, then the maximal torus of g is contained in the stabilizer of W (indeed in a subgroup $\mathrm{SO}(W)$ of the stabilizer). The same is certainly true for the stabilizer of a singular 1-space, since this is obtained as the fixed points of a connected subgroup of the algebraic group. Therefore, using Theorem 7.3.6, we see that the proportion of elements belonging to $R^\pm$ (resp. P_1) is equal to $O(1/q)$ plus the proportion of elements of the Weyl group $W(B_m)$ with product of sign $\pm$ (resp. with a positive 1-cycle), which is $1/2$ (resp. at least $(1 - 1/e)/2$ for sufficiently large m). □

7.8.3 Symplectic groups in even characteristic

The arguments are much the same as in the previous subsection.

We view $\mathrm{Sp}_{2m}(q) \cong \mathrm{SO}_{2m+1}(q) = G$, the group of isometries of a nonsingular quadratic form Q on a $(2m+1)$-dimensional space V. (Here, by nonsingular we mean that there are no nonzero vectors v of $V^\perp$ such that $Q(v) = 0$.) Under these assumptions, it turns out that $V^\perp$ is a 1-dimensional subspace of V. Through this identification, the subgroups $\mathrm{SO}^\pm_{2m}(q)$ correspond to stabilizers of nondegenerate hyperplanes of V (i.e., complements of $V^\perp$) of plus or minus type. Note that G acts trivially on $V^\perp$, since q is even and Q does not vanish on $V^\perp$.

As in the previous subsection, we denote by R^+ and R^- the union of the stabilizers of hyperplanes of plus and minus sign, respectively. It is well known that $G = R^+ \cup R^-$ (cf. [Dye79]). We will see in Lemma 7.8.10 that, also in this case, $\mathbf{P}_{\mathrm{inv}}(G, x)$ is bounded away from 1 for every $x \in G$. We can prove Theorem 7.1.1.

Proof of Theorem 7.1.1. Let $x \in G$ be as in [GK00, Table II]. The same argument given for the other classical groups in Subsection 7.8.1 applies, except that x stabilizes a unique nondegenerate hyperplane of plus or minus type. Therefore

$$\frac{|\mathscr{M}(x)|}{|G|} = \frac{|R^\pm|}{|G|} + O(n^{-0.005}).$$

Since $|R^\pm|/|G|$ is bounded away from 1 by [FG17, Theorem 9.15] (it is at most 0.86 for n sufficiently large), Theorem 7.1.1 follows. □

We can also prove the lower bound to $\mathbf{P}_{\mathrm{inv}}(G, A_\ell)$ in Theorem 7.1.4.

Proof of the lower bound to $\mathbf{P}_{\mathrm{inv}}(G, A_\ell)$. Let $x = x_1$ be as in Table [GK00, Table II]. In case m is odd (recall $G = \mathrm{Sp}_{2m}(q)$) for convenience we modify x_1 as follows. If $m \equiv 1 \bmod 4$, we choose x_1 acting on the symplectic module as $(m-1)/2 \perp (m+3)/2 \perp (m-1)$; and if $m \equiv 3 \bmod 4$ we choose x_1 acting as $(m+1)/2 \perp (m-3)/2 \perp (m+1)$. We let x_1 have order $q^b + 1$ on each block of dimension $2b$. Similarly to the proof of the lower bound to $\mathbf{P}_{\mathrm{inv}}(G, A_\ell)$ in the previous subsection, we can easily prove that x_1 does not belong to subgroups of classes $\mathcal{C}_2$ and $\mathcal{C}_3$ if m is large. (Subgroups of class $\mathcal{C}_3$ are ruled out since the element has nondegenerate irreducible modules whose dimensions differ by 2; recall Lemma 7.5.6.) In this way, our element x_1 belongs to $\mathrm{SO}^-_{2m}(q)$ in all cases, both for m even and m odd.

Let moreover $x_2 \in G$ act as follows: if m is odd, it acts as $(m-1) \perp (m+1)$; if $m \equiv 0 \bmod 4$, it acts as $(m-2) \perp (m+2)$; if $m \equiv 2 \bmod 4$, it acts as $(m-4) \perp (m+4)$. Assume moreover x_2 has order $q^b + 1$ on each block of dimension $2b$. Except for stabilizers of subspaces, the only maximal overgroup of x_2 is a conjugate of $\mathrm{SO}^+_{2m}(q)$ (see [BH19, Lemma 6.2]; in fact, since we are allowed to consider only classes $\mathcal{C}_2$ and $\mathcal{C}_3$, a simpler argument suffices).

