Growth in finite groups and the Graph Isomorphism Problem

Dissertation

for the award of the degree

"Doctor rerum naturalium" (Dr. rer. nat.)

of the Georg-August-Universität Göttingen

within the doctoral program

"Mathematical Sciences"

of the Georg-August University School of Science (GAUSS)

submitted by

Daniele Dona

from Torino

Göttingen, 2020

Thesis committee

Harald Andrés Helfgott
Mathematisches Institut
Georg-August-Universität Göttingen

Valentin Blomer
Mathematisches Institut
Universität Bonn

Members of the Examination Board

First reviewer:
Harald Andrés Helfgott
Mathematisches Institut
Georg-August-Universität Göttingen

Second reviewer:
Laurent Bartholdi
Mathematisches Institut
Georg-August-Universität Göttingen

Further members of the Examination Board

Valentin Blomer
Mathematisches Institut
Universität Bonn

Axel Munk
Institut für Mathematische Stochastik
Georg-August-Universität Göttingen

Matthew Tointon
Department of Pure Mathematics and Mathematical Statistics
University of Cambridge

Péter Varjú
Department of Pure Mathematics and Mathematical Statistics
University of Cambridge

Date of the oral examination: 17 July 2020.

Contents

Preface

The present thesis embraces two major areas of mathematics, namely group theory (especially growth in finite groups) and graph theory (especially the graph isomorphism problem).

Chapter 1 serves as a somewhat lengthy introduction to both, with §1.1-1.2-1.3-1.4 focusing on growth in groups and §1.5-1.6 on graph isomorphisms.

The next two chapters are mostly graph-theoretic, although they bear many connections to growth in groups as well. Chapter 2 is based on the author's published article [Don19c]; its main results are Theorem 2.1.6, Theorem 2.1.7 and Theorem 2.1.8. Chapter 3 is based on the author's preprint [Don18]; its main result is Theorem 3.2.1.

The two chapters that follow are entirely on the topic of growth in groups. Chapter 4 is based on the author's preprint [Don19b]; its main results are Theorem 4.1.3 and Theorem 4.1.4. Chapter 5 is based on the author's preprint [Don19a]; its main result is Theorem 5.1.1.

Finally the last one, Chapter 6, is firmly rooted into both areas at once: more precisely, graph-theoretic tools intervene in group-theoretic problems; its main result is Theorem 6.3.6, dependent on Conjecture 6.3.4.

Notation. Any and all notations hold unless otherwise stated.

We adopt the big O notation for describing orders of magnitude. If f, g are some real-valued functions, we say $f(x) = O(g(x))$ to mean that there exists a constant $C > 0$ such that $|f(x)| \leq Cg(x)$ for all x in the intersection of the domains of f, g; since we are almost always considering f to have domain $\mathbb{N}$ and codomain inside $\mathbb{R}_{\geq 0}$, in those cases it suffices to say that $f(x) \leq Cg(x)$ for all x large enough. We also use $f(x) = o(g(x))$ to mean that for all $C > 0$ and all x large enough (depending on C) we have $f(x) < Cg(x)$. Finally, $f(x) = \Omega(g(x))$ means that there exists $C > 0$ such that $f(x) \geq Cg(x)$ for all x large enough: in this, we follow Knuth's definition of the symbol instead of Hardy and Littlewood's convention (see Knuth's own letter to the editor [Knu76], where "O" is incidentally revealed to be an omicron!). If we want to emphasize that the constant C in the notations above depends on other parameters (say n, k), we write them as indices to the symbol (say $O_{n,k}(g(x))$). Many other authors, especially of the number theory school, use also Vinogradov's $\ll$ and $\gg$ notation: the author appreciates the fact that essentially the same symbol facing two directions can do the job of both $O(\cdot)$ and $\Omega(\cdot)$, but he also needs to write things like $e^{O(x)}$, for which Vinogradov provides no solution; thus, no $\ll$ will be used.

For the set $\{1, 2, \ldots, n\}$ of natural numbers from 1 to n, we often write $[n]$ for brevity, as is common in the literature regarding permutation groups; to be clear, the author subscribes to the convention that $0 \in \mathbb{N}$, but 0 has usually little space in the context of permutations. For us, p denotes a prime number, and q denotes a prime power.

If X is a finite set, the set of permutations of X is denoted by $\mathrm{Sym}(X)$, and the set of even permutations by $\mathrm{Alt}(X)$; in particular, we write $\mathrm{Sym}(n), \mathrm{Alt}(n)$ for $\mathrm{Sym}([n]), \mathrm{Alt}([n])$ (we will not use the notations S_n, A_n and $\mathrm{Sym}_n, \mathrm{Alt}_n$ that frequently occur elsewhere). As for algebraic groups, say the special linear group, we use the notation $\mathrm{SL}_n(\mathbb{F}_q)$ instead of the equally widespread $\mathrm{SL}(n, \mathbb{F}_q)$ and $\mathrm{SL}(n, q)$.

We use $\log_a x$ to denote the logarithm of x in base a, and $\log^a x$ to denote $(\log x)^a$. Since we will not be using longer expressions than $\log \log x$, there is no need to use either notation for the iterated logarithm, as some authors do (and for good reasons need to do).

About identity elements, notation varies with the context. For general groups, like in §1 and §5, we use e to denote the group's identity; for permutation groups of degree n like in §3 we use Id_n, while for the matrix groups in §4 we use Id without index since we work only with 2×2 matrices. In §2, where several identities coexist, we try and use distinct notations: e for general groups, Id_n for $n \times n$ matrices, $\mathrm{Id}_{\mathfrak{X}}$ for automorphisms on the object $\mathfrak{X}$. In §6, where we work with permutation groups but their identity elements are encountered only in their quality of group identities, we use e.

Finally, since we abundantly use several terms describing orders of magnitude, which may be unfamiliar to the readers, we collect them here:

- $f(x)$ is *quasipolynomial* in x when $f(x) \leq e^{C \log^k x}$ for some absolute constants $C, k > 0$;
- $f(x)$ is *polylogarithmic* in x when $f(x) \leq C \log^k x$ for some absolute constants $C, k > 0$.

Acknowledgements

One can always aspire to be like the Laplacian *intelligence*, embracing the whole universe in her thought, for whom nothing would be uncertain, the future as well as the past always present in her eyes[1]. And one can always delude oneself about having an arm strong enough to pull Leviathan out of the water[2]; and being able to produce a book on which everything is contained as a whole, and from it judge the world[3], as if Protagoras had actually meant a very specific man to be the measure of all things[4]. These solitary God-worthy feats one can surely believe to be within reach, and fancy oneself the fixed terminus of eternal counsel[5], the skies and the stars revolving and revolving under the push of this one finger[6]; and expect that all that is in the world bends slowly and surely to one's will, weighed by the doom of one's thought[7].

But *mit der Dummheit kämpfen Götter selbst vergebens*[8]. And even as I pile quotations in the hope of building a stand high enough for my sense of self-importance, and lose myself in my own mighty station and my stupendous brain[9] until disillusion comes too late to evade my own ruinous shipwreck[10], I must recognize the many other people around me that have contributed to making this achievement of mine possible, people without whom I would not be standing where I am, and to whom I owe in different ways all that I have today. Half-joking ramblings aside, I often hesitate, get discouraged, and doubt that the position that I am in is well-deserved[11]; that is why people who supported me throughout this whole endeavour, and continue to do so even now, deserve all the gratitude that I can muster and should get their due credit.

[1] Pierre-Simon Laplace, *Essai philosophique sur les probabilités*, p. 4 (1840 edition).
[2] Anonymous, *'Iyyōḇ*, 41:1.
[3] Tommaso da Celano (attrib.), *Dies Irae*, 5.
[4] Prōtagóras, as quoted in Plátōn, *Θεαίτητος*, 152a.
[5] Dante Alighieri, *Paradiso*, XXXIII, 3.
[6] Momotsuki Gakuen Ichinen Shigumi, *Rūretto Rūretto*, 6-7.
[7] J. R. R. Tolkien, Christopher Tolkien (ed.), *The Children of Húrin*, §3.
[8] Friedrich Schiller, *Die Jungfrau von Orleans*, III, 6.
[9] W. S. Gilbert, *H.M.S. Pinafore*, II, 4.
[10] A double reference! First, to sir Joseph Porter himself in *H.M.S. Pinafore*, who blinded by his own sense of superiority inadvertently convinces his fiancée to love someone else. Second, to Robert Terwilliger, who victim of his own histrionism wastes his moment of triumph by singing the whole *H.M.S. Pinafore* and literally shipwrecks into his own incarceration.
[11] Impostor syndrome is strong in this one. To add one last reference, I guess one could still take solace in that quote attributed to Tolstoj about the fraction representing the true worth of an individual.

First, from a mathematical perspective. The greatest thanks must indubitably go to H. A. Helfgott, who has accompanied me throughout my doctorate and has been an inexhaustible source of questions, answers, pastry, and delightful conversations; in the course of four years he displayed the charts of many unknown territories all around, and while not neglecting to guide me through the best paths he encouraged me to run free (or some would say run amok) as I pleased. I must also thank V. Blomer, another professor whose guidance I have had the privilege to have, and who never refused a helping hand whenever I came to him.

Together with them, many postdocs and fellow doctorands have come and gone; some are or will hopefully be my coauthors, and all aided me in becoming more of a mathematician, whether they answered my doubts or posed some of their own, proved or disproved my extemporaneous claims, or just engaged in thoughtful discussions among themselves while I was luckily around to listen: J. Bajpai, H. Bradford, V. Finkelshtein, L. Guan, K. Müller, S. Myerson, A. Sedunova, S. Zúñiga Alterman, and certainly others.

Finally, for the environment builds a person as much as single individuals do, I wholeheartedly thank the Georg-August-Universität Göttingen for creating a fertile one. I must also thank the European Research Council, which financed the research of Prof. Helfgott and thus mine as well. And going back in time, I can only acknowledge in my thought the many, oh so many places and institutions and people that eventually contributed to define me and my professional persona: I hope my own contribution to mathematics will not shame any of your efforts.

On a personal level, again many would justly vie for a position in these scarce lines, as there was nigh anyone of the individuals I have met that did not impress their footprint in me, in one way or another; for brevity and forgetfulness, and not for anything else, I will restrict myself to a handful of names.

Among my friends, first I thank Fefe: you have been my closest companion through thick and thin in almost everything that I have experienced, and one could not overestimate the depth of my feelings for you. A twin peak of affection goes to Simo: from philosophy to science, from wars to yellow people, you have enriched me at every corner of my soul. Both of you have given me so much, and we have shared laughter and tears, knowledge and deeds, days and nights for many years: I share one ring with you both, for Frodo wouldn't have got far without Sam.

An even longer friendship I have enjoyed with Vale: from my childhood to this day, I have followed you up the trees and into abandoned houses, and seen you grow and marry and start a family; you are a beloved sister to this only child. Edo has also been always there, through light and tough moments, a valuable and solid friend: no matter how far I go, I always return to you.

Last in order, but definitely not least, I must thank my parents, who have always loved, encouraged and supported me, and by whose pride in everything I do I have drawn my strength: only one of them lived long enough to see this moment of my life, but I hope both have known how much I appreciated them.

And after having celebrated all the people that made me better, let it be said that if some mistake has been done along the road, only two are responsible: Titivillus and I.

Καὶ καταλιπὼν αὐτοὺς ἐξῆλθεν ἔξω τῆς πόλεως
εἰς Βηθανίαν, καὶ ηὐλίσθη ἐκεῖ.
(Bibl.Vat., Vat.gr.1209/1267, col.2, 13-17)

E lasciatili, uscì dalla città
e se ne andò a Betania ove passò la notte.
(Mt 21:17, Italian transl. by T. Lovejoy)

Chapter 1

General introduction

1.1 Growth in groups, Cayley graphs, diameters

These first introductory sections on growth in groups (from §1.1 to §1.4) owe much to two surveys by Helfgott [Hel15] [Hel19a]. The reader can find more information about the topics presented here in them and in their references.

Let G be a group. For any two subsets $A, B \subseteq G$, define

$$AB = \{g \in G | \exists a \in A, b \in B (g = ab)\}.$$

As a shortcut, we can also define recursively for any $k \in \mathbb{N}$ the set $A^{k+1} = A^k A$, starting with $A^0 = \{e\}$.

The very general problem we are going to talk about is the following: what is the behaviour of the size $|A^k|$ of the set A^k as the exponent k grows? This question can be addressed in a myriad of particular situations, and from many different points of view. To tackle them all would be a monumental task beyond the aim, and possibly beyond the strength, of the author, so let us immediately restrict our attention to the case of G finite.

The first thing we observe is that, when G is finite, there will be a certain k such that $|A^k|$ stabilizes from that point onwards, in the sense that $|A^{k'}| = |A^k|$ for all $k' \geq k$. It is always true in fact that $|A^{k+1}| \geq |A^k|$, simply because $|A^k a| = |A^k|$ for any $a \in A$. Moreover, when equality occurs, it must be that $A^k a = A^k a' = A^{k+1}$ for any $a, a' \in A$: therefore $a'' A^k a = a'' A^k a'$ as well, and taking the union of all such sets among all $a'' \in A$ we get $A^{k+1} a = A^{k+1} a'$ as well and $|A^{k+2}| = |A^{k+1}|$. Given this scenario, to interpret the original question as a problem on the asymptotic behaviour $\lim_{k \to \infty} |A^k|$ (which would be, and is, natural for G infinite) would amount to produce only trivial answers, for all finite groups and all their finite subsets would have a constant as the limit above. On the contrary, the problem of what is the least k such that $|A^k|$ stabilizes is the correct and interesting question we would want to examine.

The set $H = \bigcup_{i=0}^{\infty} A^i$ is a subgroup of G: in fact we have that $a^{-1} = a^{|G|-1}$, since G is finite. When $|A^k|$ stabilizes, all $A^{k'}$ with $k' \geq k$ are cosets of the

same subgroup of H. If we have not only $|A^{k+1}| = |A^k|$ but also $A^{k+1} = A^k$, or equivalently A^k becomes eventually the subgroup H defined before, then we say that A is a set of generators of such a subgroup, which we denote by $\langle A \rangle$. In that case, the least k such that $A^k = \langle A \rangle$ is called the diameter: the reason behind the name is that we can reframe these concepts in a graph-theoretic language, as we are going to do now.

Definition 1.1.1. *Let G be a finite group, and let A be a set of generators of G. The Cayley graph* $\mathrm{Cay}(G, A)$ *is the graph (V, E) with set of vertices $V = G$ and set of edges $E = \{(g, ag) | g \in G, a \in A\}$.*

Other notations are also used in the literature, most prominently $\Gamma(G, A)$, as in [Hel15].

The concept of Cayley graph dates back to 1878 [Cay78]. In the definition above, we have arbitrarily chosen the edges to be defined by left multiplication: there is nothing special about this choice, and we could have used right multiplication without hampering our progress to any of the results that follow, so long as we are consistent about our decision.

Since $\langle A \rangle = G$, $\mathrm{Cay}(G, A)$ is strongly connected: using the identity $(g'g^{-1})g = g'$, there exists a directed path from the vertex g to the vertex g' determined by a finite sequence of generators $a_i \in A$ such that $a_1 a_2 \dots a_m = g'g^{-1}$ (this path is not unique in general, of course). The set A is allowed (and often encouraged) to contain e, so the set of edges E may contain loops; in some contexts it is also useful to consider $\mathrm{Cay}(G, A)$ as a labelled graph, where the labelling of E is given by A (i.e. (g, ag) is labelled a: this is unambiguous, as $ag = a'g$ obviously implies $a = a'$).

As is commonplace with graphs, we can define the length of a walk as the number of edges involved in the definition of the walk itself, and then for any two vertices $v, w \in G$ we can define the distance $d_A(v, w)$ as the length of the shortest walk from v to w in the graph $\mathrm{Cay}(G, A)$ (if there is no risk of confusion, the index in the notation d_A can be dropped). This allows us to give the following definition.

Definition 1.1.2. *Let G be a finite group, and let A be a set of generators of G. The (directed) diameter of* $\mathrm{Cay}(G, A)$ *is*

$$\mathrm{diam}^+(\mathrm{Cay}(G, A)) = \max\{d_A(v, w) | v, w \in G\}.$$

The (directed) diameter of G is

$$\mathrm{diam}^+(G) = \max\{\mathrm{diam}^+(\mathrm{Cay}(G, A)) | A \subseteq G, \langle A \rangle = G\}.$$

Other authors use the notation $\overrightarrow{\mathrm{diam}}(G)$, as in [HS14].

It is clear that $\mathrm{diam}^+(\mathrm{Cay}(G, A))$ is the same as the diameter considered, a bit differently from before, as the least k with $\bigcup_{i=0}^{k} A^i = \langle A \rangle$. Such k is the same as the maximum distance from the identity $e \in G$ to any other vertex of G (so that we have the inequality in one direction), and the maximum distances from all vertices are the same in a Cayley graph: as a matter of fact, Cayley graphs are vertex-transitive, meaning that there is an automorphism sending v_1 to v_2 for any

given pair of vertices $v, w \in G$, given by right multiplication (as $av_1 = v_2$ if and only if $av_1w = v_2w$)[1], therefore the existence of a vertex at a certain distance d from v_i must be preserved by that transformation.

The questions we are trying to answer in this context become: given a finite group G and a set of generators A, what is the diameter of $\mathrm{Cay}(G, A)$? Is there some important difference between classes of sets A that is reflected in different diameters? Given a finite group G, what is its diameter? Are there general lower or upper bounds on $\mathrm{diam}(G)$ that hold for all groups, or for ample classes of groups, as either $|G|$ or other parameters relevant to the class of groups one considers tend to infinity?

Before we go into the more particular and interesting cases, around which most of today's research revolves, let us establish some very basic facts. As we are talking about diameters of (strongly connected) graphs, the most trivial bound one could think of would be $\mathrm{diam}^+(\mathrm{Cay}(G, A)) \leq |G| - 1$. If we are taking into consideration the fact that the edges of a Cayley graph are directed, the equality is achieved by some groups: for example, we can take $G = \mathbb{Z}/n\mathbb{Z}$ and $A = \{0, 1\}$, and the element -1 would be reached only in $|G| - 1$ steps.

In reality, as many authors do (and as we will do as well in the future), one can focus on undirected Cayley graphs[2], or equivalently on sets of generators A such that $A = A^{-1}$, since one can see an undirected edge $\{v, w\}$ as two edges $(v, w), (w, v)$. In that case, the trivial upper bound becomes $\left\lfloor \frac{1}{2}|G| \right\rfloor$: in fact for every such graph Γ and every vertex v there cannot be only one vertex with distance $0 < d < \mathrm{diam}(\Gamma)$ from v, or else such a vertex would become a separating set on its own and, by the vertex-transitivity of Cayley graphs, all vertices would be (which is impossible). Again, the bound is actually achieved, for example by $G = \mathbb{Z}/n\mathbb{Z}$ and $A = \{-1, 0, 1\}$. A second common simplification is to consider only sets A with $e \in A$: in this case we have $A^{k'} \supseteq A^k$ for $k' \geq k$ and then $\bigcup_{i=0}^{k} A^i = A^k$; in particular, every set A is a set of generators of some $\langle A \rangle \leq G$, and there is a minimum k with $A^k = \langle A \rangle$ that is equal to $\mathrm{diam}^+(\mathrm{Cay}(\langle A \rangle, A))$.

Our desire to adopt these two simplifications, namely $e \in A = A^{-1}$, prompt us to write the following definition.

Definition 1.1.3. *Let G be a finite group, and let A be a set of generators of G. The (undirected) diameter of* $\mathrm{Cay}(G, A)$ *is*

$$\mathrm{diam}(\mathrm{Cay}(G, A)) = \max\{d_{A \cup A^{-1}}(v, w) | v, w \in G\}.$$

The (undirected) diameter of G is

$$\mathrm{diam}(G) = \max\{\mathrm{diam}(\mathrm{Cay}(G, A)) | A \subseteq G, \langle A \rangle = G\}.$$

[1] Evidently, the automorphism even preserves labels. Conversely, right multiplications are obviously the only automorphisms of Cayley graphs, if we must respect the labels; however, if we are required to respect the labelling but not the labels themselves (i.e. we can send an edge labelled a_1 to another labelled a_2, as long as *all* edges labelled a_1 are sent to a_2), or if we can ignore the labelling entirely, the question is more complicated. It is not known in general what the group of automorphisms of a Cayley graph is in those cases, although some partial results exist (see [God81] [BG82] [DSV16] [PSV17]). This is quite an interesting topic, which unfortunately we are not going to explore.

[2] For instance, Babai [Bab06] refers to directed Cayley graphs as "Cayley digraphs", and to undirected Cayley graphs as "Cayley graphs".

From now on, whenever we talk about Cayley graphs and diameters, we always refer to undirected ones unless otherwise stated. It is clear that if $e \in A = A^{-1}$ then diam and diam^+ of the corresponding Cayley graph coincide, and that they are the same as the least k with $A^k = \langle A \rangle$.

Our questions on diameters remain unchanged, although now they refer to undirected diameters. This does not entail a significant loss of generality, as long as we are not particular about small factors: in fact we know that

$$\mathrm{diam}^+(\mathrm{Cay}(G, A)) \leq C \cdot (\mathrm{diam}(\mathrm{Cay}(G, A)))^2 \cdot \log^3 |G|,$$
$$\mathrm{diam}^+(G) \leq (3 + o(1)) \cdot \mathrm{diam}(G) \cdot \log^2 |G|,$$

where $C > 0$ is some absolute constant (see [Bab06, Thm. 1.4] and [Bab06, Cor. 2.3] respectively).

The fact that considering A as symmetric is not a big difference form the non-symmetric case can also be seen, from another point of view, by looking more closely at the growth of A itself (intended as the ratio between the size of the powers of A and the size of A) instead of counting the number of steps that are necessary to fill the group. This assertion is based on the fact that the growth of the set $B = A \cup A^{-1} \cup \{e\}$ can be controlled by the growth of A itself: more precisely, we have

$$\frac{|B^3|}{|B|} \leq \left(\frac{3|A^3|}{|A|}\right)^3 \tag{1.1.1}$$

for any finite G and any $A \subseteq G$. This bound is retrieved through rather elementary means starting from the ideas of Ruzsa: see for example the slightly stronger statement in [Hel15, (3.2)]. Notice also that we measure growth in terms of the cube of the set: again, we do so without loss of generality, since one can show that

$$\frac{|A^k|}{|A|} \leq \left(\frac{|A^3|}{|A|}\right)^{k-2} \tag{1.1.2}$$

for all $k \geq 3$ and all sets $A = A^{-1}$ (a weaker version of which appears in [RT85, Thm. 3] for $G = \mathbb{Z}$). Both (1.1.1) and (1.1.2) descend in particular from the fundamental Ruzsa triangle inequality [Ruz96, Thm. 4.2], which dates at least as back as 1976 (see [Ruz79]) and whose arguments generalize even outside groups (see [GHR15, Lemma 4.2]). Many of the arguments of Ruzsa (and Plünnecke [Plü70] before him) were originally framed in the context of abelian groups, but they generalize without much difficulty to the non-abelian case. It has to be noted however that the growth in abelian groups can be measured in terms of $|A^2|$ instead of $|A^3|$; this does not happen in general, for example in the case of $A = H \cup \{g\}$ with $H \leq G$ and $g \in G \setminus H$ and HgH much larger than A (see for instance [Hel19a, Ex. 2.1]).

Returning to the question that we were considering before, upper bounds of the form $o(|G|)$ for the diameter of G are less trivial: as we have observed, there are groups for which it would not be possible to prove such a statement, but it may be (and is) possible for more restricted but still interesting classes of groups. It is clear that we cannot do better than $O(\log |G|)$, at least for groups that are generated by

a number of elements not larger than a given constant: clearly $|A^k| \leq |A|^k$, with equality realized only when all products of k elements of A are effectively distinct, therefore we can have $A^k = G$ only for $k \geq \frac{\log |G|}{\log |A|}$; examples of groups that we are going to see later and that are generated by a constant number of elements are $\mathrm{Sym}(n)$ and $\mathrm{SL}_n(\mathbb{F}_q)$. By the fact that any group has a set of generators of size $\leq \log_2 |G|$ (because if $H \leq G$ and $g \in G \setminus H$ then $|\langle H \cup \{g\}\rangle| \geq 2|H|$), we could also reach at most $\mathrm{diam}(G) = O\left(\frac{\log |G|}{\log \log |G|}\right)$, for any finite group.

Actually such strong bounds are not even true in many interesting cases: see Example 5.4.1, which shows that $\mathrm{diam}(\mathrm{Alt}(n)^2) = \Omega(n^2)$ whereas $\log |\mathrm{Alt}(n)^2| \leq 2n \log n$ (with the same methods we can prove the same lower bound for $\mathrm{Alt}(n)$ itself). Another example would be $G = (\mathbb{Z}/2\mathbb{Z})^n$, which cannot be generated by less that n elements: the set $A = \{0^n\} \cup \{0^i 10^{n-i-1} | 0 \leq i < n\}$ however generates G in n steps, so that $\mathrm{diam}(G) = \Omega(\log |G|)$. In the latter example, we are hindered by the abelianness of G, which stifles the growth of A since not all products of k elements are distinct; case in point, we have seen $\mathbb{Z}/n\mathbb{Z}$ having linear diameter.

Many finite groups are, however, not abelian. In fact, most of them are not: the number of abelian groups of order $\leq n$ is linear in n (see [Ivi85, Thm. 14.6]), while the number of groups of order p^n is at least $p^{\left(\frac{2}{27} - o(1)\right)n^3}$ for p fixed prime and $n \to \infty$, as was already known to Higman [Hig60]. For non-commutative groups, one can hope and often expect A to cover G in a shorter number of steps, and the Cayley graph $\mathrm{Cay}(G, A)$ or even G itself to have a relatively small diameter. In the next sections, we are going to introduce the concept of *simple group* and see what has been proved or conjectured about their diameters.

1.2 Finite simple groups

Let G be a group, not necessarily finite. Among all its subgroups, *normal subgroups* occupy a special place: these are the subgroups N such that $gN = Ng$ for any $g \in G$, and the property of being normal is denoted by the notation $N \trianglelefteq G$. Normal subgroups enjoy many nice properties, first and foremost the fact that we do not have to be careful about left and right multiplication, which is used to show that the set of (left or right) cosets of N behaves like a group as well, the *quotient group* G/N.

One could expect to be able to understand G, at least to a certain extent, by studying N and G/N instead: a spectacularly appropriate example of this is provided by Lemma 6.2.5. Hence, we could and should aim to reduce ourselves to the smallest possible unit of study.

Definition 1.2.1. *Let G be a group. G is said to be simple if there are no normal subgroups $N \trianglelefteq G$ other than $N = \{e\}$ and $N = G$.*

We have claimed somewhat vaguely that simple groups are the smallest objects worth studying in some contexts. Their role is often compared to the one played by prime numbers in the context of integers: we can split a number n into two factors a, b such that $n = ab$, and then keep going until we end up with a product

of prime numbers. Of course, we are making use of the fact that $\mathbb{Z}$ is a UFD, i.e. that there is an essentially unique way to split n into its prime factors; if we want to keep up this *giuoco delle parti*, we need an analogous result in the context of groups. Fortunately, at least for finite groups, we do have such a result.

Definition 1.2.2. *Let G be a finite group. A composition series is a finite chain of proper normal subgroups*

$$\{e\} = H_0 \lhd H_1 \lhd H_2 \lhd \ldots \lhd H_{n-1} \lhd H_n = G$$

such that the quotient H_i/H_{i-1} is simple for every $1 \leq i \leq n$. The quotients H_i/H_{i-1} are called composition factors.

Theorem 1.2.3 (Jordan-Hölder theorem)**.** *Let G be a finite group. Then, any two composition series of G are equivalent, i.e. they have the same length and the composition factors that appear in the two series are the same up to isomorphism and up to permutation of their position in the series.*

This theorem is named after Jordan, who proved that the quotients have the same size up to permutation [Jor70, §55][3], and Hölder, who proved that they are actually isomorphic [Höl89]; see [Bau06] for a short, modern proof. It is worth noting that the theorem is generalizable to infinite groups and transfinite series, as long as they are ascending and not descending: see [Bir34, Thm. 1] for a proof of the ascending case and [Bir34, Thm. 2] for a counterexample of the descending one.

Thanks to Jordan-Hölder, the study of finite groups can often reduce to the study of finite simple groups instead. The first question that comes to mind then is: which are the finite simple groups?

The history of the search for a definitive answer to this question is quite articulate, and in some sense still ongoing. Solomon [Sol01] offers a good overview; we will give here only a handful of highlights. The concepts of normal subgroup and of simple group go back to Galois, who famously proved that $\mathrm{Alt}(n)$ is simple for $n \geq 5$ in order to show that general equations of degree ≥ 5 are not soluble through radicals [Gal46b]; he later proved that $\mathrm{PSL}(2, p)$ is also simple for primes $p > 3$ [Gal46a]. An actual conscious search for all finite simple groups is conventionally believed to have started with a question by Hölder [Höl92]. The work of classifying all such groups went on for almost a century after that.

The period of most intense advancement is generally considered to have begun in 1955, when the Brauer-Fowler theorem [BF55] showed a concrete way of attacking the problem through the study of centralizers of involutions: on this same track lies the Feit-Thompson theorem, whose proof appeared eight years later [FT63]. In the years between 1976 and 1983 the classification project was essentially wrapping up, and in this period it was declared to be near completion or completed by several mathematicians, like Brauer [Bra79], Collins [Col80] and

[3] In 150 years, the language has changed: a "substitution" is an element of the group [Jor70, §23], and "permutable" means normal [Jor70, §35]. The groups that Jordan was studying were permutation groups, whence the terminology; the proof itself does not restrict only to these groups, though.

Gorenstein [Gor82]; it was chiefly Gorenstein's announcement that marked the moment when the project was believed to have reached its completion. However, the case of quasithin groups had been solved only partially, in an unpublished manuscript by Mason that still had gaps to be filled. Aschbacher and Smith fixed this last important missing piece only in 2004 [AS04a] [AS04b].

From that year onwards, the Classification of the Finite Simple Groups (CFSG) has generally been accepted to be proved. Here is the statement.

Theorem 1.2.4 (CFSG)**.** *For any group G, G is a finite simple group if and only if it is one of the following:*

(a) a cyclic group $\mathbb{Z}/p\mathbb{Z}$ with p prime (the only abelian groups of the list);

(b) an alternating group $\mathrm{Alt}(n)$ for $n \geq 5$;

(c) a group of Lie type among the following 16 families (in all of them, q is a prime power): $\mathrm{PSL}_n(q)$ (with $n \geq 2$ and $(n,q) \neq (2,2),(2,3)$), $\mathrm{PSU}_n(q)$ (with $n \geq 3$ and $(n,q) \neq (3,2)$), $\mathrm{PSp}_{2n}(q)$ (with $n \geq 2$ and $(n,q) \neq (2,2)$), $\mathrm{P}\Omega_{2n+1}(q)$ (with $n \geq 3$ and q odd), $\mathrm{P}\Omega^+_{2n}(q)$ (with $n \geq 4$), $\mathrm{P}\Omega^-_{2n}(q)$ (with $n \geq 4$), $G_2(q)$ (with $q \neq 2$), $F_4(q)$, $E_6(q)$, ${}^2E_6(q)$, ${}^3D_4(q)$, $E_7(q)$, $E_8(q)$, ${}^2B_2(2^{2n+1})$ (with $n \geq 1$), ${}^2G_2(3^{2n+1})$ (with $n \geq 1$), ${}^2F_4(2^{2n+1})$ (with $n \geq 1$);

(d) one of 26 sporadic groups (M_{11}, M_{12}, M_{22}, M_{23}, M_{24}, Co_1, Co_2, Co_3, McL, HS, Suz, J_2, Fi_{22}, Fi_{23}, Fi'_{24}, $\mathbb{M}$, $\mathbb{B}$, Th, HN, He, J_1, J_3, J_4, O'N, Ly, Ru) or the Tits group ${}^2F_4(2)'$.

For the notation, see [Wil09, §1.2]; mind that there is a finite number of repetitions in points (b) and (c) of the list.

The proof of Theorem 1.2.4, as it stands today, is distributed across hundreds of articles that total around 10000 pages: this is chiefly the reason why most people refer to CFSG as "widely accepted" instead of saying straight up "a theorem", for its unwieldy proof is not fit for human consumption. The truth is, mathematics is still on some extent based on trusting the community of mathematicians: while in principle it is a game of absolute rigour, humans are humans and may make mistakes in writing and proofreading depending on whether they have skipped lunch a certain day[4]. The author eventually learned to accept this fact[5], and length is for sure not a sufficient reason for making a proof not a proof, certainly not when the proofreading machines that we are last at most 122 years and 164 days[6]: thus, for us CFSG is Theorem 1.2.4, emphasis on "Theorem".

[4] *[E]t idem / indignor, quandoque bonus dormitat Homerus: / verum operi longo fas est obrepere somnum.* (Quintus Horatius Flaccus, *Ars Poetica*, 358-360)

[5] The healthiest attitude towards this problem on the part of a scientist, as far as the author has encountered, is expressed by Stephen Jay Gould, the paleontologist: faced with a change of the consensus on a particular scientific issue which was outside his expertise but affected his own conclusions, he had to "acknowledge, and [...] provisionally accept" (Gould, *The Structure of Evolutionary Theory*, §9.3.2). If "provisionally", for whatever reason, extends until we reach death or retirement or other invalidating circumstances, a problem that Gould does not address, the author (in a beautiful Italian turn of phrase that is coincidentally appropriate on multiple levels) accepts *con filosofia*.

[6] Jeanne Calment.

In any case, there is an ongoing process of writing a second-generation proof: as of the time of writing, 8 volumes out of the planned 13 have been published [GLS94] [GLS96] [GLS98] [GLS99] [GLS02] [GLS05] [GLS18a] [GLS18b]. Moreover, there has been a computer verification of an important part of CFSG, namely the Feit-Thompson theorem has been proved using Coq, a theorem-proving software (see [GAA+13]).

Let us leave the topic of CFSG itself, however interesting its history and philosophical implications may be, and move to one of its consequences that will be important to us. It is a classification of primitive subgroups of $\mathrm{Sym}(n)$; the version below is due to Maróti [Mar02], but the original result comes from Cameron [Cam81]. For the definition of k-transitive, primitive and wreath product, see §3.1.

Theorem 1.2.5. *Let $n \geq 1$ and let $G \leq \mathrm{Sym}(n)$ be primitive. Then, one of the following alternatives holds:*

(a) there are integers m, r, k such that $\mathrm{Alt}(m)^r \leq G \leq \mathrm{Sym}(m) \wr \mathrm{Sym}(r)$, where $\mathrm{Alt}(m)$ acts on k-subsets of $\{1, 2, \ldots, m\}$ and the wreath product action is the primitive one (so that in particular $n = \binom{m}{k}^r$);

(b) G is one of the sporadic groups $\mathrm{M}_{11}, \mathrm{M}_{12}, \mathrm{M}_{23}, \mathrm{M}_{24}$ with their 4-transitive action;

(c) $|G| \leq n \prod_{i=0}^{\lfloor \log_2 n \rfloor - 1} (n - 2^i) < n^{1+\log_2 n}$.

Even the history of this particular result is quite involved. The first version of a classification of primitive permutation subgroups like the one above appeared in 1981 and was due to Cameron [Cam81, Thm. 6.1]; the proof depends on CFSG (whose statement was already known and considered likely to be correct at the time, and it is referenced to as a "hypothesis" in [Cam81, §1]) and on the O'Nan-Scott theorem. The latter result appeared first in an article for a 1979 conference by Scott [Sco80], who stated in a footnote that O'Nan had also independently obtained it: the theorem, which does not depend on CFSG, offers a classification of maximal permutation subgroups. However, the O'Nan-Scott theorem itself was incorrectly proved: one case, the "twisted wreath action" case (in the language of [LPS88]), was omitted; this has no consequence on the validity of the statement though, as the groups that arise from this case are not maximal. The proof was first corrected by Aschbacher and Scott [AS85], after Cameron's article had already appeared (as the authors themselves point out)[7]. Cameron's original theorem thus is in the unusual position of having been deduced from two major results whose statements were both correct but whose proofs had both an undiscovered gap at the time.

After Cameron's version, another appeared due to Liebeck [Lie84]: this version is closer to the kind of result we will need to use, and it already acknowledges both

[7]Technically the first *published* correction is in [CPSS83], which was received in 1982 and appeared the following year, but the authors of this article make reference to Aschbacher and Scott's paper "to appear" (it was received in 1983 and it appeared in 1985). Liebeck [Lie84] refers to the theorem as being corrected in [CPSS83], but adds "for instance": this was in 1984, after all.

CFSG as a "theorem" (this was after Gorenstein's announcement, but before the quasithin gap had been truly acknowledged) and the correction of the O'Nan-Scott theorem. Even later, Maróti offered the version stated before (see [Mar02, Thm. 1.1]), which is in some sense the furthest possible refinement of Liebeck's theorem: if we were to tighten (c) inside Theorem 1.2.5 even further, an infinite family of exceptions as in (b) would emerge.

Now that we have a list of what a finite simple group can be, let us move to the next problem we face, namely what we can state about the diameter of such a group.

1.3 Babai's conjecture

As we have already said, investigating the properties of finite simple groups is often a good step for describing the properties of finite groups as a whole. This is true in particular for diameters. A consequence of Schreier's lemma [Sch27], which we state as Lemma 5.2.1, is that the diameter of a finite group can essentially be bounded linearly in terms of the product of the diameters of its composition factors (see in fact Lemma 6.2.5, in the same spirit). When the group is the direct product of finite simple groups, the dependence is even nicer: its diameter can be bounded linearly by the maximum diameter of its factors, as shown in the author's preprint [Don19a] based on previous results by Babai and Seress [BS92] and Helfgott [Hel18] (see §5).

Our objective then becomes to estimate, as accurately as we are able to, the diameter of finite simple groups. By CFSG, we only need to do so for a handful of well-described families, for which we desire to give bounds depending on their size (or on the parameters that determine them, such as n and q). The sporadic groups in point (d) of Theorem 1.2.4 do not pose a problem at all: they are finitely many, albeit possibly very big ($\mathbb{M}$, the largest one, approaches size 10^{54}), so their diameter is just a constant that is even computable in principle. We have already seen in §1.1 that $\operatorname{diam}(\mathbb{Z}/p\mathbb{Z}) = \left\lfloor \frac{p}{2} \right\rfloor$, i.e. we have a linear dependence on $|G|$ for G finite simple abelian. We are thus left with two cases to examine: alternating groups and groups of Lie type.

Conjecture 1.3.1 (Babai's conjecture)**.** *Let G be a finite simple non-abelian group. Then, there is an absolute constant $C > 0$ such that*

$$\operatorname{diam}(G) \leq \log^C |G|.$$

From what we discussed in §1.1, this is essentially best possible: already for $G = \operatorname{Alt}(n)$, a lower bound is known with $C = 2 - \varepsilon$ for any $\varepsilon > 0$ and $|G|$ large enough. In fact, for any $\varepsilon > 0$, given any finite group G large enough with respect to ε and given a non-redundant set A of generators of G (meaning that there is no proper subset of A that still generates G), we must have $|A| \leq \log_2 |G|$ and then $\operatorname{diam}(\operatorname{Cay}(G, A)) \geq \log^{1-\varepsilon} |G|$. Therefore, even a bound with $o(1)$ instead of C would be false for any infinite class of finite groups.

Babai's conjecture was stated in the literature for the first time in 1988 by Babai and Seress [BS88, Conj. 1.7], and in its full generality it is still unsolved.

There have been however numerous weaker or partial results, which we are going to illustrate here.

As we showed, two infinite classes of finite non-abelian simple groups exist, the alternating groups and the groups of Lie type, and results often apply to only one of the two. Let us start with $\mathrm{Alt}(n)$. The oldest nontrivial bound appears in the same paper by Babai and Seress [BS88] in which Babai's conjecture is first reported: their "modest first step", as they called it, was to show that

$$\mathrm{diam}(\mathrm{Sym}(n)), \mathrm{diam}(\mathrm{Alt}(n)) \leq e^{(1+o(1))\sqrt{n \log n}}. \tag{1.3.1}$$

Soon after, they went on to show that the case of transitive permutation groups reduces to $\mathrm{Alt}(n)$, by claiming that for any $G \leq \mathrm{Sym}(n)$ transitive we have

$$\mathrm{diam}(G) = e^{O(\log^3 n)} \mathrm{diam}(\mathrm{Alt}(m)) \tag{1.3.2}$$

where $\mathrm{Alt}(m)$ is the largest alternating group that is a composition factor of G (see [BS92, Thm. 1.4]). The proof however contains a bookkeeping mistake (as pointed out by Pyber, see [Hel18, §1]); the correct statement should be

$$\mathrm{diam}(G) = e^{O(\log^2 n)} \prod_i \mathrm{diam}(\mathrm{Alt}(m_i)) \tag{1.3.3}$$

with $\prod_i m_i \leq n$, as shown in [Hel18, Prop. 4.15] (known to Pyber). Setting aside the correctness of (1.3.2), results of this kind are in any case particularly significant for us, because combining them with (1.3.1) we obtain diameter bounds for all transitive groups; however, while (1.3.1) does not rely on CFSG (it has a purely combinatorial proof), both (1.3.2) and (1.3.3) follow from Cameron's theorem which in turn, as we mentioned in §1.2, follows from CFSG: it will be the objective of §6 to try (and only partially succeed) to give a CFSG-free version of (1.3.3).

The next step, or giant leap, came at the hands of Helfgott and Seress [HS14]: they proved that

$$\mathrm{diam}(\mathrm{Sym}(n)), \mathrm{diam}(\mathrm{Alt}(n)) \leq e^{O(\log^4 n \log\log n)}. \tag{1.3.4}$$

This quasipolynomial bound in n is much closer to Babai's conjecture than (1.3.1): since $|\mathrm{Alt}(n)| = \frac{1}{2}n!$, a polylogarithmic bound in $|G|$ as in Conjecture 1.3.1 corresponds to a polynomial bound in n. In particular, combining (1.3.4) and (1.3.3) (in fact they used the incorrect (1.3.2), but one can replace it with (1.3.3) without changing the end result, as stated in [Hel18, §1]), Helfgott and Seress settled a different conjecture stated in [BS88, Conj. 1.6], which claimed that the diameter of transitive groups is bounded quasipolynomially in n (a stronger polynomial conjecture has also been formulated [KMS84]). The upper bound for $\mathrm{diam}(\mathrm{Alt}(n))$ in (1.3.4) is the best to date.

Let us turn now to groups of Lie type. Here, the advancements went mostly hand in hand with generalizations of the following theorem.

Theorem 1.3.2. *Let p be a prime, let $G = \mathrm{SL}_2(\mathbb{F}_p), \mathrm{PSL}_2(\mathbb{F}_p)$, and let A be a set of generators of G. Then there exist absolute constants $\delta > 0$ and $k \geq 1$ such that at least one of the following alternatives holds:*

(a) $|A^3| \geq |A|^{1+\delta}$;

(b) $(A \cup A^{-1} \cup \{e\})^k = G$.

The theorem above, in this particular form, is due to Helfgott [Hel08]. It describes the behaviour of a set of generators as being subject to an alternative: any such set either has large growth or quickly fills the entire group. Notice that, as we pointed out in §1.1, "growth" is measured by the ratio $\frac{|A^3|}{|A|}$. Any theorem displaying the same kind of dichotomy is commonly referred to as a *product theorem* (as in [Bre14] [Raz14] [Hel18], to name a few instances). The importance of product theorems in the context of diameters is obvious: applying repeatedly Theorem 1.3.2 to A, A^3, A^9, etc... until we fall into case (b), the number of steps (a) that we can pass through at most is bounded by $\frac{\log\log|G| - \log\log|A|}{\log(1+\delta)}$, and the diameter is bounded by $k + (\log|G|)^{\log 3/\log(1+\delta)}$; hence, Babai's conjecture holds for the family of simple groups $\mathrm{PSL}_2(\mathbb{F}_p)$ (p prime) [Hel08].

Helfgott's result, which dates back to 2005, was the first of a series of increasingly general proofs of Babai's conjecture for classes of finite simple groups of Lie type. First of all, Theorem 1.3.2 holds with $A^3 = G$ replacing case (b) [NP11], with $|A^3| = \Omega(|A|^{1+\frac{1}{20}})$ replacing case (a) [RS18], and *cannot* hold with $\delta \geq \frac{1}{6}(\log_2 7 - 1)$ [BRD15]. The theorem was generalized for $\mathrm{PSL}_2(\mathbb{F}_q)$ with any prime power q by Dinai [Din11], and for $\mathrm{PSL}_3(\mathbb{F}_p)$ by Helfgott [Hel11]; moreover, case (a) was shown to hold for all sets of generators of $\mathrm{PSL}_n(\mathbb{F}_p)$ that are not too large [GH11]. Afterwards, a product theorem that holds for all finite simple groups of Lie type of bounded rank (or equivalently, for all such groups but where δ depends on the rank) was proved independently by Breuillard, Green and Tao [BGT11] and by Pyber and Szabó [PS16]: Babai's conjecture thus holds in this case as well.

There are however limitations to product theorems. In the sense of Theorem 1.3.2, a product theorem cannot hold for groups of Lie type with no condition on the rank, as the counterexample in [PS16, Ex. 77] shows. The same is true for $\mathrm{Alt}(n)$, with counterexamples in [Spi12, §4] and [PPSS12, Thm. 17]: a proof of Babai's conjecture for $\mathrm{Alt}(n)$ therefore cannot pass directly through a statement as strong as Theorem 1.3.2; there exists however a weaker version of a product theorem that does hold in the alternating group case, and thanks to which one can prove a diameter bound almost as strong as (1.3.4): the result is [Hel18, Thm. 1.4], and we will see more of it in §6. A diameter bound for finite simple groups of Lie type of bounded base field (and unbounded rank) also exists: for such groups G, Biswas and Yang [BY17] proved that the diameter is at most $e^{O(\sqrt{\log|G|}(\log\log|G|)^3)}$, and the exponent 3 has been further reduced to 2 in [HMPQ19].

Existing proofs of the two main cases of Babai's conjecture, $G = \mathrm{Alt}(n)$ and G of Lie type, do not mingle much, as all the results cited above show. However, there are some deep-running similarities that pop their head out of the water here and there: for example, the counterexamples to strong product theorems are structurally analogous, as recognized for instance in [BGH$^+$14, §1]; the authors of [BY17] acknowledge that their search for a matrix of small degree and close to the identity reminds of the search for a permutation of small support in [BS88].