Set now $A_\ell = \{x_1, x_2\}$. By Theorem 7.8.1 and [FG18, Theorem 2.4], the proportion of elements lying in conjugates of overgroups of both x_1 and x_2 is $|R^+ \cap R^-|/|G| + O(n^{-0.005})$.

By [GL01, Theorem 2.3], the proportion of regular semisimple elements in G is at least $1 - 6/q$. A regular semisimple element does not have eigenvalue 1 on the symplectic module (or, in other words, centralizes only $V^\perp$ on the orthogonal module V). It follows that a regular semisimple element g fixes only one nondegenerate hyperplane, namely $[g, V]$. Then $|R^+ \cap R^-|/|G| \leqslant 6/q$, which shows that $\mathbf{P}_{\mathrm{inv}}(G, A_\ell) \geqslant 1 - 6/q + O(n^{-0.005})$. □

Now we prove the upper bound to $\mathbf{P}_{\mathrm{inv}}(G, G)$. We first observe that, if $g \in G$ is semisimple and centralizes a 2-space, then g fixes hyperplanes of both signs (in case q odd we could exploit the discriminant to see this; here we use a different argument).

Lemma 7.8.7. *Assume $g \in G$ is semisimple and $\dim \mathrm{C}_V(g) \geqslant 2$ on the orthogonal module. Then g fixes nondegenerate hyperplanes of both signs.*

Proof. Assume $V^\perp = \langle v \rangle$. Since every element of $\mathbf{F}_q$ is a square, by rescaling we may assume $Q(v) = 1$. Assume now g is semisimple and fixes a nondegenerate hyperplane W; we want to show that g fixes also a hyperplane of opposite sign.

Since $V = W \perp V^\perp$, by assumption there exists $0 \neq e \in W$ such that $eg = e$. Write $W = \langle e \rangle \oplus T$, with T fixed by g. Assume first $Q(e) \neq 0$. Consider $e' := Q(e)^{-1/2} e + v$. Clearly $e'g = e'$ and $Q(e') = 0$. Moreover, g fixes $W' := \langle e' \rangle \oplus T$, which is a complement of $V^\perp$, i.e., a nondegenerate hyperplane. If W' has opposite sign with respect to W, the proof is finished. Hence, replacing W by W' and e by e', we may assume from the beginning that $Q(e) = 0$.

Since g is semisimple, g centralizes a nondegenerate 2-subspace $\langle e, f \rangle$ of W, where $Q(f) = 0$ and $(e, f) = 1$. Write now $W = \langle e, f \rangle \perp U$, with U fixed by g. Pick $\xi \in \mathbf{F}_q$ such that the polynomial $X^2 + X + \xi^2$ is irreducible over $\mathbf{F}_q$. Then set $e' := e + v$, $f' := f + \xi + v$ and $W' := \langle e', f' \rangle \perp U$. A straightforward computation shows that $\langle e', f' \rangle$ is a nondegenerate anisotropic space, i.e., $Q(x) \neq 0$ for every $0 \neq x \in \langle e', f' \rangle$ (cf. [KL90b, p. 26]). It follows now from [KL90b, Propositions 2.5.3 and 2.5.11] that W' has opposite sign with respect to W. This concludes the proof. □

Theorem 7.8.8. *[FG17, Theorem 7.11] The $m \to \infty$ proportion of regular semisimple elements of $\mathrm{Sp}_{2m}(q)$ is at least* 0.283.

Theorem 7.8.9. *If m is sufficiently large, the proportion of elements of G which act (on the symplectic module) as the identity on a nondegenerate 2-space, and which are regular semisimple on the orthogonal complement, is at least $1/4q^3$. These elements fix nondegenerate hyperplanes of both signs.*

Proof. The last part of the statement follows from Lemma 7.8.7. The first part is exactly the same as in Theorem 7.8.5 (one essentially replaces Ω by Sp throughout, and we use $|\mathrm{Sp}_2(q)| \leqslant q^3$). □

At this point we can prove the upper bound to $\mathbf{P}_{\mathrm{inv}}(G,G)$ in Theorem 7.1.4. We already recalled that $G = R^+ \cup R^-$, and by Theorem 7.8.9 we have that $|R^+ \cap R^-|/|G| \geqslant 1/4q^3$ for sufficiently large m. Therefore $\mathbf{P}_{\mathrm{inv}}(G,G) \leqslant 1 - 1/4q^3$ by Lemma 7.1.7.

We conclude by showing that we cannot have $|A_\ell| = 1$ in Theorem 7.1.4.

Lemma 7.8.10. *R^+ and R^- have proportion in G bounded away from zero absolutely. In particular, for every $x \in G$, $\mathbf{P}_{\mathrm{inv}}(G,x)$ is bounded away from 1 absolutely.*

Proof. The last part follows from the first, Lemma 7.1.7 and $G = R^+ \cup R^-$. We now prove the first. For q fixed, we proved a stronger statement in Theorem 7.8.9. Now we deal with large q. By Theorem 7.3.3, the proportion of regular semisimple elements in G is $1 - O(1/q)$. If a regular semisimple element g fixes a nondegenerate hyperplane W, then the maximal torus of g is contained in the stabilizer of W. By Theorem 7.3.6, we deduce that the proportion of elements belonging to $R^\pm$ is equal to $O(1/q)$ plus the proportion of elements of $W(B_m)$ with product of sign $\pm$, which is $1/2$. □

Bibliography

[AG84] M. Aschbacher and R. M. Guralnick. Some applications of the first cohomology group. *Journal of Algebra*, 90(2):446–460, 1984.

[AK14] P. Apisa and B. Klopsch. A generalization of the Burnside basis theorem. *Journal of Algebra*, 400:8–16, 2014.

[Apo76] T. M. Apostol. *Introduction to analytic number theory.* Springer-Verlag, New York-Heidelberg, 1976. Undergraduate Texts in Mathematics.

[Asc84] M. Aschbacher. On the maximal subgroups of the finite classical groups. *Invent. Math.*, 76(3):469–514, 1984.

[BBDW17] A. Blokhuis, A. Brouwer, and B. De Weger. Binomial collisions and near collisions. *Integers*, 17:paper a64, 8, 2017.

[BBE06] A. Ballester-Bolinches and L. M. Ezquerro. *Classes of finite groups*, volume 584. Springer Science & Business Media, 2006.

[BCP97] W. Bosma, J. Cannon, and C. Playoust. The Magma algebra system. I. The user language. *J. Symbolic Comput.*, 24(3-4):235–265, 1997. Computational algebra and number theory (London, 1993).

[BGH20] T. C. Burness, R. Guralnick, and S. Harper. The spread of a finite group. *arXiv preprint arXiv:2006.01421*, 2020.

[BH19] T. C. Burness and S. Harper. On the uniform domination number of a finite simple group. *Trans. Amer. Math. Soc.*, 372(1):545–583, 2019.

[BHRD13] J. N. Bray, D. F. Holt, and C. M. Roney-Dougal. *The maximal subgroups of the low-dimensional finite classical groups.*, volume 407. Cambridge: Cambridge University Press, 2013.

[BL02] D. Bubboloni and M. S. Lucido. Coverings of linear groups. *Communications in Algebra*, 30(5):2143–2159, 2002.

[BLW06] D. Bubboloni, M. S. Lucido, and T. Weigel. Generic 2-coverings of finite groups of Lie type. *Rendiconti del Seminario Matematico della Università di Padova*, 115:209–252, 2006.

[BLW11] D. Bubboloni, M. S. Lucido, and T. Weigel. 2-coverings of classical groups. *arXiv preprint arXiv:1102.0660*, 2011.

[BM13] J. Britnell and A. Maróti. Normal coverings of linear groups. *Algebra & Number Theory*, 7(9):2085–2102, 2013.

[BPS13] D. Bubboloni, C. Praeger, and P. Spiga. Normal coverings and pairwise generation of finite alternating and symmetric groups. *Journal of Algebra*, 390:199–215, 2013.

[BT71] A. Borel and J. Tits. Éléments unipotents et sous-groupes paraboliques de groupes réductifs. I. *Invent. Math.*, 12:95–104, 1971.

[Bub10] D. Bubboloni. Coverings of the symmetric and alternating groups. *arXiv preprint arXiv:1009.3866*, 2010.

[Bur19] T. C. Burness. Simple groups, generation and probabilistic methods. *Proceedings of Groups St Andrews 2017*, pages 200–229, 2019.