The use itself of a weakened product theorem in [Hel18] is a deliberate effort in bringing the alternating and the Lie type case closer together, "towards a unified perspective". There is however one nontrivial diameter bound that holds for all finite simple groups: Breuillard and Tointon [BT16] have proved that for every $\varepsilon > 0$ there is a constant C_ε such that

$$\operatorname{diam}(G) \leq \max\{|G|^\varepsilon, C_\varepsilon\} \tag{1.3.5}$$

for all finite non-abelian simple groups G (the bound in [BT16, Cor. 1.4] is for Cayley graphs with symmetric generating sets, to be precise). The proof has also the remarkable characteristic of being CFSG-free: for $\mathrm{Alt}(n)$, the best CFSG-free bound is (1.3.1), which is stronger than (1.3.5) but not as general, and in any case much weaker than (1.3.4). A CFSG-free result much closer in strength to (1.3.4) (and in spirit to [Hel18]) would be Theorem 6.3.6, if Conjecture 6.3.4 were to be true (see §6 for details).

Another possible way to unify the treatment of the two cases, yet to bear fruits, could be through $\mathbb{F}_1$, the field with one element. There are ways to define such an object (which in any case is neither a field nor a one-element object: an astounding feat, lying twice in the short string of symbols "$\mathbb{F}_1$") and more importantly to define objects over it, in the sense of algebraic geometry: what is relevant to us is that generally one shows that "$\mathrm{GL}_n(\mathbb{F}_1) = \mathrm{SL}_n(\mathbb{F}_1) = \mathrm{Sym}(n)$", whatever that means (see for instance [Lor18, §2.1.2]).

Yet another way to close the gap between alternating and Lie type case, subject of recent research, could be to try and put into practice the following suggestion by Pyber, based on [BBS04]. One might be able to prove Babai's conjecture for all finite simple non-abelian groups by performing three steps: 1) finding quickly an element of "support" $n(1-\varepsilon)$ for some $\varepsilon > 0$; 2) using this first element, finding quickly a second element of smallest "support"; 3) using this second element, concluding the proof. In [BBS04], an element as in (1) is already inside our set of generators of $\mathrm{Alt}(n)$, and the other two steps follow; one would hope to do the same for groups of Lie type, although even defining what "support" means in this situation is not obvious. There has been some very recent progress on this front, due to Halasi [Hal20] and Eberhard and Jezernik [EJ20].

1.4 Other results on growth and diameter

There are many other results that, albeit not sitting directly under the umbrella of Babai's conjecture, are closely related to it, either historically or in methods or purposes or otherwise. Any account the author would make of them would be incomplete, his perspective skewed by personal interests and general ignorance; nevertheless, here is an unavoidably incomplete account.

Growth of A. First of all, one could focus more on the information about the growth of the set A itself, without concerns about the diameter: in this scenario, the finiteness of G is often not important. The results run mostly on the dichotomy: either A has large growth or A has structure. This is a point of view that comes from additive combinatorics, and that originated from the famous Freiman-Ruzsa

theorem: the result states that any finite set $A \subset \mathbb{Z}$ with growth $K = \frac{|A+A|}{|A|}$ (in the abelian setting, we use + and need only $A^2 = A+A$ instead of A^3) is contained in a generalized arithmetic progression of size $O_K(|A|)$ and rank $O_K(1)$ [Fre73], so that both parameters are bounded in terms of $|A|, K$ only and are independent from the particular choice of A itself. There have been generalizations to torsion-free abelian groups [Ruz94], abelian groups [GR07], nilpotent groups [Toi14], and solvable groups [Tao10, Thm. 1.17], with particular focus on subsets of $\mathrm{GL}_n(K)$ [Hru12, Cor. 5.11] [GH14]; a particularly general theorem by Breuillard, Green and Tao [BGT12, Thm. 1.6] proves in particular the Helfgott-Lindenstrauss conjecture [Hel15, Conj. 1]. See [BGT13] for a more extensive survey on the subject.

In this context, obtaining quantitatively good results is often a challenge. To this day, even for abelian groups we have yet to prove that a set A with $|A+A| = K|A|$ can be covered by $\log^{O(1)} K$ many translates of a generalized arithmetic progression of rank $\log^{O(1)} K$ and size $K^{O(1)}|A|$: this is known as polynomial Freiman-Ruzsa conjecture, in analogy with the Freiman-Ruzsa theorem, although Ruzsa [Ruz99] attributes its formulation to Marton; the word "polynomial" refers to the $K^{O(1)}$, which is the prominent feature of the statement. The best known result is due to Sanders [San12], who proved as an application of the methods of Croot and Sisask [CS10] that we can have a quasipolynomial bound in K (see also [San13] for a detailed survey).

Similar issues occur in more general scenarios: for instance, the Helfgott-Lindenstrauss conjecture was proved in [BGT12] "in an impressively general but quantitatively very weak sense" (to quote [Hel15, §4.2]). The conjecture, in a few words, asks one to show that for a set $A \subseteq G = \mathrm{GL}_n(\mathbb{F})$ with $|A^3| = K|A|$ and $\mathbb{F}$ an arbitrary field we have two normal subgroups H_1, H_2 of $\langle A \rangle$ with $H_1 \subseteq A^{O_n(1)}$, with H_2/H_1 nilpotent, and with A covered by $K^{O_n(1)}$ translates of H_2. Now, consider these three known results: [BGT12, Thm. 1.6] proves the conjecture even for G an arbitrary group, but only with $O_K(1)$ instead of $K^{O_n(1)}$; [PS14, Thm. 8] proves it with H_2/H_1 solvable instead of nilpotent; and [GH14, Thm. 2] (using [PS16]) proves it only for $\mathbb{F} = \mathbb{F}_p$.

In the last example, one could also extend the result rather easily to $\mathbb{F} = \mathbb{F}_q$, but then we would have to use $K^{O_{n,e}(1)}$ translates of H_2, where $e = \log_p q$. The contrast is stark with the situation of simple groups, where structural theorems (meaning, diameter bounds) depend on n alone. It is to be hoped that generalizing [GH14] to arbitrary finite fields is only a question of time, in a similar fashion as with simple groups leading up to [BGT11] (such hope is expressed in [GH14, §1.3] and [Hel15, §3.3.2], for instance)[8]. In a sense, solvable and simple groups are at the opposite ends of a spectrum, the ones having many nested normal subgroups whose quotients are abelian (the derived series), the others having no normal subgroups at all and thus having only one large and complicated quotient; the only finite simple solvable groups are the prime cycles in Theorem 1.2.4(a).

Focusing on the small-scale example offered by the *affine group* on a finite field

[8] *Note added in proof:* A very recent preprint by Murphy and Wheeler [MW20] offers a first step in this direction, treating the case of 2×2 upper triangular matrices over $\mathbb{F}_q$. The authors seem also to be able to deal with the general case, and prove a result analogous to [GH14] for $\mathbb{F}_q$ (personal communication).

$\mathbb{F}$ (see (4.0.1) for its definition), the structural statements about slow-growing sets A can take stronger forms; this is helped by the fact that we have access to whole new toolboxes. First of all, the sum-product theorem, i.e. results on the growth of finite fields under both addition and multiplication: this is a rich terrain of investigation, in which we search for lower bounds on $\max\{|A+A|, |A\cdot A|\}$ for A inside a field $\mathbb{F}$, in terms of powers $|A|^c$ with $c>1$. The first results are due to Erdős and Szemerédi [ES83] for $c = 1+\varepsilon$ and $\mathbb{F}=\mathbb{R}$, and to Bourgain, Katz and Tao [BKT04] for $c=1+\varepsilon$, $\mathbb{F}=\mathbb{F}_p$ provided that $|A|\geq p^\delta$ for some $\delta>0$. Today there have been many small improvements for c, trying to get closer and closer to $c=2-\varepsilon$, which is widely considered to be the correct value, and it has been generalized to arbitrary fields $\mathbb{F}$ provided that A is not stuck in some subfield; see for instance [RS20] for the current record $c=\frac{4}{3}+\frac{2}{1167}-\varepsilon$ for $\mathbb{F}=\mathbb{R}$. Techniques for sum-product estimates transfer to results on the affine group: this happened with [Hel15, Prop. 4.8] (based on [GK07]) and with [Mur17, Thm. 27], which give theorems for $\mathrm{Aff}(\mathbb{F}_p)$. Another set of tools at our disposal comes from geometric arguments on finite planes: [Mur17, Thm. 27] relies also on this second point of view, and a more recent proof for $\mathrm{Aff}(\mathbb{F}_p)$ has been given in [RS18, Thm. 5] (based on [Sző99]) without the use of sum-product techniques at all. In all these cases, the authors focused only on the case of prime fields; in §4 we will give a structural result for $\mathrm{Aff}(\mathbb{F}_q)$ as well, with q a generic prime power, essentially adopting the strategy of [RS18].

Probabilistic results. Going back to questions about diameters, another avenue of research would be on probabilistic results, namely results that hold for random (read: most) Cayley graphs of G. In this situation, random walks are examined and results are obtained not only on the diameter but on the mixing time, i.e. the number of steps necessary for a (lazy) random walk to approach equidistribution on the graph. Particularly beneficial are results on expander graphs: a family of graphs $\{\Gamma_i\}_i$ is an *expander family* if there is a constant $\varepsilon>0$ such that for all i and for all subsets S of vertices of Γ_i of size $\leq \frac{1}{2}|\Gamma_i|$ we have $|S\cup\Delta S|\geq(1+\varepsilon)|S|$, where ΔS is the set of vertices that are linked to at least one vertex of S by an edge[9]; by their nature, expander graphs have small mixing time (of the form $O(\varepsilon^{-1}\log|G|)$, see [Hel19a, Ex. 6.1]), which in turn implies small diameter.

Results on expansion of some Cayley graphs predate [Hel08]: it was already known that the family of graphs

$$\{\mathrm{Cay}(\mathrm{SL}_2(\mathbb{F}_p), A) | p \text{ prime} \geq 5\}, \qquad A = \left\{ \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \right\}$$

was an expander family [Lub94, Thm. 4.4.2(i)], a fact that uses Selberg's $\frac{3}{16}$-theorem [Sel65]. There have been some fascinating results holding in much more

[9] Definitions may vary slightly across the literature. We note that the definition given here is for a *vertex expander*; there are also *edge expanders*, but most importantly *spectral expanders*, which are defined using the eigenvalues of the adjacency operator of the graph. For regular symmetric graphs, the three concepts are equivalent up to renaming ε, a fact proved by Alon and Milman [AM85] and going under the name of Cheeger inequality in analogy with the case of manifolds [Che70]; see [Hel19a, §1.1].

generality: building upon Helfgott [Hel08] and Tao [Tao08], Bourgain and Gamburd showed that the same holds for all fixed sets A not contained in proper subgroups [BG08, Thm. 1] and for randomly chosen A of fixed size [BG08, Thm. 2] ("fixed" here means "the same for all p"); moreover, the family of all Cayley graphs of $\mathrm{SL}_2(\mathbb{F}_p)$ for almost all p is also an expander family [BG10, Cor. 1.1]. However, it is not known whether all Cayley graphs for all p form an expander family. There are also results about $\mathrm{SL}_n(\mathbb{F}_p)$: for example, choosing their sets of generators appropriately, there is an expander family of Cayley graphs of $\mathrm{SL}_{3n}(\mathbb{F}_p)$ for all p, n [Kas07b, Thm. 8(a)]. For more discussion on the topic, see [Hel15, §2.2] [Hel19a, §6.1].

Going to permutation groups, expanders have been known to exist for $\mathrm{Sym}(n)$ and $\mathrm{Alt}(n)$ since [Kas07a, Thm. 2]. The fact that other graphs related to $\mathrm{Sym}(n)$ (*Schreier graphs*, see Definition 2.1.1) are almost always expanders [FJR$^+$98, Thm. 2.2] plays a role in the proof of [HSZ15, Thm. 1.1], which states that for almost all $g, h \in \mathrm{Sym}(n)$ the Cayley graph $\mathrm{Cay}(G, \{g, h\})$ (where $\langle g, h\rangle = G$) has diameter at most $n^2 \log^C n$ for some absolute C: in other words, almost all Cayley graphs of $\mathrm{Alt}(n)$ would satisfy Babai's conjecture; the non-obvious fact that two random g, h generate almost always either $\mathrm{Alt}(n)$ or $\mathrm{Sym}(n)$ is a classical theorem by Dixon [Dix69]. Previous results on the same wavelength are [BH05, Thm. 2.2], the first polynomial bound for almost all pairs of elements $g, h \in \mathrm{Sym}(n)$, and [BBS04, Thm. 1.1], a polynomial bound for all $A \subseteq \mathrm{Sym}(n)$ having an element with support of size $\leq \delta n$ for some fixed $\delta < \frac{1}{3}$. The proof of the latter statement also uses a general expansion result that goes back to Landau and Odlyzko [LO81] (no, not that Landau), and the constant δ has been improved over time, with $\delta = 0.63$ obtained in [BGH$^+$14], with some small margin for improvement (see [BGH$^+$14, §5]).

The strategies involved especially in [BBS04] and [HSZ15] are of an algorithmic nature: thus, it is in some measure possible to even give constructive procedures to determine a word in g, h for any $k \in \mathrm{Sym}(n)$, in a relatively short time. The proof in [BBS04] yields a Las Vegas polynomial-time construction[10]; for [HSZ15], a running time of $O(n^2 \log^C n)$ is almost always possible, and the authors speculate that $O(n \log^C n)$ might be reachable as well (see [HSZ15, App. A]).

Girth. Finally, let us mention that problems related to the *girth* of Cayley graphs can also contribute to diameter bounds; the girth of a graph is the length of the shortest cycle contained in it. Obviously, if the girth of a graph Γ is g, we have $\mathrm{diam}(\Gamma) \geq \left\lfloor \frac{g}{2} \right\rfloor$: we can show as much in the same way as we have done for cyclic groups in §1.1; however, we can use the girth to give upper bounds as well.

A straightforward application is the following: by [GHS$^+$09, Thm. 8], the girth of random Cayley graphs of $\mathrm{SL}_2(\mathbb{F}_p)$ with bounded set of generators is at least

[10] Apposing "Las Vegas" to a running time of an algorithm is to say that the *expectation* of the running time is as described (for instance, polynomial). The terminology was introduced by Babai [Bab79], although a "Las Vegas N time" is most commonly described in terms of an algorithm able to either output the correct solution in time N with probability $> \frac{1}{2}$ or recognize a failure: this reduces to our deterministic description, once we note that making the same algorithm run k times reduces the probability of failure to less than $\frac{1}{2^k}$. For Las Vegas polynomial time specifically, another nomenclature is ZPP (see for example [AB09, §24.3]).

$\Omega(\log p)$. This, as observed for instance in [Hel08, Proof of Cor. 6.3] and [BGT10, §7], yields a diameter bound of $O(\log |G|)$: in fact, a girth of $\Omega(\log p)$ means that all words up to that length are distinct, so that in $\log p$ steps we reach p^ε elements for some $\varepsilon > 0$, and then Theorem 1.3.2 tells us that in finitely many iterations we fill the whole $\mathrm{SL}_2(\mathbb{F}_p)$. Actually, [BG08, Thm. 3] shows even more: a girth of $\Omega(\log p)$ implies a spectral gap for the eigenvalues of the family of Cayley graphs we are referring to, which in turn translates into the graphs being expanders.

A more recent result is contained in [LS19, Prop. 3], where in particular the diameter of Cayley graphs of $\mathrm{Alt}(n)$ with sets of generators of fixed size are polynomially bounded almost always by their girth.

1.5 The graph isomorphism problem

Let us seemingly change subject, and move to algorithmic graph theory. First of all, let us start with a small introduction to complexity, which may be unnecessary depending on the reader's background.

In an algorithm, *complexity* refers to the dependence of the running time of the algorithm on the size of its input: as it should be clear, an algorithm does not only need to solve a problem, but also solve it in a reasonable amount of time; as a consequence, algorithms are placed in different complexity classes depending chiefly (but not only) on their running times. We shall refer mostly to algorithm runtimes themselves, with respect to the size of the input, without using classes. However, let us mention at least the following, which needs more description. A problem is NP when verifying whether a solution is correct takes polynomial time; among the NP problems, the NP-complete are the "hardest" ones, in the following sense: producing a solution to every other NP problem reduces in polynomial time to producing a solution to an NP-complete problem. For more details, see the standard textbook [CLRS01, §34].

Now, back to graphs. If one tries to imagine what the first important problem regarding graph algorithms would be, this hypothetical person will have a good chance of answering: "To give an algorithm that finds out, in the shortest possible amount of time, whether or not two arbitrary graphs are actually the same".

That is essentially the *graph isomorphism problem*, often shortened to GIP: given two graphs $(V_1, E_1), (V_2, E_2)$ with n vertices, say as subsets of all the unordered pairs of numbers from 1 to n (where an arbitrary ordering is given to V_i, and the pairs represent the elements of E_i as defined by this ordering), find the fastest algorithm that produces the set of isomorphisms from one graph to the other, where an isomorphism is in this case just a permutation of $[n]$ that respects the property of pairs being in E_i.

Before we give a brief history of the problem, let us note here that GIP reduces to the *string isomorphism problem* (SIP): given two strings of length k, namely two functions from $[k]$ to a finite alphabet Σ, and a subgroup $G \leq \mathrm{Sym}(k)$, find the fastest algorithm that produces the set of their isomorphisms inside G, i.e. the permutations in G forming a commutative diagram with the two functions. We can in fact take $k = \binom{n}{2}$ and describe a graph of n vertices as a string of k

letters with alphabet $\Sigma = \{0, 1\}$, where the m-th letter is 1 if the m-th pair of vertices has an edge between them in the graph and it is 0 otherwise; the group G is then the natural embedding of $\mathrm{Sym}(n)$ inside $\mathrm{Sym}(k)$. In particular, other variations of GIP still reduce to SIP, for example considering directed graphs (for which $k = n^2 - n$) or coloured edges (for which $|\Sigma| > 2$).

We remark that, in the algorithmic context in which we find ourselves now, when we say that a group G or a coset $G\tau$ thereof is "given" at the input (or "found" at the output) we mean that what we are given or what we find is in fact an explicit set of generators for G, and additionally an element τ for S. Consequently, the complexity of our routines may be affected by the choice (see also the observations after Proposition 3.3.3).

Small cases of GIP have been known for decades: that the isomorphism set of trees is computable in polynomial time is a result by Zemlyachenko [Zem70], and a polynomial-time algorithm for planar graphs has been given by Hopcroft and Tarjan [HT71] building on Weinberg [Wei66]. One of the oldest algorithms to treat GIP from a group-theoretic point of view is contained in [Bab79]: it gives a Las Vegas polynomial time for the isomorphism problem on vertex-coloured graphs having only a bounded number of vertices of any given colour. On the other hand, in 1980 no algorithm for the general GIP was known to run deterministically even in time $e^{o(n \log n)}$; in fact, nothing essentially more clever than brute force was known at the time (as observed in [Bab80]). The question was: is it possible to give a polynomial-time algorithm for solving GIP?

Today we still do not know, although we are much closer to a positive answer than 40 years ago. Some evidence suggests that, at the very least, GIP is not NP-complete: if that were the case, the polynomial-time hierarchy would collapse to some finite level [GMW86] [Sch88]. It has also been clear for a long time that, if we only ask for probabilistic results, actually almost all graphs are quickly discriminated under very simple algorithms: [BES80] offers one made of just sixteen lines from "Input" to "End", finding whether two graphs Γ_1, Γ_2 are isomorphic for almost all choices of Γ_i, and running in time $O(n^2)$ too for such choices. Moreover, following a preprint version of [BES80], [BK79] gave a procedure with $O(n^2)$ expected time for all pairs of graphs[11].

Going back to the actual runtime of GIP, the first important step was Luks's polynomial-time algorithm for graphs of bounded degree [Luk82]: not only it uses techniques upon which even modern results build, like efficient algorithms for working with permutation groups [FHL80] derived from the Schreier-Sims algorithm [Sch27] [Sim67] [Luk82, §1.2], but combined with a valence reduction technique by Zemlyachenko [ZKT85] it also yields an algorithm for GIP that runs in time $e^{O(\sqrt{n \log n})}$. To be more precise, if we have graphs whose vertices have all degree d, Luks's algorithm takes time $n^{O(d/\log d)}$ and the GIP algorithm based upon it takes time $e^{O(\sqrt{n \log n / \log d})}$; then we can use this case to deal with general graphs (see [BKL83, §9] for details). Both SIP and the related coset intersection problem are

[11] [BK79, Thm. 1.2] says "linear expected time", followed by the remarkable line "Linear means $O(n^2)$": the truth is, graphs are represented by inputs of size $\Omega(n^2)$, be they strings as described before or adjacency matrices, therefore $O(n^2)$ is indeed linear in the input. These nomenclature problems obviously do not arise when we say "polynomial" or "quasipolynomial".

also solvable in the same time, as proved by using some additional group-theoretic reasoning (see [Bab83] and [BKL83, §9-10]).

The algorithm for GIP derived from Luks was undefeated until 2015, when Sun and Wilmes classified coherent configurations with large automorphism groups and managed to prove that uniprimitive permutation groups are of size at most $e^{O(n^{1/3}\log^C n)}$ for some absolute C except for a family of known exceptions ([SW16], see the extended abstract in [SW15]). The proof is purely combinatorial and is interesting on its own, as it gives a CFSG-free upper bound on such a size: [Bab81] and [Bab82] had shown it with $\frac{1}{2}$ instead of $\frac{1}{3}$, and further improvements to Sun-Wilmes go through Cameron's theorem (see §1.2); moreover, we can use this to produce a GIP algorithm that also runs in time $e^{O(n^{1/3}\log^C n)}$. Results on primitive permutation groups, in particular Cameron's classification and similar ones, are very useful for GIP.

Then, an algorithm for GIP (and in fact for SIP) with quasipolynomial runtime, i.e. $e^{O(\log^C n)}$ for some absolute C, came at the hands of Babai ([Bab16a], see the extended abstract in [Bab16b]). This marked a stark improvement over the previous record by Sun and Wilmes, and it appeared only weeks later, as recognized in [Bab16a, Rem. 6.1.3] and [Bab16a, Acknowl.].

Theorem 1.5.1. *There is an algorithm that, given two strings $\mathbf{x}, \mathbf{y}$ of length n and a group $G \leq \mathrm{Sym}(n)$, outputs the set $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ of isomorphisms from $\mathbf{x}$ to $\mathbf{y}$ in time $e^{O(\log^C n)}$, where $C > 0$ is an absolute constant.*

Shortly after, Helfgott showed that with minor modifications one can take $C = 3$ in the aforementioned runtime bound [Hel19b] [HBD17]; we will also show in §3 that the implicit constant in the big O notation can be taken to be $103 + o(1)$ (Theorem 3.2.1).

We will see Babai's algorithm (in Helfgott's formulation) in more detail in §1.6 and §3; for now, we observe that it draws from both Luks and Sun-Wilmes, and that it depends on CFSG, mostly because the reduction process on which it is based goes through Cameron's classification. However, a CFSG-free version of Babai's algorithm exists: Babai himself observed how to replace Cameron [Bab16a, §13.1] [Hel19b, §3.1] with a procedure that descends from a CFSG-free classification of 2-transitive groups due to Pyber [Pyb93]; then, Pyber [Pyb16] [Hel19b, §4.1] removed the last *theoretical* dependence on CFSG (meaning that the algorithm did not change, but its runtime had not yet been assured to be quasipolynomial without the use of CFSG). In its CFSG-free version we have to ask for $C = 6+\varepsilon$ for any $\varepsilon > 0$: the algorithm has practically only rearranged its subroutines without any serious modification (see §3.5.2 especially), but the theoretical independence from CFSG makes us lose efficiency. The implicit constant becomes $26e^{1/\varepsilon^2} + o(1)$ for ε small enough.

Among the developments that followed Babai's result, we highlight that his techniques can be used to improve the algorithm for graphs of degree at most d to achieve a runtime of $n^{O(\log^C d)}$ [GNS18]. Also, with some extra arguments, Babai [Bab19] seems to have achieved a quasipolynomial-time algorithm for the graph (and string) *canonization* problem, i.e. for the problem of choosing a canonical

representative in the isomorphism class of a given graph: GIP reduces to this problem very quickly, because once we have found the two canonical representatives we can just check whether they are equal or not instead of isomorphic ([BES80] and [BK79], which we have mentioned before, actually feature canonization algorithms), but it is not known if the inverse reduction also holds. In any case, historically speaking a canonization algorithm has always quickly followed its GIP counterpart (for instance [BKL80] followed [Bab79] [FHL80], and [FSS83] [BL83] followed [Luk82]).

1.6 Babai's algorithm

Here we offer a brief description of Babai's quasipolynomial algorithm for GIP; this will be useful as a bare-bones reference for the work in §3. Rather than Babai's original work [Bab16a], we follow Helfgott's version [Hel19b] (and its English translation [HBD17]).

We remind that the algorithm in fact solves SIP, which in turn yields GIP as a particular case. We start with a group $G \leq \mathrm{Sym}(n)$ and two strings $\mathbf{x}, \mathbf{y} : [n] \to \Sigma$. At any time, polynomial-time procedures allow us to gain some insight into the structure of G, particularly its size and its systems of orbits and blocks. Group-theoretic arguments known since Luks [Luk82] let us reduce to the case of G transitive (Proposition 3.5.2) and to its primitive action on a minimal system of blocks, up to quotienting by the system stabilizer.

Now, the classification theorems kick in [Hel19b, §3.1]. Using Cameron (Theorem 1.2.5, or the handier Theorem 3.5.6), we know that either G is small enough to be treated case-by-case (Proposition 3.5.3) or it acts on the largest blocks as $\mathrm{Alt}(m)$ acts on k-subsets of $[m]$: the latter case, which entails a situation of considerable symmetry, was the bottleneck on which research was stuck before Babai's breakthrough (as pointed out in [Bab16a, §1.1.3]). In the CFSG-free algorithm, we similarly reduce ourselves to either a small case or a giant case or a not-doubly-transitive case (Theorem 3.5.14).

At this point, the novelties introduced by Babai intervene in the process. First is the method of *local certificates* [Bab16a, §10.1] [Hel19b, §6.1]: very succintly, if G modulo the system stabilizer acts like $\mathrm{Alt}(m)$, we consider all sets $T \subseteq [m]$ of size t (for some fixed t) and manage to determine whether the set of automorphisms of $\mathbf{x}$ inside G that preserve T contains $\mathrm{Alt}(T)$ or not. The theorems invoked here to make the procedure work require t to be large enough with respect to n (at least $\Omega(\log n)$ for the original algorithm, $\Omega(\log^5 n)$ for the CFSG-free version), and the cost involved is at least quasipolynomial also for this reason. Then we gather all the pieces T that are guaranteed to give $\mathrm{Alt}(T)$ inside the automorphisms of $\mathbf{x}$ (we *aggregate the certificates* [Bab16a, §10.2] [Hel19b, §6.2]): if the resulting group acts as an alternating group on a large chunk of $[m]$, this alternating piece will be part of the automorphisms of $\mathbf{x}$ themselves[12] and we can completely extract it and

[12] That we gain this information not on G, but on the actual set of automorphisms of $\mathbf{x}$ inside G (which we do not know at all, and in fact is the objective of the whole algorithm), is astonishing to the author of these lines and is the reason why everything we are doing works out.

put it aside; if not, up to fixing some arbitrary logarithmic choice of elements of m (which also entails a quasipolynomial cost, see Remark 3.5.10), we can recover enough asymmetry to colour tuples of elements of $[m]$ in different ways according to orbits they fall into under the action of the automorphisms of $\mathbf{x}$.

Then, it is time for the second new piece in Babai's algorithm, the *Split-or-Johnson* routine [Bab16a, §7] [Hel19b, §5]. Take the coloured tuples of elements of $[m]$ as described before: to start, we can apply a well-known algorithm due to Weisfeiler and Leman [WL68] that splits canonically the elements of $[m]$ (thus bringing us back to the intransitive case) unless some deep symmetry still exists inside a large part of $[m]$, thus forming a *coherent configuration* (see Definition 2.1.3). If such symmetry indeed exists, a routine that involves another logarithmic choice of elements either partitions the coherent configuration or finds an even more rigid structure in it, a Johnson scheme: in the first case, we have found an even coarser system of blocks on which we can act with G; in the second, by the definition of Johnson scheme, the surjection of G onto $\mathrm{Alt}(m)$ reduces to a surjection onto a certain $\mathrm{Alt}(m')$ with $m' = O(\sqrt{m})$. Both processes can be performed only a logarithmic number of times, so that reduction must eventually occur in some other form.

The cost analysis is performed in [Bab16a, §11.2], then more precisely in [Hel19b, App. A], and even more so in §3.6. It seems clear that, given the types of recursion and the group-theoretic tools required in the algorithm, such as for example the local certificate procedure whose consequences we have sketched, some essentially new idea would be necessary to solve GIP in polynomial time. However, it might still be feasible to tweak the present algorithm to make it work in time $e^{O(\log^2 n)}$ instead of $e^{O(\log^3 n)}$.

Of the subroutines involved in the procedure, the Weisfeiler-Leman algorithm [WL68] ranks among the most interesting and most studied. It is an old and widely used algorithm designed (in the context we need) to provide a colouring that encodes structural information about the initial graph, or more generally the initial collection of coloured tuples of a set: furthermore it runs in polynomial time in its classical form, and the same goes for its k-ary generalizations, where however the exponent in the polynomial time depends on k. The Weisfeiler-Leman algorithm is not sufficient to crack GIP on its own (see [CFI92] [EP99] for theoretical discussions, and [Shr59] [KZA17] for small examples), nevertheless it has ample applications even today, enough to warrant a dedicated 50th anniversary conference; see [Gro17, §3.5] for an introduction, and [EP09] for a modern survey on the topic of coherent configurations. In §2, we will provide a connection between the number of iterations of the Weisfeiler-Leman algorithm and the diameter of Cayley and Schreier graphs (published as [Don19c]).

Finally, let us draw another connection between GIP and the diameter problem, which we will explore more deeply in §6. The group-theoretic classifications that provide the core engine necessary to make Babai's algorithm work, namely Cameron's theorem in the original version and Pyber's in the CFSG-free one, are reduction tools that can potentially be plugged in other contexts where a descent of the same kind can become effective. An example of this versatility at work is the proof of the diameter bound for $\mathrm{Alt}(n)$ contained in [Hel18]: in it, Cameron's

theorem is used to justify the statement that at every step the quotients of a composition series we want to work with are either small groups or alternating groups, and then other arguments intervene to control the overall diameter by the diameter of the factors (see §6.1 and [Hel18, §4]); also, this is the only real dependence on CFSG of the whole proof. It may be possible then to replace this piece with the equivalent procedure in the CFSG-free GIP algorithm: as one can observe, the descent that the new process entails is compatible with our needs in Helfgott's proof, and the fact that the final diameter bound is not affected too deeply by the substitution is essentially a consequence of the cost analysis of Babai's algorithm.

The one decisive drawback is that the subgroups involved in the descent are not normal, as they would be under Cameron, so that the machinery involved in [Hel18, §4] is not suitable anymore[13]. In §6.3 we partially fix the problem, and show that under some hypothesis (Conjecture 6.3.4) we can retrieve a diameter bound of $e^{e^{\frac{1}{\log 2}(\log\log n)^2}}$, which is worse than [Hel18, Thm. 6.1] and (1.3.4) but better than the currently known CFSG-free bounds in (1.3.1) and (1.3.5). The author hopes that §6 will point towards a new direction of investigation in the search for CFSG-free diameter bounds, even if he himself got only halfway there.

[13]In particular, it also means that the CFSG-free workaround replacing Cameron in Babai's algorithm cannot work as a CFSG-free version of Cameron's classification in a more general sense. The author thanks L. Pyber for a private communication that made him realize the importance of this fact. The author also takes this occasion to point out that in the past he had written down an arXiv preprint announcing a CFSG-free diameter bound for $\mathrm{Alt}(n)$ (in consequence of which the aforementioned communication took place): the proof was wrong, being based on a mistake in the original analysis performed here in §3, and the preprint was withdrawn. The correction resulted in §6, with a worse bound and depending on a conjecture: alas, this is also mathematics.

Chapter 2

The Weisfeiler-Leman algorithm and the diameter of Schreier graphs

The content of this chapter is essentially taken from [Don19c].

In this chapter, we work with Schreier graphs (see Definition 2.1.1) of finite groups with the natural colouring given by the set of generators from which they are defined, akin to the labelling discussed in §1.1 for Cayley graphs. To them, we apply the Weisfeiler-Leman algorithm [WL68] (see §1.6), which refines their colourings and encodes information about the structure of the graph itself in the new colouring: as it turns out, the number of iterations taken before stopping is tightly related to the diameter of the graphs themselves.

An upper bound for the number of iterations is found in the case of general Schreier graphs: this is the content of Theorem 2.1.6. A lower bound also holds for some interesting particular cases, such as for Schreier graphs with $G = \mathrm{SL}_n(\mathbb{F}_q)$ (with $q > 2$) acting on k-tuples of vectors in $\mathbb{F}_q^n$: the result is expressed in Theorem 2.1.7. We underline that the upper bound depends only on the diameter, and is independent from the group and the set of generators it is defined from; the lower bound is similarly independent from n and q, as long as $\mathrm{char}(\mathbb{F}_q) > 2$.

In the case of Cayley graphs, instead of bounds we will be able to find an exact expression for the number of iterations as a function of the diameter: see Theorem 2.1.8.

2.1 Introduction

Let G be a finite group and let S be a set of generators of G such that $S = S^{-1}$ and $e \in S$: the *Cayley graph* $\mathrm{Cay}(G, S)$ is defined as the graph having G as its set of vertices and $\{(g, sg) | g \in G, s \in S\}$ as its set of edges (see Definition 1.1.1). We have already introduced Cayley graphs in §1.1 and discussed some very basic characteristics. Unlike what we said therein, in this chapter we will consider often

edges as being directed: although it does not make any practical difference thanks to our choice of S symmetric, in our discussion we will colour *ordered* pairs of vertices, so our mindset should be the one for directed graphs.

Cayley graphs are special cases of a more general class of graphs that we have mentioned a couple of times. As a matter of fact, we will define them as directed multigraphs (i.e. directed graphs that may have multiple edges starting and ending in the same vertices).

Definition 2.1.1. *Let G be a finite group, and let S be a set of generators of G; let V be a set on which G acts transitively (on the left). The Schreier graph* $\mathrm{Sch}(V, S)$ *is the graph having V as its set of vertices and $\{(v, sv) | v \in V, s \in S\}$ as its set of edges.*

As in Definition 1.1.1, acting on the left is just a convention: as long as the two definitions agree, we are fine. A Schreier graph depends of course also on G and its action, which is not reflected in the notation, but we assume that it is implicit in the choice of working with V; we will never use two different groups in the same context, so there is no risk of confusion.

Since S generates G and G acts transitively on V, the graph is strongly connected: if $gv = v'$ for some $g \in G$, there exists a directed path from the vertex v to the vertex v' determined by those $s_i \in S$ such that $s_1 s_2 \dots s_m = g$. A Cayley graph is just a Schreier graph where $V = G$ and the action is the usual group multiplication.

The construction of these graphs and the choice of a symmetric set of generators containing e allow us to see $\mathrm{Cay}(G, S)$ and $\mathrm{Sch}(V, S)$ as a different type of structure.

Definition 2.1.2. *A (classical) configuration $\mathfrak{X}$ is a pair $(\Gamma, c : \Gamma^2 \to \mathcal{C})$ (where Γ is a finite set of vertices and $\mathcal{C}$ is a finite set of colours) with the following properties:*

(i) for any $c \in \mathcal{C}$, if for some $v \in \Gamma$ we have $c(v, v) = c$, then for all $v_1, v_2 \in \Gamma$ such that $c(v_1, v_2) = c$ we have $v_1 = v_2$;

(ii) for any $c \in \mathcal{C}$ there exists a $c^{-1} \in \mathcal{C}$ such that for any $v_1, v_2 \in \Gamma$ with $c(v_1, v_2) = c$ we have $c(v_2, v_1) = c^{-1}$.

The addendum "classical" comes from the fact that a more general definition is often used, where the colouring is $c : \Gamma^k \to \mathcal{C}$ and consequently with some differences in how to define conditions (i) and (ii). It is also to be noted that this is a "weak" version of the definition of configuration, as provided in [SW16, §1] and [Bab16a, Def. 2.3.4], as opposed to the "strong" version that can be found in [Hel19b, Déf. 2.5]: in that paper, it is also required inside condition (ii) that there exist c', c'' such that $c(v_1, v_1) = c'$ and $c(v_2, v_2) = c''$, which was needed to prove properties of non-classical configurations that here are not needed (see [Hel19b, Ex. 2.7]). By property (i), in a configuration we can distinguish between *vertex colours* (colours coming from $c(v, v)$) and *edge colours* (colours coming from $c(v_1, v_2)$ with $v_1 \neq v_2$): these names come from the natural observation that we can think of a configuration as a particular colouring of the complete graph of $|\Gamma|$

vertices, giving to v the colour $c(v,v)$ and to the directed edge (v_1,v_2) the colour $c(v_1,v_2)$ and noticing that by condition (i) a vertex and an edge will always have different colours in this situation.

There is a natural way to define a configuration $\mathfrak{X}_C$ from $\mathrm{Cay}(G,S)$: Γ can be chosen to be the group G, while the colouring is given by $c(g_1,g_2)=g_2g_1^{-1}$ if $g_2g_1^{-1}\in S$ and $c(g_1,g_2)=\emptyset$ otherwise; in this case then $\mathcal{C}=S\cup\{\emptyset\}$ (or $\mathcal{C}=S$ in the trivial case $S=G$). $\mathfrak{X}_C$ is a configuration (in the weak sense): thanks to $e\in S$ the only vertex colour is e and all the others are edge colours, while thanks to $S=S^{-1}$ the inverse of $s\in S$ as a colour is exactly s^{-1} (and the inverse of $\emptyset$ is $\emptyset$); since we have only one vertex colour $\mathfrak{X}_C$ is also a configuration in the strong sense, but it does not make any difference. Notice the similarity between the colouring and the edge labelling naturally defined on Cayley graphs (see §1.1): all edges have as colour exactly the label s, and we add two more colours, e and $\emptyset$, to cover the rest of the pairs $(g_1,g_2)\in G^2$.

In a similar fashion we can define a configuration $\mathfrak{X}_S$ from $\mathrm{Sch}(V,S)$: Γ can be defined to be the set V, while the colouring is given by $c(v_1,v_2)=\{s\in S|sv_1=v_2\}$; in this case then $\mathcal{C}\subseteq\mathcal{P}(S)$. $\mathfrak{X}_S$ is a classical configuration (in the weak sense, but not in the strong sense): to prove that it satisfies (i), notice that if for a colour c we have $c(v,v)=c$ then $e\in c$, so that for any other two vertices v_1,v_2 with $c(v_1,v_2)=c$ we have $v_1=v_2$ (in other words, vertex colours are exactly those who contain e); to prove that it satisfies (ii), observe that for all $c\in\mathcal{C}$ we have a natural definition $c^{-1}=\{s^{-1}|s\in c\}$, thanks to $S=S^{-1}$. If we see Cayley graphs as particular Schreier graphs, the configurations $\mathfrak{X}_C$ and $\mathfrak{X}_S$ built on the same $\mathrm{Cay}(G,S)$ are clearly isomorphic, with each colour $s\neq\emptyset$ in $\mathfrak{X}_C$ corresponding to $\{s\}$ in $\mathfrak{X}_S$.

As mentioned before, we now introduce a more refined type of structure.

Definition 2.1.3. *A (classical) coherent configuration is a pair $\mathfrak{X}=(\Gamma,c:\Gamma^2\to\mathcal{C})$ that satisfies* (i) *and* (ii) *and such that*

(iii) for every $c_0,c_1,c_2\in\mathcal{C}$ there is a constant $\gamma=\gamma(\mathfrak{X},c_0,c_1,c_2)\in\mathbb{N}$ such that, for every $v_1,v_2\in\Gamma$ with $c(v_1,v_2)=c_0$, the number of $w\in\Gamma$ with $c(v_1,w)=c_1$ and $c(w,v_2)=c_2$ is γ (independently from the choice of v_1,v_2).

The colouring of a coherent configuration contains much more information about its structure than the one coming from a usual configuration. Especially important to us is the following result.

Proposition 2.1.4. *Let $\mathfrak{X}=(\Gamma,c:\Gamma^2\to\mathcal{C})$ be a coherent configuration, and let $c_0,c_1,\dots,c_k$ be a sequence of colours with $k\geq 2$. Then there is a constant $\gamma=\gamma(c_0,c_1,\dots,c_k)\in\mathbb{N}$ such that for every $v_1,v_2\in\Gamma$ with $c(v_1,v_2)=c_0$ the number of k-tuples $(w_1,\dots,w_{k-1})\in\Gamma^k$ with $c(v_1,w_1)=c_1$, $c(w_{i-1},w_i)=c_i$ for all $1<i<k$ and $c(w_{k-1},v_2)=c_k$ is γ (independently from the choice of v_1,v_2).*

So an edge colour $c(v_1,v_2)$ in a coherent configuration not only knows by definition about colourings of triangles v_1,w,v_2, but knows also about colourings of walks $v_1,w_1,\dots,w_{k-1},v_2$ of any length.

Proof. This is [Hel19b, Ex. 2.16(a)]; we give the same proof as in [HBD17, App. B]. We proceed by induction on k: for $k = 2$ the statement is exactly condition (iii) of coherent configurations, so there is nothing to prove.

Suppose now that this is true for k. We are given v_1, v_2 with $c(v_1, v_2) = c_0$ and we have to find the number of walks of colours $c_1, c_2, \dots, c_{k+1}$ from v_1 to v_2: such a walk however is merely the composition of two walks $c_1, c_2, \dots, c_{k-1}$ and c_k, c_{k+1}, so we can just consider any walk $c_1, c_2, \dots, c_{k-1}, c'$ of length k from v_1 to v_2 (for all c') and for each of them any triangle of colours c', c_k, c_{k+1} built on (w_{k-2}, v_2) of colour c'; the composition of these two structures will give us the desired walk of length $k+1$. The constants γ for walks of length 2 and k are independent from the choice of initial vertices, thus the same will occur for $k+1$: we have

$$\gamma(c_0, c_1, \dots, c_k, c_{k+1}) = \sum_{c' \in \mathcal{C}} \gamma(c_0, c_1, \dots, c_{k-1}, c')\gamma(c', c_k, c_{k+1}),$$

and the inductive step is complete. □

A configuration, and in particular the configurations $\mathfrak{X}_C, \mathfrak{X}_S$ that we are going to study, is not necessarily coherent. There is a natural way to refine a configuration into a coherent configuration, through the Weisfeiler-Leman algorithm, which is given as follows.

(1) At the 0-th iteration, we define $\mathcal{C}^{(0)} = \mathcal{C}$.

(2) If $\mathcal{C}^{(h)}$ is the colouring at the h-th iteration, we can define $\mathcal{C}^{(h+1)}$ by calling $c^{(h+1)}(v_1, v_2)$ the tuple

$$\left(c^{(h)}(v_1, v_2), \left(\left| \left\{ w \in V | c^{(h)}(v_1, w) = c_1, c^{(h)}(w, v_2) = c_2 \right\} \right| \right)_{c_1, c_2 \in \mathcal{C}^{(h)}} \right). \tag{2.1.1}$$

In truth, in practical applications we do not actually define $\mathcal{C}^{(h+1)}$ using the whole (2.1.1): it is sufficient to give the pairs (v_1, v_2) colours that are different if and only if their tuples (2.1.1) are different, so as not to make the algorithm's required runtime and memory space blow up. Obviously there can only be at most $|V|^2$ colours, so we are safe on that front.

(3) If we reach an iteration where there is no refinement, meaning that for a certain h every time that $c^{(h+1)}(v_1, v_2) \neq c^{(h+1)}(v_3, v_4)$ we also had $c^{(h)}(v_1, v_2) \neq c^{(h)}(v_3, v_4)$, the Weisfeiler-Leman algorithm stops.

The colouring $c^{(h+1)}$ is more refined than $c^{(h)}$; notice that $c^{(h+1)}(v_1, v_2)$ contains as information, for each choice of $c_1, c_2 \in \mathcal{C}^{(h)}$, the number of vertices w as in condition (i) (which is not yet independent from the choice of v_1, v_2). Once the Weisfeiler-Leman algorithm stops, it means that all these numbers are the same for each pair (v_1, v_2) with the same colour, i.e. the configuration has become coherent. That we must stop eventually is clear, since we can refine a colouring only as many times as the total number of pairs of vertices.

One last observation is necessary with regard to Weisfeiler-Leman.

Proposition 2.1.5. *Define $\mathfrak{X}^{(h)}$ as the configuration at the h-th step of Weisfeiler-Leman. Then* $\mathrm{Aut}(\mathfrak{X}^{(h)}) = \mathrm{Aut}(\mathfrak{X}^{(h+1)})$.

Proof. As $c^{(h+1)}$ is a refinement of $c^{(h)}$, we have already $\mathrm{Aut}(\mathfrak{X}^{(h)}) \supseteq \mathrm{Aut}(\mathfrak{X}^{(h+1)})$. On the other side, if $\sigma \in \mathrm{Aut}(\mathfrak{X}^{(h)})$, then for any pair $v_1, v_2 \in V$ and any pair $c_1, c_2 \in \mathcal{C}^{(h)}$ each vertex w with $(c(v_1, w), c(w, v_2)) = (c_1, c_2)$ is sent to a vertex $\sigma(w)$ such that $(c(\sigma(v_1), \sigma(w)), c(\sigma(w), \sigma(v_2))) = (c_1, c_2)$: this implies that the numbers in (2.1.1) are also preserved by σ, therefore $\sigma \in \mathrm{Aut}(\mathfrak{X}^{(h+1)})$ too. □

This explains why Weisfeiler-Leman is so interesting in the context of GIP (the *graph isomorphism problem*, see §1.5). After using the algorithm we have a more refined colouring, which means that we have more possibility to exploit the subtle differences between two graphs; at the same time the algorithm is designed to preserve all automorphisms, which in turn implies that all isomorphisms between two graphs are preserved as well: the set of isomorphisms from a graph to another is a coset of the set of automorphisms of the first graph (see for instance Remark 3.3.2), and permuting vertices and their colouring does not affect (2.1.1). Hence, when we prove that a certain bijection σ between the vertices of the graphs is not an isomorphism for the final coherent configurations, we have proved that it is also not an isomorphism for the original graphs.

We state now our main results of the chapter.

Theorem 2.1.6. *Let G be a finite group and let S be a set of generators of G with $e \in S = S^{-1}$. Suppose that G acts transitively on a set V, and consider the configuration $\mathfrak{X}_S$ coming from* $\mathrm{Sch}(V, S)$. *Then the number* $\mathrm{WL}(\mathfrak{X}_S)$ *of nontrivial iterations of the Weisfeiler-Leman algorithm satisfies*

$$\mathrm{WL}(\mathfrak{X}_S) \leq \log_2 \mathrm{diam}(\mathrm{Sch}(V, S)) + 3.$$

By counting nontrivial iterations we merely want to ignore the last one with no colour refinement.

Together with this upper bound, lower bounds also hold in some more limited but still very interesting cases. The scope of the lower bound is explicitly stated later (see Theorem 2.3.1), but here we specialize it to a more interesting group-theoretic situation, which is arguably especially relevant in the context of Babai's conjecture (Conjecture 1.3.1).

Theorem 2.1.7. *Let $G = \mathrm{SL}_n(\mathbb{F}_q)$ with $q > 2$ and let S be a set of generators of G with $\mathrm{Id}_n \in S = S^{-1}$; for any $0 < k < n$, let V be the set of linearly independent k-tuples of vectors of $\mathbb{F}_q^n$, with the action of G on V defined as $A(v_1, \ldots, v_k) = (Av_1, \ldots, Av_k)$. Consider the configuration $\mathfrak{X}_S$ coming from* $\mathrm{Sch}(V, S)$. *Then, if p is the smallest prime such that $p|(q-1)$, the number* $\mathrm{WL}(\mathfrak{X}_S)$ *of nontrivial iterations of the Weisfeiler-Leman algorithm satisfies*

$$\mathrm{WL}(\mathfrak{X}_S) \geq \log_2 \mathrm{diam}(\mathrm{Sch}(V, S)) - \log_2(p-1) - 3.$$

Notice that this result does not depend on n, and we have dependence on q only when $\mathrm{char}(\mathbb{F}_q) = 2$ (since otherwise $p = 2$ and $\log_2(p-1) = 0$).

Finally, we state a result that gives an exact expression for the number of iterations for any Cayley graph.

Theorem 2.1.8. *Let G be a finite group and let S be a set of generators of G with $e \in S = S^{-1}$. Consider the configuration $\mathfrak{X}_C$ coming from* $\mathrm{Cay}(G,S)$*. Then the number* $\mathrm{WL}(\mathfrak{X}_C)$ *of nontrivial iterations of the Weisfeiler-Leman algorithm satisfies*

$$\mathrm{WL}(\mathfrak{X}_C) = \begin{cases} \lceil \log_2(\mathrm{diam}(\mathrm{Cay}(G,S)) - 1) \rceil & \text{if } \forall g\ \exists! g' \text{ with } d(g,g') \text{ the diameter,} \\ \lceil \log_2 \mathrm{diam}(\mathrm{Cay}(G,S)) \rceil & \text{otherwise.} \end{cases}$$

We remark that Theorem 2.1.6 holds in particular for Cayley graphs too, and so does and a bound analogous to the one in Theorem 2.1.7, although as expected they are weaker than the theorem above.

2.2 The upper bound

We first prove the upper bound in Theorem 2.1.6. In fact, we prove the same upper bound for a more general class of configurations, of which the configurations $\mathfrak{X}_C, \mathfrak{X}_S$ that we defined from $\mathrm{Cay}(G,S)$ and $\mathrm{Sch}(V,S)$ are just particular cases.

Theorem 2.2.1. *Let $\mathfrak{X}$ be a configuration; call an edge colour $c \in \mathcal{C}$ nonempty if for every $v \in \Gamma$ there is at most one $w \in \Gamma$ with $c(v,w) = c$, and call it empty otherwise. Suppose that the coloured graph $\Gamma_{\mathfrak{X}} = (\Gamma, \{(v_1,v_2) \in \Gamma^2 | c(v_1,v_2) \text{ nonempty}\})$ is connected. Then the number* $\mathrm{WL}(\mathfrak{X})$ *of nontrivial iterations of the Weisfeiler-Leman algorithm satisfies*

$$\mathrm{WL}(\mathfrak{X}) \leq \log_2 \mathrm{diam}(\Gamma_{\mathfrak{X}}) + 3.$$

Here and everywhere else, $\log_2$ denotes the logarithm in base 2.

Proof that Thm. 2.2.1 $\Rightarrow$ Thm. 2.1.6. Consider the configuration $\mathfrak{X}_S$ coming from the graph $\mathrm{Sch}(V,S)$: the only possible empty colour is $\emptyset$ (hence the name) and all the other edge colours are nonempty, because for any $v \in V$ and $s \in S$ there is evidently only one v' with $sv = v'$, so by our definition for any colour $c \subseteq S$, $c \neq \emptyset$ at most one v' would realize $c(v,v') = c$ for any fixed v. If $\emptyset$ is indeed empty, by construction a nonempty-coloured pair in $\Gamma_{\mathfrak{X}_S}$ corresponds to an edge (or multiedge) in $\mathrm{Sch}(V,S)$: thus connectedness of $\mathrm{Sch}(V,S)$ implies the same for $\Gamma_{\mathfrak{X}_S}$, and the two have the same diameter. If $\emptyset$ is nonempty, i.e. in the extreme case when S sends any v to any v' except at most one, we have $\mathrm{diam}(\Gamma_{\mathfrak{X}_S}) = 1$ and $\mathrm{diam}(\mathrm{Sch}(V,S)) = 2$, so the bound still holds; if $\emptyset$ is not a colour at all, $\mathrm{diam}(\Gamma_{\mathfrak{X}_S}) = \mathrm{diam}(\mathrm{Sch}(V,S)) = 1$. □

From now on, for the sake of clarity, colours that appear during the iterations of the Weisfeiler-Leman algorithm as refinements of empty (resp. nonempty) colours are still called empty (resp. nonempty), even if some new empty colours could now satisfy the nonemptyness criterion: as a matter of fact, it is a key point in the success of the argument that eventually *all* new colours will be nonempty according to our definition of the word; however our need to refer ourselves to the origin of the intermediate colours is more pressing than highlighting the acquisition of the

property of nonemptyness. Observe that, by the construction of the new colours in (2.1.1), each of them contains information about every past colour from which it descended, so that recognizing whether a given intermediate colour is empty or nonempty does not create any problem.

Inside $\mathfrak{X}$, call *walk* (of length l) any sequence of consecutive pairs of vertices $(w_0, w_1), (w_1, w_2), \dots, (w_{l-1}, w_l)$ (with their respective colours) which corresponds to a walk of length l in $\Gamma_{\mathfrak{X}}$; equivalently, a walk in $\mathfrak{X}$ is any sequence of consecutive pairs with only nonempty colours.

Throughout the rest of the paper, we will colloquially say multiple times that a colour c *knows* something. With this expression, we want to convey that for any pair of vertices (v, w) for which $c(v, w) = c$ that particular bit of information is true: in particular, when we say that c knows a certain walk (intended as a sequence of nonempty colours) we mean that for every two vertices v, w with $c(v, w) = c$ there is a walk from v to w made of pairs of vertices having exactly those colours. Moreover, when we say that c knows all walks, or all walks up to a certain length, we mean that for every pair (v, w) with $c(v, w) = c$ walks (or walks of a certain length) from v to w made of a given sequence of colours occur the same number of times regardless of the choice of v, w.

In order to prove Theorem 2.2.1, we need the next lemma.

Lemma 2.2.2. *For every $v_1, v_2 \in \Gamma$, at the k-th iteration, $c^{(k)}(v_1, v_2)$ knows all walks of length $\leq 2^k$ from v_1 to v_2.*

Proof. We proceed by induction on k. When $k = 0$ the statement is trivial, since the only walk of length 1 from v_1 to v_2 that could possibly exist is the edge (v_1, v_2) provided that its colour is nonempty, which the colour $c^{(0)}(v_1, v_2)$ evidently knows.

Suppose that the statement has been proved for k, and consider any walk of length $\leq 2^{k+1}$ from v_1 to v_2: for any such walk, there exists a w that splits the original walk into two walks (from v_1 to w and from w to v_2) of length $\leq 2^k$. The existence of these walks is an information contained inside $c^{(k)}(v_1, w)$ and $c^{(k)}(w, v_2)$ respectively, so at the $(k+1)$-th iteration $c^{(k+1)}(v_1, v_2)$ will know about the existence of this pair of walks (and consequently of the original long walk). □

Compare the statement of Lemma 2.2.2 with that of Proposition 2.1.4: according to the latter, colours of a coherent configuration know all walks of any length, while the former describes how, iteration after iteration, the colours of a configuration arrive to gain knowledge of longer and longer walks, whose length doubles at every step. The origin of the $\log_2$ of the diameter in all our results clearly resides in this "learning" process; proving the lower and the upper bounds involves making sure that such process is, in a sense, more or less respectively necessary and sufficient to make the configuration $\mathfrak{X}$ coherent.

As for the upper bound, it turns out that, given a vertex v, knowing which walks starting from v reach the same endpoint is enough information (in a graph like $\Gamma_{\mathfrak{X}}$) to reconstruct a piece of the graph around v.

Lemma 2.2.3. *Fix any vertex $v \in \Gamma_{\mathfrak{X}}$ and suppose that we know, for any two sequences of nonempty colours of length $\leq k$, whether the walks starting from v defined by these sequences exist and have the same endpoint. Then it is possible*

to reconstruct in a unique way the subgraph $\Gamma^{(k)}(v) \subseteq \Gamma_{\mathfrak{X}}$ *given by the edges at distance* $\leq k$ *from* v*. In other words, for any other coloured graph* Γ' *whose walks of length* $\leq k$ *from a certain vertex* w *satisfy the same conditions of existence and equality of endpoints, there exists a unique graph isomorphism* $\Gamma^{(k)}(v) \to \Gamma'^{(k)}(w)$ *sending* v *to* w*.*

Proof. The information provided to us is the following:

- a collection of statements of the form "the walk from v consisting of consecutive edges of colour $c_1, c_2, \ldots, c_l$ exists", for some strings of nonempty colours $c_1c_2 \ldots c_{l\leq k}$;
- a collection of statements of the form "the two walks from v consisting of consecutive edges of colour $c_1, c_2, \ldots, c_l$ and $c'_1, c'_2, \ldots, c'_{l'}$ end in the same v'", for some pairs of strings of nonempty colours $(c_1c_2 \ldots c_{l\leq k}, c'_1c'_2 \ldots c'_{l'\leq k})$.

With these strings and pairs of strings in our hands, we will manage to rebuild what the graph $\Gamma_{\mathfrak{X}}$ looks like up to distance k.

We proceed by induction on k. For $k = 1$, walks of length ≤ 1 are just single edges: there is no possible equality of two endpoints, so $\Gamma^{(1)}(v)$ is just a star whose internal vertex is v and whose leaves are all the edges (v, v_i) of $\Gamma_{\mathfrak{X}}$. There will obviously be then a unique isomorphism to any other star with internal vertex w and leaves with the same colour as $\Gamma^{(1)}(v)$.

Suppose now that the statement is true for k, and consider for $k + 1$ the two graphs $\Gamma^{(k+1)}(v), \Gamma'$: we already have a partial isomorphism from $\Gamma^{(k)}(v) \subseteq \Gamma^{(k+1)}(v)$ to the subgraph given by the edges at distance at most k from w in Γ', and we just need to extend this isomorphism to the edges and vertices at distance $k+1$. The already existing isomorphism already covers walks of length $\leq k$, so we need to consider only relations involving a walk that includes an edge at distance $k + 1$ from v or w (which will necessarily be the last edge).

Each edge at distance $k + 1$ in $\Gamma^{(k+1)}(v)$ is (v_1, v_2) where either $d(v, v_1) = d(v, v_2) = k$ or $d(v, v_1) = k$ and $d(v, v_2) = k + 1$. In the first case, the two vertices v_1, v_2 are already sent in a unique way to $w_1, w_2 \in \Gamma'$. Moreover, for any string $c_1c_2 \ldots c_kc_{k+1}$ whose last edge would be (v_1, v_2) in $\Gamma_{\mathfrak{X}}$ there are pairs of the type $(c_1c_2 \ldots c_kc_{k+1}, c'_1c'_2 \ldots c'_k)$ (where the second corresponds to a walk of length k from v to v_2). Inside Γ', the walks $c_1c_2 \ldots c_k, c'_1c'_2 \ldots c'_k$ of length k have endpoints w_1 and w_2, so the information gathered from the pairs above is: "the edge starting from w_1 of colour c_{k+1} ends in w_2". This implies that the edge (v_1, v_2) has the same colour as the edge (w_1, w_2), and the isomorphism can be extended to all such edges.

In the second case, consider all pairs $(c_1c_2 \ldots c_kc_{k+1}, c'_1c'_2 \ldots c'_kc'_{k+1})$ of walks of length $k + 1$ whose last edge (v_1, v_2) have $d(v, v_1) = k$ and $d(v, v_2) = k + 1$. We can group these walks into equivalence classes corresponding to their endpoints v_2 inside $\Gamma^{(k+1)}(v) \setminus \Gamma^{(k)}(v)$. For all strings $c^i_1c^i_2 \ldots c^i_kc^i_{k+1}$ inside a given class, their subwalks of length k end in vertices $v^i_1 \in \Gamma_{\mathfrak{X}}$, which uniquely correspond to vertices $w^i_1 \in \Gamma'$: therefore the information they carry, when seen in Γ', signifies that "the edges starting from w^i_1 of colour c^i_{k+1} all end in the same vertex" (say w_2): we extend then the isomorphism by sending each v_2 to the corresponding w_2

and each (v_1^i, v_2) to the corresponding (w_1^i, w_2). This covers all vertices at distance $k+1$ and all edges of the second type. The isomorphism has been extended in a unique way to the whole $\Gamma^{(k+1)}(v)$, so we are done. □

Notice how important it is that we are working with nonempty colours, i.e. in a situation where every vertex has at most one adjacent edge of any given colour: only with this condition we can talk about *the* edge starting from v and having colour c, or about *the* walk starting from v and having colours $c_1, c_2, \dots, c_l$. Without it, there would be no guarantee that there is an isomorphism as the one we constructed, i.e. we would not be able to deduce the shape of the subgraph $\Gamma^{(k)}(v)$ only by looking at the information about the walks, because more than one graph could satisfy the same conditions: this is crucial in the proof of the upper bound.

We are now ready to prove Theorem 2.2.1.

Proof of Thm. 2.2.1. Define $k = \lceil \log_2(\mathrm{diam}(\Gamma_{\mathfrak{X}}) + 1) \rceil$: for any $v_1, v_2 \in \Gamma_{\mathfrak{X}}$, the colour $c^{(k)}(v_1, v_2)$ knows all walks of length $\leq 2^k$ from v_1 to v_2 by Lemma 2.2.2. Since $d(v_1, v_2) \leq \mathrm{diam}(\Gamma_{\mathfrak{X}}) < 2^k$, there exists at least one of these walks: this implies that we have $c^{(k)}(v_1, v_2) \neq c^{(k)}(v_1, v_2')$ for any $v_2 \neq v_2'$, because the walk from a given vertex defined by a given sequence of nonempty colours is unique (if it exists).

At the next iteration, the colour $c^{(k+1)}(v_1, v_1)$ knows the number of vertices w such that $(c^{(k)}(v_1, w), c^{(k)}(w, v_1)) = (c_1, c_2)$ for any choice of colours $c_1, c_2 \in \mathcal{C}^{(k)}$, so in particular it knows whether there is a colour c_1 containing one or two given walks (with the expression "the colour c_1 contains a given walk" we mean that two vertices v, v' with $c(v, v') = c_1$ would have this walk going from v to v') such that there is one w with $c^{(k)}(v_1, w) = c_1$: in other words, the colour $c^{(k+1)}(v_1, v_1)$ knows if a walk starting from v_1 of length $\leq 2^k$ given by a certain sequence of nonempty colours exists, thanks to the information about colours containing one given walk, and it knows also if two walks starting from v_1 of length $\leq 2^k$ defined by given sequences of nonempty colours have the same endpoint, thanks to the information about colours containing two given walks. So we are in the situation described in the statement of Lemma 2.2.3, where at the $(k+1)$-th iteration from the colour $c^{(k+1)}(v_1, v_1)$ we can reconstruct the subgraph $\Gamma_{\mathfrak{X}}^{(2^k)}(v_1)$.

Now, $\Gamma_{\mathfrak{X}}^{(2^k)}(v_1)$ is actually the whole $\Gamma_{\mathfrak{X}}$ for any v_1, since every edge is at distance at most $\mathrm{diam}(\Gamma_{\mathfrak{X}}) + 1$ from v_1: this means that at the $(k+1)$-th iteration we can look at any colour of any vertex and reconstruct the whole graph $\Gamma_{\mathfrak{X}}$ around it. A consequence of the above is that for any two vertices that still have the same colour after $k+1$ iterations there must be an automorphism of $\Gamma_{\mathfrak{X}}$ sending one to the other. We still need to ensure the same for the edges.

After one more iteration, also each pair of distinct vertices (v_1, v_2) will have a colour capable of reconstructing the subgraph $\Gamma_{\mathfrak{X}}^{(2^k)}(v_1)$, because $c^{(k+2)}(v_1, v_2)$ will in turn contain $c^{(k+1)}(v_1, v_1)$ in a way that allows us to identify it unequivocally: remember, in a configuration vertex colours and edge colours are distinct, so $c^{(k+1)}(v_1, v_1)$ can be defined as the unique vertex colour c such that the number of w with $\big(c^{(k+1)}(v_1, w), c^{(k+1)}(w, v_2)\big) = \big(c, c^{(k+1)}(v_1, v_2)\big)$ is nonzero. Therefore,

at the $(k+2)$-th iteration two vertices or two edges with the same colour can be sent to each other via an automorphism of $\Gamma_{\mathfrak{X}}$.

We still have to check that this extends to an automorphism of $\mathfrak{X}$. We remark that this passage is not necessary to prove the special cases $\mathfrak{X}_C, \mathfrak{X}_S$ that we are interested in, since in that case there is at most one empty colour $\emptyset$ and any automorphism preserving all nonempty colours (i.e. an automorphism of $\Gamma_{\mathfrak{X}}$) would preserve all colours (i.e. it would be an automorphism of $\mathfrak{X}$).

For any pair (v_1, v_2) and any $c(v, w)$ empty, there is a sequence of consecutive pairs that goes first through a walk from v_1 to v, then includes the pair (v, w), then through another walk goes from w to v_2; if $c^{(k+2)}(v_1, v_2) = c^{(k+2)}(v_1', v_2')$, the resulting automorphism of $\Gamma_{\mathfrak{X}}$ sends the walk $v_1 \to v$ to a walk $v_1' \to v'$ (and analogously for the second walk), and all these walks can be chosen to be at most of length $\operatorname{diam}(\Gamma_{\mathfrak{X}}) < 2^k$. The colour $c^{(k)}(v_1, v)$ knows the walk $v_1 \to v$, therefore $c^{(k+1)}(v_1, w)$ knows the colour of the whole sequence composed by $v_1 \to v$ and the pair (v, w); on the other side, $c^{(k)}(w, v_2)$ knows the walk $w \to v_2$, so that $c^{(k+2)}(v_1, v_2)$ knows the whole long sequence $v_1 \to v \to w \to v_2$: $c^{(k+2)}(v_1, v_2) = c^{(k+2)}(v_1', v_2')$ then implies $c(v, w) = c(v', w')$ (as always, we have repeatedly used the uniqueness of nonempty-coloured walks from a given vertex: otherwise we would have not been able only from colours to identify (v', w') as the image of (v, w)). We have extended the automorphism of $\Gamma_{\mathfrak{X}}$ to the empty colours, i.e. it is in fact an automorphism of the whole configuration $\mathfrak{X}$.

By Proposition 2.1.5, the existence of this automorphism means that if two pairs $(v_1, v_2), (v_1', v_2')$ have the same colour at the $(k+2)$-th iteration they will always have the same colour, i.e.

$$\mathrm{WL}(\mathfrak{X}) \leq k + 2 \leq \log_2 \operatorname{diam}(\Gamma_{\mathfrak{X}}) + 3,$$

and the theorem is proved. □

2.3 The lower bound

We now prove the lower bound in Theorem 2.1.7. Again, we prove the same for a more general class of configurations than the ones given by Cayley and Schreier graphs.

Theorem 2.3.1. *Let $\mathfrak{X}$ be a configuration with only one empty colour, and let $\Gamma_{\mathfrak{X}}$ be defined as in Theorem 2.2.1; suppose that there exists a $\varphi \in \mathrm{Aut}(\mathfrak{X}), \varphi \neq \mathrm{Id}_{\mathfrak{X}}$ with the property that each nontrivial $\varphi^i \in \langle\varphi\rangle$ (where $\langle\varphi\rangle$ is the cyclic subgroup of $\mathrm{Aut}(\mathfrak{X})$ generated by φ) has no fixed points, and consider such a φ with minimal $|\langle\varphi\rangle|$. Then the number $\mathrm{WL}(\mathfrak{X})$ of nontrivial iterations of the Weisfeiler-Leman algorithm satisfies*

$$\mathrm{WL}(\mathfrak{X}) \geq \log_2 \operatorname{diam}(\Gamma_{\mathfrak{X}}) - \log_2(|\langle\varphi\rangle| - 1) - 3.$$

Proof that Thm. 2.3.1 ⇒ Thm. 2.1.7. As we already said in the previous section, the only possible empty colour is $\emptyset$; the extreme cases of $\emptyset$ nonempty and $\emptyset$ not a colour give a trivially true bound in Theorem 2.1.7 since $\operatorname{diam}(\mathrm{Sch}(G, S)) = 2$

and 1 respectively, while when $\emptyset$ is empty we have $\mathrm{diam}(\Gamma_{\mathfrak{X}_S}) = \mathrm{diam}(\mathrm{Sch}(G,S))$. As for the automorphism condition, if $q > 2$ the elements $h \in \mathbb{F}_q \setminus \{0,1\}$ induce automorphisms φ_h defined by $h(v_1,\dots,v_k) = (hv_1,\dots,hv_k)$: they are really automorphisms since $Av = v' \Leftrightarrow Ahv = hv'$ and they obviously do not fix any point since $hv \neq v$ for $h \neq 1$ ($v = 0$ is not considered linearly independent even on its own). The multiplicative group $\mathbb{F}_q^*$ is just a cycle of order $q-1$, so there will be an element of order p as defined in the statement of Theorem 2.1.7, and p will provide an upper bound for $\min_\varphi\{|\langle\varphi\rangle|\}$. □

Again, notice that Theorem 2.3.1 would also apply to Cayley graphs: the same reasoning applies with regard to the colour $\emptyset$, and right multiplication by any $h \in G$ gives an automorphism φ_h that behaves as required.

The key idea to prove Theorem 2.3.1 is the following: for any φ with the property described in the statement, the various pairs $(v, \varphi^i(w))$ despite being distinct from one another will all have the same colour until we manage to "cover the distance" between v and at least one of the $\varphi^i(w)$, i.e. to encode inside the colours information about walks of length $d(v,\varphi^i(w))$. This is what we meant when we said that the process of bestowing upon our colours knowledge of the walks was necessary in order to reach coherence through Weisfeiler-Leman: we prove that, since the surroundings of the $\varphi^i(w)$ are indistinguishable, from the point of view of v we will not be able to differentiate among them until we manage to touch at least one of them.

This idea translates to the following result.

Lemma 2.3.2. *For any pair of vertices $v, w \in \Gamma_{\mathfrak{X}}$ and for any integer $k \geq 0$ such that $d(v,\varphi^i(w)) > 2^k$ for all i, we have that $c^{(k)}(v,\varphi^i(w))$ is the same for all i.*

Proof. We proceed by induction on k. For $k = 0$ the statement is obvious, since by hypothesis $d(v,\varphi^i(w)) > 1$ for all i and $c^{(0)}$ is the original colouring of $\mathfrak{X}$, so that $c^{(0)}(v,\varphi^i(w)) = \emptyset$ for each of these pairs, where $\emptyset$ is the unique empty colour.

Now suppose that the statement is true for k, i.e. for every two elements $v, w \in \Gamma_{\mathfrak{X}}$ such that $\forall i \left(d(v,\varphi^i(w)) > 2^k\right)$ the colour $c^{(k)}(v,\varphi^i(w))$ is the same for all i. We fix now elements v, w that satisfy the condition $\forall i \left(d(v,\varphi^i(w)) > 2^{k+1}\right)$ and we prove for them the statement for $k+1$.

The colour $c^{(k+1)}(v,w)$ consists of the previous colour $c^{(k)}(v,w)$ and of the number of v' such that $\left(c^{(k)}(v,v'), c^{(k)}(v',w)\right) = (c_1,c_2)$ for each choice of $c_1, c_2 \in \mathcal{C}^{(k)}$: since by hypothesis we already have the same $c^{(k)}(v,\varphi^i(w))$ for all i, we only need to prove that for each pair (c_1,c_2) the numbers are the same for all $(v,\varphi^i(w))$, i.e. for each i there exists a bijection $\sigma_i : V \to V$ such that

$$\left(c^{(k)}(v,v'), c^{(k)}(v',w)\right) = \left(c^{(k)}(v,\sigma_i(v')), c^{(k)}(\sigma_i(v'),\varphi^i(w))\right). \tag{2.3.1}$$

Define $A^{(k)} = \left\{v' \in V \,|\, \min_i d(\varphi^i(v),v') \leq 2^k\right\}$: we show that the bijections then can be defined to be

$$\sigma_i(v') = \begin{cases} v' & \text{if } v' \in A^{(k)}, \\ \varphi^i(v') & \text{if } v' \notin A^{(k)}. \end{cases}$$

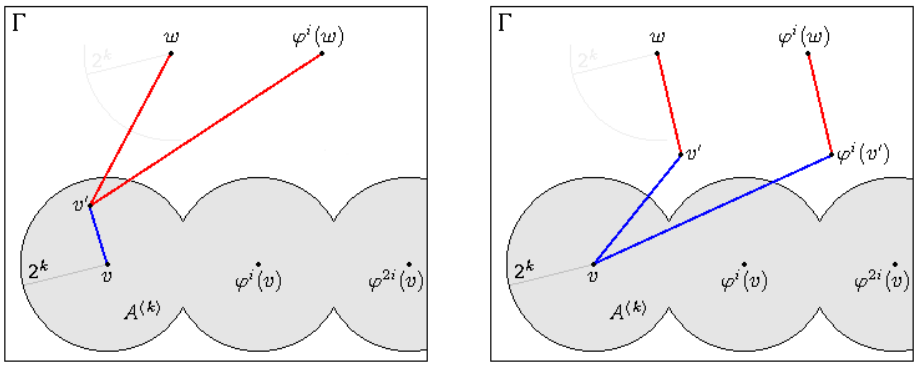

Figure 2.1: Colours $c^{(k)}$ for the two cases $v' \in A^{(k)}$ and $v' \notin A^{(k)}$.

These are really bijections, since the fact that $\varphi \in \mathrm{Aut}(\mathfrak{X})$ preserves distances implies $v' \notin A^{(k)} \Leftrightarrow \varphi^i(v') \notin A^{(k)}$. Figure 2.1 shows what happens in our situation: since w is at least 2^{k+1} far from all the $\varphi^i(v)$ and the walks that the colours know have length at most 2^k, either we are close to one of the $\varphi^i(v)$ and then far away from w (which allows us to "confuse" w and $\varphi^i(w)$ from the point of view of v') or we are far away from the $\varphi^i(v)$ (and we can confuse v' and $\varphi^i(v')$ from the point of view of v).

When $v' \in A^{(k)}$, the first components in (2.3.1) are already identical. By definition, there is a j such that $d(\varphi^j(v), v') \leq 2^k$; then using $d(\varphi^j(v), \varphi^i(w)) = d(v, \varphi^{i-j}(w))$ we have

$$d(v', \varphi^i(w)) \geq d(\varphi^j(v), \varphi^i(w)) - d(\varphi^j(v), v') > 2^{k+1} - 2^k = 2^k,$$

thus obtaining equality for the second components in (2.3.1) by inductive hypothesis.

When $v' \notin A^{(k)}$, we have $d(v, \varphi^i(v')) = d(\varphi^{-i}(v), v') > 2^k$ for all i by definition: this gives us already $c^{(k)}(v, v') = c^{(k)}(v, \varphi^i(v'))$ by inductive hypothesis, and the first components in (2.3.1) are equal; equality of the second components is also clear, because φ^i is an automorphism of $\mathfrak{X}$ that sends (v', w) to $(\varphi^i(v'), \varphi^i(w))$, which means that their colours will always be equal to each other at every iteration (as consequence of Proposition 2.1.5).

We have thus shown that the σ_i are bijections satisfying (2.3.1): this means that the colours $c^{(k+1)}(v, \varphi^i(w))$ are all the same for every i, proving the inductive step. □

It is now easy to prove Theorem 2.3.1 from this lemma.

Proof of Thm. 2.3.1. Suppose that we have reached the end of the Weisfeiler-Leman algorithm, i.e. we have reached an iteration k that does not refine any colour in $\mathcal{C}^{(k-1)}$. For any two elements $v, w_1 \in \Gamma_{\mathfrak{X}}$ there is a walk of nonempty colours from v to w_1 that is unique to w_1 (by definition of nonempty colour) in the sense that no other w_2 will have the same walk going from v to w_2. This means in

particular that the final colours $c^{(k)}(v, \varphi^i(w))$ must be all distinct for fixed v, w, now that we have reached the end of Weisfeiler-Leman and the configuration has become coherent: in fact for each of them there is a walk that is not shared by the others, so by Proposition 2.1.4 these walks will give different colours to the pairs $(v, \varphi^i(w))$. By Lemma 2.3.2 however these pairs have all the same colour when $\forall i \left(d(v, \varphi^i(w)) > 2^k\right)$, which implies that for any fixed $v \in \Gamma_{\mathfrak{X}}$ and for any other $w \in \Gamma_{\mathfrak{X}}$ there exists an i such that $\varphi^i(w)$ has distance at most 2^k from v.

We now suppose that the cyclic group generated by φ in $\mathrm{Aut}(\mathfrak{X})$ is of minimal size, the only property of φ required in the statement that we have not used up to this point. Fix now a v such that $\mathrm{diam}(\Gamma_{\mathfrak{X}})$ is realized as a distance by some pair (v, v'), and fix a w such that $d(v, w) = 2^k + 1$ (of course we are supposing that the diameter is larger than 2^k, otherwise the bound of the theorem is already true): we must have $d(v, \varphi^i(w)) \leq 2^k$ for some i, so changing the name of the generator of this shortest cycle $\langle\varphi\rangle$ we can suppose that $d(v, \varphi(w)) \leq 2^k$. For a v' such that $d(v, v') = \mathrm{diam}(\Gamma_{\mathfrak{X}})$ we must also have $d(v, \varphi^j(v')) \leq 2^k$ for some j; therefore

$$\begin{aligned}
\mathrm{diam}(\Gamma_{\mathfrak{X}}) = d(v, v') \;&=\; d(\varphi^j(v'), \varphi^j(v)) \\
&\leq d(\varphi^j(v'), v) + \sum_{i=1}^{j} \left(d(\varphi^{i-1}(v), \varphi^i(w)) + d(\varphi^i(w), \varphi^i(v))\right) \\
&= d(v, \varphi^j(v')) + j(d(v, \varphi(w)) + d(v, w)) \\
&\leq 2^k + j(2^k + 2^k + 1) \;\leq\; 2^{k+2} j \;\leq\; 2^{k+2}(|\langle\varphi\rangle| - 1),
\end{aligned}$$

and we have the bound on the last nontrivial iteration $k - 1$ as

$$\mathrm{WL}(\mathfrak{X}) \;=\; k - 1 \;\geq\; \log_2 \mathrm{diam}(\Gamma_{\mathfrak{X}}) - \log_2(|\langle\varphi\rangle| - 1) - 3.$$

The proof is concluded. □

2.4 The case of Cayley graphs

Now we prove Theorem 2.1.8. The case of a Cayley graph $\mathrm{Cay}(G, S)$ is particularly nice with respect to general Schreier graphs: the reason behind this fact is ultimately the existence of automorphisms (namely, right multiplications by elements of G) that can send any g to any g', as we will be able to observe during the proof of the result[1].

Before we move to the actual proof, here is an idea (without details) that shows why the result is reasonable: the author thanks L. Bartholdi (personal communication) for the observation. If at a certain iteration the algorithm has assigned a colour to two elements s, s', in the sense that $c(g, sg)$ has been established as a nonempty colour, then at the next step we would have a colour for ss' too: this is because the pair $(g, ss'g)$ has a triangle with sides $(g, s'g), (s'g, ss'g)$ built on it. Therefore, at the k-th iteration the configuration we obtain from $\mathrm{Cay}(G, S)$ is

[1] To turn a Russian quote inside out, Cayley vertices being all alike make us happy, and Schreier vertices being all in their own way make us unhappy.

substantially the same as the initial configuration coming from $\mathrm{Cay}(G, S^{2^k})$. Thus it is natural to expect that the algorithm would stop when S^{2^k} becomes the whole G (or possibly G except for just one g), and the number of steps would become essentially $\log_2 \mathrm{diam}(\mathrm{Cay}(G, S))$.

Now, let us turn to filling the necessary details. We start with a lemma that provides an upper bound for the number of iterations; more specifically, we prove that the right hand side of the equality in Theorem 2.1.8 provides at least a sufficient number of iterations to make the configuration $\mathfrak{X}_C$ coherent.

Lemma 2.4.1. *The inequality with $\leq$ in Theorem 2.1.8 holds.*

Proof. Right multiplication by any $h \in G$ gives an automorphism φ_h of $\mathfrak{X}_C$, since $sg = g' \Leftrightarrow sgh = g'h$ for any $g, g' \in G$, $s \in S$; by Proposition 2.1.5, at the end of the algorithm φ_h will still be an automorphism of the final coherent configuration, which means in particular that every vertex will have the same colour and that for any three vertices $g, g', h \in G$ there exists an $h' \in G$ such that $c(g, h) = c(g', h')$. Therefore, the maximum possible colour refinement that we can expect to obtain from Weisfeiler-Leman is the one where all sets $\{c(g, h) | h \in G\}$ are equal for every g but where any two colours $c(g, h), c(g, h')$ are distinct for $h \neq h'$.

Indeed, this is the colouring that we reach at the end: as we have already observed in the general Schreier case, any coloured walk from g to h (corresponding to a certain product of generators that represents $g^{-1}h$) is the unique walk from g consisting of that sequence of colours, and by Proposition 2.1.4 the final colour $c(g, h)$ knows it. Thus, every $c(g, h)$ knows a walk that all other $c(g, h')$ do not know, and the colouring described above is achieved.

When every colour $c(g, h)$ for all $h \in G$ has become distinct from the others, we have undoubtedly reached the end of the algorithm; this happens when the colour of every pair (g, h) knows at least one walk connecting them, or at the very least when for a fixed g all but one of them know such a walk (so that the remaining pair (g, h) has the unique colour that knows no walk: this corresponds, informally speaking, to the situation where even the emptiest descendant of the colour $\emptyset$ is nonempty according to the definition given in Theorem 2.2.1). By Lemma 2.2.2, this happens when at the k-th iteration we have 2^k at least as large as the diameter, or at least $\mathrm{diam}(\mathrm{Cay}(G, S)) - 1$ if for any g there is only one h whose distance from g is the diameter; the result follows. □

To prove an inequality in the other direction, we make use of the abundance of automorphisms in the Cayley graph to prove a stronger version of Lemma 2.3.2.

Lemma 2.4.2. *For any four vertices $g, h, g', h' \in G$ and for any integer $k \geq 0$ such that $d(g, h), d(g', h') > 2^k$, we have $c^{(k)}(g, h) = c^{(k)}(g', h')$.*

Proof. Again, we proceed by induction on k. For $k = 0$ the statement is obvious, because all pairs (g, h) of vertices with distance > 1 have the same colour $\emptyset$ at the 0-th step.

Now suppose that the statement is true for k. First, we are going to prove that for any three vertices $g, h, h' \in G$ with $d(g, h), d(g, h') > 2^{k+1}$ the two pairs $(g, h), (g, h')$ will still have the same colour at the $(k + 1)$-th step; the idea is,

as in Lemma 2.3.2, to construct a suitable bijection $\sigma : G \to G$ with a property analogous to (2.3.1), namely

$$\left(c^{(k)}(g, g'), c^{(k)}(g', h)\right) = \left(c^{(k)}(g, \sigma(g')), c^{(k)}(\sigma(g'), h')\right). \tag{2.4.1}$$

Define

$$\sigma(g') = \begin{cases} g' & \text{if } d(g', h), d(g', h') > 2^k, \\ g'h^{-1}h' & \text{if } d(g', h) \leq 2^k, \\ \tau(g') & \text{if } d(g', h) > 2^k, d(g', h') \leq 2^k, \end{cases}$$

where τ is an arbitrary bijection from the set $\{g' \in G | d(g', h) > 2^k, d(g', h') \leq 2^k\}$ to the set $\{g' \in G | d(g', h) \leq 2^k, d(g', h') > 2^k\}$.

In the first case we have obviously a bijection, whose image is the set of all vertices of distance $> 2^k$ from both h and h'; also, by inductive hypothesis in this set we have $c^{(k)}(g', h) = c^{(k)}(g', h')$, so (2.4.1) is satisfied. In the second case, right multiplication by $h^{-1}h'$ is an automorphism (hence a bijection) from the ball around h to the one around h' both of radius 2^k: this descends trivially from the fact that if we have $s_i \in S$ such that $(\prod_i s_i)\, g' = h$ then also $(\prod_i s_i)\, g'h^{-1}h' = h'$, so that in particular $d(g', h) = d(g'h^{-1}h', h')$; moreover for the same reason we must have $c^{(k)}(g', h) = c^{(k)}(g'h^{-1}h', h')$, because for every walk given by a sequence of colours s_i from g' to h the same walk exists from $g'h^{-1}h'$ to h': as observed in the proof of Lemma 2.4.1, for our three vertices $g', h, g'h^{-1}h'$ there must be a fourth x that will have $c(g', h) = c(g'h^{-1}h', x)$ at the end, and $x = h'$ is the only possible candidate. We also have $c^{(k)}(g, g') = c^{(k)}(g, g'h^{-1}h')$ by inductive hypothesis, since g' and $g'h^{-1}h'$ are at distance $\leq 2^k$ from h and h' (both at distance $> 2^{k+1}$ from g); thus, (2.4.1) is satisfied again. In the third case, τ is a bijection because the balls around h' and h of radius 2^k have the same number of vertices, and domain and codomain of τ are these two balls minus their intersection; in addition, the colours of the four pairs $(g, g'), (g', h), (g, \tau(g')), (\tau(g'), h')$ are all the same since each of their distances is $> 2^k$, so we have (2.4.1) for this case too. Given that the codomains in the three cases are disjoint, σ is indeed a bijection satisfying (2.4.1), which implies that $c^{(k+1)}(g, h) = c^{(k+1)}(g, h')$.

Now we know that for any vertex g there is a colour c_g such that all vertices h at distance $> 2^{k+1}$ will have $c^{(k+1)}(g, h) = c_g$: but then it is obvious that c_g does not depend on g, since any g is sent to any g' by some automorphism that will preserve distances in the graph, so that (g, h) of distance $> 2^{k+1}$ is sent to some (g', h') of same distance (and consequently $c_g = c_{g'}$). This proves that $c^{(k+1)}(g, h) = c^{(k+1)}(g', h')$ whenever $d(g, h), d(g', h') > 2^{k+1}$, concludes the inductive step and proves the lemma. □

Now we can easily prove Theorem 2.1.8.

Proof of Thm. 2.1.8. We have already shown the $\leq$ direction in Lemma 2.4.1. On the other hand, proving the $\geq$ direction means proving that if at the k-th iteration there are three vertices g, h, h' such that $d(g, h), d(g, h') > 2^k$ then there are more iterations to come; by Lemma 2.4.2, however, in this situation $c^{(k)}(g, h) = c^{(k)}(g, h')$ and we know that at the end of the algorithm we will have $c(g, h) \neq c(g, h')$ (ultimately coming from Proposition 2.1.4), so the statement holds. □

2.5 Concluding remarks

As we have noted, the Weisfeiler-Leman algorithm behaves especially well on Cayley graphs: Schreier graphs, which in general have not as many automorphisms as Cayley graphs, satisfy less stringent upper and lower bounds. In one direction, having vertices with different stabilizers can make pairs of vertices of large distance receive different colours early on, so that the colouring at the k-th step of the algorithm conveys more information than the mere colouring coming from the choice of S^{2^k} as generators. Consider for example the set $V = \mathbb{Z}/n\mathbb{Z}$ and the group $G = \mathrm{Sym}(n)$ generated by the set S consisting of all transpositions of the type $(i\ i+1)$ and the identity: the diameter of $\mathrm{Sch}(V, S)$ is $\lfloor n/2 \rfloor$, but since every $i \in V$ is stabilized by a different subset of the generators (i.e. all of them except $(i-1\ i)$ and $(i\ i+1)$) after the first step all the pairs of colour $\emptyset$ are differentiated immediately; therefore the number of iterations in this case will be 1, for any choice of n (the bound given by Theorem 2.3.1 also fails, because there are no non-trivial automorphisms of the coloured graph).

On the other hand, the fact that the maximum possible colour refinement that we can expect from Weisfeiler-Leman is more than the one described during the proof of Lemma 2.4.1 could mean that more steps are necessary than just the ones needed to reach the end of the graph: the information to reconstruct the whole graph (to put it in the language of Lemma 2.2.3) exists already but it could be scattered among the various pairs of vertices of the graph and it could take a few more steps to make sure that every single pair knows everything about the graph. Consider for example the set $V = \{1, 2, \dots, 14\}$ and the set $S = \{e, \sigma^{\pm 1}, \tau^{\pm 1}\}$ (the group G has little importance here: for the sake of simplicity, think of it as the free group F_2, or as a suitable subgroup of $\mathrm{Sym}(14)$) acting on V as follows:

$$\begin{aligned}
\sigma : & 1 \mapsto 2 \mapsto 3 \mapsto 4 \mapsto 5 \mapsto 6 \mapsto 7 \mapsto 1, \\
& 8 \mapsto 9 \mapsto 10 \mapsto 11 \mapsto 12 \mapsto 13 \mapsto 14 \mapsto 8; \\
\tau : & 1 \mapsto 8 \mapsto 9 \mapsto 2 \mapsto 1,\ 3 \mapsto 10 \mapsto 3,\ 4 \mapsto 11 \mapsto 4, \\
& 5 \mapsto 12 \mapsto 5,\ 6 \mapsto 13 \mapsto 6,\ 7 \mapsto 14 \mapsto 7.
\end{aligned}$$

In this situation, $\mathrm{Sch}(V, S)$ looks like two coloured heptagons whose corresponding vertices are linked, so that its diameter is 4; from the reasoning in Lemma 2.2.2 and Lemma 2.2.3, after the second iteration of the algorithm there is enough information to reconstruct the whole graph, and this would be in accord with a hypothetical estimate as in Theorem 2.1.8. Nevertheless, it is possible to verify that we have $c^{(2)}(5,5) = c^{(2)}(12,12)$ and $c^{(3)}(5,5) \neq c^{(3)}(12,12)$, so that the number of iterations for this configuration is > 2 (it is 3 indeed).

***

Theorems 2.1.6-2.1.7-2.1.8 establish a rather strong correlation between the number of iterations of Weisfeiler-Leman and the diameter of Cayley and Schreier graphs. In particular, Theorem 2.1.8 allows us to describe the diameter of Cayley graphs as a function of $\mathrm{WL}(\mathfrak{X}_C)$; it is natural, in the context of Babai's conjecture,

to ask ourselves whether it is possible that the number of Weisfeiler-Leman iterations could be reflected in another way in the construction of the graph: usually, determining the runtime of the algorithm would involve from the beginning the actual construction of the graph, thus making it useless for the solution of the conjecture. In light of this, it would be interesting to find results that express $\mathrm{WL}(\mathfrak{X}_C)$ as a more intrinsic feature of the construction of Cayley graphs.

On the other side, to the best of our knowledge the results established here are the first ones that determine nontrivial bounds for the number of iterations of the Weisfeiler-Leman algorithm on configurations, either general or of a specific form (the trivial bound on a generic classical configuration being $|\Gamma|^2-|\mathcal{C}|$). In this direction, it would be interesting to find results in the style of Theorems 2.2.1-2.3.1 with different initial conditions: a case that appears to be particularly appealing is the case of non-coloured graphs, for which one can wonder whether it could be possible to bound $\mathrm{WL}(\mathfrak{X})$ from above by some function of the diameter of the graph, as we have done here for the particular coloured graphs described in the statement of Theorem 2.2.1.

Chapter 3

Short expressions for cosets of permutation subgroups

The content of this chapter is essentially based on [Don18].

In this chapter, we analyze more closely Babai's quasipolynomial algorithm for the graph and string isomorphism problems (GIP and SIP): see §1.5 for the history of GIP in general, and §1.6 for a short introduction to the algorithm. As said therein, the reduction and recursion process in the algorithm has at its core a theorem by Cameron [Cam81] (and later Liebeck [Lie84] and Maróti [Mar02]) that describes all the primitive permutation groups as either having relatively small size or being very close to a wreath product of alternating groups, and such a classification makes the whole result operationally depend on CFSG (quoted here as Theorem 1.2.4, essentially from [Wil09]). On the other hand, it is possible to slightly modify Babai's proof to make it independent from CFSG, a feat due to Babai himself [Bab16a, §13.1] and Pyber [Pyb16]: in particular, the algorithm can avoid the use of Cameron by resorting to another result by Pyber [Pyb93] that describes doubly transitive permutation subgroups.