[Car78] R. W. Carter. Centralizers of semisimple elements in finite groups of Lie type. *Proc. London Math. Soc. (3)*, 37(3):491–507, 1978.

[Car93] R. W. Carter. *Finite groups of Lie type*. Wiley Classics Library. John Wiley & Sons, Ltd., Chichester, 1993. Conjugacy classes and complex characters, Reprint of the 1985 original, A Wiley-Interscience Publication.

[CC02] P. J. Cameron and P. Cara. Independent generating sets and geometries for symmetric groups. *Journal of Algebra*, 258(2):641–650, 2002.

[CCN+85] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson. *Atlas of finite groups*. Oxford University Press, Eynsham, 1985. Maximal subgroups and ordinary characters for simple groups, With computational assistance from J. G. Thackray.

[CL13a] E. Crestani and A. Lucchini. The generating graph of finite soluble groups. *Israel Journal of Mathematics*, 198(1):63–74, 2013.

[CL13b] E. Crestani and A. Lucchini. The non-isolated vertices in the generating graph of a direct powers of simple groups. *Journal of Algebraic Combinatorics*, 37(2):249–263, 2013.

[Coo78] B. N. Cooperstein. Minimal degree for a permutation representation of a classical group. *Israel Journal of Mathematics*, 30(3):213–235, 1978.

[DF91] D. I. Deriziotis and A. P. Fakiolas. The maximal tori in the finite Chevalley groups of type E_6, E_7 and E_8. *Comm. Algebra*, 19(3):889–903, 1991.

[Dix92] J. D. Dixon. Random sets which invariably generate the symmetric group. *Discrete Mathematics*, 105(1-3):25–39, 1992.

[DL03] E. Detomi and A. Lucchini. Crowns and factorization of the probabilistic zeta function of a finite group. *Journal of Algebra*, 265(2):651–668, 2003.

[DL15] E. Detomi and A. Lucchini. Invariable generation with elements of coprime prime-power orders. *Journal of Algebra*, 423:683–701, 2015.

[DL16] E. Detomi and A. Lucchini. Invariable generation of prosoluble groups. *Israel Journal of Mathematics*, 211(1):481–491, 2016.

[DSC98] P. Diaconis and L. Saloff-Coste. Walks on generating sets of groups. *Inventiones mathematicae*, 134(2):251–299, 1998.

[dW97] B. M. de Weger. Equal binomial coefficients: Some elementary considerations. *J. Number Theory*, 63(2):373–386, 1997.

[Dye79] R. H. Dye. Interrelations of symplectic and orthogonal groups in characteristic two. *Journal of Algebra*, 59(1):202–221, 1979.

[EFG16] S. Eberhard, K. Ford, and B. Green. Permutations fixing a k-set. *Int. Math. Res. Not. IMRN*, (21):6713–6731, 2016.

[EFG17] S. Eberhard, K. Ford, and B. Green. Invariable generation of the symmetric group. *Duke Math. J.*, 166(8):1573–1590, 2017.

[EFK16] S. Eberhard, K. Ford, and D. Koukoulopoulos. Permutations contained in transitive subgroups. *Discrete Anal.*, page Paper No. 12 (34 pages), 2016.

[Eno72] H. Enomoto. The characters of the finite symplectic group $\mathrm{Sp}(4, q)$, $q = 2^f$. *Osaka J. Math.*, 9:75–94, 1972.

[FG03] J. Fulman and R. M. Guralnick. Derangements in simple and primitive groups. In *Groups, combinatorics & geometry (Durham, 2001)*, pages 99–121. World Sci. Publ., River Edge, NJ, 2003.

[FG12] J. Fulman and R. M. Guralnick. Bounds on the number and sizes of conjugacy classes in finite Chevalley groups with applications to derangements. *Trans. Amer. Math. Soc.*, 364(6):3023–3070, 2012.

[FG17] J. Fulman and R. M. Guralnick. Derangements in subspace actions of finite classical groups. *Trans. Amer. Math. Soc.*, 369(4):2521–2572, 2017.

[FG18] J. Fulman and R. M. Guralnick. Derangements in finite classical groups for actions related to extension field and imprimitive subgroups and the solution of the Boston-Shalev conjecture. *Trans. Amer. Math. Soc.*, 370(7):4601–4622, 2018.

[FNP05] J. Fulman, P. M. Neumann, and C. E. Praeger. A generating function approach to the enumeration of matrices in classical groups over finite fields. *Mem. Amer. Math. Soc.*, 176(830):vi+90, 2005.