Our analysis here, on a first superficial level, provides a more explicit runtime for Babai's algorithm, both in the CFSG and the CFSG-free case. We will follow Helfgott's description of Babai's result given in [Hel19b] [HBD17], instead of Babai's original formulation in [Bab16a]: Helfgott makes the algorithm more explicit and proves that the procedure actually takes time $n^{O(\log^2 n)}$ when CFSG is available; we will make it even more explicit and determine the constants in front of the logarithm. Also, in [Hel19b] the reader's attention is justifiably focused on the proof of the single steps that are involved in the procedure, while the interstitial reasoning that details the recursion is only sketched: in [Hel19b], this part is contained mostly in §3, §5.3, §6.2 and Appendix A; conversely, we will concentrate on the jumping between the main processes to delineate what the flow of the algorithm is, while using its individual theorems and subroutines as black boxes whose validity and well-functioning is taken for granted (we will mention the most important ones in §3.4). This will give us the control we need to determine

the runtime with the desired accuracy.

On a deeper level, the way we achieve the goal described above is interesting on its own. Babai's algorithm is combinatorial in nature, although it is based on group-theoretic results; on the other hand, the combinatorial techniques developed by Babai have also been used before to deduce consequences for permutation subgroups, such as in [Bab81]. It turns out that this is possible also in the case of Babai's quasipolynomial algorithm: since the procedure described by him is closely translatable to the CFSG-free case, it is possible to give a description of permutation subgroups that shares some characteristics of Cameron's result even when CFSG is not available, simply by making a subgroup pass through the algorithm, in a way that will be clarified in the next section; in brief, the use of the algorithm reveals structural information about permutation subgroups that we translate in the language of Theorem 3.2.1 as being able to write them as short expressions made of "easy" or "atomic" subgroups, where shortness here is just another face of the quasipolynomiality of the whole process.

That all of this can be useful, and that Theorem 3.2.1 can potentially do a job qualitatively similar to Cameron's theorem despite its different language, can be witnessed in §6. A decomposition similar to what we achieve in Theorem 3.2.1, but based directly on Cameron, makes its appearance in [Hel18, Prop. 4.6] and is fundamental in proving a diameter bound for $\mathrm{Alt}(n)$ that goes through a sort of product theorem, like Theorem 1.3.2. Passing through our decomposition instead, we will achieve the more modest and conditional result laid out in Theorem 6.3.6, which however shows already the potential power of this chapter's analysis.

3.1 Standard definitions

Before we start, let us recall here some standard terms and properties, coming from permutation group theory, that we have already mentioned multiple times and are long due an explanation.

Definition 3.1.1. *Let $n \geq 1$, and let $G \leq \mathrm{Sym}(n)$ be a permutation subgroup. G is said to be transitive if for any two elements $x, y \in [n]$ there exists a $g \in G$ with $g(x) = y$. G is intransitive if it is not transitive.*

Let $d \geq 1$. G is said to be d-transitive if for any two d-tuples of distinct elements $(x_1, \dots, x_d), (y_1, \dots, y_d) \in [n]^d$ there is a $g \in G$ with $g(x_i) = y_i$ for each $1 \leq i \leq d$. A 2-transitive subgroup is also referred to as doubly transitive.

The group $G \leq \mathrm{Sym}(n)$ is a giant if either $G = \mathrm{Sym}(n)$ or $G = \mathrm{Alt}(n)$.

Transitive subgroups of $\mathrm{Sym}(n)$ have only one orbit for their natural action on $[n]$. There is another action of permutation subgroups that we will have to consider, namely the one on the set of k-subsets of $[n]$, denoted by $\binom{[n]}{k}$ (in obvious analogy with the binomial coefficients); in particular, the action of a d-transitive group on $\binom{[n]}{d}$ has only one orbit too. The same abstract group G can be embedded into symmetric groups of different degrees, and thus be transitive or intransitive depending on the situation, therefore we will always specify "$G \leq \mathrm{Sym}(n)$" or similar notations to indicate that G is considered to be of degree n; one of the

reductions we operate, the one we call "fourth action" in §3.6, is a passage to a smaller degree without changing G, so it is an important detail to keep in mind.

Let us see another important characteristic of the action of permutation groups.

Definition 3.1.2. *Let $G \leq \mathrm{Sym}(n)$ be transitive. A system of blocks of (the action of) G is a partition $\mathcal{B} = \{B_1, B_2, \dots, B_r\}$ of $[n]$ such that for every $g \in G$ and every $1 \leq 1, j \leq r$ either $B_i = g(B_j)$ or $B_i \cap g(B_j) = \emptyset$. A trivial system of blocks is either the system $\mathcal{B} = \{[n]\}$ or the system $\mathcal{B} = \{\{1\}, \{2\}, \dots, \{n\}\}$.*

G is primitive if the only systems of blocks it has are the trivial ones; G is imprimitive if it is not primitive. G is uniprimitive if it is primitive and not 2-transitive.

By transitivity, all the blocks of the same system have the same size. Every 2-transitive group is primitive, but not vice versa: in other words, there exist uniprimitive groups, for example $\mathrm{Alt}(n)$ acting on $\binom{[n]}{2}$, provided that n is large enough ($n = 6$ is sufficient[1]). Similarly, there are transitive but imprimitive groups: an example of minimal size in terms of $|G| + n$ is $\langle (1\ 2), (1\ 3)(2\ 4) \rangle$ acting on $\{1, 2, 3, 4\}$.

Finally, let us not miss an opportunity to describe the following action, since it plays a central role in Cameron.

Definition 3.1.3. *Let G, H be finite groups acting on finite sets V, W respectively. Then $G \wr H$, the wreath product of G by H, is defined to be the semidirect product $G^{|W|} \rtimes H$; in other words, $G \wr H$ is the group whose underlying set is $G^{|W|} \times H$ and whose group operation is*

$$(g_{w_1}, \dots, g_{w_{|W|}}, h) \cdot (g'_{w_1}, \dots, g'_{w_{|W|}}, h') = (g_{w_1} g'_{h^{-1}(w_1)}, \dots, g_{w_{|W|}} g'_{h^{-1}(w_{|W|})}, hh').$$

The primitive action of $G \wr H$ on $V^{|W|}$ is defined to be

$$(g_{w_1}, \dots, g_{w_{|W|}}, h) \cdot (v_{w_1}, \dots, v_{w_{|W|}}) = (g_{h^{-1}(w_1)} v_{h^{-1}(w_1)}, \dots, g_{h^{-1}(w_{|W|})} v_{h^{-1}(w_{|W|})}).$$

There are several wreath products in more general contexts, but for us this will be sufficient. The primitive action of $G \wr H$ is also called product action or exponentiation in the literature [Cam99, §4.3] [DM96, §2.7] [JK81, §4.1]; there is also another natural action of the wreath product, the *imprimitive action* on $V \times W$, but we will not encounter it.

In Theorem 1.2.5(a), we are using the definition above with $G = \mathrm{Sym}(m)$ and $H = \mathrm{Sym}(r)$ and their natural actions on $V = \binom{[m]}{k}$ and $W = [r]$ respectively.

3.2 Main theorem: statement

Let us start with a permutation subgroup $G \leq \mathrm{Sym}(n)$. How "easy" is it to describe? Or rather, what are the "easy" permutation subgroups and how can we obtain all subgroups by building them out of the easy ones?

[1] Double transitivity fails because if a is stabilized then all pairs containing a are sent to each other; let us sketch the argument for primitivity. There is 1 such that $\{1, 2\}$ is in a block B and $\{1, 3\}$ is not: if there is $\{4, 5\} \in B$, use (2 3 6). If on the contrary all pairs in B touch either 1 or 2, we can have either $\{1, 3\} \in B$ and $\{1, 4\} \notin B$ (and use (3 4 5)) or every $\{1, x\}, \{2, y\} \in B$ (and use (1 2 3)); if there are none at all, $|B| = 1$.

The easiest kind of subgroup that one can imagine would likely be a product of symmetric groups: given a partition $\{[n_i]\}_i$ of $[n]$, in the sense that $\sum_i n_i = n$, the subgroup corresponding to $\prod_i \mathrm{Sym}(n_i)$ (provided that we fix a way to partition $[n]$ into these $[n_i]$) is very easily describable, in terms of generators, size, membership, etc...; we are curious about the way in which we can assemble groups of this sort to create G, or more generally a coset of G if possible. Specifically, given a certain $H = \prod_j \mathrm{Sym}(n_j)$ with $\sum_j n_j = n$ and a general $G \leq \mathrm{Sym}(n)$, we are going to give a description of cosets of the form $G \cap H\sigma$ in terms of easy subgroups; note that this does not include all the possible permutation cosets: for example, $G'\eta$ with G' transitive is of the form $G \cap H\sigma$ only if $H = \mathrm{Sym}(n)$, which implies that η is the identity permutation. On the other hand, by the same reasoning we promptly see that any subgroup G' falls into this class of cosets. The reason why we restrict to these cosets will lie in our use of Babai's result (see Definition 3.3.1).

Let us define now more rigorously what it means to build an expression for $G \cap H\sigma$ starting from easy building blocks. Our *atomic* elements are:

($\mathcal{A}$) cosets $G\sigma$ of permutation subgroups G of the form $\mathrm{Alt}(\bigsqcup_i A_i) \cap \prod_i \mathrm{Sym}(A_i)$ (where the A_i are disjoint sets).

So the atoms are defined to be the cosets of the even permutation part of the aforementioned "easiest subgroups". In particular, the trivial subgroup $\{\mathrm{Id}_{|\Omega|}\}$ is an atom, being simply $\mathrm{Sym}(1)^{|\Omega|}$, and so are all singletons $\{\sigma\}$, being its cosets.

We declare the atoms to be *well-formed*. We can combine well-formed expressions to form more complex ones; the legitimate ways to do it are the following three.

($\mathcal{C}1$) *Paste cosets of a subgroup to get the whole group.*

Let $G' \leq G \leq \mathrm{Sym}(A)$ with $\{\sigma_i\}_i$ a set of representatives of G' in G, and let $H = \prod_j \mathrm{Sym}(A_j)$ for some partition $\{A_j\}_j$ of A; suppose that for some fixed $\sigma \in \mathrm{Sym}(A)$ the cosets $G' \cap H\sigma\sigma_i^{-1}$ are all well-formed: then $G \cap H\sigma = \bigcup_i (G' \cap H\sigma\sigma_i^{-1})\sigma_i$ is also well-formed.

($\mathcal{C}2$) *Paste disjoint domains to get a group acting on both.*

Let $G \leq \mathrm{Sym}(A_1) \times \mathrm{Sym}(A_2)$; for $i = 1, 2$, let $\pi_i : G \to \mathrm{Sym}(A_i)$ be the natural projections, let $H_i = \prod_j \mathrm{Sym}(A_{ij})$ for some partition $\{A_{ij}\}_j$ of A_i, and let $\sigma_i \in \mathrm{Sym}(A_i)$. Suppose that $\pi_1(G) \cap H_1\sigma_1 = K\tau$ is well-formed, and suppose that $\pi_2(\pi_1^{-1}(K)) \cap H_2\sigma_2\pi_2(\pi_1^{-1}(\tau))$ is well-formed too: then $G \cap (H_1 \times H_2)(\sigma_1, \sigma_2)$ is well-formed.

($\mathcal{C}3$) *Paste a group fixing a set of blocks with an alternating group permuting them.*

Let $G \leq \mathrm{Sym}(A)$ be a well-formed subgroup, contained in $\prod_i \mathrm{Sym}(A_i)$ for some partition $\{A_i\}_i$ of A into equally sized parts; let $\sigma_1, \sigma_2, \sigma'$ be three permutations of A and suppose that $\langle\{\sigma_1, \sigma_2\}\rangle$ permutes the A_i in the same way as $\mathrm{Alt}(\Gamma)$ permutes $\binom{\Gamma}{k}$ for some Γ, k: then $\langle G \cup \{\sigma_1, \sigma_2\}\rangle\sigma'$ is also well-formed.

Since the trivial subgroup is an atom, all subgroups G could be written as a well-formed expression by ($\mathcal{C}1$), choosing $G' = \{\mathrm{Id}_{|\Omega|}\}$, $H = \mathrm{Sym}(\Omega)$ and any σ.

That is uninteresting, though, since we need $|G|$ atoms to perform such a task: the point is to use as few of them as possible. Our main theorem gives a way to build a well-formed expression of small length for G, and even for any $G \cap H\sigma$.

Theorem 3.2.1. *Let $n \geq 1$, let $G \leq \mathrm{Sym}(n)$, let $H = \prod_i \mathrm{Sym}(\Sigma_i)$ for some partition $\{\Sigma_i\}_i$ of $[n]$, and let $\sigma \in \mathrm{Sym}(n)$.*

Then, we can write a well-formed expression for $G \cap H\sigma$, starting from atomic elements ($\mathcal{A}$) and combining them using ($\mathcal{C}1$)-($\mathcal{C}2$)-($\mathcal{C}3$), such that the number of atomic elements involved in the construction is bounded by $n^{K \log^c n}$, where $(K, c) = (103, 2)$ if we assume CFSG and $(K, c) = (26e^{1/\varepsilon^2}, 5 + \varepsilon)$ otherwise for any $\varepsilon > 0$ small enough.

The time necessary to find such an expression is bounded by $O(n^{11+K \log^c n})$.

One can verify that $\varepsilon < \frac{1}{100}$ is indeed small enough.

The similarities with [Hel18, Prop. 4.6] are important, as they are exactly of the nature that we would need to free the bound on $\mathrm{diam}(\mathrm{Alt}(n))$ proved therein from the use of CFSG: the descent to smaller cosets (or ascent to larger ones, for us) works in the same way, and the quasipolynomial bound is fundamental for the diameter. The only difference that prevents a direct substitution is the fact that ($\mathcal{C}1$) allows for any subgroup, instead of restricting to normal subgroups like we would need for other procedures given in the course of such a proof. This whole discussion will take place again in §6.

The runtime claimed in Theorem 3.2.1 is in reality a bound on the runtime for Babai's algorithm: the construction process of the well-formed expression, as illustrated in the following sections, is part of the description process necessary to solve the string isomorphism problem; in the proof we will calculate the cost for the latter, thus retrieving a bound for the former as well.

Setting aside the time issue, this theorem does not surprise us if we assume CFSG. Cameron implies in its stronger form that any primitive permutation subgroup either is small enough to be expressed as the union of $\leq n^{O(\log^2 n)}$ singletons through ($\mathcal{C}1$) or it has as large subgroup a wreath product $\mathrm{Alt}(\Gamma) \wr \mathrm{Alt}(s)$ where $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$ (see Definition 3.1.3), so that it is susceptible of being described using repeatedly ($\mathcal{C}3$); if the subgroup is not primitive, it is not difficult to reduce to this case by working on each block separately and then uniting and glueing together the pieces with ($\mathcal{C}1$) and ($\mathcal{C}2$).

Without assuming CFSG however, the situation changes. It is true that, for doubly transitive permutation subgroups, Theorem 3.2.1 would be a consequence of Pyber's result: either such a group is $\mathrm{Sym}(n)$ or $\mathrm{Alt}(n)$, or it has size $\leq n^{O(\log^2 n)}$; the discussion goes basically as above. Pyber's result does not however say anything about subgroups that are transitive but not doubly transitive; in this sense, Theorem 3.2.1 extends this CFSG-free description to this class of permutation subgroups as well (and [Hel18, Prop. 4.6] is needed for all transitive groups).

One last note: the computation of K in the main theorem, and many of the intermediate results leading to it, have been performed with SageMath, version 8.9. The calculations are elementary enough to be easily reproducible with any

software, but SageMath is open-source and can be embedded into LaTeX, which is why the author chose to use it.

3.3 Elementary routines

Let us define the fundamental objects in the study of SIP.

Definition 3.3.1. *Let Ω be a finite set, let $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings. The set of isomorphisms from $\mathbf{x}$ to $\mathbf{y}$ in G is defined as*

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \{g \in G | \mathbf{x}^g = \mathbf{y}\} = \{g \in G | \forall r \in \Omega(\mathbf{x}(r) = \mathbf{y}(g(r)))\}.$$

The group of automorphisms of $\mathbf{x}$ in G is defined as $\mathrm{Aut}_G(\mathbf{x}) = \mathrm{Iso}_G(\mathbf{x}, \mathbf{x})$.

The sets of isomorphisms $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ are precisely the intersections $G \cap H\sigma$, H being a product of smaller symmetric groups, that are featured in Theorem 3.2.1: in fact, a permutation of Ω is in such a set if and only if it is in G and for every letter of Σ it sends the preimage of that letter in $\mathbf{x}$ to its preimage in $\mathbf{y}$. H is therefore $\prod_{\alpha \in \mathbf{x}(\Omega)} \mathrm{Sym}(\mathbf{x}^{-1}(\alpha))$, and vice versa, given a product of symmetric groups and a σ, it is possible to define $\mathbf{x}$ as being piecewise constant with a letter for each symmetric group and then define $\mathbf{y} = \mathbf{x}^\sigma$.

This also reveals how to find an expression for any permutation subgroup $G \leq \mathrm{Sym}(\Omega)$: this corresponds to finding $\mathrm{Aut}_G(\alpha^{|\Omega|})$, where $\alpha^{|\Omega|}$ is the constant string consisting of one letter repeated $|\Omega|$ times, or in other words to making the algorithm run "in neutral" on a trivial string so as to capture only G.

Remark 3.3.2. Every time we describe $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ as a coset $G'\tau$, where $G' \leq \mathrm{Sym}(\Omega)$ and $\tau \in \mathrm{Sym}(\Omega)$, G' is actually $\mathrm{Aut}_G(\mathbf{x})$ and τ is an element of G sending $\mathbf{x}$ to $\mathbf{y}$.

In fact, since G' is a subgroup of $\mathrm{Sym}(\Omega)$ it contains the trivial permutation, so that $\tau \in \mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$: this proves what we claimed about τ. If $g \in G'$ (so that $g\tau$ sends $\mathbf{x}$ to $\mathbf{y}$) then g fixes $\mathbf{x}$ since permutations are bijections and any $\mathbf{x}' \neq \mathbf{x}$ will not be sent to $\mathbf{y}$ by τ; therefore by definition g is also an element of $\mathrm{Aut}_G(\mathbf{x})$. On the other hand, if $\sigma \in \mathrm{Aut}_G(\mathbf{x})$ then $\sigma\tau \in \mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = G'\tau$ and $\sigma \in G'$; this proves also that $G' = \mathrm{Aut}_G(\mathbf{x})$.

We begin by providing several simple results on computations that we have to constantly perform throughout the whole procedure. Before that, a couple of definitions; if $G \leq \mathrm{Sym}(\Omega)$ and $\Delta \subset \Omega$, the *setwise stabilizer* and the *pointwise stabilizer* of Δ are respectively

$$G_\Delta = \{g \in G | g(\Delta) = \Delta\},$$
$$G_{(\Delta)} = \{g \in G | \forall r \in \Delta(g(r) = r)\}.$$

We also write $G_{(r_1,\dots,r_i)}$ for $G_{(\{r_1,\dots,r_i\})}$. Trying to find the setwise stabilizer for a generic Δ is a task of difficulty comparable to producing $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ itself; on the other hand, producing pointwise stabilizers is much easier (see Corollary 3.3.4(e)), and we can walk down this route to obtain basic but useful algorithms.

Proposition 3.3.3 (Schreier-Sims algorithm)**.** *Let $\Omega = \{x_1, x_2, \ldots, x_n\}$ and let $G \leq \mathrm{Sym}(\Omega)$ be provided with a set of generators A. Then there is an algorithm that finds in time $O(n^5 + n^3|A|)$ a set C of generators of G of size $\leq n^2$ such that for every $0 \leq i \leq n-2$ and for every coset of $G_{(x_1,\ldots,x_i,x_{i+1})}$ inside $G_{(x_1,\ldots,x_i)}$ there exists a unique $\gamma \in C$ that is a representative of that coset.*

Proof. See [Luk82, §1.2] or [Hel19b, Alg. 1]. □

We will see that in our base cases corresponding to the atoms ($\mathcal{A}$) the number of generators will be polynomial in n, so that we will not have problems supposing that the Schreier-Sims algorithm takes polynomial time in n; from now on, when we talk about polynomial time (or size, or cost) we mean polynomial in n, the length of the strings involved. It also happens at some point that we take the union of several cosets, and the process produces sets of generators of size comparable to the number of cosets (as described in Proposition 3.5.3); in that case, the time will be more conspicuous: for instance, Corollary 3.5.7(a) and Proposition 3.5.15 entail a cost of order $m^{O(\log^2 n)}n^{O(1)}$ for the filtering of generators through Schreier-Sims.

In any case, every time a G is already "given", or has been "described" or "determined", or other similar locutions, we will suppose that it has a quadratic number of generators thanks to Schreier-Sims (unless explicitly stated otherwise).

Proposition 3.3.3 provides us with many useful polynomial-time procedures, as shown below.

Corollary 3.3.4. *Let $|\Omega| = n$ and let $G \leq \mathrm{Sym}(\Omega)$ be provided with a set of generators A of polynomial size. Then the following tasks can be accomplished in polynomial time:*

(a) determine $|G|$;

(b) determine whether a certain $g \in \mathrm{Sym}(\Omega)$ is in G;

(c) given a subgroup $H \leq G$ with index $[G : H]$ of polynomial size and given a polynomial-time test that determines whether a certain $g \in G$ is in H, determine H and a representative of each coset of H in G;

(d) given a homomorphism $\varphi : G \to \mathrm{Sym}(\Omega')$ with Ω' of polynomial size and given a subgroup $H \leq \mathrm{Sym}(\Omega')$, determine $\varphi^{-1}(H)$, or given an element $\tau \in \mathrm{Sym}(\Omega')$, determine an element of $\varphi^{-1}(\tau)$;

(e) given a set $S \subseteq \Omega$, determine $G_{(S)}$;

(f) provided that G acts transitively imprimitively on Ω and given a system of blocks of its action on Ω, determine the stabilizer of this system;

Moreover, we can explicitly write in time $O(n^5 + n^3|A| + n^2|G|)$ all the elements of G.

Proof. For parts (a)-(b)-(c) see [Hel19b, Ex. 2.1(a)-2.1(c)], based on [FHL80, Cor. 1] and [Luk82, Lemma 1.2]; the representatives in part (c) are the elements

of C_{-1} in the solution of [Hel19b, Ex. 2.1(c)] given in [HBD17, App. B][2]. Part (d) is similar to (c), see [Hel19b, Ex. 2.1(b)]; finding an element of the preimage of a generator is a passage inside the proof of the procedure that finds $\varphi^{-1}(H)$, so to solve the second issue we can take $H = \langle\tau\rangle$. Finding pointwise stabilizers $G_{(S)}$ is a byproduct of Schreier-Sims itself, so we simply have to order Ω so that $S = \{x_1, \dots, x_{|S|}\}$ and Proposition 3.3.3 will solve part (e) directly. Part (f) is an application of (d): Ω' will be the system of blocks (which means that $|\Omega'| < n$) and $H = \{\mathrm{Id}_{|\Omega'|}\}$.

The last statement is a consequence of the particular structure of the set of generators C found through Schreier-Sims: C is divided into sets $C_0, \dots, C_{n-2}$, each consisting of the generators $\gamma \in G_{(x_1,\dots,x_i)} \setminus G_{(x_1,\dots,x_{i+1})}$, and each element of G is written uniquely as a product $\gamma_0\gamma_1 \dots \gamma_{n-2}$ with $\gamma_i \in C_i$. There are $|G|$ such products, and a product of ttwo permutations is computable in time $O(n)$, whence the result.

Let us include here the runtimes of the other items, too. Parts (a)-(b)-(e) consist in using the Schreier-Sims algorithm at most twice with at most one more generator, so the runtime is $O(n^5 + n^3|A|)$. In Schreier-Sims, the time is more explicitly of order $n \cdot (n^2 \cdot n^2 + n^2 \cdot |A|)$, where n comes from the use of the subroutine FILTER in [Hel19b, Alg. 1] and n^2 is the bound on the size of the final C; by this analysis, part (c) employs time $O(n^{2i+t} + n^{i+t}|A|)$, where i is the maximum between 2 and the exponent of the index $[G : H]$ and t is the maximum between 1 and the exponent of the test time for H. For part (d), we use Schreier-Sims first on G, then on each preimage of $\mathrm{Sym}(\Omega')_{(x'_1,\dots,x'_i)}$, then we express each generator of H as product of images of generators of G: this takes time $O(n^{5s} + n^3|A| + n^{h+2s})$, where s is the maximum between 2 and the exponent of $|\Omega'|$ and h is the exponent of the number of generators of H. Using (d), part (f) takes time $O(n^{10} + n^3|A|)$. □

All these polynomial costs will not be particularly relevant: in the course of our reasoning we will not encounter an exponent of a polynomial cost that is larger than 14, and this is negligible against the $n^{K \log^e n}$ we have at the end. The constants hidden in the big O notation are only depending on the cost of procedures like reading, writing, comparing elements, etc...: we will not care about them, but just carry them around inside the O.

Another important polynomial-time algorithm is the one illustrated in the following lemma: recalling the definition of transitivity and primitivity for permutation subgroups, it is clear that being able to quickly determine respectively orbits and blocks of the actions of groups that do not present these two properties is a beneficial skill for us to possess.

Lemma 3.3.5. *Let $|\Omega| = n$ and $G \leq \mathrm{Sym}(\Omega)$. Then the orbits of the action of G on Ω can be determined in time $O(n^3)$; also, if G is transitive but imprimitive,*

[2]Between [Hel19b] and [HBD17], Exercise 2.1(b) in one corresponds to Exercise 2.1(c) in the other. The author apologizes, but that was the order in which he proved things during the translation process: if he had respected the original order, part (b) would have depended on part (c).

a system of minimal blocks for the action of G on Ω can be determined in time $O(n^4)$.

Proof. To determine the orbits, we follow [HBD17, Ex. B.2]. Let A be a set of generators of G, which by Schreier-Sims we can suppose is of size $\leq n^2$: the sets $A_x = \{x^a | a \in A\}$ for every $x \in \Omega$ can be determined in time $O(n^3)$. After that, we follow this procedure: we start with any fixed $x_0 \in \Omega$ and set $\Delta_{x_0} = \{x_0\} \cup A_{x_0}$; we divide the elements of Δ_{x_0} in "examined" (at this stage, only x_0) and "unexamined" (the other elements of Δ_{x_0}). Then at every step we take an unexamined $x \in \Delta_{x_0}$ and we update Δ_{x_0} by adding the elements of A_x to it: the newly added elements are marked as unexamined, while x now is examined; the procedure stops when Δ_{x_0} becomes the orbit $\{x_0^g | g \in G\}$. If there is an element x_1 that has not yet been considered, we define $\Delta_{x_1} = \{x_1\} \cup A_{x_1}$ and go through the whole procedure again, until we have considered all the elements of Ω: the final sets $\Delta_{x_0}, \Delta_{x_1}, \dots, \Delta_{x_m}$ are the orbits of the action of G on Ω; this part takes time $O(n)$, so the runtime of the whole algorithm is $O(n^3)$.

Suppose now that G is transitive imprimitive: to determine the blocks we follow [Hel19b, §2.1.2], which is based on an idea by Higman (through Sims and then Luks). The idea in the previous case was basically to follow the edges of the Schreier graph of G with set of generators A on Ω: we will do the same with different graphs now. Our preparatory work this time consists in considering all the pairs $\{x, x'\} \subseteq \Omega$ and constructing the sets $A_{x,x'} = \{\{x^a, x'^a\} | a \in A\}$ in time $O(n^4)$, forming a first graph; then we fix $x_0 \in \Omega$ and for every other $x \in \Omega$ we build the following graph: the set of vertices is Ω and the edges are the pairs contained in the connected component of $\{x_0, x\}$ of the first graph (finding the connected component takes linear time in the number of vertices, so $O(n^2)$ here). In the newly formed graphs, the connected components containing $\{x_0, x\}$ are the smallest blocks containing $\{x_0, x\}$ (see [Sim67, Prop. 4.4]; again, finding the connected components is a $O(n)$ routine): once we find among the blocks constructed from each x a block that is properly contained in Ω, which exists for G imprimitive, we can find a whole system by taking the other components of the graph given by the same x. The system may not be minimal, but we have only to repeat the whole process working with the set of blocks instead of Ω; since at each iteration the blocks are at least twice the size of the ones at the previous step, eventually we reach a system that has blocks of maximal size, i.e. a minimal system. The whole process works in time $O\left(n^4 + \left(\frac{n}{2}\right)^4 + \left(\frac{n}{2^2}\right)^4 + \dots\right) = O(n^4)$. □

Finally, we illustrate several equalities among different sets of isomorphisms (employed here in a slightly more flexible way than Definition 3.3.1) that will allow us to pass from difficult problems to easier ones, or to break down problems into smaller ones.

Lemma 3.3.6. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$, $\sigma \in \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings. For $\Delta \subseteq \Omega$ invariant under G, σ, define the set of partial isomorphisms $\mathrm{Iso}^{\Delta}_{G\sigma}(\mathbf{x}, \mathbf{y})$ as in Definition 3.3.1 with $g \in G\sigma$ and $\mathbf{x}(r) = \mathbf{y}(g(r))$ necessary only for $r \in \Delta$.*

(a) We can pass from cosets to groups using

$$\mathrm{Iso}^{\Delta}_{G\sigma}(\mathbf{x},\mathbf{y}) = \mathrm{Iso}^{\Delta}_{G}(\mathbf{x},\mathbf{y}^{\sigma^{-1}})\sigma.$$

(b) We can split unions of cosets using

$$\mathrm{Iso}^{\Delta}_{G\sigma_1\cup G\sigma_2}(\mathbf{x},\mathbf{y}) = \mathrm{Iso}^{\Delta}_{G\sigma_1}(\mathbf{x},\mathbf{y}) \cup \mathrm{Iso}^{\Delta}_{G\sigma_2}(\mathbf{x},\mathbf{y}).$$

(c) We can split unions of windows using

$$\mathrm{Iso}^{\Delta_1\cup\Delta_2}_{G\sigma}(\mathbf{x},\mathbf{y}) = \mathrm{Iso}^{\Delta_2}_{G_1}(\mathbf{x},\mathbf{y}^{\sigma_1^{-1}})\sigma_1,$$

where $\mathrm{Iso}^{\Delta_1}_{G\sigma}(\mathbf{x},\mathbf{y}) = G_1\sigma_1$.

(d) For every $g \in G$, call $g|_\Delta$ its restriction to Δ, defined by simply forgetting what happens in $\Omega \setminus \Delta$ (since G leaves Δ invariant, this is well-defined); define $S|_\Delta, H|_\Delta, \mathbf{x}|_\Delta$ for any $S \subseteq G$, $H \leq G$, $\mathbf{x} : \Omega \to \Sigma$ analogously. For any $h \in G|_\Delta$, let $\overline{h}$ be any element of G whose restriction to $G|_\Delta$ is h; if $H \leq G|_\Delta$, define $\overline{H}$ analogously as the subgroup of G whose restriction to $G|_\Delta$ is H (since G leaves Δ invariant, $\overline{H}$ is indeed a subgroup).

We can eliminate windows using

$$\mathrm{Iso}^{\Delta}_{G}(\mathbf{x},\mathbf{y}) = \overline{G'}\overline{\sigma},$$

where $\mathrm{Iso}_{G|_\Delta}(\mathbf{x}|_\Delta,\mathbf{y}|_\Delta) = G'\sigma$*; this is independent from the choice of* $\overline{\sigma}$.

Proof. (a) It is easy from the definition: inside Δ, the permutation $g = g'\sigma \in G\sigma$ sends $\mathbf{x}$ to $\mathbf{y}$ if and only if g' sends $\mathbf{x}^\sigma$ to $\mathbf{y}$, i.e. if and only if it sends $\mathbf{x}$ to $\mathbf{y}^{\sigma^{-1}}$.

(b) It is obvious from the definition, since both sides mean the exact same thing, allowing in both cases g to be either in $G\sigma_1$ or in $G\sigma_2$.

(c) First, we obtain $\mathrm{Iso}^{\Delta_1\cup\Delta_2}_{G\sigma}(\mathbf{x},\mathbf{y}) = \mathrm{Iso}^{\Delta_2}_{G_1\sigma_1}(\mathbf{x},\mathbf{y})$ easily by examining the definitions: both sides simply mean that $g \in G\sigma$ has to respect both windows Δ_1, Δ_2. Then we get $\mathrm{Iso}^{\Delta_2}_{G_1\sigma_1}(\mathbf{x},\mathbf{y}) = \mathrm{Iso}^{\Delta_2}_{G_1}(\mathbf{x},\mathbf{y}^{\sigma_1^{-1}})\sigma_1$ from part (a).

(d) $G'\sigma$ is the collection of permutations of Δ that send $\mathbf{x}$ to $\mathbf{y}$ as far as Δ is able to perceive. Passing to the whole Ω by considering $\overline{G'}$ and $\overline{\sigma}$, the result is the definition itself of $\mathrm{Iso}^{\Delta}_{G}(\mathbf{x},\mathbf{y})$. □

Remark 3.3.7. In the future we are going to need to differentiate the cases of n large and n small. This will come in the form of $C\log^c n \leq n$, for certain $C, c > 0$: if such an inequality is true, which would allow us to have an intermediate integer m between them when needed, then n is considered large. Let us make now this choice.

Assuming CFSG, we suppose that largeness means $102\log^2 n < m \leq n$, which implies $m, n \geq 8308$. See (3.6.8) inside the proof of the main theorem, which is the final quantity to optimize. Without assuming CFSG we suppose instead that largeness means $25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon} < m \leq n$, which implies in particular $m, n \geq 25e^{1/\varepsilon^2}$. For ε small (say $\varepsilon < \frac{1}{10}$), the CFSG-free condition is a stronger restriction.

3.4 Major routines

Before we turn to the algorithm itself, let us describe separately a couple of major routines that were introduced for the first time by Babai. We will not prove their validity here: both [Bab16a] and [Hel19b] do that for us. What we want is to sum up their contribution to the runtime.

We start with a theoretical result, needed to differentiate between the CFSG and the CFSG-free case.

Lemma 3.4.1. *Let $G \leq \mathrm{Sym}(n)$ be primitive, and let $\phi : G \to \mathrm{Alt}(a)$ be an epimorphism.*

(a) Assuming CFSG, if $a > \max\{8, 2 + \log_2 n\}$ then ϕ is an isomorphism.

(b) Not assuming CFSG, if $a > \max\{e^{1/\varepsilon^2}, (\log_2 n)^{4+\varepsilon}\}$ then ϕ is an isomorphism, for any $\varepsilon > 0$ small enough.

Proof. For (a) see [Bab16a, Lemma 8.2.4] or [Hel19b, Lemme 4.1]. For (b) see [Pyb16, Lemma 12], which states that it's sufficient to take $a > \max\{C, \log_2^5 n\}$ for some constant C (his "log" is our "$\log_2$"). Let us compute our version of the bound.

Using [MR96, Thm. C(v)], the sum of the first s primes for $s \geq 6$ is bounded by $\frac{1}{2}s^2(\log s + \log_2 s)$, so $\mathrm{Alt}(a)$ contains a cyclic subgroup whose order is the product of the first $\left\lceil \frac{5}{4}\sqrt{\frac{a}{\log a}} \right\rceil$ primes for $a \geq 10000$ (say). From [Pyb16, Thm. 7] and $a > (\log_2 n)^{4+\varepsilon}$, if ϕ were not an isomorphism we would get $\left\lceil \frac{5}{4}\sqrt{\frac{a}{\log a}} \right\rceil < 2a^{\frac{2}{4+\varepsilon}} < 2a^{\frac{1}{2}-\frac{\varepsilon}{10}}$ (for small ε), which can be true only if $a \leq e^{1/\varepsilon^2}$ (again for small ε). □

A short verification shows that $\varepsilon < \frac{1}{100}$ is plenty enough for the result above to hold.

We are using Lemma 3.4.1 in the computation of the runtime of the following routine. The production and aggregation of *local certificates* (see [Bab16a, §10] or [Hel19b, §6]) is an important part of the algorithm.

Proposition 3.4.2. *Let $G \leq \mathrm{Sym}(n)$, and let $\phi : G \to \mathrm{Alt}(m)$ be an epimorphism; let $\mathbf{x}$ be a string of length n. Then we can find the group $F \leq \mathrm{Aut}_G(\mathbf{x})$ generated by the certificates of fullness in the time taken by $\frac{1}{2}m^{2a}naa!$ calls of the whole algorithm for strings of length $\leq \frac{n}{a}$, where*

(a) $a \in (1.66431, 1.77512) \cdot \log n$, for $102 \log^2 n < m \leq n$ (assuming CFSG), or

(b) $a \in (6.24999, 6.25) \cdot e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$, for $25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon} < m \leq n$ for any $\varepsilon > 0$ small enough (without assuming CFSG),

and in both cases $a \leq \frac{m}{4}$, plus some additional time $O(m^{2a}n^{11})$.

Proof. The proof is contained in [Hel19b, §6.1]. We are going to discuss the details of the runtime.

Let T, T' be two ordered a-tuples of elements of Γ, where a is as in Lemma 3.4.1. Updating one window $A(W)$ relative to the production of the certificate for (T, T') one time takes $\frac{1}{2}aa!$ calls for strings of length $\leq \frac{n}{a}$, and we need to apply also some of the routines in Corollary 3.3.4, which take time $O(n^{10})$ at most. This can happen at most n times for each window (see the end of [Hel19b, §6.1.1]), and the number of windows to update is $\leq m^{2a}$ (see [Hel19b, §6.1.2]), so we obtain the claimed runtime for producing the certificates of fullness. Then, we need to generate F: we simply take the union of the generators of all certificates, but we do it one certificate at a time and we apply Schreier-Sims at every step, so that the number of generators stays quadratic in n (see the observation after Proposition 3.3.3). The certificates of fullness are at most m^{2a}, so this cost is absorbed in the additional time already.

Finally, we need to justify the bounds on a given in the statement. First, by the restrictions on m, n we must have $m, n \geq X$, where $X = 8308$ in the CFSG case and $X = 25e^{1/\varepsilon^2}$ (say) in the CFSG-free case: these are the choices we made in Remark 3.3.7. The conditions then follow, noticing that for our choice of n the two a respect all bounds in Lemma 3.4.1 and for ε small the two intervals are large enough to contain an integer. □

Again, $\varepsilon < \frac{1}{100}$ is plenty enough.

Let us also insert here a short lemma that we will use as part of the aggregation of certificates: it is a classical bound on d-transitivity for non-giants.

Lemma 3.4.3. *Let $G \leq \mathrm{Sym}(n)$ be d-transitive and $G \neq \mathrm{Sym}(n), \mathrm{Alt}(n)$. Then*

(a) $d \leq 5$ (assuming CFSG), or

(b) $d \leq 3 \log n$ (without assuming CFSG).

Proof. See [Cam99, Thm. 4.11] for the CFSG result and [Wie34, Satz C] for the CFSG-free result. □

Then we estimate the cost of another major routine, the one represented by the *Design Lemma* and *Split-or-Johnson* (see [Bab16a, §§6-7] or [Hel19b, §5][3]).

Proposition 3.4.4. *Let $\mathfrak{X}$ be a b-ary coherent configuration on Γ, with $|\Gamma| = m \geq 8308$ and $2 \leq b \leq \frac{1}{2}m$, such that there is no twin class with $> \frac{1}{2}m$ elements. Then we can find either*

(a) a coloured $\frac{2}{3}$-partition of Γ, or

(b) a Johnson scheme of size $\geq \frac{2}{3}m$ inside Γ,

at a multiplicative cost of $m^{b+55 \log m}$ and at an additive cost of $O(m^{b+14})$.

Again, the condition on m is the largeness condition of Remark 3.3.7 (regardless of our position on CFSG).

[3] In parts of the next proof, we use terms from the English version [HBD17] instead of the original French ones. The author thinks the reader is better served by this choice, considering also that Babai's original article is in English.

Proof. As in Proposition 3.4.2, we are going to discuss only the runtime here. The proof of the rest of the statement is contained in [Hel19b, §§5.1-5.2]. The "multiplicative cost" we incur here is the cost of fixing images of a certain number of points of Γ (or parts of a partition of Γ, but fixing the image of a point in the part implies fixing the image of the whole part): arbitrarily fixing a point x in a configuration (or in a graph) in an isomorphism problem translates to trying all possible images of that point, consequently multiplying its contribution. See also Remark 3.5.10.

First, we plug the configuration $\mathfrak{X}$ into the Design Lemma, so that we can pull out a classical configuration to use inside Split-or-Johnson: this involves a multiplicative cost of m^{b-1} at most, and a time of $O(m^b)$ to find the right tuple to use (see [Hel19b, §5.1]). Then, either we terminate by fixing 1 more point (i.e. another multiplicative cost of m) if the new configuration is not primitive, or we call Split-or-Johnson (SoJ, [Hel19b, Thm. 5.3]).

SoJ itself fixes 1 element and then, if it does not terminate, calls Bipartite Split-or-Johnson (BSoJ, [Hel19b, Prop. 5.7]). Call $T(m, v)$ the number of elements fixed by BSoJ when $|V_2| = v$. The base case is $v \leq (6 \log m)^{\frac{3}{2}}$, and here the multiplicative cost is at most $v!$; we use Robbins's bound [Rob55] for factorials,

$$v! < \sqrt{2\pi} v^{v+\frac{1}{2}} e^{-v+\frac{1}{12v}}$$

(the latter being an increasing function), and the cost is in turn bounded by

$$\sqrt{2\pi}(6 \log m)^{\frac{3}{2}\left((6 \log m)^{\frac{3}{2}}+\frac{1}{2}\right)} e^{-(6 \log m)^{\frac{3}{2}}+\frac{1}{12}(6 \log m)^{-\frac{3}{2}}} = m^{f(m) \log m},$$

where

$$f(m) = \frac{3\sqrt{6}(3 \log \log m + 3 \log 6 - 2)}{\sqrt{\log m}} + \frac{3 \log \log m + \log(6^3 \cdot 4\pi^2)}{4 \log^2 m} + \frac{(72\sqrt{6})^{-1}}{\log^{\frac{7}{2}} m}.$$

Now suppose we are outside the base case; first, we apply the Design Lemma again, for a cost of at most

$$v^{6\left\lceil \frac{\log m}{\log v} \right\rceil} < v^{12 \frac{\log m}{\log v}} = m^{12}.$$

Then we fall again into two subcases: either we recur to a new v that is $\leq \frac{2}{3}$ times the old v, with no other cost along the way, or we pass through Coherent Split-or-Johnson (CSoJ, [Hel19b, Prop. 5.8]) and recur to $\leq \frac{1}{2}$ times the old v, with 1 more element fixed in the process (in both cases, it might also happen that we exit the recursion, which is even better). The two situations lead to bounds $T(m, v) \leq m^{12} T\left(m', \frac{2}{3} v\right)$ and $T(m, v) \leq m^{13} T\left(m', \frac{1}{2} v\right)$ respectively, where m' may be smaller than m but still $> \frac{2}{3} m$, or we would exit the recursion again. Since $v < m$ and given the bound in the base case, we obtain in the end

$$T(m, v) \leq m^{f(m) \log m} \cdot \max\left\{m^{12 \log_{3/2} m}, m^{13 \log_2 m}\right\} = m^{\left(f(m) + \frac{12}{\log 3/2}\right) \log m}.$$

As for the additive time incurred during the procedure, the heaviest costs come from the use of the Weisfeiler-Leman algorithm inside BSoJ ([Hel19b, Alg. 3],

see also [WL68] or §2.1), which is performed on a c-ary configuration of V_2 with $c \leq 6 \left\lceil \frac{\log m}{\log v} \right\rceil$, entailing spending $O(c^2 v^{2c+1} \log v) \leq O(m^{13} \log^3 m)$ time for each encounter we have with Weisfeiler-Leman: by what we described before, we call BSoJ at most $O(\log m)$ times, so that we can safely bound the runtime by $O(m^{14})$. All other costs inside SoJ and its relatives (finding twins, colours, etc...) can also be bounded by $O(m^{14})$.

Hence, at the end we incurred a multiplicative cost of $m^{b+\left(f(m)+\frac{12}{\log 3/2}\right)\log m}$ and an additive cost of $O(m^{b+14})$. For $m \geq 8308$ we have $f(m) \leq 24.44853$, and we obtain the bound in the statement. □

3.5 The algorithm

During the whole process, we are working with a pair of strings of the same length $|\Omega|$ and with a group G that respects a system of blocks in Ω; every time we go through the various steps, we are going to either decrease the length of Ω, increase the size of the blocks or decrease the degree of G (in the sense that G will not vary but we will decrease m where $G \leq \mathrm{Sym}(m)$ as abstract groups).

Remark 3.5.1. The case of n small is trivial to examine, and could work as a base case for our algorithm (although we actually follow another path): if $n \leq C$ for some fixed constant C, then we can determine $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ in constant time with constant number of generators.

To achieve this, just try all the permutations of G: we can write all its elements in constant time by Corollary 3.3.4, then check whether each of them sends $\mathbf{x}$ to $\mathbf{y}$. If we do not find one, $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ is empty, otherwise after we find the first one (call it τ) we check which elements of G fix $\mathbf{x}$; the collection of all those that pass the test are all the elements of $\mathrm{Aut}_G(\mathbf{x})$, and they also trivially form a set of generators of $\mathrm{Aut}_G(\mathbf{x})$: since $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \mathrm{Aut}_G(\mathbf{x})\tau$ by Remark 3.3.2 (or by Lemma 3.3.6(a) and $G\tau = G$), we are done.

As we already mentioned, the base case of the atoms ($\mathcal{A}$) will be treated in a different way, as presented in Proposition 3.5.8. Here we need only to cover $n = 1$, which is trivial: this is also an atom, as $\mathrm{Sym}(1) = \mathrm{Alt}(1) = \{\mathrm{Id}_1\}$; from now on we can suppose $n > 1$.

Let us start now with the simplest of recursions, the one with G intransitive.

Proposition 3.5.2. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings. If G is intransitive, we can reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to determining sets $\mathrm{Iso}_{G_i}(\mathbf{x}_i, \mathbf{y}_i)$ such that $\sum_i |\mathbf{x}_i| = \sum_i |\mathbf{y}_i| = n$ and each G_i is transitive. The reduction takes time $O(n^{11})$ and no multiplicative cost.*

Proof. Let Δ be an orbit induced by the action of G on Ω, nonempty and properly contained in Ω since G is intransitive; we can find orbits in time $O(n^3)$ by Lemma 3.3.5. We call $G_1 = G|_\Delta, \mathbf{x}_1 = \mathbf{x}|_\Delta, \mathbf{y}_1 = \mathbf{y}|_\Delta$ the restriction of $G, \mathbf{x}, \mathbf{y}$ to Δ, as in Lemma 3.3.6(d); we suppose that we can compute the set $\mathrm{Iso}_{G_1}(\mathbf{x}_1, \mathbf{y}_1) = H_1\tau_1$. As in Lemma 3.3.6(d), we will use $\overline{\alpha}$ to indicate the object

(or *an* object) whose restriction to a subset of Ω is α: this subset will be either Δ or $\Omega \setminus \Delta$, depending on α; by Corollary 3.3.4(d) with $s = h = 2$, finding $\overline{\alpha}$ from α takes time $O(n^{10})$.