[Gal70] P. X. Gallagher. The number of conjugacy classes in a finite group. *Mathematische Zeitschrift*, 118(3):175–179, 1970.

[GAP19] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.10.2*, 2019.

[Gar20a] D. Garzoni. The invariably generating graph of the alternating and symmetric groups. *Journal of Group Theory*, 1(ahead-of-print), 2020.

[Gar20b] D. Garzoni. A note on the invariably generating graph of a finite group. *arXiv preprint arXiv:2009.14536*, 2020.

[Gas59] W. Gaschütz. Die Eulersche funktion endlicher aufloÈsbarer gruppen. *Illinois Journal of Mathematics*, 3(4):469–476, 1959.

[Gel15] T. Gelander. Convergence groups are not invariably generated. *International Mathematics Research Notices*, 2015(19):9806–9814, 2015.

[GK00] R. M. Guralnick and W. M. Kantor. Probabilistic generation of finite simple groups. volume 234, pages 743–792. 2000. Special issue in honor of Helmut Wielandt.

[GKS94] R. M. Guralnick, W. M. Kantor, and J. Saxl. The probability of generating a classical group. *Comm. Algebra*, 22(4):1395–1402, 1994.

[GL01] R. M. Guralnick and F. Lübeck. On p-singular elements in Chevalley groups in characteristic p. In *Groups and computation, III (Columbus, OH, 1999)*, volume 8 of *Ohio State Univ. Math. Res. Inst. Publ.*, pages 169–182. de Gruyter, Berlin, 2001.

[GL15] M. Garonzi and A. Lucchini. Covers and normal covers of finite groups. *Journal of Algebra*, 422:148–165, 2015.

[GL20] D. Garzoni and A. Lucchini. Minimal invariable generating sets. *Journal of Pure and Applied Algebra*, 224(1):218–238, 2020.

[GLSS99] R. M. Guralnick, M. W. Liebeck, J. Saxl, and A. Shalev. Random generation of finite simple groups. *J. Algebra*, 219(1):345–355, 1999.

[GLT12] R. M. Guralnick, M. Larsen, and P. H. Tiep. Representation growth in positive characteristic and conjugacy classes of maximal subgroups. *Duke Math. J.*, 161(1):107–137, 2012.

[GM12a] R. M. Guralnick and G. Malle. Products of conjugacy classes and fixed point spaces. *J. Amer. Math. Soc.*, 25(1):77–121, 2012.

[GM12b] R. M. Guralnick and G. Malle. Simple groups admit Beauville structures. *J. Lond. Math. Soc. (2)*, 85(3):694–721, 2012.

[GM13] R. M. Guralnick and A. Maróti. On the non-coprime $k(GV)$-problem. *Journal of Algebra*, 385:80–101, 2013.

[GM15] M. Garonzi and A. Maróti. On the number of conjugacy classes of a permutation group. *Journal of Combinatorial Theory, Series A*, 133:251–260, 2015.

[GM17] T. Gelander and C. Meiri. The congruence subgroup property does not imply invariable generation. *International Mathematics Research Notices*, 2017(15):4625–4638, 2017.

[GM20] D. Garzoni and E. McKemmie. On the probability of generating invariably a finite simple group. *arXiv preprint arXiv:2008.03812*, 2020.

[GMPS15] S. Guest, J. Morris, C. E. Praeger, and P. Spiga. On the maximum orders of elements of finite almost simple groups and primitive permutation groups. *Trans. Am. Math. Soc.*, 367(11):7665–7694, 2015.

[Gor07] D. Gorenstein. *Finite groups*, volume 301. American Mathematical Soc., 2007.

[GT05] R. M. Guralnick and P. H. Tiep. The non-coprime $k(GV)$ problem. *Journal of Algebra*, 293(1):185–242, 2005.

[Gut72] M. M. Guterman. A characterization of the groups $F_4(2^n)$. *J. Algebra*, 20:1–23, 1972.

[Hal36] P. Hall. The Eulerian functions of a group. *The Quarterly Journal of Mathematics*, (1):134–151, 1936.

[Hei06] H. Heineken. On groups all of whose elements have prime power order. In *Mathematical Proceedings of the Royal Irish Academy*, pages 191–198, 2006.