First, by Lemma 3.3.6(d) we have $\mathrm{Iso}_G^{\Delta}(\mathbf{x}, \mathbf{y}) = \overline{H_1}\overline{\tau_1}$; then, by Lemma 3.3.6(c),

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \mathrm{Iso}_{\overline{H_1}}^{\Omega\setminus\Delta}(\mathbf{x}, \mathbf{y}^{\overline{\tau_1}^{-1}})\overline{\tau_1}. \tag{3.5.1}$$

If we can compute

$$\mathrm{Iso}_{\overline{H_1}|_{\Omega\setminus\Delta}}(\mathbf{x}|_{\Omega\setminus\Delta}, \mathbf{y}^{\overline{\tau_1}^{-1}}|_{\Omega\setminus\Delta}) = K_1 \upsilon_1, \tag{3.5.2}$$

we can use again Lemma 3.3.6(d) to plug (3.5.2) inside (3.5.1) and obtain that $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \overline{K_1}\overline{\upsilon_1}\overline{\tau_1}$. The whole process reduces in time $O(n^{10})$ the determination of $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to the determination of Iso sets on the shorter pieces $\Delta, \Omega \setminus \Delta$.

We can repeat the same procedure on the Iso in (3.5.2): notice that the group and the strings are all defined on $\Omega \setminus \Delta$, so if the group $\overline{H_1}|_{\Omega\setminus\Delta}$ is intransitive we again have a $\Delta' \subsetneq \Omega \setminus \Delta$, a group $G_2 = \overline{H_1}|_{\Delta'}$ and strings $\mathbf{x}_2 = \mathbf{x}|_{\Delta'}, \mathbf{y}_2 = \mathbf{y}^{\overline{\tau_1}^{-1}}|_{\Delta'}$ and we continue as before. This happens at most n times.

In the end, we have spent time $O(n^{11})$ and computed sets $\mathrm{Iso}_{G_i}(\mathbf{x}_i, \mathbf{y}_i)$: each G_i is defined in a way that makes it transitive, because we always restrict to an orbit, and each $\mathbf{x}_i, \mathbf{y}_i$ is the restriction of strings $\mathbf{x}, \mathbf{y}^{\sigma}$ to a different part of Ω, so that the sum of their lengths is n. □

The partition of Ω into the orbits of the action of G, and the reduction of the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to problems on shorter strings, corresponds (in reverse, so to speak) to the glueing process of cosets on disjoint sets featured in ($\mathcal{C}2$).

Then, let us continue tackling the next route to recursion, the case of G imprimitive.

Proposition 3.5.3. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings. If G is transitive but imprimitive, call N the stabilizer of a minimal set of blocks: then we can reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to computing the elements of G/N and determining $|G/N|$ sets $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}_i)$ (where N is intransitive). The reduction takes time $O(|G/N|n^{10})$ and no multiplicative cost.*

Proof. Let $\{B_j\}_j$ be a minimal system of blocks for G (it is not a trivial partition since G is imprimitive), which we can retrieve in time $O(n^4)$ by Lemma 3.3.5. Let N be the stabilizer of this system: by Corollary 3.3.4(f), we can compute it in time $O(n^{10})$.

Write $G = \bigcup_i N\sigma_i$, where each σ_i is a representative of a coset of N, so that the number of elements σ_i is $|G/N|$; if we know all the elements of G/N, we can determine each σ_i in time $O(n^{10})$ by Corollary 3.3.4(d) with $s = h = 2$. By Lemma 3.3.6(a)-3.3.6(b),

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \bigcup_i \mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\sigma_i^{-1}})\sigma_i,$$

so we only have to compute the $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\sigma_i^{-1}})$ now; after having done so, we have a description of those sets as $H\tau_i$ where $H = \mathrm{Aut}_N(\mathbf{x})$ is generated by a certain set S, and

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \bigcup_i H\tau_i\sigma_i = \langle S \cup \{\tau_i\sigma_i\sigma_1^{-1}\tau_1^{-1}\}_i\rangle\tau_1\sigma_1.$$

Finally, we can filter the set $S \cup \{\tau_i\sigma_i\sigma_1^{-1}\tau_1^{-1}\}_i$ using the Schreier-Sims algorithm in time $O(n^5+n^3(n^2+|G/N|))$ to obtain a description of $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ with quadratically many generators, and the claim is proved. □

This process, which essentially reduces the problem to a case-by-case examination, corresponds in reverse to the union of cosets featured in ($\mathcal{C}1$). Proposition 3.5.3 cannot be used directly, as a case-by-case reduction is very expensive in general: nevertheless, seeing this reduction process is useful, as it is used when G/N is especially small (Corollary 3.5.7(a), Proposition 3.5.15).

Before going to the key steps of the main algorithm, we introduce a couple of combinatorial lemmas that will be useful in the future. The spirit behind them is to be able to start with the set $\binom{\Gamma}{k}$ of all the k-subsets of some Γ and:

(a) in one case, after finding a partition of Γ, transfer the partition to $\binom{\Gamma}{k}$ itself (Lemma 3.5.4);

(b) in the other case, after identifying Γ with another $\binom{\Gamma'}{k'}$, use this identification to partition $\binom{\Gamma}{k}$ (Lemma 3.5.5).

In the following, a *coloured partition* of a set is a partition in which each part is assigned a colour. A permutation subgroup *respects* a coloured partition if it respects both the partition and the colouring: in other words, for any permutation in the group, the image of any part of a given colour is another part of the same colour.

Lemma 3.5.4. *Let $|\Gamma| = m$ and let $\mathcal{B} = \binom{\Gamma}{k}$, with $k \leq \sqrt{\frac{m}{\log m}}$; suppose that $G \leq \mathrm{Sym}(\Gamma)$ acts on Γ in such a way that there is a coloured partition $\mathcal{C}$ of Γ respected by G and whose parts are of size $\leq \alpha|\Gamma|$ (for some $\alpha \leq \frac{2}{3}$). Then either $m \leq 1045$ or $\mathcal{B}$ has a coloured partition $\mathcal{C}'$, respected by the natural action of G on $\mathcal{B}$, whose parts are of size $\leq \frac{2}{3}|\mathcal{B}|$.*

Proof. Starting from the partition $\mathcal{C}$ of Γ, we can naturally construct the following partition $\mathcal{C}'$ of $\mathcal{B}$: each part of $\mathcal{C}'$ collects the elements of $\mathcal{B}$ (i.e. the k-subsets of Γ) that intersect each part of $\mathcal{C}$ with a specific intersection size; $\mathcal{C}'$ is also naturally a coloured partition: if in a given part $A' \in \mathcal{C}'$ the ordered tuple of intersection sizes with parts $A_i \in \mathcal{C}'$ is $(k_i)_i$, we can give to A' the colour given by the ordered tuple of unordered tuples of intersection sizes for all parts of the same colour for every colour of $\mathcal{C}$ (remember, the fact that G respects $\mathcal{C}$ means that different colours will not mix but different parts of the same colour can be sent to each other).

Now we must prove the claim about the size of the parts $A'_j \in \mathcal{C}'$. Fix any part $A'_0 \in \mathcal{C}'$: from what we said above, all the k-subsets belonging to A'_0 are intersecting the parts of $\mathcal{C}$ in the same number of points, so fix a part $A_0 \in \mathcal{C}$

whose intersection with them is of a certain size $a > 0$. The number of k-subsets of Γ intersecting A_0 in a points is $\binom{|A_0|}{a}\binom{m-|A_0|}{k-a}$, so this is an upper bound for $|A_0'|$: we just have to prove that this number is at most $\frac{2}{3}\binom{m}{k}$ (for m large enough).

If $k = 1$ the task is already accomplished: in this case in fact we also have $a = 1$ and then $|A_0| \leq \alpha m \leq \frac{2}{3}m$. From now on, $k > 1$.

Let us call $|A_0| = \beta m$, where $\beta \leq \alpha \leq \frac{2}{3}$. Then

$$\binom{\beta m}{a}\binom{(1-\beta)m}{k-a} = \frac{1}{k!}\binom{k}{a}\beta m(\beta m - 1)\dots(\beta m - a + 1)\cdot \\ \cdot (1-\beta)m((1-\beta)m-1)\dots((1-\beta)m - k + a + 1).$$

First, since $\beta < 1$ we have obviously $\beta m - i \leq \beta(m-i)$ for all $0 \leq i < a$. On the other hand, for $0 \leq i < k - a$,

$$\frac{(1-\beta)m - i}{(1-\beta)(m-i-a)} = 1 + \frac{a - (i+a)\beta}{(1-\beta)(m-i-a)} \;\leq\; 1 + \frac{a}{m-i-a} \\ \leq 1 + \frac{k}{m-k} \;<\; 1 + \frac{2k}{m},$$

so that

$$\binom{\beta m}{a}\binom{(1-\beta)m}{k-a} < \frac{1}{k!}\binom{k}{a}\beta^a m(m-1)\dots(m-a+1)\cdot \\ \cdot (1-\beta)^{k-a}(m-a)\dots(m-k+1)\left(1+\frac{2k}{m}\right)^{k-a} \\ = \binom{m}{k}\binom{k}{a}\beta^a(1-\beta)^{k-a}\left(1+\frac{2k}{m}\right)^{k-a}. \tag{3.5.3}$$

The last factor can be easily bounded in the following way:

$$\left(1+\frac{2k}{m}\right)^{k-a} < \left(1+\frac{2}{\sqrt{m\log m}}\right)^{k} \\ = \left(1+\frac{2}{\sqrt{m\log m}}\right)^{\frac{\sqrt{m\log m}}{2}\cdot\frac{2k}{\sqrt{m\log m}}} \\ < e^{\frac{2}{\log m}}.$$

Let us treat the rest now. We are going to prove that

$$\binom{k}{a}\beta^a(1-\beta)^{k-a} \leq \frac{1}{2}. \tag{3.5.4}$$

First, we start with the case $k \geq 5$ and $2 \leq a \leq k-2$, implying that $a \geq 2$, $k - a \geq 2$ with at least one being a strict inequality. We have

$$\frac{(1-\beta)a}{\beta(k-a+1)} + \frac{\beta(k-a)}{(1-\beta)(a+1)} = \frac{(1-\beta)^2 a(a+1) + \beta^2(k-a)(k-a+1)}{\beta(1-\beta)(a+1)(k-a+1)}$$

$$> \frac{(1-\beta)^2a^2+\beta^2(k-a)^2}{\beta(1-\beta)a(k-a)} \cdot \frac{a}{a+1}\frac{k-a}{k-a+1}.$$

The first fraction is of the form $\frac{x^2+y^2}{xy}$, which is equal to $\frac{(x-y)^2}{xy}+2 \geq 2$; as for the other two, they are both $\geq \frac{2}{3}$ and at least one is $\geq \frac{3}{4}$: therefore the whole product is ≥ 1. This means that

$$\begin{aligned}
1 = (\beta+1-\beta)^k &= \sum_{a'=0}^{k} \binom{k}{a'} \beta^{a'}(1-\beta)^{k-a'} \\
&> \sum_{a' \in \{a-1,a,a+1\}} \binom{k}{a'} \beta^{a'}(1-\beta)^{k-a'} \\
&= \binom{k}{a} \beta^a (1-\beta)^{k-a} \left(\frac{(1-\beta)a}{\beta(k-a+1)} + 1 + \frac{\beta(k-a)}{(1-\beta)(a+1)} \right) \\
&\geq 2\binom{k}{a} \beta^a (1-\beta)^{k-a},
\end{aligned}$$

and (3.5.4) is proved in this case. For $k=4$ and $a=2$,

$$\frac{2(1-\beta)}{3\beta} + \frac{2\beta}{3(1-\beta)} = \frac{2}{3}\frac{(1-\beta)^2+\beta^2}{\beta(1-\beta)} \geq \frac{4}{3} > 1,$$

and we are done as before. Now, let $a=1$ or $a=k-1$: we can suppose $a=k-1$ by exchanging the role of β and $1-\beta$ if necessary (although we cannot use the bound $\beta \leq \frac{2}{3}$ anymore); $k\beta^{k-1}(1-\beta)$ has a maximum in $\beta = 1-\frac{1}{k}$, in which it is equal to $\frac{k}{k-1}\left(1-\frac{1}{k}\right)^k$. The factor $\left(1-\frac{1}{k}\right)^k$ is bounded from above by $\frac{1}{e}$, so for $k \geq 4$ we obtain the bound $< \frac{1}{2}$; for $k = 2, 3$ we just check directly obtaining $\frac{1}{2}, \frac{4}{9}$ respectively. Finally, let $a = k$: then we have just β^k, which is $\leq \beta^2 \leq \frac{4}{9}$, and (3.5.4) is proved for all cases.

Plugging our results into (3.5.3),

$$\binom{\beta m}{a}\binom{(1-\beta)m}{k-a} < \binom{m}{k}\frac{1}{2}e^{\frac{2}{\log m}},$$

and for $m \geq 1046$ we obtain $\frac{1}{2}e^{\frac{2}{\log m}} < \frac{2}{3}$. □

Given our choice of large m, n inside Remark 3.3.7, Lemma 3.5.4 applies any time we are assuming $m > C\log^e n$ for the appropriate C, e.

Lemma 3.5.5. *Let Γ' be a set, let $\Gamma = \binom{\Gamma'}{k'}$ for some $2 \leq k' \leq \frac{|\Gamma'|}{2}$, and let $\mathcal{B} = \binom{\Gamma}{k}$ for some $2 \leq k \leq \frac{|\Gamma|}{2}$; suppose that $|\Gamma'| = m' \geq 12$. Let any permutation of Γ' induce the natural permutations of Γ and $\mathcal{B}$; then any $H \leq \mathrm{Sym}(\Gamma')$ divides $\mathcal{B}$ into a system of orbits and blocks such that each part is $\leq \frac{1}{2}|\mathcal{B}|$.*

Proof. Let Δ be any orbit of $\mathcal{B}$ under the action given in the statement. Any element $x \in \Delta$ is a k-set of k'-sets of elements of Γ': since every $x' \in \Delta$ can be

sent to x by some permutation induced by some $h \in H$, all the elements of Δ are constructed respecting the same equalities among the elements of their elements (for example, if there are $a_1, a_2 \in x$ with $b_1, b_2, b_3 \in a_1 \cap a_2$, then any x' also has a'_1, a'_2 with $b'_1, b'_2, b'_3 \in a'_1 \cap a'_2$, and so on). Every orbit Δ is therefore contained in the subset $\mathcal{B}_r \subseteq \mathcal{B}$ of elements of $\mathcal{B}$ respecting some given set of relations r; if we prove that either $\mathcal{B}_r$ is of size $\leq \frac{1}{2}|\mathcal{B}|$ or can be divided into blocks with the same property, the same will hold for Δ and we would be done.

For any $x \in \mathcal{B}_r$, let $A(x) \subseteq \Gamma'$ be the set of the elements of all the elements of x, with $|A(x)| = a$ (a does not depend on x since it is determined by the relations r); we divide $\mathcal{B}_r$ into blocks, where each of them collects all the x with the same $A(x)$: these are really blocks, in the sense that the elements of $\mathcal{B}_r$ inside them move together under the action of H since this movement depends ultimately on where $A(x)$ is moved inside Γ'. We have to exclude that the so formed block system is trivial, i.e. that either the blocks have size 1 or that the whole $\mathcal{B}_r$ is a block: if we do it, we are done.

Having blocks of size 1 means that each x already collects all the possible k'-subsets of its own $A(x)$, so that x is its own only permutation under $\mathrm{Sym}(A(x))$: this means that $k = \binom{a}{k'}$ and that $\mathcal{B}_r$ has $\binom{m'}{a}$ elements, one for each $A(x)$. $\mathcal{B}$ has $\binom{|\Gamma|}{k}$ elements, where $|\Gamma| = \binom{m'}{k'}$, so to prove the statement in this case it is sufficient to prove that

$$\binom{m'}{a} \leq \frac{1}{2}\binom{\binom{m'}{k'}}{\binom{a}{k'}}, \tag{3.5.5}$$

and we would have shown that $\mathcal{B}_r$ is small.

Since $k \geq 2$ there are at least two distinct k'-subsets of Γ' participating in the formation of $A(x)$, so $a > k'$ and then $a \leq \binom{a}{k'}$; we also recall the easy bounds $\left(\frac{x}{y}\right)^y \leq \binom{x}{y} \leq \left(\frac{ex}{y}\right)^y$. Then, since $m' \geq 12$, $2 \leq k' \leq \frac{m'}{2}$, $k \leq \frac{|\Gamma|}{2}$ and $\left(\frac{11}{2e}\right)^a > 2$, we obtain

$$\binom{\binom{m'}{k'}}{\binom{a}{k'}} \geq \binom{\binom{m'}{k'}}{a} \geq \left(\frac{\binom{m'}{k'}}{a}\right)^a \geq \left(\frac{\frac{11}{2}m'}{a}\right)^a > 2\left(\frac{em'}{a}\right)^a \geq 2\binom{m'}{a}, \tag{3.5.6}$$

and (3.5.5) is proved.

Having $\mathcal{B}_r$ as a whole block means that all the $x \in \mathcal{B}_r$ are coming from the same $A(x)$; as $\mathcal{B}_r$ just collects all elements of $\mathcal{B}$ with the same relations, with no other discriminating condition, $A(x)$ must be the whole Γ'. For each $x \in \mathcal{B}_r$ and $\gamma \in \Gamma'$, call $N(\gamma, x)$ the number of elements of x that contain γ: the multiset $\{N(\gamma, x) | \gamma \in \Gamma'\}$ is independent from x, since it is a reflection of the relations of $\mathcal{B}_r$.

Suppose first that such multiset has all equal elements, i.e. every γ is contained in the same number N of k'-subsets of Γ' belonging to a fixed x (or to any x, given our hypotheses): this is a rather constraining condition in $\mathcal{B}$, so we will show that $\mathcal{B}_r$ is small. Consider the set $\mathcal{C}_1 \subseteq \mathcal{B}$ of all x with multiset $\{N, N, N, \ldots, N\}$ (m' times), so that $\mathcal{B}_r \subseteq \mathcal{C}_1$, and consider the set $\mathcal{C}_2 \subseteq \mathcal{B}$ of all x with multiset $\{N+1 \ldots, N+1, N-1 \ldots, N-1, N, \ldots, N\}$, where the number k'' of $N+1$ is equal to the number of $N-1$ and runs among all $1 \leq k'' \leq k'$: construct the bipartite

graph $\mathcal{C}_1 \cup \mathcal{C}_2$ where $\{x_1, x_2\}$ is an edge if and only if we can change exactly one k'-subset inside x_1 to obtain x_2. Every $x_1 \in \mathcal{C}_1$ has $k\left(\binom{m'}{k'} - k\right) \geq \binom{m'}{k'}$ neighbours, since we can move each of the k'-subsets of x_1 to any of the k'-subsets that are not already in x_1 and obtain some (distinct) element of $\mathcal{C}_2$; on the other hand, the number of neighbours of a given x_2 is at most $\binom{m'-2k''}{k'-k''}$: in fact, each k'-subset that contains all the γ with $N+1$ can be moved only in one way to produce an element of $\mathcal{C}_1$, namely by replacing the γ with $N+1$ with the γ with $N-1$ and fixing the other ones, and the number of such subsets is bounded by $\binom{m'-2k''}{k'-k''}$. Provided that $b_i \leq \frac{1}{2}a_i$, $a_1 \leq a_2$ and $b_1 \leq b_2$ imply $\binom{a_1}{b_1} \leq \binom{a_2}{b_2}$; therefore

$$\binom{m'}{k'}|\mathcal{C}_1| \leq |\{\text{edges of } \mathcal{C}_1 \cup \mathcal{C}_2\}| \leq \binom{m'-2k''}{k'-k''}|\mathcal{C}_2| \leq \binom{m'}{k'}|\mathcal{C}_2|,$$

and since $\mathcal{C}_1$ and $\mathcal{C}_2$ are disjoint we obtain $|\mathcal{B}_r| \leq \frac{1}{2}|\mathcal{B}|$.

Now suppose that the multiset $\{N(\gamma, x) | \gamma \in \Gamma'\}$ has at least two distinct elements; take the least frequent of these elements (or the smallest of the least frequent ones, if more than one exists), say that there are k'' of them with $k'' \leq \frac{m'}{2} < \frac{kk'}{2}$: the second inequality comes from the fact that $A(x) = \Gamma'$, implying that $kk' \geq m'$, and that equality is excluded because it would imply $N(\gamma, x) = 1$ regardless of γ. Call $A'(x)$ the set of γ with this specified N for x; $A'(x)$ is properly contained in Γ', so there must exist elements x with different $A'(x)$: we collect elements $x \in \mathcal{B}$ based on their $A'(x)$, and as we said before for $A(x)$ this forms a system of blocks, which are not the whole $\mathcal{B}$ since $A'(x) \neq \Gamma'$. We have to exclude that this system has blocks of size 1.

Assume that these blocks have indeed size 1, which means that $|\mathcal{B}_r| = \binom{m'}{k''}$ (one element for each $A'(x)$); as before, we have to prove that

$$\binom{m'}{k''} \leq \frac{1}{2}\binom{\binom{m'}{k'}}{k}.$$

When $k'' \leq k'$ we have $\binom{\binom{m'}{k'}}{k} \geq \binom{\binom{m'}{k''}}{k} > 2\binom{m'}{k''}$, and when $k'' \leq k$ we can say $\binom{\binom{m'}{k'}}{k} \geq \binom{\binom{m'}{k'}}{k''}$ and continue as in (3.5.6), so we can assume $k'' > k, k'$; this also excludes the cases $k = 2$ and $k' = 2$, using $k'' < \frac{kk'}{2}$. Let us start with the case $\frac{m'}{k'} > 4$; using $k \geq \lceil \frac{m'}{k'} \rceil$, the bounds on binomial coefficients and $m' \geq 12$,

$$\binom{\binom{m'}{k'}}{k} \geq \binom{\binom{m'}{k'}}{\lceil \frac{m'}{k'} \rceil} \geq \left(\frac{m'}{k'}\right)^{\frac{(k'-1)m'}{k'}} > 4^{\frac{2}{3}m'} > 2(\sqrt{2e})^{m'} \geq 2\binom{m'}{\lfloor \frac{m'}{2} \rfloor} \geq 2\binom{m'}{k''}.$$

Similarly, for $3 < \frac{m'}{k'} \leq 4$ (implying $k \geq 4$),

$$\binom{\binom{m'}{k'}}{k} \geq \binom{\binom{m'}{k'}}{4} \geq 4^{m'-4} \geq 4^{\frac{2}{3}m'} > 2\binom{m'}{k''}.$$

For $2 \leq \frac{m'}{k'} \leq 3$ and $k \geq 4$,

$$\binom{\binom{m'}{k'}}{k} \geq \binom{\binom{m'}{k'}}{4} \geq \frac{3^{\frac{4}{3}m'}}{4^4} \geq 4^{\frac{2}{3}m'} > 2\binom{m'}{k''}.$$

Finally, for $2 \leq \frac{m'}{k'} \leq 3$ and $k = 3$, we can first check directly that

$$\binom{\binom{m'}{k'}}{3} \geq \binom{\binom{m'}{\lceil \frac{m'}{3} \rceil}}{3} \geq 2\binom{m'}{\lfloor \frac{m'}{2} \rfloor} \geq 2\binom{m'}{k''}$$

for each $12 \leq m' < 16$, while for $m' \geq 16$

$$\binom{\binom{m'}{k'}}{3} \geq 3^{m'-3} > 2(\sqrt{2e})^{m'} \geq 2\binom{m'}{k''}.$$

Since $\frac{m'}{k'} \geq 2$ is always true, this covers all cases and concludes the proof. □

We are now at a point where we must introduce the cornerstone of the algorithm, the group-theoretic result thanks to which the branching into different cases starts and the recursion is performed. Actually, as anticipated, we have two of them: Theorem 3.5.6 assumes CFSG and Theorem 3.5.14 does not; consequently, henceforth we split our reasoning into two different parts, according to our attitude towards CFSG: the two approaches present many points of contact with each other nonetheless, enough to make the proof of the main theorem virtually the same both times.

3.5.1 The algorithm, assuming CFSG

Let us start immediately with our theoretic main tool.

Theorem 3.5.6. *Let $|A| = a$ and let $G \leq \mathrm{Sym}(A)$. Assume CFSG. If G is primitive, then one of the following alternatives holds:*

(a) $|G| \leq C(a) = \max\{C_0, a^{1+\log_2 a}\}$ *for* $C_0 = 244823040$*;*

(b) there is a system $\mathcal{A}$ of (possibly size 1) blocks of A with $|\mathcal{A}| = \binom{b}{t} \leq a$ and there is a $G' \trianglelefteq G$ with $[G : G'] \leq a$ and preserving $\mathcal{A}$, such that we can construct in time $O(n^{10})$ a bijection φ between $\mathcal{A}$ and the set $\binom{B}{t}$ of t-subsets of a b-set B in a way that makes G' isomorphic to $\mathrm{Alt}(B)$, with the action of G' on $\mathcal{A}$ agreeing with the natural action induced by $\mathrm{Alt}(B)$ on $\binom{B}{t}$.

Proof. This theorem is a consequence of Cameron's classification of primitive permutation groups in its formulation due to Maróti, which we have already displayed as Theorem 1.2.5. Case (a) in the present result collects cases (b) and (c) in the other one, and C_0 is the size of the largest of the four Mathieu groups that appear in (b), namely M_{24}. The other alternative is realized by taking Theorem 1.2.5(a) and choosing G' to be the power of $\mathrm{Alt}(B)$ that is guaranteed to exist as a subgroup of G; the rest of the structure is retrieved by creating the partition $\mathcal{A}$ with one block for each of the possible values of the first coordinate (say) in the formulation of the wreath product in Definition 3.1.3, and forgetting the structure coming from all other coordinates, so that we see only one $\mathrm{Alt}(B)$ among all the ones that compose G'.

As for the polynomial-time construction of φ, it is described in [BLS87, §4] (see also [Hel19b, §2.8]). The procedure NATURAL_ACTION thereby described

produces a set D divided into blocks $\{B_i | i \in I\}$ such that the elements of A correspond to subsets of D of a certain form; our B is any of the B_i (say B_1) and if $\pi : G \to \mathrm{Sym}(I)$ is the map describing how G permutes the B_i then our G' is $\pi^{-1}(\mathrm{Sym}(I)_{(1)})$. All the passages involved in finding B and G' and constructing φ come from Corollary 3.3.4 and Lemma 3.3.5 (on sets of size at most $|A|^2$): together, they cost at most time $O(n^{10})$ as claimed. □

When we start the whole algorithm to compute $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$, we can divide G into its orbits and blocks (if G is intransitive or imprimitive) in time $O(n^4)$ by Lemma 3.3.5, and then treat the intransitive case thanks to Proposition 3.5.2: therefore we can suppose that G is transitive and acts primitively on some system of blocks $\mathcal{B}$ that we are able to assume to be known.

Corollary 3.5.7. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks of Ω with $1 < |\mathcal{B}| = r \leq n$, on which G acts primitively: call N the stabilizer of the system $\mathcal{B}$, and suppose that there are a set Γ of size m and a bijection between $\mathcal{B}$ and $\binom{\Gamma}{k}$ (for some k) such that the action of G/N on $\mathcal{B}$ corresponds to the action of some transitive subgroup $H \leq \mathrm{Sym}(\Gamma)$ on $\binom{\Gamma}{k}$. Assume CFSG. Then we can reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to one of the following problems:*

(a) determining $\leq m^{102 \log^2 n}$ sets of isomorphisms $\mathrm{Iso}_M(\mathbf{x}, \mathbf{y}_i)$, where $M \trianglelefteq N$ stabilizes all blocks, in time $O(m^{102 \log^2 n} n^{10})$ and at no multiplicative cost;

(b) determining $\leq m$ sets of isomorphisms $\mathrm{Iso}_{G'}(\mathbf{x}, \mathbf{y}_i)$, where G' respects a system of orbits and/or blocks $\mathcal{B}'$ strictly coarser than $\mathcal{B}$ and whose parts are of size $\leq \frac{2}{3}|\Omega|$, in time $O(n^{10})$ and at no multiplicative cost;

(c) determining $\leq m$ sets of isomorphisms $\mathrm{Iso}_{G'}(\mathbf{x}, \mathbf{y}_i)$, where G'/N acts on $\mathcal{B}$ in the same way as $\mathrm{Alt}(\Gamma')$ acts on $\binom{\Gamma'}{k'}$ (where $|\Gamma'| = m' > 102 \log^2 n$), in time $O(n^{10})$ and at no multiplicative cost.

Proof. Before we start, we point out that we hypothesize the existence of Γ in the statement (or, from another perspective, the fact that k may be ≥ 2) because we want to leave open the possibility that we are returning to this situation after having already been through this step before and found a bijection as in Theorem 3.5.6(b) (using the theorem itself or by other means) that we have then carried forth until this moment, as it may happen. In any case, either we are provided with such $\Gamma, k, \mathcal{B}, N$ from past procedures, or in their absence we can determine $\mathcal{B}, N$ in time $O(n^{10})$ by Lemma 3.3.5 and Corollary 3.3.4(f) (setting $\mathcal{B} = \{\{x\} | x \in \Omega\}$ if G is primitive) and then impose $\Gamma = \mathcal{B}$ and $k = 1$.

As it can be imagined, we want to use Theorem 3.5.6 on $A = \Gamma$. First, H must be primitive: if it were not, then its action on $\binom{\Gamma}{k}$ would also be imprimitive (even intransitive, if $k > 1$) and this contradicts our hypothesis on G; hence we can actually use the theorem. The generators of G (at most n^2 in number) can be seen as generators of $G/N \simeq H$ and can be processed through Schreier-Sims to determine $|H|$ in time $O(n^5)$ by Corollary 3.3.4(a), so that we are able to determine whether we are in case (a) or (b) of Theorem 3.5.6.

If we are in case (a), we can write all the elements of H in time $O(n^5+C(m)n^2)$ by Corollary 3.3.4 and we are exactly in the situation described in Proposition 3.5.3 (with the computation of all the elements of $H \simeq G/N$ already taken care of). This falls into case (a) of the present corollary: we have $N = M$ for the subgroup; also, for $n \leq 3$ obviously $|G/N| \leq m! \leq m^{102\log^2 n}$, while for $n \geq 4$ both $C_0 < 2^{102\log^2 4} \leq m^{102\log^2 n}$ and $m^{1+\log_2 m} < m^{102\log^2 n}$, so the bound on the number of problems holds. The runtime, in light of the previous reasoning on $C(m)$, is also $O(m^{102\log^2 n}n^{10})$ as required.

If we are in case (b), there is some $H' \trianglelefteq H$ with $[H : H'] \leq m$ acting on a partition Γ^o of Γ as $\mathrm{Alt}(\Gamma')$ acts on $\binom{\Gamma'}{k'}$ for some $|\Gamma'| = m'$ and some $k' \geq 1$: Γ^o, Γ', k' and the action are all found in time $O(n^{10})$, as we already said. First, suppose that $m \leq 102\log^2 n$: then $|G/N| < m^m \leq m^{102\log^2 n}$, and repeating what we did before we retrieve again case (a).

Now suppose that $m > 102\log^2 n$ and that Γ^o is a nontrivial partition: as observed in Remark 3.3.7 we have $m \geq 8308$, and the hypothesis on Γ^o makes it into a coloured partition (with only one colour) whose parts are of size $\leq \frac{1}{2}|\Gamma|$; to use Lemma 3.5.4, we still have to prove that $k \leq \sqrt{\frac{m}{\log m}}$. For $k = 1$ this is true for any m, so suppose that $k > 1$. Obviously we can assume that $m \geq 2k$: in fact there is a natural identification between $\binom{\Gamma}{k}$ and $\binom{\Gamma}{|\Gamma|-k}$, just by taking the complement of each of their elements; therefore

$$n \geq \binom{m}{k} \geq \left(\frac{m}{k}\right)^k \geq 2^k \implies k \leq \frac{1}{\log 2}\log n \implies m > k^2 \cdot 102\log^2 2 > k^2,$$

and, using this new bound again,

$$n \geq \binom{m}{k} \geq \left(\frac{m}{k}\right)^k > k^k \implies \log n > k\log k.$$

The function $f(y) = \frac{y}{\sqrt{\log y}}$ is increasing and $f(k\log k) > k$ for $k > 1$, therefore using $k\log k < \log n \leq \sqrt{\frac{1}{102}m}$ we get $k < \sqrt{\frac{m}{102\log\sqrt{\frac{1}{102}m}}} < \sqrt{\frac{m}{\log m}}$ (where $m \geq 8308$ is amply sufficient to satisfy the second inequality). Now we are free to use Lemma 3.5.4, which makes us fall into case (b) of the present corollary.

Finally, let us have $m > 102\log^2 n$ and $\Gamma^o = \Gamma$: since $m \geq 8308$ and $m = \binom{m'}{k'}$, we have $m' \geq 12$ regardless of our choice of k'. If both k and k' are > 1, we can use Lemma 3.5.5 and we fall again into case (b). If $k' = 1$, then $\Gamma' = \Gamma$ and H' acts as $\mathrm{Alt}(\Gamma)$ on Γ itself, thus acting as $\mathrm{Alt}(\Gamma)$ on $\binom{\Gamma}{k} \simeq \mathcal{B}$. If $k = 1$, then $\Gamma = \mathcal{B}$ and H' acts as $\mathrm{Alt}(\Gamma')$ on $\binom{\Gamma'}{k'} \simeq \mathcal{B}$; if $m' \leq 102\log^2 n$ we reduce again to case (a) exactly as before, so $m' > 102\log^2 n$. In both cases, whether $k' = 1$ or $k = 1$, we can take the pullback G' of H' in G (in time $O(n^{10})$ by Corollary 3.3.4(d)) and $G'/N \simeq H'$ will satisfy the requirements of case (c) of this corollary: in fact $[G : G'] = [H : H']$ and we can obtain (a preimage of) all the elements of G/G' in time $O(n^{10})$, continuing then with $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \bigcup_i \mathrm{Iso}_{G'}(\mathbf{x}, \mathbf{y}^{\sigma_i^{-1}})\sigma_i$ as in Proposition 3.5.3. □

We point out that [Hel19b] uses actually a bound on m of the form $m > C \log n$ for the case equivalent to our case (c). In order to follow our line of thought we need a stronger bound, quadratic in $\log n$, because otherwise we obtain a weaker inequality than $k \leq \sqrt{\frac{m}{\log m}}$ and then Lemma 3.5.4 does not work: the issue is with the last factor in (3.5.3), which needs to decrease with the growth of m; the problem is treated incorrectly in [Hel19b, §4.2]. A bound $m > C \log n$ is more than we need to obtain the bound on the runtime of the form $n^{O(\log^2 n)}$ anyway: as observed in [Hel19b, §3.1], it is consistent even with a $n^{O(\log n)}$ runtime, to this day unproven.

After we have reached case (a) in the previous corollary, we can simply go through Proposition 3.5.2 and reduce to examinate each block singularly: this makes n decrease, and we return to the top of this corollary. After case (b), Ω is divided into orbits and blocks that are coarser than the original $\mathcal{B}$: this makes n decrease or the block size increase (or both). Case (c) is the one we will examine in the following results.

Proposition 3.5.8. *Let $|\Omega| = n$, and let the action of $G \leq \mathrm{Sym}(\Omega)$ on Ω be as in Corollary 3.5.7*(c)*, i.e. there is a system of blocks $\mathcal{B}$ with stabilizer N such that G/N acts on it as $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$, where $|\Gamma| = m$ and $|\mathcal{B}| = \binom{m}{k}$. If $k = 1$ and the blocks have size 1, the set $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ can be determined in time $O(n^6)$ with at most n^2 generators.*

Proof. Having $k = 1$ means that $\Gamma = \Omega$, and having block size 1 means that $G = G/N \simeq \mathrm{Alt}(\Gamma) = \mathrm{Alt}(\Omega)$. This is a trivial case: if $\mathbf{x}$ and $\mathbf{y}$ do not send the same number of elements of Ω to the same letter of the alphabet Σ, the set is empty.

Otherwise, we first obtain $\mathrm{Aut}_{\mathrm{Sym}(\Omega)}(\mathbf{x})$ as a product $\prod_i \mathrm{Sym}(\Delta_i)$, where the Δ_i are the parts of Ω whose elements are sent by $\mathbf{x}$ to the same letter: more precisely, for each generator of $\mathrm{Sym}(\Delta_i)$ we find the corresponding element in $\mathrm{Sym}(\Omega)_{(\Omega \setminus \Delta_i)}$, and then we take the union of these preimages for all i; each $\mathrm{Sym}(\Delta_i)$ can be described by two generators, a transposition and a cycle of length $|\Delta_i|$, therefore up until now we are working with $\leq \frac{2}{3}n$ generators. Then, we find $H = \mathrm{Aut}_{\mathrm{Alt}(\Omega)}(\mathbf{x})$: by Corollary 3.3.4(c), since the index is ≤ 2 and the test to prove whether a permutation is even is linear-time (just by computing the length of the cycles), we obtain polynomially many generators of H in time $O(n^5)$; more precisely, the number of generators is at most $\left(\frac{2}{3}n + 1\right)^3$ by Schreier's lemma ([Sch27], see for example [Ser03, Lemma 4.2.1]) and we can reduce it to $\leq n^2$ using Schreier-Sims and spending time $O(n^6)$ by Proposition 3.3.3.

Finally we take any bijection $\pi : \Omega \to \Omega$ sending elements sent to each letter of Σ by $\mathbf{x}$ to the elements sent to the same letter by $\mathbf{y}$. If this bijection is in $\mathrm{Alt}(\Omega)$ we have $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = H\pi$; if it is not, there are two possibilities: if there is a letter that appears twice in the strings (say $\mathbf{x}(r_1) = \mathbf{x}(r_2)$) we have $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = H\tau\pi$ where τ is the transposition $(r_1\ r_2)$, otherwise the set is empty again. □

The situation described in Proposition 3.5.8 (apart from the case of $n = 1$ taken care of in Remark 3.5.1) is the only true base case of the whole algorithm; the rest of

the time, the procedure either stops and gives $\emptyset$ as a result or it reduces to simpler cases, until we arrive to the one given above. Proposition 3.5.8 corresponds to the case of the atom ($\mathcal{A}$) in the main theorem.

Let us see what happens aside from the base case.

Theorem 3.5.9. *Let $|\Omega| = n$, let $\mathbf{x} : \Omega \to \Sigma$ be a string, and let the action of $G \leq \mathrm{Sym}(\Omega)$ on Ω be as in Corollary 3.5.7(c), i.e. there is a system of blocks $\mathcal{B}$ such that G acts on it as $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$, where $|\Gamma| = m$ and $|\mathcal{B}| = \binom{m}{k}$. Assume CFSG; suppose also that $m > 102 \log^2 n$. Then we can reduce to one of the following cases:*

(a) Γ has a canonical coloured partition in which each part has size $\leq \frac{1}{2}|\Gamma|$;

(b) there is a canonical set $S \subseteq \Gamma$ of size $> \frac{1}{2}|\Gamma|$ such that for any $\sigma \in \mathrm{Alt}(S)$ there is an element of $\mathrm{Aut}_G(\mathbf{x})$ that induces σ on S;

(c) at a multiplicative cost of at most $m^{57 \log n}$, either:

(c1) Γ has a coloured partition in which each part has size $\leq \frac{2}{3}|\Gamma|$, or

(c2) there are two disjoint sets $V_1, V_2 \subseteq \Gamma$, with V_2 divided into a system of blocks $\mathcal{G}$ with $\binom{|\mathcal{G}|}{k'} = |V_1| \geq \frac{2}{3}|\Gamma|$ for some $k' \geq 2$, and there is a bijection between V_1 and $\binom{\mathcal{G}}{k'}$ such that if a $g \in G$ induces a permutation $\sigma \in \mathrm{Sym}(\mathcal{G})$ of the blocks then it also induces the corresponding permutation of V_1 through the identification of its elements with the k'-subsets of $\mathcal{G}$.

The time necessary for this reduction is the cost of $\frac{1}{2}m^{2a}naa!$ calls of the whole algorithm for strings of length $\leq \frac{n}{a}$ where $a \in (1.66431, 1.77512) \cdot \log n$, plus some additional time $O(m^{3a}n^{11})$.

Proof. We are in the scenario of Proposition 3.4.2: the construction in our hypothesis yields in particular a surjective map $\phi : G \to \mathrm{Alt}(\Gamma)$. After $\frac{1}{2}m^{2a}naa!$ calls of the algorithm for strings of length $\leq \frac{n}{a}$ and an additional time of $O(m^{2a}n^{11})$, we have obtained the group $F \leq \mathrm{Aut}_G(\mathbf{x})$ generated by all certificates of fullness. Now we follow the case subdivision in [Hel19b, §6.2].

- "Cas 1" is the case of $|\mathrm{supp}(\phi(F))| \geq \frac{1}{2}m$ and no orbit of $\phi(F)$ of length $> \frac{1}{2}m$.
- "Cas 2a" is the case of $|\mathrm{supp}(\phi(F))| \geq \frac{1}{2}m$, an orbit Φ of $\phi(F)$ of length $> \frac{1}{2}m$, and $\mathrm{Alt}(\Phi) \leq \phi(F)|_\Phi$.
- "Cas 2b" is the case of $|\mathrm{supp}(\phi(F))| \geq \frac{1}{2}m$, an orbit Φ of $\phi(F)$ of length $> \frac{1}{2}m$, and $\mathrm{Alt}(\Phi) \not\leq \phi(F)|_\Phi$.
- "Cas 3" is the case of $|\mathrm{supp}(\phi(F))| < \frac{1}{2}m$.

In "Cas 1" we colour each element of Γ by the length of its orbit (in time $O(m^3)$ by Lemma 3.3.5) and we are in our case (a). "Cas 2a" is our case (b) for $S = \Phi$.

"Cas 2b" starts by arbitrarily fixing some points of Γ, precisely $d-1$ many for d as in Lemma 3.4.3(a), and then feeds the resulting configuration to the Split-or-Johnson procedure (without passing through the Design Lemma). In "Cas 3", the information we already have at hand after the production of the local certificates lets us have a colouring of $(\Gamma \setminus \mathrm{supp}(\phi(F)))^a$ with less than half twins (as long as $a \leq \frac{m}{4}$): we can make it into an a-ary configuration and refine it through Weisfeiler-Leman at a cost of $O(a^2 m^{2a+1} \log m)$ for the runtime, and then invoke the Design Lemma plus Split-or-Johnson.

In both cases, we can apply Proposition 3.4.4: the two alternatives (a) and (b) therein correspond respectively to cases (c1) and (c2) here. We have explicitly written in our statement what the sentence "nous pouvons trouver [...] un schéma de Johnson plongé sur [...] Γ" means in the statement of [Hel19b, Thm. 5.3]: in particular, the fact that the objects that when permuting induce a permutation of V_1 may be the parts of $\mathcal{B}'$ (instead of being directly the elements of V_2) is due to the use of [Hel19b, Ex. 2.18] inside CSoJ, where from a graph made of elements of V_2 we pass to a contracted graph made of its parts.

The multiplicative cost of "Cas 2b" and "Cas 3" is bounded by $m^{a+55\log m}$ (certainly $d \leq a$, so the "Cas 2b" expense is subsumed by the "Cas 3" expense), and their additive cost is safely absorbed into the $O(m^{3a}n^{11})$. For our choice of a, we obtain the cost featured in (c). □

Remark 3.5.10. The multiplicative cost described in case (c) of Theorem 3.5.9 means the following: since a permutation in G induces also an even permutation of Γ, for any choice of s points $x_1, \dots, x_s \in \Gamma$ each isomorphism from $\mathbf{x}$ to $\mathbf{y}$ falls into a particular coset of the stabilizer of these points; these cosets are one for each possible choice of images of the points in Γ.

Call N the preimage in G of $\mathrm{Alt}(\Gamma)_{(x_1,\dots,x_s)}$, found in time $O(n^5)$ by Corollary 3.3.4(e) (N need not be normal in G: we call it N in analogy to Proposition 3.5.3); $[G:N] \leq m^s$, so again by Corollary 3.3.4(c) we can write an element σ_i of each coset of N in time $O(n^5 + m^s n^3)$. Thus the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ reduces to $\leq m^{s'}$ problems of determining $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}_i)$, because

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \bigcup_i \mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\sigma_i^{-1}})\sigma_i$$

exactly as in Proposition 3.5.3. It is important to consider that s' as above, the exponent of the multiplicative cost, is not the same as s (despite them being certainly related) and is indeed smaller: the fact is that the elements of Γ are not all indistinguishable (due to the presence of V_1, V_2), so many possibilities for the choice of $x_1, \dots, x_s$ are as a matter of fact forbidden; seen in a different light, many of the Iso_N that emerge are known to be empty without the need for computing them, as they do not make V_1, V_2 correspond in $\mathbf{x}$ and $\mathbf{y}^{\sigma_i^{-1}}$.

Now that the situation described in the hypothesis of Theorem 3.5.9 has been split into its various cases, we show how to treat each of them while making at least one among our parameters $n, |\mathcal{B}|, m$ decrease.

Corollary 3.5.11. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks such that G acts on it as $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$, where $|\Gamma| = m$ and $|\mathcal{B}| = \binom{m}{k}$. Suppose that $m > 102 \log^2 n$; suppose also that, fixing the images $(y_i)_{i=1}^s$ of some elements $(x_i)_{i=1}^s \subseteq \Gamma$, we can find a coloured partition of Γ in which each part has size $\leq \alpha|\Gamma|$ (with $\alpha \leq \frac{2}{3}$).*

Then, if N is the preimage of $\mathrm{Alt}(\Gamma)_{(x_1,\dots,x_s)}$ inside G, N divides Ω into a system $\mathcal{B}'$ of orbits and blocks (at least as coarse as $\mathcal{B}$) of size $\leq \frac{2}{3}|\Omega|$. Moreover, for any orbit Δ with $|\Delta| > \frac{2}{3}|\Omega|$, $\mathcal{B}'|_\Delta$ is nontrivial and strictly coarser than $\mathcal{B}|_\Delta$ and its elements are k-subsets of blocks of $\mathcal{B}$ all contained in the same colour Γ_0 of Γ of size $> \frac{2}{3}|\Gamma|$; also, the stabilizer of blocks of $\mathcal{B}'|_\Delta$ coincides with the stabilizer of blocks of Γ_0.