[Hig57] G. Higman. Finite groups in which every element has prime power order. *Journal of the London Mathematical Society*, 1(3):335–342, 1957.

[Jam13] S. Jambor. The minimal generating sets of $\mathrm{PSL}(2, p)$ of size four. *Lond. Math. Soc. J. Comput. Math*, 16:419–423, 2013.

[Kan79] W. M. Kantor. Permutation representations of the finite classical groups of small degree or rank. *J. Algebra*, 60:158–168, 1979.

[Kan07] D. M. Kane. Improved bounds on the number of ways of expressing t as a binomial coefficient. *Integers*, 7(1):paper a53, 7, 2007.

[KL90a] W. M. Kantor and A. Lubotzky. The probability of generating a finite classical group. *Geom. Dedicata*, 36(1):67–87, 1990.

[KL90b] P. B. Kleidman and M. W. Liebeck. *The subgroup structure of the finite classical groups*, volume 129 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1990.

[Kle88] P. B. Kleidman. The maximal subgroups of the Chevalley groups $G_2(q)$ with q odd, the Ree groups ${}^2G_2(q)$, and their automorphism groups. *J. Algebra*, 117(1):30–71, 1988.

[KLS11] W. M. Kantor, A. Lubotzky, and A. Shalev. Invariable generation and the Chebotarev invariant of a finite group. *J. Algebra*, 348:302–314, 2011.

[KLS15] W. M. Kantor, A. Lubotzky, and A. Shalev. Invariable generation of infinite groups. *Journal of Algebra*, 421:296–310, 2015.

[KNN71] L. G. Kovács, J. Neubüser, and B. H. Neumann. On finite groups with 'hidden' primes. *Journal of the Australian Mathematical Society*, 12(3):287–300, 1971.

[KR93] L. G. Kovács and G. R. Robinson. On the number of conjugacy classes of a finite group. *Journal of Algebra*, 160:441–460, 1993.

[KZ12] E. Kowalski and D. Zywina. The Chebotarev invariant of a finite group. *Exp. Math.*, 21(1):38–56, 2012.

[Law99] R. Lawther. The action of $F_4(q)$ on cosets of $B_4(q)$. *J. Algebra*, 212(1):79–118, 1999.

[LM09] A. Lucchini and A. Maróti. Some results and questions related to the generating graph of a finite group. In *Ischia group theory 2008*, pages 183–208, 2009.

[LMS05] M. W. Liebeck, B. M. S. Martin, and A. Shalev. On conjugacy classes of maximal subgroups of finite simple groups, and a related zeta function. *Duke Math. J.*, 128(3):541–557, 2005.

[LP93] T. Luczak and L. Pyber. On random generation of the symmetric group. *Combin. Probab. Comput.*, 2(4):505–512, 1993.

[LP97] M. W. Liebeck and L. Pyber. Upper bounds for the number of conjugacy classes of a finite group. *J. Algebra*, 198(2):538–562, 1997.

[LPS88] M. W. Liebeck, C. E. Praeger, and J. Saxl. On the O'Nan-Scott theorem for finite primitive permutation groups. *J. Aust. Math. Soc., Ser. A*, 44(3):389–396, 1988.

[LS96] M. W. Liebeck and A. Shalev. Maximal subgroups of symmetric groups. *Journal of combinatorial theory, Series A*, 75(2):341–352, 1996.

[LS98] M. W. Liebeck and G. M. Seitz. On the subgroup structure of classical groups. *Invent. Math.*, 134(2):427–453, 1998.

[LS99] M. W. Liebeck and A. Shalev. Simple groups, permutation groups, and probability. *J. Amer. Math. Soc.*, 12(2):497–520, 1999.

[LSS92] M. W. Liebeck, J. Saxl, and G. M. Seitz. Subgroups of maximal rank in finite exceptional groups of Lie type. *Proc. London Math. Soc. (3)*, 65(2):297–325, 1992.

[LT17] A. Lucchini and G. Tracey. An upper bound on the Chebotarev invariant of a finite group. *Israel J. Math.*, 219(1):449–467, 2017.

[Luc13a] A. Lucchini. The largest size of a minimal generating set of a finite group. *Archiv der Mathematik*, 101(1):1–8, 2013.

[Luc13b] A. Lucchini. Minimal generating sets of maximal size in finite monolithic groups. *Archiv der Mathematik*, 101(5):401–410, 2013.

[Luc18] A. Lucchini. The Chebotarev invariant of a finite group: a conjecture of Kowalski and Zywina. *Proc. Amer. Math. Soc.*, 146(11):4549–4562, 2018.

[Mac81] I. G. Macdonald. Numbers of conjugacy classes in some finite classical groups. *Bull. Aust. Math. Soc.*, 23:23–48, 1981.

[Mar05] A. Maróti. Bounding the number of conjugacy classes of a permutation group. *Journal of Group Theory*, 8(3):273–289, 2005.

[McK19] E. McKemmie. Invariable generation of finite classical groups. *arXiv preprint arXiv:1910.03623*, 2019.

[MT11] G. Malle and D. Testerman. *Linear algebraic groups and finite groups of Lie type*, volume 133 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2011.

[Pak01] I. Pak. What do we know about the product replacement algorithm. *Groups and computation, III (Columbus, OH, 1999)*, 8:301–347, 2001.

[Pel13] M. A. Pellegrini. 2-coverings for exceptional and sporadic simple groups. *Archiv der Mathematik*, 101(3):201–206, 2013.

[PPR16] R. Pemantle, Y. Peres, and I. Rivin. Four random permutations conjugated by an adversary generate S_n with high probability. *Random Structures & Algorithms*, 49(3):409–428, 2016.

[Pri09] W. Pribitkin. Simple upper bounds for partition functions. *The Ramanujan Journal*, 18(1):113–119, 2009.

[PS80] C. E. Praeger and J. Saxl. On the orders of primitive permutation groups. *Bull. Lond. Math. Soc.*, 12:303–307, 1980.

[Rob55] H. Robbins. A remark on Stirling's formula. *Am. Math. Mon.*, 62:26–29, 1955.

[Rob83] G. Robin. Estimation de la fonction de Tchebychef θ sur le k-ième nombre premier et grandes valeurs de la fonction $\omega(n)$ nombre de diviseurs premiers de n. *Acta Arithmetica*, 42(4):367–389, 1983.

[Rob12] D. J. S. Robinson. *A Course in the Theory of Groups*, volume 80. Springer Science & Business Media, 2012.

[SB81] H. P. Sankappanavar and S. Burris. *A course in universal algebra*, volume 78. Graduate Texts Math, 1981.

[Sha98] A. Shalev. A theorem on random matrices and some applications. *J. Algebra*, 199(1):124–141, 1998.

[Sin71] D. Singmaster. How often does an integer occur as a binomial coefficient? *Am. Math. Monthly*, 78:385–386, 1971.

[SL96] P. Stevenhagen and H. W. Lenstra. Chebotarëv and his density theorem. *The Mathematical Intelligencer*, 18(2):26–37, 1996.

[Ste62] R. Steinberg. Generators for simple groups. *Canadian Journal of Mathematics*, 14:277–283, 1962.

[Suz62] M. Suzuki. On a class of doubly transitive groups. *Ann. Math. (2)*, 75:105–145, 1962.

[Suz82] M. Suzuki. *Group theory*, volume 247. Springer, 1982.

[Tar75] A. Tarski. An interpolation theorem for irredundant bases of closure structures. *Discrete Mathematics*, 12(2):185–192, 1975.

[Tra19] G. Tracey. Invariable generation of permutation and linear groups. *Journal of Algebra*, 524:250–289, 2019.

[Wal63] G. E. Wall. On the conjugacy classes in the unitary, symplectic and orthogonal groups. *J. Aust. Math. Soc.*, 3:1–62, 1963.

[Wal75] G. E. Wall. Secretive prime-power groups of large rank. *Bulletin of the Australian Mathematical Society*, 12(3):363–369, 1975.

[Wei92] T. S. Weigel. Generation of exceptional groups of Lie-type. *Geom. Dedicata*, 41(1):63–87, 1992.

[Whi00] J. Whiston. Maximal independent generating sets of the symmetric group. *Journal of Algebra*, 232(1):255–268, 2000.

[Wie76] J. Wiegold. Transitive groups with fixed-point free permutations. *Archiv der Mathematik*, 27(1):473–475, 1976.

[Wie77] J. Wiegold. Transitive groups with fixed-point-free permutations II. *Archiv der Mathematik*, 29(1):571–573, 1977.

[Wil09] R. A. Wilson. *The finite simple groups*, volume 251 of *Graduate Texts in Mathematics*. Springer-Verlag London, Ltd., London, 2009.