Proof. This corollary covers cases (a) and (c1) of Theorem 3.5.9. The focus on N is due to the reduction to the problem of determining $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\sigma^{-1}})$ featured in Remark 3.5.10, where $\sigma \in G$ is an element that sends each x_i to y_i.

We have a coloured partition $\mathcal{C}$ on Γ with parts of size $\leq \alpha|\Gamma|$ (with $\alpha \leq \frac{2}{3}$); we can repeat the same reasoning as in Corollary 3.5.7 (the case $m > 102 \log^2 n$ and Γ^o nontrivial) and show that the hypotheses of Lemma 3.5.4 hold here. By this lemma, Ω itself has a coloured partition $\mathcal{C}'$ that is at least as coarse as $\mathcal{B}$ and whose parts are also of size $\leq \frac{2}{3}|\Omega|$: the fact that N respects the colours of $\mathcal{C}'$ means that elements with different colours will not be sent to each other, i.e. they sit in different orbits, while respecting the parts with the same colours translates to sending all the elements of one part to the same part, i.e. moving them as a block.

If we are in an orbit Δ of size $> \frac{2}{3}|\Omega|$, it means that inside $\mathcal{C}'$ we are in a colour of size $> \frac{2}{3}|\Omega|$, so that it will also have to be divided into smaller parts with the same colour: therefore, $\mathcal{B}'|_\Delta$ is nontrivial and strictly coarser than $\mathcal{B}|_\Delta$, since each part will contain not all blocks and at least two blocks of $\mathcal{B}$. Using the reasoning in Lemma 3.5.4, Δ must come from a Γ_0 as in our statement, and by our description of $\mathcal{C}'$ in that lemma the block stabilizer of $\mathcal{B}'|_\Delta$ contains the block stabilizer of Γ_0; the other direction also holds: in fact, the only case in which a σ permutes blocks of Γ_0 without permuting anything in $\mathcal{B}'|_\Delta$ is when Δ represents k-subsets of Γ_0 intersecting all parts of Γ_0 equally, but then there would be only one block in $\mathcal{B}'|_\Delta$ itself in contradiction with the fact that $|\Delta| > \frac{2}{3}|\Omega|$. □

This corollary divides Ω into orbits and blocks that are coarser than the original $\mathcal{B}$: this makes n decrease or the block size increase, or both.

Corollary 3.5.12. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks such that G acts on it as $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$, where $|\Gamma| = m$ and $|\mathcal{B}| = \binom{m}{k}$. Suppose also that there exist sets $S_\mathbf{x}, S_\mathbf{y} \subseteq \Gamma$ of size $> \frac{1}{2}|\Gamma|$, canonical for $\mathbf{x}, \mathbf{y}$ respectively, such that for any $\sigma \in \mathrm{Alt}(S_\mathbf{x})$ there is an element of $\mathrm{Aut}_G(\mathbf{x})$ inducing σ on $S_\mathbf{x}$ (and similarly for $\mathbf{y}$).*

Then in time $O(n^{10})$ we can reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to determining 4 sets $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}_i)$, where N induces orbits of size $\leq \frac{2}{3}|\Omega|$.

Proof. This corollary covers case (b) of Theorem 3.5.9.

If π is the map going from G to $\mathrm{Alt}(\Gamma)$ mentioned in the statement, define $N = \pi^{-1}(\mathrm{Alt}(\Gamma)_{(S_{\mathbf{x}})})$: we can find N in time $O(n^{10})$ by Corollary 3.3.4(d)-3.3.4(e). Also, define $N' = \pi^{-1}(\mathrm{Alt}(\Gamma)_{S_{\mathbf{x}}})$: since $S_{\mathbf{x}}$ is canonical for $\mathbf{x}$, $\mathrm{Aut}_G(\mathbf{x})$ stabilizes $S_{\mathbf{x}}$ setwise, which means that it is contained inside N'. For any even permutation of Γ sending $S_{\mathbf{x}}$ to $S_{\mathbf{y}}$, we can find a preimage $\tau \in G$ in time $O(n^{10})$ by Corollary 3.3.4(d); we have

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \mathrm{Iso}_G(\mathbf{x}, \mathbf{y}^{\tau^{-1}})\tau = \mathrm{Iso}_{N'}(\mathbf{x}, \mathbf{y}^{\tau^{-1}})\tau = \mathrm{Aut}_{N'}(\mathbf{x})\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\tau^{-1}})\tau,$$

using Lemma 3.3.6(a), the fact that $G\tau = G$, and (by canonicity) the fact that any string isomorphism between $\mathbf{x}$ and $\mathbf{y}^{\tau^{-1}}$ must stabilize $S_{\mathbf{x}}$.

Now we have to describe $\mathrm{Aut}_{N'}(\mathbf{x})$: by the canonicity of $S_{\mathbf{x}}$, it is equal to $\mathrm{Aut}_G(\mathbf{x})$. Since by hypothesis $\mathrm{Alt}(S_{\mathbf{x}})$ is contained in $\mathrm{Aut}_G(\mathbf{x})$, there exist two elements in $\mathrm{Aut}_G(\mathbf{x})$ that induce two generators of $\mathrm{Alt}(S_{\mathbf{x}})$; to find them, we can take preimages σ_1, σ_2 of these two generators in G (again in time $O(n^{10})$ by Corollary 3.3.4(d)) and then determine the sets $\mathrm{Aut}_{N\sigma_i}(\mathbf{x}) = \mathrm{Iso}_N(\mathbf{x}, \mathbf{x}^{\sigma_i^{-1}})\sigma_i$ for $i = 1, 2$: any two elements τ_1, τ_2 inside them will give us the whole $\mathrm{Aut}_{N'}(\mathbf{x})$, since this is $\langle A \cup \{\tau_1, \tau_2\}\rangle$ for any set A of generators of $\mathrm{Aut}_N(\mathbf{x})$. We have reduced the problem to the four problems $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}_i)$ with $\mathbf{y}_1 = \mathbf{x}$, $\mathbf{y}_2 = \mathbf{y}^{\tau^{-1}}$, $\mathbf{y}_3 = \mathbf{x}^{\sigma_1^{-1}}$, $\mathbf{y}_4 = \mathbf{x}^{\sigma_2^{-1}}$.

We still have to prove that N has the property described in the statement. The partition $\{S_{\mathbf{x}}, \Gamma \setminus S_{\mathbf{x}}\}$ can be seen as a coloured partition where $S_{\mathbf{x}}$ and $\Gamma \setminus S_{\mathbf{x}}$ are two parts of different colours (if $S_{\mathbf{x}} = \Gamma$ then the second part is empty, but this will not be a problem): examining the proof of Lemma 3.5.4, we see that each subset Ω_a collecting (the elements contained in blocks corresponding to) the k-subsets of Γ containing $a > 0$ elements of $\Gamma \setminus S_{\mathbf{x}}$ is of size $\leq \frac{2}{3}|\Omega|$; on the other hand, the blocks corresponding to k-subsets of $S_{\mathbf{x}}$ are stabilized by N since this subgroup stabilizes $S_{\mathbf{x}}$ itself pointwise. Therefore N has only orbits of size $\leq \frac{2}{3}|\Omega|$. □

Again, this corollary makes n decrease or the block size increase (or both) by dividing Ω into orbits and blocks coarser than $\mathcal{B}$.

Corollary 3.5.13. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks such that G acts on it as $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$, where $|\Gamma| = m$ and $|\mathcal{B}| = \binom{m}{k}$. Suppose also that, fixing the images $(y_i)_{i=1}^s$ of some elements $(x_i)_{i=1}^s \subseteq \Gamma$, we can find two disjoint sets $V_1, V_2 \subseteq \Gamma$, with V_2 divided into a system of (possibly size 1) blocks $\mathcal{G}$ with $\binom{|\mathcal{G}|}{k'} = |V_1| \geq \frac{2}{3}|\Gamma|$ for some $k' \geq 2$, and a bijection between V_1 and $\binom{\mathcal{G}}{k'}$ such that each element of G, seen as a permutation in $\mathrm{Sym}(\mathcal{G})$, also induces the natural permutation of V_1 given by the previous identification.*

Then, if N is the preimage of $\mathrm{Alt}(\Gamma)_{(x_1,\dots,x_s)}$ inside G and $\Delta \subseteq \Omega$ is an orbit induced by N of size $> \frac{2}{3}|\Omega|$, $N|_\Delta$ respects a system $\mathcal{B}'$ of blocks inside Δ (at least as coarse as $\mathcal{B}|_\Delta$), and if M is the stabilizer of $\mathcal{B}'$ then $N|_\Delta/M \leq \mathrm{Sym}(\mathcal{G})$ (and $|\mathcal{G}| < 1 + \sqrt{2m}$).

Proof. This corollary covers case (c2) of Theorem 3.5.9. The focus on N is due to the reduction to the problem of determining $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\sigma^{-1}})$ featured in Remark 3.5.10, where $\sigma \in G$ is an element that sends each x_i to y_i.

We can see $\{V_1, V_2, \Gamma \setminus (V_1 \cup V_2)\}$ as a coloured partition on Γ, where the last two parts are of size $\leq \frac{1}{3}|\Gamma|$ combined. Looking at the proof of Lemma 3.5.4, each subset Ω_a collecting (the elements contained in blocks corresponding to) the k-subsets of Γ containing $a > 0$ elements of $\Gamma \setminus V_1$ is of size $\leq \frac{2}{3}|\Omega|$; thus, the orbit Δ (if it exists at all) can only be one of the orbits collecting k-subsets of Γ entirely contained in V_1.

An element $B \in \mathcal{B}|_\Delta$ corresponds to a k-subset R of V_1 and each element of R is a k_0-subset of $\mathcal{G}$; each element of $N|_\Delta$ induces a permutation of $\mathcal{G}$, so any two subsets R, R' whose elements cover the same blocks of $\mathcal{G}$ (rather, their union does) move together under the action of $N|_\Delta$, i.e. they are in a same block of Δ. A system of blocks $\mathcal{B}'$ is therefore at least as coarse as the system formed by collecting all the B corresponding to the R based on the same blocks of $\mathcal{G}$, which is in turn at least as coarse as $\mathcal{B}$; the image of a block $B' \in \mathcal{B}'$ is determined by the movement of the blocks of $\mathcal{G}$, since a permutation of $\mathcal{G}$ determines the new k_0-subsets of $\mathcal{G}$ represented in V_1, so $N|_\Delta/M \leq \mathrm{Sym}(\mathcal{G})$.

The fact that $|\mathcal{G}| < 1 + \sqrt{2m}$, which will be helpful in the recursion process, is evident from the hypotheses we made in the statement: since $V_1 \subseteq \Gamma$ is in bijection with $\binom{\mathcal{G}}{k'}$ and $k' \geq 2$ we have $m \geq \binom{|\mathcal{G}|}{2} > \frac{(|\mathcal{G}|-1)^2}{2}$, and the inequality follows. □

This corollary either decreases n or reduces the degree of the symmetric group that contains G (as an abstract group, in the sense that we do not care about the precise action). In fact, while recursing through Cameron in this circumstance, if G is not too small we will obtain a subgroup of G that is $\mathrm{Alt}(\Gamma')$ for some Γ', and $|\Gamma'| \leq 1 + \sqrt{2m}$ where m was the size of the old Γ.

3.5.2 The algorithm, not assuming CFSG

Now we examine what the algorithm looks like when we are not assuming CFSG: the result by Cameron and Maróti, which provided us with the initial crossroads to guide us in the recursion, does not hold anymore. On the other hand, the fact that the action of G/N on $\mathcal{B}$ is the same as the action of $\mathrm{Alt}(\Gamma)$ on $\binom{\Gamma}{k}$ (in Theorem 3.5.6(b), Corollary 3.5.7(c) and beyond) is not always essential: in many occasions the important fact is that each block of $\mathcal{B}$ corresponds to a k-subset of a certain Γ, but G/N may act on it as some $H \leq \mathrm{Sym}(\Gamma)$, and not necessarily as $H = \mathrm{Alt}(\Gamma)$. We will see this in the next results.

We start with our new building block, a result due to Pyber [Pyb93] that replaces Cameron and does not depend on CFSG.

Theorem 3.5.14. *Let $|\Gamma| = m$ and let $G \leq \mathrm{Sym}(\Gamma)$. Do not assume CFSG. If G is primitive, then one of the following alternatives holds:*

(a) $|G| \leq m^{8\lceil 4 \log_2 m \rceil \log_2 m}$;

(b) G is either $\mathrm{Sym}(\Gamma)$ *or* $\mathrm{Alt}(\Gamma)$;

(c) G is transitive but not doubly transitive.

Proof. See the proof of [Pyb93, Thm. A]. □

Let us tackle each of these alternatives that emerge in our determination of $\mathrm{Iso}_G(\mathbf{x},\mathbf{y})$. We start again with the case of G/N small enough to be able to effectively use Proposition 3.5.3.

Proposition 3.5.15. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x},\mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks preserved by G, and call N the stabilizer of $\mathcal{B}$: suppose that there are a set Γ of size m and a bijection between $\mathcal{B}$ and $\binom{\Gamma}{k}$ (for some k) such that the action of G/N on $\mathcal{B}$ corresponds to the action of some $H \leq \mathrm{Sym}(\Gamma)$ on $\binom{\Gamma}{k}$. Do not assume CFSG.*

If $|H| \leq m^{8\lceil 4\log_2 m\rceil \log_2 m}$, or if $m \leq 25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$, then we can reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x},\mathbf{y})$ to determining $\leq m^{25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}}$ sets of isomorphisms $\mathrm{Iso}_N(\mathbf{x},\mathbf{y}_i)$, in time $O(m^{25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}} n^{10})$ and at no multiplicative cost.

Proof. The proof is very similar to part of the proof of Corollary 3.5.7, as expected: the current proposition corresponds to the route taken by Corollary 3.5.7(a). We add that, if we know both Γ and the bijection, it is a polynomial-time task to find out whether the conditions on H are satisfied: we can calculate $|H|$ in time $O(m^5)$ by Corollary 3.3.4(a), which will tell us if either condition is true.

First, $|H|$ is always bounded by $m! \leq m^{m+\frac{1}{2}}e^{1-m}$. For $m \leq 5656$ we have $\left(m+\frac{1}{2}\right)\log m + 1 - m \leq 67\log^3 m$, while for $m \geq 5657$ we have $4\log_2 m > 49.8$ and then $\lceil 4\log_2 m\rceil \leq \frac{51}{50}\frac{4}{\log 2}\log m$; hence, for any m,

$$|H| \leq m^{8\lceil 4\log_2 m\rceil \log_2 m} \implies |H| \leq m^{\max\left\{67, \frac{51}{50}\frac{32}{\log^2 2}\right\}\log^2 m} < m^{68\log^2 m}.$$

As for $m \leq 25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$, this implies easily that $|H| < m^m \leq m^{25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}}$. Since $m \leq n$, for ε small we have $68\log^2 m \leq 25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$, so both bounds on $|H|$ can be summed up by using the unique bound $m^{25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}}$. We can conclude the proof by producing all the elements of G/N and working as in Proposition 3.5.3. □

Case (b) of Theorem 3.5.14 is extremely similar to the process followed in the CFSG case, as shown in the following proposition.

Proposition 3.5.16. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x},\mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks preserved by G, and call N the stabilizer of $\mathcal{B}$: suppose that there are a set Γ of size m and a bijection between $\mathcal{B}$ and $\binom{\Gamma}{k}$ (for some k) such that the action of G/N on $\mathcal{B}$ corresponds to the action of $H = \mathrm{Sym}(\Gamma), \mathrm{Alt}(\Gamma)$ on $\binom{\Gamma}{k}$. Do not assume CFSG.*

If $m > 25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$, then we reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x},\mathbf{y})$ to one of the following:

(a) determining ≤ 8 sets $\mathrm{Iso}_{N'}(\mathbf{x},\mathbf{y}_i)$, where N' divides Ω into orbits of size $\leq \frac{2}{3}|\Omega|$;

(b) determining $\leq m^{7e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}}$ sets $\mathrm{Iso}_{N'}(\mathbf{x},\mathbf{y}_i)$, where N' divides Ω into a system of orbits and/or blocks $\mathcal{B}'$ (at least as coarse as $\mathcal{B}$) such that if there is an orbit Δ of size $> \frac{2}{3}|\Omega|$ then either

(b1) $\mathcal{B}'|_\Delta$ *is nontrivial and strictly coarser than* $\mathcal{B}|_\Delta$*, with stabilizer of* $\mathcal{B}'|_\Delta$ *equal to the block stabilizer of the large colour of* Γ *(in the sense of Corollary 3.5.11), or*

(b2) if M *is the stabilizer of* $\mathcal{B}'|_\Delta$*,* $N'|_\Delta/M$ *acts on* $\mathcal{B}'|_\Delta$ *as some* $H' \leq \mathrm{Sym}(\Gamma')$ *acts on* $\binom{\Gamma'}{k'}$ *with* $|\Gamma'| < 1 + \sqrt{2m}$*.*

The time necessary for this reduction is the cost of $\frac{1}{2}m^{2a}naa!$ *calls of the whole algorithm for strings of length* $\leq \frac{n}{a}$ *where* $a \in (6.24999, 6.25) \cdot e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$*, plus some additional time* $O(m^{3a}n^{11})$*.*

Proof. First, in the case of $H = \mathrm{Sym}(\Gamma)$ we can reduce the problem to 2 sets with $H = \mathrm{Alt}(\Gamma)$. Now we are exactly in the case described in Corollary 3.5.7(c). We can retrace all the steps from Theorem 3.5.9 to Corollary 3.5.13, this time using the CFSG-free versions of the results in §3.4, and the results correspond to one of the final situations thereby reached: case (a) corresponds to Corollary 3.5.12 (where 4 becomes 8 because of the aforementioned reduction from Sym to Alt), case (b1) corresponds to Corollary 3.5.11, and case (b2) corresponds to Corollary 3.5.13.

We need only to justify how to obtain the action in part (b2) rather than only a bound on the degree of $N|_\Delta/M$ like in Corollary 3.5.13 (as we observed, this stronger statement is necessary for the recursion, given the unavailability of Cameron).

Let us start with the first problem. Following the reasoning up to Corollary 3.5.13, we ended up finding two disjoint sets $V_1, V_2 \subseteq \Gamma$ and a partition $\mathcal{G}$ of V_2 that respect the various hypotheses mentioned in the corollary, and in its proof we find a system of blocks $\mathcal{B}'|_\Delta$ on an orbit Δ of size $> \frac{2}{3}|\Omega|$ (if such an orbit exists) such that the action of $N|_\Delta$ is induced by the permutations of $\mathcal{G}$, up to the stabilizer of the system. If $k = 1$, $\mathcal{B}$ corresponds to Γ itself: therefore Δ of size $> \frac{2}{3}|\Omega|$ must correspond to V_1 itself, and by hypothesis the permutations of $\mathcal{G}$ induce permutations of V_1 in a way that respects the bijection $V_1 \leftrightarrow \binom{\mathcal{G}}{k'}$ ($\mathcal{G}$ is then the sought Γ'). If $k \geq 2$, we can use Lemma 3.5.5 to prove that Δ is further split into blocks that are strictly coarser than $\mathcal{B}$: in that lemma, we use $\Gamma', \Gamma, \mathcal{B}$ to refer in this situation to $\mathcal{G}, V_1, \mathcal{B}|_\Delta$ respectively; we only have to show that the bounds on $|\mathcal{G}|$ hold. If $m > 102\log^2 n$, by Remark 3.3.7 we have $m \geq 8308$; $|V_1| \geq \frac{2}{3}m$, so that $|V_1| \geq 5539$: whatever will be our choice of k', we have $5539 \leq \binom{|\mathcal{G}|}{k'} \leq \binom{|\mathcal{G}|}{\lfloor \frac{1}{2}|\mathcal{G}| \rfloor}$, hence $|\mathcal{G}| \geq 12$.

Finally, let us obtain the exponent in part (b) and the value of a. The interval of a is taken directly from Proposition 3.4.2(b). As for the exponent, we notice that exactly as in Theorem 3.5.9 we still have $d \leq a$ (with d as in Lemma 3.4.3(b)), so that the multiplicative cost is still $m^{a+55\log m}$. For our choice of a, our bounds $25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon} < m \leq n$, and ε small enough, we can bound this cost as in the statement (remember that we also have a possible multiplication by 2, from the reduction in the case of $H = \mathrm{Sym}(\Gamma)$). The additive cost is the same as in Theorem 3.5.9. □

Finally, we treat case (c) of Theorem 3.5.14, whose procedure is a somewhat shortened version of the one covered in the previous proposition.

Proposition 3.5.17. *Let $|\Omega| = n$, $G \leq \mathrm{Sym}(\Omega)$ and let $\mathbf{x}, \mathbf{y} : \Omega \to \Sigma$ be two strings; let $\mathcal{B}$ be a system of blocks preserved by G, and call N the stabilizer of $\mathcal{B}$: suppose that there are a set Γ of size m and a bijection between $\mathcal{B}$ and $\binom{\Gamma}{k}$ (for some k) such that the action of G/N on $\mathcal{B}$ corresponds to the action of some $H \leq \mathrm{Sym}(\Gamma)$ on $\binom{\Gamma}{k}$. Do not assume CFSG.*

If $m > 25e^{1/\varepsilon^2}(\log n)^{4+\varepsilon}$ and H is transitive but not doubly transitive, then in time $O(m^{14})$ we reduce the problem of determining $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ to determining $\leq m^{55 \log m}$ sets $\mathrm{Iso}_{N'}(\mathbf{x}, \mathbf{y}_i)$ where N' divides Ω into a system of orbits and/or blocks $\mathcal{B}'$ (at least as coarse as $\mathcal{B}$) such that if there is an orbit Δ of size $> \frac{2}{3}|\Omega|$ then either

(a) $\mathcal{B}'|_\Delta$ is nontrivial and strictly coarser than $\mathcal{B}|_\Delta$, with stabilizer of $\mathcal{B}'|_\Delta$ equal to the block stabilizer of the large colour of Γ (in the sense of Corollary 3.5.11), or

(b) if M is the stabilizer of $\mathcal{B}'|_\Delta$, $N'|_\Delta/M$ acts on $\mathcal{B}'|_\Delta$ as some $H' \leq \mathrm{Sym}(\Gamma')$ acts on $\binom{\Gamma'}{k'}$ with $|\Gamma'| < 1 + \sqrt{2m}$.

Proof. If H is transitive but not doubly transitive, we can determine the nontrivial orbits of the action of H on $\binom{\Gamma}{2}$ in time $O(m^6)$ by Lemma 3.3.5; giving to each orbit its own colour, we can make $\binom{\Gamma}{2}$ into a coherent configuration in time $O(m^{10} \log m)$ (mostly due to Weisfeiler-Leman, see [Hel19b, §§2.3-2.5]): the result would be a nontrivial homogeneous coherent configuration, where homogeneity is consequence of the fact that this is a canonical process and H moves every point of Γ to any other, so that we are unable to distinguish them with different colours.

Now we can use SoJ directly. We use Proposition 3.4.4, where from the costs we can remove the exponent b (since we do not perform the Design Lemma).

The shape of the action of $N'|_\Delta/M$ on $\mathcal{B}'|_\Delta$ in part (b) is again proved as in part (b2) of Proposition 3.5.16, i.e. resorting to Lemma 3.5.5. □

All these cases reduce to some sort of recursion with lower parameters, either by decreasing n or m or increasing the block size. This works exactly as in the CFSG case.

3.6 Main theorem: proof

We are at last ready to prove Theorem 3.2.1.

The group-theoretic results to which we keep returning in our recursions are Theorem 3.5.6 in the CFSG case and Theorem 3.5.14 in the CFSG-free case; we have already declared this multiple times, but we repeat it here (now with references, though): except for exiting through the base cases given in Remark 3.5.1 and Proposition 3.5.8 and for breaking down Ω into smaller orbits through Proposition 3.5.2, the only other alternatives are that on a large chunk of Ω either the system of blocks $\mathcal{B}$ on which we are working becomes coarser and coarser (the conclusion featured in Corollary 3.5.11, Proposition 3.5.16(b1) and Proposition 3.5.17(a)) or the group in which we are operating is contained in a symmetric

group of degree smaller and smaller (the conclusion featured in Corollary 3.5.13, Proposition 3.5.16(b2) and Proposition 3.5.17(b)).

Proof of Thm. 3.2.1. There are several tasks to accomplish: we need to analyze the possible passages mentioned above and see that they fit the description given in terms of ($\mathcal{C}1$)-($\mathcal{C}2$)-($\mathcal{C}3$), and that the final base cases fit ($\mathcal{A}$), and we need to estimate their contribution in terms of both the multiplicative cost (which will lead us to a bound on the number of atomic elements) and additive cost (which will yield the total runtime).

To determine the multiplicative cost of the procedure, we start in medias res. We are working on a certain orbit Δ of Ω, of size $|\Delta| = n' \leq n$, divided into a system of blocks $\mathcal{B}$, of size $|\mathcal{B}| = r \leq n'$, such that the group G/N permuting the blocks is isomorphic to a subgroup of $\mathrm{Sym}(m)$, of degree $m \leq r$. We call $M(n', r, m)$ (an upper bound on) the multiplicative cost that we incur from this moment until we manage to make each block into an orbit of its own. Call $T(n', r, m)$ the intermediate time cost, in an analogous fashion as we did with $M(n', r, m)$; we also suppose that $T(n', r, m)$ includes the cost of performing Proposition 3.5.2 on the resulting orbits, so as to cover the time spent to bridge one intermediate problem to the next one.

The proof is articulated in the following main steps.

(1) From the already known passages we delineate a handful of "actions" and the reduction they entail on $M(n', r, m)$; note that here we are using the word "action" *not* in a mathematical sense, but in the everyday meaning of "something done purposefully to accomplish a certain end". This step gives us a series of conditions that our function M must respect in order to work.

(2) We choose M and show that it is compatible with the previous conditions coming from the actions; then $M(n, n, n)$ by definition turns out to be a bound on the multiplicative cost incurred throughout the whole algorithm.

(3) We translate actions into ($\mathcal{C}1$)-($\mathcal{C}2$)-($\mathcal{C}3$) and end-cases into ($\mathcal{A}$), and use $M(n, n, n)$ to bound the number of atomic elements.

(4) We refine the computations of the second part to tackle $T(n', r, m)$.

For the sake of notation, we are going to perform our computations by bounding $\log M$ instead of M, so that the focus will be on the exponents of the quantities involved.

(1) *Description of the actions.*

The first action that is possible to perform, following from Corollary 3.5.7(a) and Proposition 3.5.15, is to directly pass to the stabilizer of the system, thus making each block into an orbit: this concludes the calculation of M with no reduction, and it costs at most $102 \log m \log^2 n'$ in the CFSG case and $25e^{1/\varepsilon^2} \log m (\log n')^{4+\varepsilon}$ in the CFSG-free case; these are direct lower bounds for $\log M(n', r, m)$, therefore

$$\log M(n', r, m) \geq K_1 \log m (\log n')^{e_1} \tag{3.6.1}$$

for $(K_1, e_1) = (102, 2), (25e^{1/\varepsilon^2}, 4 + \varepsilon)$ appropriately.

For notational simplicity, let us set $X = 8308$ for the CFSG case and $X = 25e^{1/\varepsilon^2}$ for the CFSG-free case: these are the values we have already encountered many times, and they separate small and large values of m, n (see Remark 3.3.7 in particular). If either n' or m is smaller than X we are using the first action, so for the other actions we can assume otherwise.

The second action, following from Corollary 3.5.12 and Proposition 3.5.16(a) and (in case there are only orbits of size $\leq \frac{2}{3}|\Omega|$) from Corollaries 3.5.11-3.5.13 and Propositions 3.5.16(b)-3.5.17, consists in reducing n' (and consequently r) by a fraction at least as small as $\frac{2}{3}$. This costs at most $K_2 \log m (\log n')^{e_2}$, where $(K_2, e_2) = (57, 1)$ assuming CFSG and $(K_2, e_2) = (7e^{1/\varepsilon^2}, 4 + \varepsilon)$ without CFSG: for our bounds on m, n' (and for ε small), these are the largest expenses, coming from Theorem 3.5.9(c) and Propositions 3.5.16(b) respectively. Hence

$$\log M(n', r, m) \geq K_2 \log m (\log n')^{e_2} + \log M\left(\frac{2}{3}n', \frac{2}{3}r, m\right). \tag{3.6.2}$$

The third action, following (in case there is an orbit of size $> \frac{2}{3}|\Omega|$) from Corollary 3.5.11 and Propositions 3.5.16(b1)-3.5.17(a), creates a new system of blocks strictly coarser than the original $\mathcal{B}$, at a cost of at most $K_2 \log m (\log n')^{e_2}$: (K_2, e_2) is as in the previous action, as the largest expenses originate in the same results. What happens is, we have first to work on the coarser system, then after we have stabilized each coarser block we have to work on each one of them as the new orbit and the finer blocks as the new system; since the stabilizer of coarser blocks coincides with some block stabilizer of Γ, we also get $m', \frac{m}{m'}$ instead of m in the two steps, for some $2 \leq m' \leq \frac{m}{2}$. The bound on $\log M(n', r, m)$ given by this action is

$$\log M(n', r, m) \geq K_2 \log m (\log n')^{e_2} + \log M(n', r', m') + \log M\left(\frac{n'}{r'}, \frac{r}{r'}, \frac{m}{m'}\right), \tag{3.6.3}$$

where $2 \leq r' \leq \frac{r}{2}$ is the size of the coarser system.

The fourth action, following (in case there is an orbit of size $> \frac{2}{3}|\Omega|$) from Corollary 3.5.13 and Propositions 3.5.16(b2)-3.5.17(b), reduces the degree of the minimal symmetric group containing G, at a cost of at most $K_2 \log m (\log n')^{e_2}$ ((K_2, e_2) as in the second and third actions); therefore,

$$\log M(n', r, m) \geq K_2 \log m (\log n')^{e_2} + \log M(n', r, 1 + \sqrt{2m}). \tag{3.6.4}$$

(2) *Choice of function M.*

Now let us prove that

$$\log M(n', r, m) = (\log n')^{e_2+1}(a \log m + b \log r) \tag{3.6.5}$$

satisfies the four conditions for some appropriate constants a, b.

Since $m \leq r$ and $e_1 \leq e_2 + 1$, in order to have (3.6.1) we have simply to ask $a + b \geq K_1$. Recall that for the other actions we can assume $m, n' \geq X$.

For $n' \geq X$ and $e_2 \geq 1$ we have $\left(\log\left(\frac{2}{3}n'\right)\right)^{e_2+1} < (\log n')^{e_2+1} - \frac{3}{4}(\log n')^{e_2}$ (for both values of X), so

$$\begin{aligned} &K_2 \log m (\log n')^{e_2} + \left(\log\left(\frac{2}{3}n'\right)\right)^{e_2+1}\left(a\log m + b\log\left(\frac{2}{3}r\right)\right) \\ &< (\log n')^{e_2+1}(a\log m + b\log r) + (\log n')^{e_2}\left(K_2\log m - \frac{3}{4}(a\log m + b\log r)\right), \end{aligned}$$

and since $m \leq r$ in order to have (3.6.2) it is sufficient to ask $\frac{3}{4}(a+b) > K_2$. For (3.6.3), using $\left(\log\frac{n'}{r'}\right)^{e_2+1} < (\log n')^{e_2+1} - \log r'(\log n')^{e_2}$ and $\log\frac{m}{m'} \geq \log 2$ the sufficiency of (3.6.5) in this case is implied by

$$f(\log r') = b\log^2 r' - (a\log 2 + b\log r)\log r' + K_2 \log m \leq 0. \tag{3.6.6}$$

The function $f(x)$ in the interval $[\log 2, \log r - \log 2]$ has its maximum in $x = \log 2$, being a quadratic polynomial with the minimum in $x = \frac{1}{2}\log r + \frac{a\log 2}{2b} > \frac{1}{2}\log r$; evaluating $f(\log 2)$ and recalling that $X \leq m \leq r$, (3.6.6) is in turn consequence of

$$b \geq \frac{K_2 \log m}{\log 2(\log r - \log 2)} - \frac{K_2 \log 2}{\log r - \log 2}a \iff b \geq \frac{K_2 \log X}{\log 2 \log(X/2)}. \tag{3.6.7}$$

To have (3.6.4), we notice that $1 + \sqrt{2m} < m^{0.53926}$ for $m \geq X$ (for both values of X); then,

$$\begin{aligned} &(\log n')^{e_2+1}(a\log m + b\log r) \\ &\geq K_2 \log m (\log n')^{e_2} + (\log n')^{e_2+1}(0.53926 a \log m + b\log r) \end{aligned}$$

means $a \geq \frac{K_2}{0.46074 \log n'}$, so that $a \geq 0.2405 K_2 \geq \frac{K_2}{0.46074 \log X}$ is enough to satisfy (3.6.4).

Putting together these conditions and considering our K_1, K_2, it turns out that $a = 13.7085$ and $b = 89.07486$ with CFSG and $a = b = \frac{25}{2}e^{1/\varepsilon^2}$ without CFSG are suitable choices for (3.6.5). The multiplicative cost of the whole algorithm is bounded by $M(n,n,n)$; thus we conclude that the multiplicative cost is bounded by

$$n^{102.78336\log^2 n} \text{ with CFSG}, \qquad n^{25e^{1/\varepsilon^2}(\log n)^{5+\varepsilon}} \text{ without CFSG}. \tag{3.6.8}$$

(3) *Reduction to* $(\mathcal{A})$-$(\mathcal{C}1)$-$(\mathcal{C}2)$-$(\mathcal{C}3)$.

Now that we have bounded the multiplicative cost, let us focus now on the actions themselves, in order to be able to describe the various stages as one among $(\mathcal{A})$-$(\mathcal{C}1)$-$(\mathcal{C}2)$-$(\mathcal{C}3)$ and to use $M(n,n,n)$ for the computation of the number of atomic elements.

The first action entails firstly a reduction of the problem of determining the set $\mathrm{Iso}_G(\mathbf{x},\mathbf{y})$ to a collection of $\mathrm{Iso}_N(\mathbf{x}, \mathbf{y}_i^{\sigma_i^{-1}})\sigma_i$ whose union is the original set, as seen in Proposition 3.5.3 or Remark 3.5.10: the way this union is performed corresponds

precisely to ($\mathcal{C}1$), and the number of subproblems is equal to the multiplicative cost incurred during this action; then, each stabilized block becomes an orbit of its own, in a reduction that corresponds to the situation described in ($\mathcal{C}2$) (see Proposition 3.5.2). This passage does not feature any multiplicative cost, but it does multiply the number of atomic elements at the end: however, since we have simply r blocks, the contribution of ($\mathcal{C}2$) here, and indeed the contribution of any nested series of ($\mathcal{C}2$) acting throughout the entire process of solving the intermediate problem with parameters (n', r, m), is at most r.

The second action features a reduction of Ω to orbits of size at most $\frac{2}{3}|\Omega|$; this can happen in two different ways. In the case of Corollaries 3.5.11-3.5.13 and Propositions 3.5.16(b)-3.5.17, after having fixed the image of a certain number of points at a multiplicative cost we find orbits of such size, and then we examine each orbit singularly: this is exactly as in the previous case, where each passage consists in using ($\mathcal{C}1$) and ($\mathcal{C}2$), and the bounds on the atomic element multiplication are as above. In the case of Corollary 3.5.12 and Proposition 3.5.16(a), we are in a situation where

$$\mathrm{Iso}_G(\mathbf{x}, \mathbf{y}) = \langle \mathrm{Aut}_N(\mathbf{x}), \tau_1, \tau_2 \rangle \tau' \tau,$$

where $\tau' \in \mathrm{Iso}_N(\mathbf{x}, \mathbf{y}^{\tau^{-1}})$ (to use the notation of the corollary); this corresponds to ($\mathcal{C}3$), and despite the multiplication cost being at most 4 or 8, there is no actual growth in the number of atomic elements through this case.

The third and the fourth action create respectively (on the large orbit) a strictly coarser system of blocks and a bijection on a permutation subgroup of strictly smaller degree: this happens at a certain multiplicative cost, that corresponds to a passage of the form shown in ($\mathcal{C}1$) and multiplies the atomic elements by the same quantity.

The various actions, as we already said, decrease at least one of the three parameters n, r, m, and when r, m become too small n itself diminishes through the use of the first action: hence, the procedure eventually stops when $n = 1$, the trivial case of Remark 3.5.1. There is also a second way to stop the algorithm, and that is Proposition 3.5.8: both cases correspond to the atom ($\mathcal{A}$). The reduction to ($\mathcal{A}$)-($\mathcal{C}1$)-($\mathcal{C}2$)-($\mathcal{C}3$) has been proved; the actual writing of the expression is done following the proofs of Proposition 3.5.3 (for ($\mathcal{C}1$)), Proposition 3.5.2 (for ($\mathcal{C}2$)) and Corollary 3.5.12 (for ($\mathcal{C}3$)). The number of atomic elements, by the reasoning above, is bounded by

$$\begin{aligned} n \cdot n^{102.78336 \log^2 n} &< n^{103 \log^2 n} && \text{with CFSG,} \\ n \cdot n^{25 e^{1/\varepsilon^2} (\log n)^{5+\varepsilon}} &< n^{26 e^{1/\varepsilon^2} (\log n)^{5+\varepsilon}} && \text{without CFSG,} \end{aligned}$$

since its intermediate multiplication is bounded by $rM(n', r, m)$, and we are done.

(4) *Runtime.*

Finally, let us tackle the runtime; we start at the end, this time. We have already proved that there are at most $n^{K \log^e n}$ atomic elements constituting the expression, and by Remark 3.5.1 and Proposition 3.5.8 we can treat each one in time $O(n^6)$, so the bound on the runtime covers this final stage; now we go back to the analysis of the recursion process that leads to it.

Call $T(n', r, m)$ the intermediate time cost, in an analogous fashion as we did with $M(n', r, m)$; most of the computations for M also hold for T, but we have to verify that the *added* time does not disrupt the final constants coming from our multiplicative reasoning: we also suppose that $T(n', r, m)$ includes the cost of performing Proposition 3.5.2 on the resulting orbits, so as to cover the time spent to bridge one intermediate problem to the next one. For the first action, the bound is as in Corollary 3.5.7(a) and Proposition 3.5.15, with the addition of the cost for the reduction to single orbits:

$$T(n', r, m) = O(m^{K_1(\log n')^{e_1}} n'^{10} + n'^{11}).$$

As for the other three actions, let us start by working on the additive cost first; recall that henceforth $n' \geq r \geq m \geq X$. The highest additive cost is featured in Theorem 3.5.9 and Proposition 3.5.16 and it involves the use of the runtime itself (for smaller n'); supposing that we want to show that it is sufficient to ask $T(n', r, m) = O(e^{(\log n')^{e_2+1}(a \log m + b \log r)} n'^{11})$, this cost is of order

$$\frac{1}{2} m^{2\nu} n' \nu\nu! \cdot e^{(\log \frac{n'}{\nu})^{e_2+1}(a \log m + b \log r)} \frac{n'^{11}}{\nu^{11}} + 2m^{3\nu} n'^{11}, \tag{3.6.9}$$

where $\nu = \alpha(\log n')^{e_2}$ for some $\alpha \in (1.66431, 1.77512)$ with CFSG and $\alpha \in (6.24999, 6.25) \cdot e^{1/\varepsilon^2}$ without CFSG. Notice that we write $2m^{3\nu} n'^{11}$ (i.e. with a 2 in front) in order to absorb the successive smaller costs, such as the n'^{11} from Proposition 3.5.2, the n'^{10} from Corollary 3.5.12 and the m^{14} from Proposition 3.5.17. For $a, b \geq 5$, it is easy to prove that the first addend of (3.6.9) is larger than the second: say for example $n' > 4$, $\nu\nu! > 1$ and $e^{(\log \frac{n'}{\nu})^{e_2+1}(a \log m + b \log r)} > e^{\frac{1}{3} \log^2 n' (a \log m + b)} = m^{\frac{a}{3} \log^2 n'} n'^{\frac{b}{3} \log n'} > m^{\nu} (2\nu)^{11}$. Now let us bound the first addend (without $\frac{1}{2}$); its logarithm is

$$\begin{aligned}
&2\nu \log m + \log(n' \nu\nu!) + \left(\log \frac{n'}{\nu}\right)^{e_2+1} (a \log m + b \log r) + \log \frac{n'^{11}}{\nu^{11}} \\
&< 2\alpha(\log n')^{e_2} \log m + \log n' + \log m + \alpha(\log n')^{e_2} \log m \\
&\quad + (\log n')^{e_2+1}(a \log m + b \log r) - 2.19999(\log n')^{e_2}(a \log m + b \log r) + \log n'^{11} \\
&< (\log n')^{e_2+1}(a \log m + b \log r) + \log n'^{11} - 2.19999 b (\log n')^{e_2} \log r,
\end{aligned}$$

using $\left(\log \frac{n'}{\nu}\right)^{e_2+1} < (\log n')^{e_2+1} - (\log n')^{e_2} \log \nu$ for $e_2 \geq 1$ and $2.19999 < \log \nu < \log m$, and noting that the negative $(\log n')^{e_2} \log m$ term absorbs the smaller $\log n', \log m, (\log n')^{e_2} \log m$ positive terms for $3\alpha + 2 < 2a$. Therefore for example $b \geq 5$ gives us already enough leeway:

$$e^{-2.19999 b (\log n')^{e_2} \log r} < 10^{-389}.$$

Now that the additive cost is accounted for, we continue with the multiplicative one. Since we want to prove that a quantity multiplied by n'^{11} is larger than its partial version multiplied by some fraction of n'^{11}, we can just ignore this polynomial cost. For the second action, we exploit the already existing margin left

out before: $\left(\log\left(\frac{2}{3}n'\right)\right)^{e_2+1} < (\log n')^{e_2+1} - \left(\frac{3}{4}+\frac{3}{100}\right)(\log n')^{e_2}$, and for $a+b\geq 1$ we are left with a constant of

$$e^{-\frac{3}{100}(\log n')^{e_2}(a\log m+b\log r)} < \frac{1}{4}$$

in front of this part of the runtime. For the third action, if b is as on the right side of (3.6.7), we can use $\left(1+\frac{1}{100000}\right)b$ as the new coefficient and going through (3.6.6) we can cut ourselves a margin of

$$e^{-\frac{b}{100000}(\log n')^{e_2}\log r'\log\frac{r}{r'}} \leq e^{-\frac{K_2}{100000}(\log n')^{e_2}\log X} < \frac{49}{50}.$$

The fourth action is treated in the same way: putting $\left(1+\frac{1}{100000}\right)a$ we carve out a $\frac{49}{50}$ constant as well. This shows that we can take the same coefficient a, b as before multiplied by $1+\frac{1}{100000}$, because $\frac{49}{50}+10^{-389}<1$; also, thanks to

$$n^{102.78336\left(1+\frac{1}{100000}\right)\log^2 n} < n^{103\log^2 n},$$
$$n^{25\left(1+\frac{1}{100000}\right)e^{1/\varepsilon^2}(\log n)^{5+\varepsilon}} < n^{26e^{1/\varepsilon^2}(\log n)^{5+\varepsilon}},$$

we achieve the bounds we wanted in the two cases for the runtime, too.

The theorem is proved. □

3.7 Concluding remarks

It must be noted that the difference between the exponents for the CFSG and the CFSG-free case in not a consequence of the different use of group-theoretic results to produce a suitable recursion (Theorems 3.5.6 and 3.5.14 respectively): they make the algorithm different in the two cases, that is true, but the different expense lies elsewhere. What is important in this respect is the theoretic tool that allows the recursion in Theorem 3.5.9 and Proposition 3.5.16, and that gives for us a different number of calls to the algorithm for shorter strings. In the local certificates procedure in Babai's algorithm, one important detail is that a certain epimorphism $G\to \mathrm{Alt}(k)$ for $G\leq \mathrm{Sym}(n)$ primitive is guaranteed to be an isomorphism, and this is ensured for $k=\Omega(\log n)$ with a proof relying on CFSG (see [Bab16a, Lemma 8.3.1] [Hel19b, Lemme 4.1]), but only for $k=\Omega(\log n)^{4+\varepsilon}$ without CFSG (see [Pyb16, Lemma 12], where $\Omega(\log^5 n)$ is used). Consequently the algorithm is still performing the same subroutines, but the tuples on which we want to build the certificates need to be larger, leading to the loss of efficiency that we witness.

The constants are likely improvable, if one were to analyze with greater care the routines. We have been quite accurate, but we have not really aimed at obtaining the best possible constant, especially in the CFSG-free case: as our position is to consider CFSG as a theorem (see the discussion after Theorem 1.2.4), the analysis of the CFSG-free procedure is more of a question of method, given the use we are going to do of the main theorem in §6.

***

In truth, the origin of the whole analysis performed in this chapter lay originally in trying to find whether we could easily arrive to an improvement of Babai's algorithm that would gets us to a $n^{O(\log n)}$ runtime, or, if not, to point out where exactly the bottleneck was and why.

It is clear, to the attentive reader of these pages, that the obstacle does not lie in the "interstitial reasoning" as we called it at the start. We have performed our analysis burdened with multiplicative costs of $n^{O(\log n)}$, or $n^{O(\log n)^{4+\varepsilon}}$, originating in the main subroutines in §3.4. However, if we had had at that point a polynomial cost, we could have continued with our bookkeeping until the end and obtained a $n^{O(\log n)}$ runtime: even the $n^{O(\log n)}$ that is weaved already into Cameron's theorem (Theorem 3.5.6(a), coming from Theorem 1.2.5(c)) does not pile up eventually, since (3.6.1) shows that $\max\{e_1, e_2 + 1\}$ is the correct exponent of the logarithm.

Hence, the bottleneck must be in the subroutines. The local certificates call the algorithm for strings of size $\Omega\left(\frac{n}{\log n}\right)$, for each of the $O(\log n)$-tuples inside an $O(n)$-set: thus, unless one manages to bypass the logarithmic requirement in Lemma 3.4.1, the routine of Proposition 3.4.2 is too expensive to improve the runtime under the $n^{O(\log^2 n)}$ threshold. Also Split-or-Johnson is in its current form too expensive, but in that case one might make do with reworking the recursion process that comes into play by showing for instance that the worst scenario does not actually happen in real life. It is already a common thread in the literature that distinguishing non-isomorphic graphs is actually pretty easy in general (see [BES80] [BK79]), and a handful of bad cases yields a much worse runtime: SoJ as well analyzes in its recursion hypothetical configurations where it is very difficult to break the symmetry of its vertices, even when we are given from the start that the are few twins among them. It might be feasible to prove that there are actually no such configurations, or alternatively that they are so well-structured that it is possible to describe them entirely and treat them separately as exceptional cases, as was done for instance with the "three exceptional families" in [SW16, Def. 1.3] (the first paper to break the $n^{O(\sqrt{n})}$ threshold on GIP).

Chapter 4

Slowly growing sets in $\mathrm{Aff}(\mathbb{F}_q)$

The content of this chapter is essentially taken from [Don19b].

We have already mentioned in §1.4 that, among the many different problems related to the study of growth and expansion in finite groups, the study of the affine group over finite fields has occupied a particularly interesting place. The affine group

$$\mathrm{Aff}(\mathbb{F}) = \left\{ \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} \middle| \, a \in \mathbb{F}^*, b \in \mathbb{F} \right\}, \tag{4.0.1}$$

where $\mathbb{F}$ is a finite field, is one of the smallest interesting examples of an infinite family of finite groups on which questions of growth of sets $A \subseteq \mathrm{Aff}(\mathbb{F})$ can yield nontrivial answers, and it has been used to showcase techniques applicable to more general situations, like the pivot argument; on the other hand, its shape makes its uniquely suitable to study the so-called sum-product phenomenon, related to growth of sets inside finite fields under both addition and multiplication. For both of these points of view, a remarkable example is provided in Helfgott's survey [Hel15, §4.2].

Structural theorems about growth in $\mathrm{Aff}(\mathbb{F}_p)$ (p prime) have been produced in the last few years, describing in substance what a set A with small growth must look like. Results like Helfgott's [Hel15, Prop. 4.8] and Murphy's [Mur17, Thm. 27] belong to a first generation of proofs that rely, one way or another, on sum-product estimates; they already accomplish the goal of characterizing quite well a slowly growing A: such a set must essentially either be a point stabilizer or be contained in a few vertical lines, which in addition get filled in finitely many steps if $|A| = \Omega(p)$.

Rudnev and Shkredov [RS18] have then quantitatively improved this classification in $\mathrm{Aff}(\mathbb{F}_p)$: the main attractivity of their result, however, resides in the fact that, in their own words, "the improvement [they] gain is due [...] to avoiding any explicit ties with the sum-product phenomenon, which both proofs of Helfgott and Murphy relate to", which makes their version of the characterization of slowly growing A part of a new generation of efforts. What they rely on instead is a geometric theorem by Szőnyi [Sző99, Thm. 5.2] that gives a good lower bound on

the number of directions spanned by a set of non-collinear points in the plane $\mathbb{F}_p^2$ for p prime.

Following the approach by Rudnev and Shkredov, we first produce an analogous version of Szőnyi's result for the plane $\mathbb{F}_q^2$, where q is any prime power; then we use that estimate to prove a structural theorem on slowly growing sets in $\mathrm{Aff}(\mathbb{F}_q)$ (resembling the corresponding ones for $\mathrm{Aff}(\mathbb{F}_p)$ mentioned before), which to the best of our knowledge is the first of its kind.

4.1 Introduction

We remind the reader that, at least for us (unlike in some of the works we reference), p will always denote a prime and q a power of p. Given a set A inside the plane $\mathbb{F}^2$, the set of *directions* spanned or determined by A denotes the set

$$D = \left\{ \frac{b'-b}{a'-a} \middle| (a,b),(a',b') \in A,\ (a,b) \neq (a',b') \right\} \subseteq \mathbb{F} \cup \{\infty\},$$

where conventionally ∞ corresponds to the fraction with $a'-a=0$. We make free use of the natural identification $\mathrm{Aff}(\mathbb{F}) \leftrightarrow \mathbb{F}^* \times \mathbb{F}$ given by

$$\begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} \in \mathrm{Aff}(\mathbb{F}) \qquad \longleftrightarrow \qquad (a,b) \in \mathbb{F}^* \times \mathbb{F},$$

so that we may refer to points, lines and directions even when speaking of the group $\mathrm{Aff}(\mathbb{F})$; in particular, we call $\pi : \mathrm{Aff}(\mathbb{F}) \to \mathbb{F}^*$ the map corresponding to the projection on the first component, so that the preimage of a point through this map is a vertical line. $\mathrm{Aff}(\mathbb{F})$ acts also on $\mathbb{F}$ as $(a,b) \cdot x = ax + b$, and we think of this action when we refer to $\mathrm{Stab}(x)$ (which also looks like a line when seen in $\mathbb{F}^2$); finally, U denotes the unipotent subgroup corresponding to $\{1\} \times \mathbb{F}$, again a vertical line.

As said before, one of the starting points of the new-style result for slowly growing sets in $\mathrm{Aff}(\mathbb{F}_p)$ is the following bound by Szőnyi.

Theorem 4.1.1. *Let p be a prime, and let $A \subseteq \mathbb{F}_p^2$ with $1 < |A| \leq p$. Then either A is contained in a line or A spans $\geq \frac{|A|+3}{2}$ directions.*

With that, Rudnev and Shkredov prove the following (see [RS18, Thm. 5]).

Theorem 4.1.2. *Let p be a prime and let $A \subseteq \mathrm{Aff}(\mathbb{F}_p) \leftrightarrow \mathbb{F}_p^* \times \mathbb{F}_p$ with $A = A^{-1}$ and $|A^3| = C|A|$. Then at least one of the following is true:*

(a) $A \subseteq \mathrm{Stab}(x)$ for some $x \in \mathbb{F}_p$;

(b) when $1 < |A| \leq (1+\varepsilon)p$ for some $0 < \varepsilon < 1$, we have $|\pi(A)| \leq 2C^4$;

(c) when $|A| > (1+\varepsilon)p$ for some $0 < \varepsilon < 1$, we have $|\pi(A)| = O_\varepsilon\left(\frac{1}{p}C^3|A|\right)$, and in particular for $|A| > 4p$ we have $|\pi(A)| \leq \frac{2}{p}C^3|A|$ and $A^8 \supseteq U$.

Szőnyi's bound is part of a long history of applications of results about *lacunary polynomials* (i.e. polynomials made of a small number of monomials with respect to their degree) over finite fields to finite geometry: the reader interested in similar applications can check [Sző99] and its bibliography.

Many results in this area can apply, with the appropriate modifications, to $\mathbb{F}_q$ as well. In this case, however, bounds on the number of directions spanned by a set in the finite plane appear to be messier, and understandably so: unlike in the case of $\mathbb{F}_p$, the number of directions determined by A tends to congregate around values $\frac{|A|}{p^i}$ for powers $p^i|q$; this is due to the fact that there may exist sets with multiples of p^i points on each line that are so well-structured that they sit in relatively few directions compared to the amount of points they have (see [BBB$^+$99, §5] for an example of this assertion when $|A| = q$).

The result we essentially use, on the number of directions spanned in $\mathbb{F}_q^2$ by some set with $1 < |A| \leq q$, is due to Fancsali, Sziklai and Takáts [FST13, Thm. 17]: for the lower bound they found we give here a proof that is very similar to theirs, but we also prove a different upper bound that can be more or less advantageous than theirs depending on the situation (Theorem 4.2.2). Used directly, the lower bound can only give us about $\frac{|A|}{q/p}$ directions; a tighter theorem, in the style of [BBB$^+$99, Thm. 1.1], would give not only $p^i|p^e = q$, but also $i|e$ (and therefore a much better lower bound of approximately $\frac{|A|}{\sqrt{q}}$ directions): [BBB$^+$99, Thm. 1.1] however works only for $|A| = q$, and the lack of a complete set of q points is crucial in worsening the condition on the denominator p^i during the proof.

Nevertheless, it turns out that a simple observation *can* make us achieve the bound with $\sqrt{q}$ in the denominator: at its core, we use the fact that a set of points A either sits on $\geq \sqrt{q}$ parallel lines or has a line with $\geq \frac{|A|}{\sqrt{q}}$ points on it. Our first main result then, playing the role of Szőnyi's bound in [RS18], is as follows.

Theorem 4.1.3. *Let $q = p^e$ be a prime power, and let $A \subseteq \mathbb{F}_q^2$ with $1 < |A| \leq q$. Then either A is contained in a line or A spans*

(a) $> \frac{|A|}{\sqrt{q}}$ directions for e even,

(b) $> \frac{|A|}{p^{\frac{e-1}{2}}+1}$ directions for e odd.

Observe that the theorem is only a constant away from Szőnyi's bound when we use it for $q = p$; we add that actually the proof can be easily adjusted to yield that bound exactly: we chose not to do so in order to get a cleaner statement, with case (b) valid for all e odd.

Using Theorem 4.1.3 and following more or less the same proof as in [RS18], we obtain our second main result, generalizing Theorem 4.1.2 to any $\mathbb{F}_q$.

Theorem 4.1.4. *Let $q = p^e$ be a prime power and let $A \subseteq \mathrm{Aff}(\mathbb{F}_q) \leftrightarrow \mathbb{F}_q^* \times \mathbb{F}_q$ with $A = A^{-1}$ and $|A^3| = C|A|$. Then at least one of the following is true:*

(a) $A \subseteq \mathrm{Stab}(x)$ for some $x \in \mathbb{F}_q$;

(b) when $1 < |A| \leq q$ *we have* $|\pi(A)| < (p^{\lfloor \frac{e}{2} \rfloor} + 2)C^4$*, while when* $q < |A| < (3 + 2\sqrt{2})q$ *we have* $|\pi(A)| < (4 + 2\sqrt{2})C^4$*;*

(c) when $|A| \geq (3 + 2\sqrt{2})q$ *we have* $|\pi(A)| < \frac{2}{q}C^3|A|$ *and* $A^8 \supseteq U$*.*

The statement above looks remarkably similar to Theorem 4.1.2, and is qualitatively as strong a structural theorem as in the case of $\mathrm{Aff}(\mathbb{F}_p)$. The $p^{\lfloor \frac{e}{2} \rfloor}$ in case (b) cannot be improved in general: for e even, there is a natural embedding of $\mathbb{F}_{\sqrt{q}}$ inside $\mathbb{F}_q$, and $A = \mathbb{F}_{\sqrt{q}}^* \times \mathbb{F}_{\sqrt{q}}$ has $|A| < q$, $C = 1$ and $|\pi(A)| = p^{\frac{e}{2}} - 1$. See §4.4 for further remarks.

Let us comment however on a small difference between Theorem 4.1.2 and the result for $\mathbb{F}_p$ featured in [RS18]. The case of a medium-sized A (i.e. $1 < \frac{|A|}{q} = O(1)$) has been placed into alternative (c) by Rudnev and Shkredov and into alternative (b) by us, essentially losing the $A^k \supseteq U$ implication: this has been done because the subgroup H of Kneser's theorem [Kne53] can stifle the growth of A, in a way that the Cauchy-Davenport inequality ([Cau13, Thm. VII] [Dav35], see [TV06, Thm. 5.4]) could not; asking for p large enough is innocuous in the latter, but not in the former: see also §4.3 where we use it.

We could still use Alon's bound [Alo86, (4.2)] on the number of lines in the projective plane as done in [RS18], since it holds for $\mathbb{F}_q$ as well: this would give for example $|\pi(A)| < \frac{2(\sqrt{5}+1)}{(7-3\sqrt{5})q}C^3|A|$ for $|A| \geq \frac{\sqrt{5}+1}{2}q$ (where the maximum of $\frac{\varepsilon^2(1-\varepsilon)}{2(1+\varepsilon)}$ is located) and in general $|\pi(A)| = O_\varepsilon\left(\frac{1}{q}C^3|A|\right)$ for $|A| \geq (1+\varepsilon)q$; then, upon using Kneser's theorem, one could either ask for p large enough ($p > 100$ in the first case, say, and $p = \Omega_\varepsilon(1)$ in general) or classify separately the sets A with large H (which should be possible, because having large $H = \mathrm{Stab}(A^2)$ is a rather restrictive condition to satisfy), and an additional conclusion $A^k \supseteq U$ for $k = O_\varepsilon(1)$ would be reached. It would probably be interesting to explore more deeply these medium-sized sets; however, for the purpose of obtaining a structural result like Theorem 4.1.4 whose numerical details are of secondary relevance, we deemed to be simpler and just as effective to reduce that case to alternative (b), especially as the observation behind our ability to do so (Lemma 4.2.1) is very elementary.

4.2 Number of directions in $\mathbb{F}_q^2$

In the present section we prove bounds about the number of directions determined by sets of points in the plane $\mathbb{F}_q^2$, which lead eventually to Theorem 4.1.3.

Let us start with the following simple statement: it does not concern Theorem 4.1.3, but it will allow us in the next section to deal quickly with the sets A whose size is slightly larger than q.

Lemma 4.2.1. *Any set* $A \subseteq \mathbb{F}_q^2$ *with* $|A| > q$ *spans all* $q + 1$ *directions.*

Proof. The result is immediate: by the pigeonhole principle, for any given direction, one of the q parallel lines in $\mathbb{F}_q^2$ following that direction has to contain at least two points of A. □

As a complement to Lemma 4.2.1, the following theorem deals with the number of directions spanned by sets of size at most q. As remarked before, a theorem of the same nature appears already in [FST13], and it is proved very similarly using the same techniques deriving from the study of lacunary polynomials.

Theorem 4.2.2. *Let $q = p^e$ be a prime power, let $A \subseteq \mathbb{F}_q^2$ with $1 < |A| \leq q$, and let D be the set of directions determined by A. Then either $|D| = 1$ (and A is contained in a line), $|D| = q + 1$ (and A spans all directions) or there are two integers $0 \leq l_2 \leq l_1 < e$ such that*

$$|D| \geq \frac{|A| - 1}{p^{l_2} + 1} + 2,$$
$$|D| \leq q - |A| + \max\left\{1, \frac{|A| - 1 - (q - |A|)\max\{0, |A| + p^{l_1} - q - 1\}}{p^{l_1} - 1}\right\}.$$

A little notational comment: if $l_1 = 0$ we consider the upper bound trivial (but the lower bound becomes $\frac{|A|+3}{2}$, which is quite strong, identical to Szőnyi's bound for $\mathbb{F}_p$).

Before we go to the proof, let us spend a few more words comparing this result with the one in [FST13]: their bounds are written as $\frac{|A|-1}{t+1} + 2 \leq |D| \leq \frac{|A|-1}{s-1}$, for some appropriately defined s, t. The lower bound is the same as the one presented here, as t and p^{l_2} are defined in the same way. The situation for the upper bound is more interesting: we have $s \leq t = p^{l_2} \leq p^{l_1}$, because the authors define s looking at the multiplicities in $H_y(x)$ alone (see the proof below for details) instead of the whole $x^q + g_y(x)$, which also gives a stronger geometric meaning to their s than to our l_1; however, our upper bound tends to be stronger when $|A|$ is fairly close to q and there is a gap between s and p^{l_1} (which can happen, as observed in [FST13]).

Proof. First of all, we can suppose $\infty \in D$. If this were not true, we could take any $d \in D \setminus \{0\}$ (D is nonempty for $|A| > 1$, and $D = \{0\}$ concludes the theorem) and consider A' made of points $(a - db, b)$ for any $(a, b) \in A$, which implies also that $|A'| = |A|$: such a set would span directions given by

$$\frac{b' - b}{a' - db' - a + db} = \frac{1}{\frac{a'-a}{b'-b} - d},$$

from which it is clear that the new set of directions D' is as large as D, since equalities are preserved, and that moreover $\infty \in D'$.

Define $n \geq 0$ such that $|A| = q - n$. First, define the *Rédei polynomial*

$$H_y(x) = \prod_{i=1}^{q-n} (x + ya_i - b_i) \in \mathbb{F}_q[x, y],$$

where the product is among all the $(a_i, b_i) \in A$: it is a polynomial of degree $q - n$ in two variables (some authors, like in [BBB+99], define it as a homogeneous polynomial in three variables, but by ensuring that $\infty \in D$ we do not need to do so). The usefulness of such polynomial lies in the fact that two points of A sitting

on the same line with slope y_0 yield the same $x + y_0 a - b$, so that a multiple root in $H_{y_0}(x)$ reflects the presence of a line with multiple points, i.e. a secant of A, and indicates that $y_0 \in D$. We also define another function in two variables,

$$f_y(x) = \sum_{j=0}^{n} (-1)^j \sigma_j(\mathbb{F}_q \setminus \{ya_i - b_i | (a_i, b_i) \in A\}) x^{n-j}, \tag{4.2.1}$$

where $\sigma_j(S)$ is the j-th elementary symmetric polynomial of the elements in the set S; $f_y(x)$ is itself a polynomial in two variables (see [Sző96, Thm. 4] for a recursive definition of $f_y(x)$), in which the coefficient of x^{n-j} has y-degree j: therefore we can write

$$x^q + g_y(x) = H_y(x) f_y(x) \in \mathbb{F}_q[x, y],$$

where $g_y(x)$ is a polynomial in two variables of x-degree $\leq q - 1$.

Substituting $y = y_0$ for some $y_0 \notin D$, we observe that by definition the set $\mathbb{F}_q \setminus \{y_0 a_i - b_i | (a_i, b_i) \in A\}$ has n elements and that $f_{y_0}(x)$ is simply the product of the $x - k_i$ for all the $k_i \in \mathbb{F}_q$ not counted in $H_{y_0}(x)$, so $g_{y_0}(x) = -x$: this means that the coefficients of $x^{q-1}, x^{q-2}, \ldots, x^{|D|}$ in $g_y(x)$ are polynomials of degree $\leq q - |D|$ in y that take value 0 for the $q - |D| + 1$ values $y_0 \in \mathbb{F}_q \setminus D$. Thus, these coefficients are the zero polynomial; in other words, the x-degree of $g_y(x)$ is at most $|D| - 1$.

Working with x, y has allowed us to give a bound on the degree of $g_y(x)$. From now on, for the sake of simplicity we substitute one value $y \in D \setminus \{\infty\}$ inside our polynomials and drop the index, and we will work with only one variable; this is possible unless $D = \{\infty\}$, from which $|D| = 1$ and A is contained in a vertical line.

Call l_2 the largest integer for which $g(x) \in \mathbb{F}_q[x^{p^{l_2}}]$: by the fact that any $x \mapsto x^{p^i}$ is an automorphism of $\mathbb{F}_q$, we have $g(x) = (\tilde{g}(x))^{p^{l_2}}$ for some $\tilde{g}(x) \in \mathbb{F}_q[x] \setminus \mathbb{F}_q[x^p]$. Decompose $x^q + g(x)$ into its irreducible factors, and call l_1 the largest integer for which p^{l_1} divides the multiplicity of each linear factor (hence $l_1 \geq l_2$): l_1, l_2 depend on our choice of y, so for our definition we suppose that we have chosen a y that yields the smallest l_1. We can write

$$x^{q/p^{l_2}} + \tilde{g}(x) = (R(x))^{p^{l_1 - l_2}} N(x),$$

where $R(x) \in \mathbb{F}_q[x] \setminus \mathbb{F}_q[x^p]$ is such that $(R(x))^{p^{l_1}}$ is the divisor of $x^q + g(x)$ made of its linear factors (the fully reducible part of $x^q + g(x)$) and $N(x) \in \mathbb{F}_q[x] \setminus \mathbb{F}_q[x^p]$ is such that $(N(x))^{p^{l_2}}$ is the divisor of $x^q + g(x)$ made of its nonlinear factors. Note that $(N(x))^{p^{l_2}}$ must be a divisor of $f(x)$. If $l_1 = e$ then $x^q + g(x) = (x + c)^q$ for some $c \in \mathbb{F}_q$, which means that all the points of A lie on a line of slope equal to the y we have fixed, contradicting $\infty \in D$: hence $l_2 \leq l_1 < e$.

Call $R^*(x)$ the divisor of $R(x)$ made of all the irreducible factors of $R(x)$ counted without multiplicity: $R^*(x)$ divides also $x^q - x$ by definition, so it divides $x^q + g(x) - (x^q - x) = g(x) + x \neq 0$ ($y \in D$ prevents us from having $g(x) = -x$). If an irreducible polynomial $k_1(x)$ divides another $k_2(x)$ with multiplicity m then it divides $k_2'(x)$ with multiplicity $m - 1$, so

$$\frac{(R(x))^{p^{l_1 - l_2}}}{R^*(x)} \Bigg| (x^{q/p^{l_2}} + \tilde{g}(x))' = \tilde{g}'(x) \neq 0,$$

where the last inequality is true because $\tilde{g}(x) \notin \mathbb{F}_q[x^p]$. From the reasoning above, we obtain

$$x^q + g(x) = \left(R^*(x) \cdot \frac{(R(x))^{p^{l_1 - l_2}}}{R^*(x)} \right)^{p^{l_2}} (N(x))^{p^{l_2}} \Big| (g(x) + x)^{p^{l_2}} (\tilde{g}'(x))^{p^{l_2}} f(x) \neq 0,$$

and therefore $q = \deg(x^q + g(x)) \leq p^{l_2}(\deg(g(x) + x) + \deg \tilde{g}'(x)) + \deg f(x)$; we have already determined that $\deg(g(x) + x) \leq \deg g(x) \leq |D| - 1$, and similarly $\deg \tilde{g}'(x) \leq \frac{\deg g(x)}{p^{l_2}} - 1 \leq \frac{|D|-1}{p^{l_2}} - 1$, whence from the definition of $f(x)$ we get

$$q \leq p^{l_2} \left(|D| - 1 + \frac{|D| - 1}{p^{l_2}} - 1 \right) + n \qquad \Longrightarrow \qquad |D| \geq \frac{q - n - 1}{p^{l_2} + 1} + 2,$$

settling the lower bound.

Let us focus now on the upper bound. Fix a point $(a, b) \in A$ and take a slope $y_0 \in \mathbb{F}_q$: the multiplicity of the linear factor $x + y_0 a - b$ inside $H(x)$ determines how many points of A sit on the line defined by (a, b) and y_0. We know that the multiplicity of every linear factor in the whole $H(x)f(x)$ is a multiple of p^{l_1} and that it is at least 1 for this particular linear factor, since (a, b) sits on the line; however, we need a way to keep under control the number of false positives that come from the fully reducible part of $f(x)$ (inexistent "ghost points" that make us overcount the contribution of a single line to A, and thus undercount $|D|$). The way to go is to bound the number of lines passing through (a, b) for which ghost points exist.

Let $f_y(x)$ be as in (4.2.1), call it for simplicity $f_y(x) = \sum_{j=0}^{n} \sigma_{y,j} x^{n-j}$ where the $\sigma_{y,j}$ are polynomials in y of degree j. Assume that $|D| < q + 1$: then there will be a direction $y_0 \in \mathbb{F}_q \setminus D$, as $\infty \in D$. For this y_0, $H_{y_0}(x) f_{y_0}(x) = x^q - x$ and $x + y_0 a - b$ has multiplicity 1 in it; moreover, it must come from our fixed point (a, b), which means that it must divide $H_{y_0}(x)$ and be coprime with $f_{y_0}(x)$: this fact implies that the two-variable linear polynomial $x + ya - b$ cannot divide $f_y(x)$. In other words, we cannot write

$$(x^{n-1} + \tau_{y,1} x^{n-2} + \ldots + \tau_{y,n-1})(x + ya - b) = x^n + \sigma_{y,1} x^{n-1} + \ldots + \sigma_{y,n} \quad (4.2.2)$$

for any choice of polynomials $\tau_{y,i}$; however, defining

$$\tau_{y,i} = \sum_{j=0}^{i} (-1)^j (ya - b)^j \sigma_{y,i-j}$$

(here $\sigma_{y,0} = 1$) we can ensure that the equality (4.2.2) works at least at the level of the coefficients of $x, x^2, \ldots, x^{n-1}$, which means that we must have

$$\sum_{j=0}^{n} (-1)^j (ya - b)^j \sigma_{y,n-j} \neq 0, \quad (4.2.3)$$

so as to violate (4.2.2) for the free coefficient.

Every time $f_{y_i}(x)$ has a $x + y_i a - b$ factor (or, geometrically speaking, every time the line determined by (a, b) and y_i has a ghost point), (4.2.2) is true for $y = y_i$ though, and in particular the LHS of (4.2.3) is indeed 0: that expression is a polynomial in y of degree n, so if there were $n+1$ lines with ghost points (4.2.3) would not be true, contradicting the fact that $x + ya - b$ cannot divide $f_y(x)$ as polynomials in two variables. Hence, at most n non-vertical lines through (a, b) have ghost points.

If $|D|-1 \leq n$ the upper bound stated in the theorem is already true, so suppose that the opposite holds: then there is a non-vertical line through (a, b) whose slope is in D with no ghosts. We can transform A as at the beginning of the proof to make that slope ∞, i.e. (a, b) lies on a vertical secant of A.

Each non-vertical line through (a, b) whose slope is in D has a multiple of p^{l_1} among true points of A and ghost points (l_1 has been defined so as to make that statement true for all slopes at the same time). On the ghost-free lines there are at least $p^{l_1} - 1$ true points besides (a, b), while on the ones with ghosts we can only say that there are at least $\max\{0, p^{l_1} - 1 - n\}$ of them (as the x-degree of $f_y(x)$ is n); finally, the vertical secant has at least p^{l_1} points including (a, b), as we made sure that it had no ghosts before the transformation. Combining all of this with the bound on the number of lines with ghosts, and counting all the points of A, we get

$$(|D| - 1 - n)(p^{l_1} - 1) + n \max\{0, p^{l_1} - 1 - n\} + p^{l_1} \leq q - n.$$

As we remarked after the statement of the theorem, for $l_1 = 0$ there is no upper bound. For $l_1 > 0$, the inequality above concludes the proof. □

Now that we have the lower bound provided by Theorem 4.2.2, we can proceed with the proof of the first main theorem. We retain the same notation as in the previous proof.

Proof of Thm. 4.1.3. Suppose that $|D| \notin \{1, q + 1\}$ (otherwise the theorem is already proved); fix a slope $y_0 \neq \infty$ and consider the polynomial $R^*(x)$ defined as in the proof of Theorem 4.2.2. Let $\varepsilon > 0$ be small enough, and let $q' = p^{\frac{e}{2}} - \varepsilon$ for e even and $q' = p^{\frac{e-1}{2}}$ for e odd.

If the degree of $R^*(x)$ is $\leq q'$, the set A must be contained in $\leq q'$ lines with slope y_0, which means that one of them (call it L) will have to contain $\geq \frac{|A|}{q'}$ points of A; since A is not contained in one line there must be also a point of A outside L, and each secant laid between this point and a point of $A \cap L$ has a different slope, so that $|D| \geq \frac{|A|}{q'}$: for e even it means $|D| > \frac{|A|}{\sqrt{q}}$, while for e odd it means $|D| \geq \frac{|A|}{p^{\frac{e-1}{2}}}$.

If $R^*(x)$ has degree $> q'$, then by the fact that $(R^*(x))^{p^{l_1}}$ divides $x^q + g(x)$ we must have $p^{l_2} \leq p^{l_1} < \frac{q}{q'}$: regardless of whether e is even or odd, $p^{l_2} \leq p^{\lfloor \frac{e}{2} \rfloor}$ since l_2 is an integer. Using the lower bound in Theorem 4.2.2 (which holds for our A), we have

$$|D| \geq \frac{|A| - 1}{p^{\lfloor \frac{e}{2} \rfloor} + 1} + 2 = \frac{|A|}{p^{\lfloor \frac{e}{2} \rfloor}} + 2 - \frac{|A|}{p^{\lfloor \frac{e}{2} \rfloor}(p^{\lfloor \frac{e}{2} \rfloor} + 1)} - \frac{1}{p^{\lfloor \frac{e}{2} \rfloor} + 1}.$$

For e even, the bound above implies $|D| > \frac{|A|}{\sqrt{q}}$, while for e odd we can obtain $|D| \geq \frac{|A|+3}{p^{\frac{e-1}{2}}+1}$. □

4.3 Growth in $\mathrm{Aff}(\mathbb{F}_q)$

We move now to the proof of Theorem 4.1.4. We follow closely the proof of the analogous result in [RS18] for $\mathbb{F}_p$: the only difference is that we use Theorem 4.1.3 instead of Szőnyi's bound, and that as we have already said we absorb the case of A of medium size into alternative (b), without resorting to Alon's bound to fall into (c).

We remind the reader of two well-known and by now classical results. First, an inequality, deducible in multiple ways from bounds by Ruzsa (see for instance [Ruz96]), states that for any group G and any $A = A^{-1} \subseteq G$ the equality $|A^3| = C|A|$ implies $|A^k| \leq C^{k-2}|A|$ for any $k \geq 4$. Second, Kneser's theorem [Kne53] tells us that, given any abelian G and any $A, B \subseteq G$, there is a proper subgroup H with $|A+B| \geq \min\{|G|, |A|+|B|-|H|\}$.

Theorem 4.1.3 and Lemma 4.2.1 will take care of small and medium $|A|$, respectively. For $|A|$ large we will instead make use of the following bound, due to Vinh [Vin11, Thm. 3]: the statement therein says actually something weaker, but it is based on a well-known graph-theoretic result [AS16, Cor. 9.2.5] that lets us reformulate as follows (as [RS18] does for $\mathbb{F}_p$).

Proposition 4.3.1. *Let q be a prime power, let P be a set of points in $\mathbb{F}_q^2$ and let L be a set of lines in $\mathbb{F}_q^2$; define $I(P,L)$ as the set of pairs $(p,l) \in P \times L$ s.t. $p \in l$. Then*

$$\left| I(P,L) - \frac{|P||L|}{q} \right| \leq \sqrt{q|P||L|}.$$

Let us also give here separately a lemma that will provide the upper bounds on $\pi(A)$ in Theorem 4.1.4(b)-(c).

Lemma 4.3.2. *For any $g = \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} \in \mathrm{Aff}(\mathbb{F}_q) \setminus \{\mathrm{Id}\}$, define the map*

$$\varphi_g : \mathrm{Aff}(\mathbb{F}_q) \to \mathrm{Aff}(\mathbb{F}_q), \qquad \varphi_g(h) = hgh^{-1}.$$

Then,

(a) any point in the image of φ_g has as preimage a line of slope $\frac{b}{a-1}$;

(b) if $A = A^{-1} \subseteq \mathrm{Aff}(\mathbb{F}_q)$ and $g \in A^k$ then $|\pi(A)| \leq \frac{|A^{k+3}|}{|\varphi_g(A)|}$.

Proof. (a) This is just an easy computation: as

$$\begin{pmatrix} r & s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} \begin{pmatrix} r^{-1} & -r^{-1}s \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} a & br + (1-a)s \\ 0 & 1 \end{pmatrix},$$

two elements are in the preimage of a single point if and only if $br + (1-a)s = br' + (1-a)s'$, from which all pairs of elements with $\frac{s'-s}{r'-r} = \frac{b}{a-1}$ must be sent

to the same point by φ_g (we allow $a = 1$ and a slope equal to ∞, but we avoid $(a, b) = (1, 0)$ since $g \neq \mathrm{Id}$).

(b) On one hand we have $|A\varphi_g(A)g^{-1}| = |A\varphi_g(A)| \leq |A^{k+3}|$, while on the other hand any element of $A\varphi_g(A)g^{-1}$ is of the form $a_1a_2ga_2^{-1}g^{-1} \in AU$: since

$$\begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & z \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} x & xz+y \\ 0 & 1 \end{pmatrix},$$

pairs in $A \times U$ with either distinct x or with the same x, y but distinct z will all give different products in AU; hence we can select one element of A for each value of x (therefore $|\pi(A)|$ of them) and all the $a_2ga_2^{-1}g^{-1}$ ($|\varphi_g(A)|$ of them, they are all multiplied by the same g^{-1}) and obtain the other side of the bound, namely $|A\varphi_g(A)g^{-1}| \geq |\pi(A)||\varphi_g(A)|$. □

With these tools at our disposal, we can proceed with the proof.

Proof of Thm. 4.1.4. Let us start with the case of A large: impose $|A| \geq cq$ for a constant $c > 1$ to be chosen later.

We use the bound from Proposition 4.3.1 with $P = A$ and $L = \overline{L(A)}$ (the set of lines that are not determined by A), interpreted as a lower bound on the expression inside the absolute value, and combine it with the trivial observation that $I(A, \overline{L(A)}) \leq |\overline{L(A)}|$ by the definition of $\overline{L(A)}$: this yields

$$|\overline{L(A)}| \leq q^2 \frac{c}{(c-1)^2} \qquad \Longrightarrow \qquad |L(A)| \geq q + q^2 \left(1 - \frac{c}{(c-1)^2}\right).$$

If $c \geq 1 + \frac{1+\sqrt{3-\frac{2}{p}}}{1-\frac{1}{p}}$ (or $c \geq 3 + 2\sqrt{2}$, which is an upper bound for all primes p), by the pigeonhole principle there must exist $\geq \frac{q}{2}\left(1 + \frac{1}{p}\right)$ non-vertical lines of $L(A)$ parallel to each other; call d the direction defined by such lines. Given any two elements of A sitting on one of these lines, we have

$$g = \begin{pmatrix} a_1 & b_1 \\ 0 & 1 \end{pmatrix}^{-1} \begin{pmatrix} a_2 & b_2 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} a_2a_1^{-1} & b_2a_1^{-1} - b_1a_1^{-1} \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} a' & b' \\ 0 & 1 \end{pmatrix},$$

with $\frac{b'}{a'-1} = \frac{b_2-b_1}{a_2-a_1} = d$, so by Lemma 4.3.2(a) there are $\geq \frac{q}{2}\left(1 + \frac{1}{p}\right) > \frac{q}{2}$ elements in $\varphi_g(A)$; by Lemma 4.3.2(b) and Ruzsa's inequality, this implies that $|\pi(A)| < \frac{2|A^5|}{q} \leq \frac{2}{q}C^3|A|$. Moreover, the unipotent subgroup U is isomorphic to $\mathbb{F}_q$ as an additive group, so that its largest proper subgroup is of size $\frac{q}{p}$; therefore, since $\varphi_g(A)g^{-1} \subseteq A^6 \cap U$ has $|\varphi_g(A)g^{-1}| \geq \frac{q}{2}\left(1 + \frac{1}{p}\right)$, by Kneser's theorem we must have

$$A^8 \supseteq AgAAg^{-1}A \supseteq (\varphi_g(A)g^{-1})(\varphi_g(A)g^{-1})^{-1} \supseteq U,$$

and we fall into case (c) of the theorem.

Let us deal now with A of medium size: suppose $q < |A| < cq$, so that by Lemma 4.2.1 every direction is determined by some pair of points of A. Partition

$A^2 \setminus \{\mathrm{Id}\}$ into $q+1$ subsets, collecting into each one of them elements having the same value for $\frac{b}{a-1} \in \mathbb{F}_q \cup \{\infty\}$. Every two distinct $a_1, a_2 \in A$ yield an element $a_1^{-1}a_2 \in A^2$ that is located inside the part corresponding to the slope of the line they define: by the pigeonhole principle there will be a part (identifiable with some $d \in \mathbb{F}_q \cup \{\infty\}$) with at most $\frac{|A^2|-1}{q+1}$ elements, and therefore every line of $L(A)$ in the direction d must have at most $\frac{|A^2|-1}{q+1}+1$ elements of A on it, since $a_1^{-1}a_i \neq a_1^{-1}a_j$ for $a_i \neq a_j$. We have thus given a bound on the number of points of A sent to the same element by the map φ_g for some $g \in A^2$ with $\frac{b}{a-1} = d$, which translates to

$$|\varphi_g(A)| \geq \frac{(q+1)|A|}{|A^2|+q} > \frac{|A|}{Cc+1};$$

proceeding as before, by Lemma 4.3.2(b) and Ruzsa's inequality we conclude that $|\pi(A)| < (Cc+1)C^3 \leq (4+2\sqrt{2})C^4$ and we reach case (b) of the theorem.

For A small (i.e. $1 < |A| \leq q$) we repeat essentially what we did for A medium, but instead of $|D| = q+1$ we use the bounds in Theorem 4.1.3 on the number of directions $|D|$ spanned by A. We obtain

$$|\varphi_g(A)| > \frac{|A|^2}{q'(|A^2|-1)+|A|} > \frac{|A|}{Cq'+1},$$

where $q' = \sqrt{q}$ for e even and $q' = p^{\frac{e-1}{2}}+1$ for e odd, from which we get $|\pi(A)| < (p^{\lfloor \frac{e}{2} \rfloor}+2)C^4$ and end up in case (b). Finally, we need to deal with the other alternative in Theorem 4.1.3, namely that A may be contained in one line: in other words, the elements of A are either all of the form $(a, ad+b)$ for some $b, d \in \mathbb{F}_q$, through the identification of $\mathrm{Aff}(\mathbb{F}_q)$ with $\mathbb{F}_q^* \times \mathbb{F}_q$, or all contained in U. In the latter alternative $A \subseteq U$ implies $|\pi(A)| = 1$, yielding (b); in the former, since $A = A^{-1}$ and $(a, ad+b)^{-1} = (a^{-1}, -a^{-1}b-d)$, we must have $b = -d$ and then $A \subseteq \mathrm{Stab}(-d)$. □

4.4 Concluding remarks

Theorem 4.1.4, as addressed several times, gives a structural result on $\mathrm{Aff}(\mathbb{F}_q)$; more than that, it shows that the techniques used for the case of $\mathbb{F}_p$ in [RS18] can be generalized to arbitrary finite fields. This is not a new concept, as the sum-product theorem has also followed an analogous trajectory, although the same cannot be said of the proofs about the affine group specifically. There is value in such feats, as it has often been the case that results for $\mathbb{F}_q$ have followed the $\mathbb{F}_p$ case once the machinery had been understood and streamlined; as commented on in §1.4, in some situations it is yet an ongoing process, for example with [GH14] still missing an equally strong counterpart in $\mathbb{F}_q$ (where weakened generalizations are included however in [BGT12] and [PS14]).

It is the hope of the author that the present result provides another small tessera in the mosaic of growth in matrix groups. The differences between Theorem 4.1.2 and Theorem 4.1.4 seem also to reflect the general divergence point

between prime fields and general finite fields: the obstacle presented by $\sqrt{q}$ is in substance the same as providing that we avoid proper subfields, for its origin lies in the complications of the Frobenius map $x \mapsto x^p$ in the course of the proof of Theorem 4.2.2. In this sense, the result also works as a reaffirmation of deeply rooted principles that are likely to resurface in future research on similar cases, with matrix groups having larger rank and a more complicated structure without close access to the perks of a well-understood geometric structure or sum-product phenomena.

In particular, this refers to the example (immediately following Theorem 4.1.4) of a set A whose growth is stifled by a subfield. The question of whether the example we provided is essentially the only one, in line with the structure predictions of the Helfgott-Lindenstrauss conjecture [Hel15, Conj. 1], is not answered here; we thank H. Helfgott and M. Tointon for independently raising this question. However, the methods involved in the proof of Theorem 4.1.4 seem to yield themselves to be employed in such a task: having a power of p as a factor in the first half of Theorem 4.1.4(b) translates into a condition on the polynomials describing the points of A as being polynomials in some power x^{p^l}, instead of simply polynomials in x. This in turn might provide enough information on the arrangement of the points of A in the affine plane (arrangement that defines $H_y(x)$, essentially) to say that A must necessarily be "stuck in a subfield". This avenue of inquiry would show a quantitative version of the aforementioned conjecture for the case of $\mathrm{Aff}(\mathbb{F}_q)$, and it is worth exploring.

Chapter 5

Diameter bounds for products of finite simple groups

The content of this chapter is essentially taken from [Don19a].

As addressed in §1, an important area of research in finite group theory in the last decades has been the production of upper bounds for the diameter of Cayley graphs of such groups. We have mentioned Babai's conjecture (Conjecture 1.3.1), asking for diameter bounds polylogarithmic in $|G|$ for finite simple groups, and we observed that it is still open, despite great progress towards a solution both for alternating groups and for groups of Lie type.

A more modest question is that of producing bounds for the diameter of direct products of finite simple groups, depending on the diameter of their factors. This is not an idle question, for bounds of this sort have been used more than once as intermediate steps towards the proof of bounds for simple groups themselves: Babai and Seress have done so in [BS92, Lemma 5.4], as well as Helfgott more than two decades later in [Hel18, Lemma 4.13]. Here we give a result that improves on both, in ways that we are going now to discuss.

5.1 Main theorem

The following is our main theorem of the chapter. The bounds presented therein feature explicit constants, unlike in [Hel18], since it does not take any particular effort to keep track of computations.

Theorem 5.1.1. *Let $G = \prod_{i=1}^{n} T_i$, where the T_i are finite simple groups.*

(a) If the T_i are all abelian (say $G = \prod_{j=1}^{s} (\mathbb{Z}/p_j\mathbb{Z})^{e_j}$, where the p_j are distinct

primes and $e_j \geq 1$), then

$$\operatorname{diam}(G) < \frac{2}{3} \max\{e_j | 1 \leq j \leq s\} \prod_{j=1}^{s} p_j.$$

(b) If the T_i are all non-abelian, call $d = \max\{\operatorname{diam}(T_i) | 1 \leq i \leq n\}$; then

$$\operatorname{diam}(G) < \frac{196}{243} n^3 \max\{C_A, C_L, C_S\}(4d+1) + d,$$

where

$$C_A = \begin{cases} \max\left\{3, \left\lfloor \frac{m}{2} \right\rfloor\right\} & \text{if there are alternating groups among the } T_i \\ & \text{and where } m \text{ is their maximum degree,} \\ 0 & \text{if there are no alternating groups among the } T_i, \end{cases}$$

$$C_L = \begin{cases} 8(5r+7) & \text{if there are groups of Lie type among the } T_i \\ & \text{and where } r \text{ is their maximum untwisted rank,} \\ 0 & \text{if there are no groups of Lie type among the } T_i, \end{cases}$$

$$C_S = \begin{cases} 6 & \text{if there are sporadic or Tits groups among the } T_i, \\ 0 & \text{if there are no sporadic or Tits groups among the } T_i. \end{cases}$$

(c) If there are abelian and non-abelian T_i, write $G = G_A \times G_{NA}$, where G_A collects the abelian factors and G_{NA} collects the non-abelian ones; then

$$\operatorname{diam}(G) \leq d_A + 4d_{NA},$$

where $d_A = \operatorname{diam}(G_A), d_{NA} = \operatorname{diam}(G_{NA})$.

Needless to say, the subdivision of cases that occurs in the theorem descends from CFSG (Theorem 1.2.4), which we take for granted and do not bother mentioning again in the course of the proof.

The result of part (a) is known and elementary: see [BS92, Lemma 5.2], where the constant is marginally worse only due to the fact that sets of generators are not required to be symmetric (cfr. also [Hel18, Lemma 4.14], which treats the case of $G = (\mathbb{Z}/p\mathbb{Z})^e$ under this assumption). Part (c) is quite natural, given the different (in some sense, opposite) behaviour of abelian and non-abelian factors, as it can be readily observed in its short proof.

Part (b) is where the novelty of the result resides. Dependence on the maximum of the diameter of the components, instead of dependence on their product as Schreier's lemma (see Lemma 5.2.1) would naturally give us, was already established in [BS92, Lemma 5.4]: in that case, the diameter was bounded as $O(d^2)$, where the dependence of the constant on n was polynomial as in our statement. This result was improved in [Hel18, Lemma 4.13] to $O(d)$, but only in the case of alternating groups: this was done in part to fix a mistake in the use of the previously available result in Babai-Seress, which is why only alternating groups

were considered, as permutation subgroups were the sole concern in both papers; a suggestion by Pyber, reported in Helfgott's paper, points at the results by Liebeck and Shalev [LS01] as a way to prove a bound of $O(d)$ for a product of arbitrary non-abelian finite simple groups.

Indeed, the general approach that we follow in our proof owes its validity to [LS01, Thm. 1.6], although we do not explicitly use the statement of that theorem: rather, we closely follow the proof of [Hel18, Lemma 4.13] and show that the same reasoning applies to groups of Lie type as well. The way that the lemma is related to Liebeck-Shalev is through the use of the fact that every element in $\mathrm{Alt}(m)$ is a commutator ([Hel18, Lemma 4.12], first proved in [Mil99, Thm. I]), which is essentially [LS01, Thm. 1.6] with $w = xyx^{-1}y^{-1}$ and a c that is just equal to 1 for $\mathrm{Alt}(m)$; the same can be said for all non-abelian finite simple groups (i.e., $c = 1$ in general) since Ore's conjecture [Ore51] was established to be true in [LOST10], a fact yet unproved at the time of [LS01].

5.2 Preliminaries

Before we turn to the proof of Theorem 5.1.1, we will need a certain number of group-theoretic results.

Lemma 5.2.1 (Schreier's lemma)**.** *Let G be a finite group, let $N \trianglelefteq G$, and let S be a set of generators of G with $e \in S = S^{-1}$. Then $S^{2d+1} \cap N$ generates N, where $d = \mathrm{diam}(G/N)$.*

Proof. This is a standard result dating back to Schreier [Sch27], written in various fashions across the literature according to the needs of the user; let us prove here the present version (we are going to see another one in Lemma 6.2.5).

Calling $\pi : G \to G/N$ the natural projection, by definition we have $\pi(S)^d = G/N$; this equality means that S^d contains at least one representative for each coset gN in G. For any coset gN, choose a representative $\tau(g) \in S^d$. Then, for any $h \in N$ and any way to write h as a product of elements $s_i \in S$, we have

$$\begin{aligned} h &= s_1 s_2 \dots s_k = \\ &= (s_1 \tau(s_1)^{-1}) \cdot (\tau(s_1) s_2 \tau(\tau(s_1)s_2)^{-1}) \cdot \dots \cdot (\tau(\tau(\tau(\dots)s_{k-2})s_{k-1})s_k). \end{aligned}$$

Each element of the form $\tau(x)s_i\tau(\tau(x)s_i)^{-1}$ is contained in $S^{2d+1} \cap N$, so the same can be said about the last element of the form $\tau(x)s_k$ (since h itself is in N); therefore $S^{2d+1} \cap N$ is a generating set of N. □

Proposition 5.2.2 (Ore's conjecture)**.** *Let G be a finite non-abelian simple group. Then, for any $g \in G$, there exist $g_1, g_2 \in G$ such that $g = [g_1, g_2]$.*

Proof. See [LOST10], for references to previously known results and for the proof of the final case. □

Notice that, for any finite non-abelian simple group G, any nontrivial conjugacy class C must generate the whole G (because $\langle C \rangle$ would be a normal subgroup). This observation justifies the following definition.

Definition 5.2.3. *Let G be a finite non-abelian simple group. The conjugacy diameter* $\mathrm{cd}(G)$ *is the smallest m such that $(C \cup C^{-1} \cup \{e\})^m = G$ for all nontrivial conjugacy classes C.*

We will need to have bounds for $\mathrm{cd}(G)$.

Proposition 5.2.4. *Let G be a finite non-abelian simple group.*

(a) If G is an alternating group of degree m, then $\mathrm{cd}(G) \leq \max\left\{3, \left\lfloor \frac{m}{2} \right\rfloor\right\}$.

(b) If G is a group of Lie type of untwisted rank r, then $\mathrm{cd}(G) \leq 8(5r+7)$.

(c) If G is a sporadic group or the Tits group, then $\mathrm{cd}(G) \leq 6$.

Proof. First of all, $\mathrm{cd}(G)$ is trivially bounded by definition by the *covering number* of G, which is defined as $\mathrm{cn}(G) = \min\{m | \forall C \neq \{e\}(C^m = G)\}$; therefore it suffices to give bounds for $\mathrm{cn}(G)$.

For (a), see [Dvi73, Thm. 9.1] (our specific result is credited therein to a manuscript by J. Stavi). For (b), see [LL98, Thm. 1]. To prove (c), the sporadic groups all satisfy $\mathrm{cn}(G) \leq 6$: this inequality can be checked directly from [Zis89, Table 1]; if $G = {}^2F_4(2)'$ is the Tits group, we can show the same inequality using [Zis89, Lemma 3] and the character values reported in the ATLAS of Finite Groups [CCN+85]. □

Let us also perform a side computation separately from the proof of the main theorem, so as not to bog down the exposition there.

Lemma 5.2.5. *Let $n \geq 1$. Then*

$$\sum_{i=1}^{n-1} 4^{\lceil \log_2 i \rceil} < \frac{196}{243} n^3.$$

Proof. Call $m = \lceil \log_2(n-1) \rceil$, and write $n - 1 = 2^{m-1} + l$, where $1 \leq l \leq 2^{m-1}$; $\lceil \log_2 i \rceil = j$ for all $i \in (2^{j-1}, 2^j]$, hence we can rewrite the sum in the statement as

$$\sum_{i=1}^{n-1} 4^{\lceil \log_2 i \rceil} = 1 + \sum_{j=1}^{m-1} 4^j 2^{j-1} + 4^m l = \frac{1}{2} + \frac{1}{2}\frac{8^m - 1}{7} + 4^m \left(2^{\log_2(n-1)} - 2^{m-1}\right) =$$

$$= \frac{3}{7} + 4^m 2^{\log_2(n-1)} - \frac{3}{7} 8^m = \frac{3}{7} + 2^{2m'} \left(1 - \frac{3}{7} 2^{m'}\right) (n-1)^3,$$

where $m' = m - \log_2(n-1) \in [0, 1)$. We have $x^2\left(1 - \frac{3}{7}x\right) \leq \frac{196}{243}$ for $x \in [1, 2)$, and $\frac{3}{7} < \frac{196}{243}(3n^2 - 3n + 1)$ for all $n \geq 1$, so the result is proved. □

5.3 Proof of the main theorem

Proof of Thm. 5.1.1(a). Let $G = (\mathbb{Z}/p_1\mathbb{Z})^{e_1} \times (\mathbb{Z}/p_2\mathbb{Z})^{e_2} \times \ldots \times (\mathbb{Z}/p_s\mathbb{Z})^{e_s}$, with primes $p_1 < p_2 < \ldots < p_s$; we have

$$G = A_1 A_2 \ldots A_s \tag{5.3.1}$$

(we are using multiplicative notation even if G is abelian) where the A_i are any sets such that

$$A_{i,i} = (\mathbb{Z}/p_i\mathbb{Z})^{e_i} \qquad A_{i,j} = (0)^{e_j} \quad (\forall j < i) \tag{5.3.2}$$

where $A_{i,j}$ is the projection of A_i to the j-th component of G.

Let S be a set of generators of G with $e \in S = S^{-1}$: $\{t^{p_1 \cdots p_{i-1}} | t \in S\} \subseteq S^{p_1 \cdots p_{i-1}}$ has elements that are all 0 on the first $i-1$ components of G and that still generate the i-th one since $(p_1 \dots p_{i-1}, p_i) = 1$; from now on, let us focus exclusively on the i-th component. $(\mathbb{Z}/p_i\mathbb{Z})^{e_i}$ is also a vector space over $\mathbb{Z}/p_i\mathbb{Z}$, so there must be e_i generators that also form a basis: any element of the space can be written as a linear combination of those generators with coefficients in $\left[-\left\lfloor \frac{p_i}{2} \right\rfloor, \left\lfloor \frac{p_i}{2} \right\rfloor\right]$, which corresponds to a word of length $\leq e_i \left\lfloor \frac{p_i}{2} \right\rfloor$; thus, each set A_i with the properties in (5.3.2) is covered in $e_i \left\lfloor \frac{p_i}{2} \right\rfloor p_1 \dots p_{i-1}$ steps. This fact and (5.3.1) imply that G has diameter bounded by

$$\sum_{i=1}^{s} \left(e_i \left\lfloor \frac{p_i}{2} \right\rfloor \prod_{j=1}^{i-1} p_j \right) \leq \frac{1}{2} \max\{e_j | 1 \leq j \leq s\} \prod_{j=1}^{s} p_j \cdot \sum_{i=1}^{s} \left(\prod_{j=i+1}^{s} \frac{1}{p_j} \right). \tag{5.3.3}$$

The sum in (5.3.3) is maximized when each p_j is the j-th prime number: for $s = 1$ the sum is 1 and for $s = 2$ it is bounded by $\frac{4}{3}$; for $s \geq 3$, we use $p_s \geq 5$ and $p_j \geq 3$ for all $1 < j < s$, so that the sum is bounded by $1 + \frac{1}{5}\frac{1}{1-\frac{1}{3}} = \frac{13}{10}$. The result follows. □

Proof of Thm. 5.1.1(b). Calling $G_j = \prod_{i=1}^{j} T_i$, we have natural projections $\pi_j : G = G_n \to G_j$ and $\rho_{j_1,j_2} : G_{j_1} \to T_{j_2}$ for any $j_1 \geq j_2$. As in (5.3.1), we write G as a product of subsets A_i with $\rho_{n,i}(A_i) = T_i$ and $\rho_{n,j}(A_i) = \{e\}$ for all $j < i$, and our aim is to cover each one of them.

Suppose that we have two subsets X_1, X_2 of G for which $\rho_{n,i}(X_1) = \rho_{n,i}(X_2) = T_i$ for some fixed $i \in \{1, \dots, n\}$ and that have $\rho_{n,j_1}(X_1) = \{e\} = \rho_{n,j_2}(X_2)$ for all $j_1 \in I_1, j_2 \in I_2$, where I_1, I_2 are two subsets of indices in $\{1, \dots, n\} \setminus \{i\}$: then, the set $X = \{[x_1, x_2] | x_1 \in X_1, x_2 \in X_2\}$ has $\rho_{n,i}(X) = T_i$ by Proposition 5.2.2 (Ore's conjecture) and $\rho_{n,j}(X) = \{e\}$ for all $j \in I_1 \cup I_2$. Now consider the set of indices $I = \{1, \dots, i-1\}$: if $|I| > 1$ we can partition I into two parts of size $\left\lfloor \frac{|I|}{2} \right\rfloor, \left\lceil \frac{|I|}{2} \right\rceil$, then partition each part I' with $|I'| > 1$ into two new parts again of size $\left\lfloor \frac{|I'|}{2} \right\rfloor, \left\lceil \frac{|I'|}{2} \right\rceil$, and continue until we reach a subdivision where all sets have size 1; the tree of partitions that we constructed to reach this subdivision will have exactly $\lceil \log_2 |I| \rceil$ layers. Notice that, given any two parts I_1, I_2 inside the tree, if we have two subsets X_1, X_2 (as described before) that are covered by a certain S^a, the resulting set X will be covered by S^{4a}: this observation, together with the information about the layers, tells us that if we can cover sets $X_{i,j}$ with $\rho_{n,i}(X_{i,j}) = T_i$ and $\rho_{n,j}(X_{i,j}) = \{e\}$ in a steps (for a fixed $i > 1$ and all $j < i$) then we are able to cover a set A_i defined as at the beginning of the proof in $4^{\lceil \log_2 (i-1) \rceil} a$ steps as well.

Let us start now with a generating set S with $e \in S = S^{-1}$ and fix two indices $i \geq j$: $\pi_i(S)$ is a set of generators for G_i, and the set $\pi_i(S)^{2d+1}$ contains

generators for the whole $T_1 \times \ldots \times T_{j-1} \times \{e\} \times T_{j+1} \times \ldots \times T_i = G_i \cap \ker(\rho_{i,j})$ by Lemma 5.2.1 (Schreier's lemma), where d is as in the statement. In particular, there is an element $x \in S^{2d+1}$ with $\rho_{n,i}(x) \neq e$ and $\rho_{n,j}(x) = e$; by hypothesis $\rho_{n,i}(S^d) = T_i$, which means that there is a set $S' = \{yxy^{-1} | y \in S^d\} \cup \{yx^{-1}y^{-1} | y \in S^d\} \cup \{e\} \subseteq S^{4d+1}$ with $\rho_{n,i}(S') = C \cup C^{-1} \cup \{e\}$ and $\rho_{n,j}(S') = \{e\}$, where C is the conjugacy class of $\rho_{n,i}(x)$. By Proposition 5.2.4, $\rho_{n,i}(S'^{\max\{3, \lfloor \frac{m_i}{2} \rfloor\}}) = T_i$ if $T_i = \mathrm{Alt}(m_i)$, $\rho_{n,i}(S'^{8(5r_i+7)}) = T_i$ if T_i is of Lie type of untwisted rank r_i, and $\rho_{n,i}(S'^6) = T_i$ otherwise; in all three cases, the projection to T_j is still $\{e\}$, therefore we managed to cover a set $X_{i,j}$ of the aforementioned form.

A set A_1 is reached in d steps, hence the final count for the whole G following the reasoning above is

$$\mathrm{diam}(G) \leq d + \sum_{i=2}^{n} 4^{\lceil \log_2(i-1) \rceil} x_i (4d+1),$$

where x_i is either $\max\{3, \lfloor \frac{m_i}{2} \rfloor\}$, $8(5r_i + 7)$ or 6, accordingly. The result follows by Lemma 5.2.5. □

A note on the connection between the proof given above and [LS01]. As mentioned before, Pyber pointed at [LS01] as a way to prove linear dependence on d for products of arbitrary non-abelian finite simple groups. In particular, [LS01, Thm. 1.6] seems to fit the bill: it states that for any word w that is not a law in a finite simple group T there is $c_w \in \mathbb{N}$, depending on w but not on T, such that any element of T can be written as a product of at most c_w values of w. We use this property, in disguise, when we want to pass from two subsets being indentically e at indices I_1, I_2 and filling an entire component T_i to a third subset that also fills the same component and is e for the whole $I_1 \cup I_2$: the creation of the new subset is made possible by taking c_w values of a word w, so that T_i remains filled, where w has two distinct letters x_1, x_2 and presents the same number of x_i and x_i^{-1} for $i \in \{1, 2\}$, so that when any one x_i is equal to e on a given factor of the product G the result is e on that factor; in our case, w was the shortest nontrivial word with these characteristics, namely the commutator $[x_1, x_2] = x_1 x_2 x_1^{-1} x_2^{-1}$ (not a law for any non-abelian group), and $c_w = 1$ by Ore's conjecture. In this sense $w = [x_1, x_2]$ is also computationally the best word we can expect, for it yields the lowest possible value of $|w| c_w$, the 4 that we find in Lemma 5.2.5.

Proof of Thm. 5.1.1(c). Define the two projections π_A, π_{NA} in the obvious way; for any generating set S of G, by definition there are a subset $X_A \subseteq S^{d_A}$ with $\pi_A(X_A) = G_A$ and a subset $X_{NA} \subseteq S^{d_{NA}}$ with $\pi_{NA}(X_{NA}) = G_{NA}$, and then

$$G = X_A [X_{NA}, X_{NA}] \subseteq S^{d_A + 4 d_{NA}},$$

again by the fact that $[T, T] = T$ for non-abelian finite simple groups by Ore's conjecture and $[T, T] = \{e\}$ for abelian groups (here with $[X, X]$ we mean the set of commutators, and not the commutator subgroup). □

5.4 Concluding remarks

One may wonder how tight the inequalities in Theorem 5.1.1 are. The results are essentially in line with what is generally expected from the behaviour of the diameter of finite groups. The abelian case is tight up to constant: for the group $G(x) = \prod_{p \leq x} \mathbb{Z}/p\mathbb{Z}$ (nontrivial for $x \geq 2$) one generator $s = (1, 1, \dots, 1)$ is enough, and then the diameter of $\mathrm{Cay}(G(x), \{s, s^{-1}, e\})$ is $\frac{1}{2}|G(x)|$; the fact that abelian groups behave in the worst possible way, i.e. linearly in the size of the group, should not be a surprise for anyone (this has been discussed in §1.1 as well).

The non-abelian bound of case (b) also matches what is anticipated in general. Babai's conjecture posits a polylogarithmic bound on the diameter of finite simple groups: the natural extension to direct products of such groups would suggest a bound of the form $n^k d$, which is exactly what we have obtained. Case (c) also fits into the same idea, as a product $|G| = |G_A||G_{NA}|$ becomes a sum of the corresponding diameters.

The dependence on d in Theorem 5.1.1(b) is almost best possible by definition (we cannot drop the "almost", as m, r are not independent from d). It would be more interesting to understand which power of n is the correct one: here we have proved $O_{m,r,d}(n^3)$, and we can quickly show that the bound is $\Omega_{m,r,d}(n)$, as illustrated in the following example.

Example 5.4.1. If $G = (\mathrm{Alt}(m))^n$ then $\mathrm{diam}(G) = \Omega(m^2 n)$. We prove it for $m \geq 5$ odd and n even, but the proof is analogous for the general case.

Consider the two permutations $\sigma = (1\ 2\ 3\ \dots\ m)$ and $\tau = (1\ 2\ 3\ \dots\ m-2)$; they generate $\mathrm{Alt}(m)$, and the elements

$$\begin{aligned} s_0 &= (\sigma, \sigma, \dots, \sigma, \sigma), \\ s_1 &= (\tau, \sigma, \dots, \sigma, \sigma), \\ s_2 &= (\sigma, \tau, \dots, \sigma, \sigma), \\ &\dots \\ s_n &= (\sigma, \sigma, \dots, \sigma, \tau) \end{aligned}$$

generate G. Let $S = \{e\} \cup \{s_i, s_i^{-1}\}_{0 \leq i \leq n}$: to prove the lower bound on the diameter of G, we construct a function $f : G \to \mathbb{N}$ such that there are two elements $g_1, g_2 \in G$ with $|f(g_1) - f(g_2)|$ large and such that $|f(g) - f(gs)|$ is small for any $g \in G, s \in S$; this is a known technique to prove lower bounds for the diameter of $\mathrm{Sym}(m)$, as shown for instance in [Tan11, Prop. 3.6].

Call $c(g, i, j) = (g(i))(j)$ the image of $j \in \{1, \dots, m\}$ under the i-th component of $g \in G$, for $1 \leq i \leq n$; define

$$f(g) = \sum_{j=1}^{m} \sum_{i=1}^{n} ||c(g, i+1, j) - c(g, i, j)||_{\mathbb{Z}/m\mathbb{Z}},$$

where $||a||_{\mathbb{Z}/m\mathbb{Z}} = \min\{a, m-a\}$ (in the case $i = n$, $c(g, n+1, j)$ means $c(g, 1, j)$). First, $f(e) = 0$; also, if we call e_m the identity element in $\mathrm{Alt}(m)$ and $\eta = \left(1\ \frac{m+1}{2}\right)\left(2\ \frac{m+3}{2}\right) \dots \left(\frac{m-1}{2}\ m-1\right)$, for $g \in G$ that has e_m at all odd components

and η at all even ones we have $f(g) = \frac{1}{2}(m-1)^2 n$. Finally, notice that σ simply adds 1 modulo m to all the elements of $\{1, \ldots, m\}$, so that $f(g) = f(gs_0^{\pm 1})$, while τ is defined so that it adds 1 for $m-3$ elements, adds 3 (modulo m) for one element and fixes two elements, which means that $|f(g) - f(gs_i^{\pm 1})| \leq 10$; these facts taken together imply that $\operatorname{diam}(G, S) \geq \frac{1}{20}(m-1)^2 n$.

The correct (or even expected) order of magnitude for a bound of the form $\operatorname{diam}(G) = O_{m,r}(n^k d)$ for a generic product G is not known to the author, besides knowing that $1 \leq k \leq 3$ by Theorem 5.1.1 and Example 5.4.1.

Chapter 6

Towards a CFSG-free diameter bound for $\mathrm{Alt}(n)$

As discussed at length in §1.3, Babai's conjecture (Conjecture 1.3.1) is an important open problem in the context of finite group theory. Because of CFSG (Theorem 1.2.4), we know that we need only to treat the two cases of groups of Lie type and of alternating groups, and in fact most proofs to date produce results in only one of the two classes.

The strongest known result for the alternating case is (1.3.4), which was proved by Helfgott and Seress [HS14]: it was a big improvement over the previous best bound (1.3.1), due to Babai and Seress [BS88], and is quite close to the actual order of magnitude that Babai's conjecture anticipates. Later, Helfgott [Hel18] gave a proof of slightly less tight bound for (1.3.4) (with $(\log\log n)^2$ in the exponent instead of $\log\log n$, see (6.1.1)) that made use of a weakened product theorem (like Theorem 1.3.2), so as to provide a more general framework for the problem and shrink the distance between proofs for permutation subgroups and proofs for groups of Lie type; on the significance and limitations of product theorems, see again §1.3.

Both [HS14] and [Hel18] rely in some way on CFSG: as a matter of fact they are both based at their core on the classification of primitive permutation subgroups given here as Theorem 1.2.5, in primis due to Cameron [Cam81] and refined later by Liebeck [Lie84] and Maróti [Mar02], which descends from CFSG and the O'Nan-Scott theorem. Removing the dependence on CFSG from (1.3.4) or analogous results would be in the words of Helfgott "another worthwhile goal" [Hel18, §1].

Our aim in this chapter will be to walk at least part of the way towards that goal: modulo an unproven assumption (Conjecture 6.3.4), we will be giving a CFSG-free proof of a diameter bound for $\mathrm{Alt}(n)$, and in fact for all transitive subgroups of $\mathrm{Sym}(n)$, which is not as strong as the one given by Helfgott and Seress but is a decisive improvement on the CFSG-free bound due to Babai and Seress; the main result is Theorem 6.3.6. To do so, we are going to make use of a tool that has not been adopted before in the context of Babai's conjecture, namely Babai's

quasipolynomial GIP algorithm, as already foreshadowed in §1.6: in particular, the analysis we have performed in §3 will be instrumental in accomplishing what we want; we discuss the strategy behind it in §6.1, together with overviewing what is needed from [Hel18] to understand the context of the already known result that relies on CFSG.

6.1 Background and strategy

Helfgott's result [Hel18, Thm. 6.1] on the diameter of $\mathrm{Alt}(n)$ is the following: if $G = \mathrm{Alt}(n), \mathrm{Sym}(n)$ then

$$\mathrm{diam}(G) \leq e^{O(\log^4 n (\log\log n)^2)}. \tag{6.1.1}$$

To prove the bound above, he shows that a sort of product theorem also holds in the context of permutation groups. We have seen that product theorems, like Theorem 1.3.2, are central to proofs in the Lie type case, and that a statement as strong as that cannot hold for $\mathrm{Alt}(n)$: there are specific counterexamples in [Spi12, §4] and [PPSS12, Thm. 17]. However, the weakened version below is still true (for the meaning of 3-transitive, see Definition 3.1.1).

Theorem 6.1.1 ([Hel18], Theorem 1.4)**.** *Let $G \leq \mathrm{Sym}(n)$ be 3-transitive, and let A be a set of generators of G with $e \in A = A^{-1}$. Then there are absolute constants $C, k > 0$ such that, if $|A| \geq n^{C \log^2 n}$, then at least one of the following alternatives holds:*

(a) $|A^{n^C}| \geq |A|^{1 + \frac{k \log\log |A|}{\log^2 n \log\log n}}$;

(b) there is a transitive $G' \leq \mathrm{Sym}(n')$ with $n' \leq n$ such that $\mathrm{diam}(\mathrm{Cay}(G, A)) \leq n^C \mathrm{diam}(G')$*, and either $n' \leq e^{-\frac{1}{10}} n$ or $G' \neq \mathrm{Sym}(n'), \mathrm{Alt}(n')$.*

Theorem 6.1.1 still qualifies as a sort of product theorem, in the sense that after as many instances as possible of growth of $|A|$ in case (a), like in Theorem 1.3.2(a), we reach in case (b) a bound on the diameter of the Cayley graph of G and the final power of A, which was neater for Theorem 1.3.2(b) (it was 3) whereas now it sparks a recursion process. In this sense, Theorem 6.1.1 is part of an effort to close the gap between the Lie type proofs and the alternating proofs.

What is important for us, though, is that Theorem 6.1.1 implies (6.1.1) via (1.3.3) (which is [Hel18, Prop. 4.15], and is part of the aforementioned recursion process), and that the part of the proof of (6.1.1) that depends on CFSG is contained solely in (1.3.3), whereas Theorem 6.1.1 itself is CFSG-free. Therefore, what we need is to show something resembling (1.3.3) without the help of CFSG, and then we can conclude in a way that is not different from what has already appeared in [Hel18]; the end of the proof of Theorem 6.3.6 will proceed exactly along these lines.

As for the strategy leading to that point, it is as follows. An intermediate result in Helfgott's proof, i.e. [Hel18, Prop. 4.6], produces a nicely shaped chain of normal subgroups necessary to reach the conclusion (1.3.3). Let us write it down for future reference.

Proposition 6.1.2. *Let $G \leq \mathrm{Sym}(n)$ be transitive. Then there is a composition series $\{e\} = H_0 \lhd H_1 \lhd \ldots \lhd H_\ell = G$ and a partition $\{C_1, C_2\}$ of the set of composition factors H_i/H_{i-1} with the following properties:*

(a) if $H_i/H_{i-1} \in C_1$ then $H_i/H_{i-1} \simeq M_i^{k_i}$ for some M_i simple and $k_i \leq 2n$;

(b) $\prod_{H_i/H_{i-1} \in C_1} |H_i/H_{i-1}| = n^{O(\log n)}$;

(c) if $H_i/H_{i-1} \in C_2$ then $H_i/H_{i-1} \simeq \mathrm{Alt}(m_i)^{k_i}$ for some $m_i \geq 5$ and $k_i \leq 2n$;

(d) $\prod_{H_i/H_{i-1} \in C_2} m_i \leq n$, and each $m_i \leq \frac{n}{2}$ unless G is a giant;

(e) $\ell = O(\log n)$.

We aim to replace this intermediate result with a CFSG-free analogue that would prove a counterpart of (1.3.3). Cameron's structure theorem (Theorem 1.2.5) is the backbone of the proof of [Hel18, Prop. 4.6], as it breaks down permutation groups into pieces that are either small or alternating (represented here by the factors in C_1 and C_2 respectively) and allows us to construct the chain: a CFSG-free structure theorem that albeit weaker is still capable of breaking down permutation groups into small and alternating pieces would be a good candidate for being the backbone of our own result. We find such a candidate in Theorem 3.2.1, which is based on Babai's algorithm for the string isomorphism problem (SIP) both in its CFSG and in its CFSG-free version.

We have discussed the structure of the algorithm both in §1.6 and in §3, but let us recall its salient points. Babai [Bab16a] has produced an algorithm that describes in quasipolynomial time the set $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ of all permutations in G that send the string $\mathbf{x}$ to the string $\mathbf{y}$: this algorithm is dependent on CFSG, in that it uses Cameron as a crossroad to pass from the original problem to a collection of subproblems with a smaller or more structured G. A slightly modified CFSG-free version of the same algorithm has been produced as well, a work started by Babai [Bab16a, §13.1] and concluded by Pyber [Pyb16]: this new version avoids the use of Cameron, but broadly speaking retains the same idea of a "crossroad through structure theorem" using a result by Pyber [Pyb93, Thm. 3.15] and the Split-or-Johnson routine of the original algorithm [Bab16a, §7].

It is possible to take a general permutation group G and make it pass through Babai's algorithm: after all, as observed after Definition 3.3.1, $\mathrm{Iso}_G(\mathbf{x}, \mathbf{y})$ is none other than $G \cap H\sigma$ where H is a product of symmetric groups (one for each distinct letter of $\mathbf{x}$) and σ is any one permutation sending $\mathbf{x}$ to $\mathbf{y}$; this means that we can choose $\mathbf{x}_0$ to be a constant string (equivalently, choose H to be the whole symmetric group) and we can obtain $G = \mathrm{Aut}_G(\mathbf{x}_0) = \mathrm{Iso}_G(\mathbf{x}_0, \mathbf{x}_0)$ as a result. Of course from the standpoint of the string isomorphism algorithm this process is utterly useless, since $\mathbf{x}_0$ is trivial and the algorithm outputs G having been given G as input; nevertheless, the algorithm is still making G pass through the whole process of reducing it into smaller subgroups, identifying alternating factor, etc...: this is exactly what we want, i.e. finding structure inside G, and the modifications by Babai and Pyber allow us to do precisely that without resorting to CFSG.

The key observation is that Babai's algorithm takes only quasipolynomial time in n, which implies that the information that we retrieve about the structure of the group is also simple enough. For example, the number of floors of the structure tree with which we are going to replace the chain in [Hel18, Prop. 4.6] (see Proposition 6.3.1) will be polylogarithmic too.

Remark 6.1.3. In using Babai's algorithm to determine the structure of a group $G \leq \mathrm{Sym}(n)$, i.e. determine $G = \mathrm{Iso}_G(\mathbf{x}_0, \mathbf{x}_0)$ with $\mathbf{x}_0 = \alpha^n$, we always reduce to subproblems that also involve only strings of the form $\mathbf{x}'_0 = \alpha^{n'}$. In fact, the only manipulations of the strings themselves that occur in the algorithm in §3.5 are restrictions $\mathbf{x} \mapsto \mathbf{x}|_\Omega$ and preimages $\mathbf{x} \mapsto \mathbf{x}^{\sigma^{-1}}$, both of which do not change the property of being a constant string. Hence, all subproblems descending from the original problem on G are also problems on some G', and not on a more general coset $G' \cap H'\sigma'$ with $H' \leq \mathrm{Sym}(n')$; in other words, in the language and notation of Theorem 3.2.1, since the first H is $\mathrm{Sym}(n)$, all intermediate H' are $\mathrm{Sym}(n')$ and the final atoms themselves are $\mathrm{Alt}(n')$.

In truth, this does not mean that we never use nonconstant strings in Babai's algorithm, even when starting with $\mathbf{x}_0$ constant: some routines feature auxiliary strings made of different letters, such as the "glauque" letter in [Hel19b, §6.1.2], but they are used only to gather structural information and the actual $\mathbf{x}_0$ does not reduce to them.

There are of course some important disadvantages in adopting this new path towards a reduction like in (1.3.3): they are going to be due mainly to the fact that the subgroups involved in the descent process are not necessarily normal, as they were in the situation where Cameron's theorem was a viable route. We will discuss them later in more depth; for now we limit ourselves to observe that the fact that Theorem 6.3.6 has a weaker final bound and depends on Conjecture 6.3.4 is exactly what we have to pay for this weakening of the intermediate result.

6.2 Tools

Let us start with a couple of easy lemmas, describing the structure of $\mathrm{Alt}(n)$.

Lemma 6.2.1. *The group* $\mathrm{Alt}(n)$ *is generated by the set of* 3*-cycles.*

Proof. This is elementary. Any element of $\mathrm{Alt}(n)$ is the product of an even number of transpositions τ_i, or equivalently a product of $\tau_{2i-1}\tau_{2i} = (a\,b)(c\,d)$: if $b = c$ then $\tau_{2i-1}\tau_{2i} = (a\,d\,b)$ is already a 3-cycle, and if $b \neq c$ then $\tau_{2i-1}\tau_{2i} = (a\,c\,b)(b\,d\,c)$ is the product of two 3-cycles. □

Lemma 6.2.2. *For any* $n \geq 5$*, any proper subgroup of* $\mathrm{Alt}(n)$ *has index* $\geq n$*.*

Proof. This is a standard result that uses the fact that $\mathrm{Alt}(n)$ is simple for all $n \geq 5$ (see for instance [DF03, §4.6, Ex. 1]). A whole classification of maximal permutation subgroups exists, the O'Nan-Scott theorem [Sco80] already mentioned in §1.2, but we do not need such a powerful tool here.

Let $G < \mathrm{Alt}(n)$: in particular, $\mathrm{Alt}(n)$ acts by permuting the cosets of G (left cosets, say), so that there is a natural group homomorphism

$$\varphi : \mathrm{Alt}(n) \to \mathrm{Sym}([\mathrm{Alt}(n) : G]).$$

Since $\mathrm{Alt}(n)$ is simple, the normal subgroup $\ker(\varphi)$ is either $\{e\}$ or $\mathrm{Alt}(n)$; however, there exists an element $\sigma \in \mathrm{Alt}(n) \setminus G$, and then σ induces a nontrivial partition of the cosets of G, so that $\ker(\varphi) \neq \mathrm{Alt}(n)$. Hence, φ is injective, and since we have $(n-1)! = \frac{1}{n-1} n! < \frac{1}{2} n!$ we can conclude that $[\mathrm{Alt}(n) : G] \geq n$. □

Thanks to the previous lemmas, we can show the following result, which will prevent the arising of large alternating factors when G itself is not giant (i.e. not equal to $\mathrm{Sym}(n)$ or $\mathrm{Alt}(n)$, see Definition 3.1.1). We also recall the notations $G_A, G_{(A)}$ for setwise and pointwise stabilizers respectively, introduced in §3.3, and $G|_A$ for the restriction of the group G to A (when it is possible to do so, namely when G already stabilizes A), introduced in Lemma 3.3.6(d).

Proposition 6.2.3. *Let $G \leq \mathrm{Sym}(n)$ be a transitive permutation subgroup, with $n \geq 5$. Consider a set $A \subseteq [n]$ with $|A| = \alpha n$ for some $\frac{2}{3} \leq \alpha < 1$, and let $H \leq G_A$. Suppose that $H|_A = \mathrm{Alt}(A)$. Then $G \geq \mathrm{Alt}(n)$.*

This is the kind of proposition that likely can be proved in several different fashions. If we were allowed to use CFSG for example, we could argue that G must be not only transitive but primitive, because an alternating group inside of it permuting more than half of the vertices prevents the formation of a nontrivial block system, and then we could use Cameron's theorem to exclude the possibility of G not being a giant given that by hypothesis $|G| \geq \frac{1}{2}(\alpha n)!$. For our purposes, however, we will need to provide a proof that does not rely on CFSG.

We remark that there is no particular reason to use $\frac{2}{3}$ as a lower bound for α: as one can readily check, we can prove the same for any constant arbitrarily close to $\frac{1}{2}$, as long as we choose n to be large enough.

Proof. By hypothesis we have that $H|_A = \mathrm{Alt}(A)$; the main idea is to prove that $H_{(\bar{A})}|_A = \mathrm{Alt}(A)$ as well, where $\bar{A} = [n] \setminus A$.

Consider an arbitrary $x \in \bar{A}$. By the isomorphism theorems we first have that $[\mathrm{Alt}(A) : H_x|_A] \leq [H : H_x]$, and in turn we also have that $[H : H_x] = [H|_{\bar{A}} : H_x|_{\bar{A}}]$ following the same reasoning and using moreover the fact that H_x contains the kernel of the restriction map to $\bar{A}$. The subgroup $H_x|_{\bar{A}}$ cannot have more than $|\bar{A}|$ cosets inside $H|_{\bar{A}}$ (one can see this as an instance of the orbit-stabilizer theorem); hence

$$[\mathrm{Alt}(A) : H_x|_A] \leq (1-\alpha)n < \alpha n = |A|, \tag{6.2.1}$$

and by Lemma 6.2.2 we must have $H_x|_A = \mathrm{Alt}(A)$. Now we can redefine H to be H_x acting on $n \setminus \{x\}$, and we can repeat the whole process with a new x': notice that α increases, so that the second inequality inside (6.2.1) is still valid. Iterating the process for all points of the original $\bar{A}$, we obtain in the end $H_{(\bar{A})}|_A = \mathrm{Alt}(A)$.

At this point it is easy to conclude. In fact, H (and therefore G) contains all the 3-cycles $(a\,b\,c)$ formed by elements $a, b, c \in A$, so we just have to use them

to get all the 3-cycles in $[n]$ and we could conclude by Lemma 6.2.1. Take any $x \in \bar{A}$: since G is transitive there exists a $g \in G$ that sends x to a given element $y \in A$, and since $\alpha \geq \frac{2}{3}$ and $n \geq 5$ there exist two elements $r, s \in A \setminus g(\bar{A})$; then $g(r\,s\,y)g^{-1}$ is the 3-cycle $(g^{-1}(r)\,g^{-1}(s)\,x)$, which contains two elements of A and one element of $\bar{A}$. Using

$$\begin{aligned}(a\,b\,c) &= (b\,c\,a) = (c\,a\,b),\\ (a\,c\,b) &= (a\,b\,c)^2,\\ (b\,c\,x) &= (a\,c\,b)(a\,b\,x)(c\,a\,b),\end{aligned}$$

we can then reorder elements and insert elements from other cycles as we please, and get all the 3-cycles of $[n]$. □

For any two groups $H \leq G$, let us denote by $\mathcal{L}(G, H)$ the set of left cosets of H inside G: to prevent confusion we would rather avoid using the notation G/H for such a set, unless we are dealing with a normal subgroup H and G/H is the quotient group[1]. We are going to work with a class of Schreier graphs (see Definition 2.1.1) arising from the action on the cosets of a subgroup; incidentally, this was the context in which Schreier graphs were originally conceived [Sch27].

Definition 6.2.4. *Let G be a group and let $H \leq G$. We define* $\mathrm{diam}(G, H)$*, the diameter of the pair (G, H), to be the maximum among the (undirected) diameters of all the Schreier graphs* $\mathrm{Sch}(\mathcal{L}(G, H), S)$*, where S runs through all sets of generators of G and the action $\eta : G \times \mathcal{L}(G, H) \to \mathcal{L}(G, H)$ defining the graphs is the left multiplication $\eta(g, g'H) = gg'H$.*

The diameter of a group $\mathrm{diam}(G)$ is then the same as $\mathrm{diam}(G, \{e\})$, and if H is normal in G then $\mathrm{diam}(G/H) = \mathrm{diam}(G, H)$. Of course, there is nothing special about our choice of "left": we could as well define $\mathcal{R}(G, H)$, and act on it through right multiplication.

We use here Schreier's lemma (in a slightly altered form than in Lemma 5.2.1) so as to be able to use a chain of subgroups as a way to bound diameters. The use we make of it is identical to what happens with [Hel18, Lemma 4.7].

Lemma 6.2.5. *Let G be a finite group, let $H \leq G$ be a proper nontrivial subgroup, and let S be a set of generators of G with $e \in S = S^{-1}$. Then*

$$\begin{aligned}\mathrm{diam}(G) &\leq (2\mathrm{diam}(G, H) + 1)\mathrm{diam}(H) + \mathrm{diam}(G, H)\\ &\leq 4\mathrm{diam}(G, H)\mathrm{diam}(H).\end{aligned}$$

Proof. First, if $d = \mathrm{diam}(G, H)$ then $S^{2d+1} \cap H$ generates H: this is almost the statement of Lemma 5.2.1. In that case however the subgroup of G was assumed to be normal, so that we could use the usual definition of diameter for the quotient group, while here we resort to Definition 6.2.4 and the set $\mathcal{L}(G, H)$; apart from that, the proof with general H is identical to the proof of Lemma 5.2.1.

[1]The author is embarrassingly prone to get confused by the notation and assume that H is normal whenever G/H is written on paper. May the reader be indulgent with him.

The result is now easy: if $S^{2d+1} \cap H$ generates H, then $(S^{2d+1})^{\mathrm{diam}(H)} \supseteq H$ and since S^d contains by definition representatives of all the left cosets of H inside G we have $S^d H = G$, thus concluding the proof. □

The condition of H being proper nontrivial is really only needed for the second inequality, since by definition $\mathrm{diam}(\{e\}) = 0$. For ease of notation, we can use the second inequality anyway and conventionally establish that $\mathrm{diam}(\{e\}) = 1$ (which we are going to do).

6.3 Main theorem

Now we begin our path towards the main result (Theorem 6.3.6). First, let us rewrite Theorem 3.2.1 in a form that suits us more.

Proposition 6.3.1. *Let $n \geq 1$ and let $G_0 \leq \mathrm{Sym}(n)$ acting on a set Ω_0 of size n. Then we can build a rooted tree $T(G_0, \Omega_0)$ (oriented away from the root, say) with the following properties:*

(a) the vertices are pairs (G, Ω) and the edges are coloured either "$(\mathcal{C}1)$", "$(\mathcal{C}2)$" or "$(\mathcal{C}3)$";

(b) the root is (G_0, Ω_0) and the leaves are $(\mathrm{Alt}(\Omega_i), \Omega_i)$ for a partition $\{\Omega_i\}_i$ of Ω_0;

(c) for any non-leaf vertex (G, Ω), either:

  (1) there is only one edge departing from it, coloured "$(\mathcal{C}1)$", and its endpoint is (G', Ω) for some $G' \leq G$, or

  (2) there are only edges coloured "$(\mathcal{C}2)$" departing from it, and their endpoints are $(G|_{\Omega_i}, \Omega_i)$ for some nontrivial partition $\{\Omega_i\}_i$ of Ω, or

  (3) there is only one edge departing from it, coloured "$(\mathcal{C}3)$", and its endpoint is (G', Ω) for some $G' \lhd G$ with G/G' isomorphic to an alternating group of degree ≥ 5;

(d) if a vertex (G', Ω) has an incoming edge coloured "$(\mathcal{C}3)$" coming from a vertex (G, Ω), then it has departing edges coloured "$(\mathcal{C}2)$" whose endpoints $(G'|_{\Omega_i}, \Omega_i)$ are such that $|\Omega| \geq m|\Omega_i|$ for all i, where $G/G' \simeq \mathrm{Alt}(m)$;

(e) every index $[G : G']$ coming from an edge $((G, \Omega), (G', \Omega))$ coloured "$(\mathcal{C}1)$" is bounded by $n^{O(\log^5 n)}$, and for any path from the root to a leaf the number of edges coloured "$(\mathcal{C}1)$" lying on the path is bounded by $O(\log^2 n)$;

(f) for any path from the root to a leaf, the product $\prod_i m_i$ of the degrees of the alternating groups coming from all the edges coloured "$(\mathcal{C}3)$" lying on the path and from the final leaf is bounded by n.

If we were to compare the tree above with the chain in Proposition 6.1.2, the two parts C_1, C_2 would correspond here to $(\mathcal{C}1)$ and $(\mathcal{C}3)$ respectively, and the bounds in Proposition 6.1.2(b)-(d)-(e) would correspond to those in Proposition 6.3.1(e)-(f). The normality of the subgroups involved adds the following perk: all composition factors fit into one chain, by making them into direct product of simple groups; instead, in Proposition 6.3.1 we are forced to deal with a tree, with bifurcations labelled $(\mathcal{C}2)$.

Proof. The construction of $T(G_0, \Omega_0)$ comes as we said from the use of Theorem 3.2.1 in the case of $H = \mathrm{Sym}(n)$. Its definition is similar to that, widely used, of a *structure tree* as in [Hel18, §4.1] and a *structure forest* in [LM88, §3] [BS92, §3.4], although it is more refined to suit our needs.

The root is the starting point of the algorithm, i.e. the input made of the group G_0 and the set Ω_0 on which the group acts, while the leaves are the atoms that are reached at the end of the procedure; as we said before in Remark 6.1.3, starting with G_0 makes us reach simple alternating groups instead of the more general possibilities described in $(\mathcal{A})$. The edges leaving a vertex represent the three possibilities $(\mathcal{C}1)$-$(\mathcal{C}2)$-$(\mathcal{C}3)$ in which an expression can break down to smaller expressions as described in the theorem; however, the construction is not exactly like giving to each vertex its smaller expressions as children.

In the case of $(\mathcal{C}1)$, in $T(G_0, \Omega_0)$ we pass from G to a subgroup G' as its only child. By Remark 6.1.3, the group H in this intermediate step is still $\mathrm{Sym}(\Omega)$, so there is no loss of information: we are simply writing $G = \bigcup_i G'\sigma_i$ for a set of representatives $\{\sigma_i\}_i$ of G' in G, so that the various subproblems (the smaller well-formed expressions in the language of $(\mathcal{C}1)$ inside §3.2) are all on the subgroup G'.

In the case of $(\mathcal{C}2)$, following its exact wording we would reduce from (G, Ω) to $(\pi_1(G), \Omega_1)$ and $(\pi_2(G), \Omega_2)$ for a partition $\Omega = \Omega_1 \sqcup \Omega_2$ respected by G: this is because, as in $(\mathcal{C}1)$, $H = \mathrm{Sym}(\Omega)$ by Remark 6.1.3. However, for simplicity we can reduce directly to subdividing Ω into its orbits[2].

The case of $(\mathcal{C}3)$ is as described in the theorem: the only child of G is a G' such that $\langle G' \cup \{\sigma_1, \sigma_2\}\rangle = G$ and the group generated by σ_1, σ_2 is some alternating group; let us prove the stronger claims that are present in our statement. The only time $(\mathcal{C}3)$ emerges in the CFSG-free algorithm of §3.5 is in Proposition 3.5.16(a), where G acts on Ω preserving a system of blocks $\mathcal{B}$ on which it acts as $\mathrm{Alt}(\Gamma)$ acts on $\binom{\Gamma}{k}$ for some Γ, k; in general we have some large set $S_{\mathbf{x}} \subseteq \Gamma$, canonical with respect to the string $\mathbf{x}$, such that for any $\sigma \in \mathrm{Alt}(S_{\mathbf{x}})$ there is an element of $\mathrm{Aut}_G(\mathbf{x})$ inducing σ on $S_{\mathbf{x}}$, and that set would be the origin of our alternating quotient (see Corollary 3.5.12, which traces in more detail the steps we are describing): however for us $\mathbf{x}$ is constant by Remark 6.1.3 and $\mathrm{Aut}_G(\mathbf{x}) = G$, so we can assume $S_{\mathbf{x}} = \Gamma$. Then our G' is the preimage of $\{e\} = \mathrm{Alt}(\Gamma)_{(S_{\mathbf{x}})}$ and our G is the preimage of $\mathrm{Alt}(\Gamma) = \mathrm{Alt}(\Gamma)_{S_{\mathbf{x}}}$ (by definition); hence $G' \lhd G$ and $G/G' \simeq \mathrm{Alt}(\Gamma)$, and since the algorithm passes through $(\mathcal{C}3)$ only under the condition $|\Gamma| = m > 22\log^2 n$ we have also $|\Gamma| \geq 5$.

[2]The order in which we subdivide Ω is relevant only when starting with nonconstant strings in the original algorithm.

To prove (d), observe that from what we just said in the case of ($\mathcal{C}3$) we have that G' stabilizes the blocks of $\mathcal{B}$ and G permutes them as $\mathrm{Alt}(\Gamma)$ permutes $\binom{\Gamma}{k}$: therefore, since G' is intransitive, the next step will be the restriction to the orbits of the action, i.e. ($\mathcal{C}2$), and each new orbit will be of the same size $\binom{m}{k}^{-1}|\Omega| \leq \frac{|\Omega|}{m}$ where $|\Gamma| = m$.

To see (e), let us turn to the proof of Theorem 3.2.1: for each use of ($\mathcal{C}1$), the number of subproblems to which the original problem reduces is bounded as $n^{O(\log^5 n)}$, as stated in Propositions 3.5.15-3.5.16-3.5.17; furthermore the number of subproblems is the same as the index $[G : G']$, since the reduction we are performing each time is as in Proposition 3.5.3. On the other hand, let us examine the four actions we are allowed to do as described in the course of the proof (§3.6): the first two involve at most one instance of use of ($\mathcal{C}1$) followed by a reduction through ($\mathcal{C}2$) from Ω to orbits of size $\leq \frac{2}{3}|\Omega|$; the third involves one instance of ($\mathcal{C}1$) in exchange for a coarser block system in Ω; the fourth involves one ($\mathcal{C}1$) for a reduction of the degree of the smallest symmetric group (that we know of) containing G, from m to $1 + \sqrt{2m}$. The last two actions can happen at most $O(\log n)$ and $O(\log \log n)$ times respectively on the same Ω, and the first two (which shrink Ω by a fraction) can happen at most $O(\log n)$ times on a single path of the tree: thus, at most $O(\log^2 n)$ edges coloured "($\mathcal{C}1$)" can exist on such a path.

Finally, (f) is a consequence of (d): every time we use ($\mathcal{C}3$) with some $\mathrm{Alt}(m)$ associated to it, we are also dividing the orbit size by at least m, so that on a path we must have $\prod_i m_i \leq n$. □

Remark 6.3.2. A language note: when talking informally about the tree, we will figure the root on top and the paths departing from the root to be vertically descending[3]. Thus, expressions like "descending the tree" mean for us "walking along its paths while moving away from the root", and anything "horizontal" is on the contrary something that singles one element out of a path across multiple paths. We also refer to elements (i.e. vertices or edges) preceding, following or being between others, or also being closer or farther away than others: all of them refer to their distance from the root of the tree in the usual graph metric.

As we mentioned in §6.1, this new route going through Proposition 6.3.1 has some important disadvantages, descending from this one fact: the reduction process may involve subgroups with small index that are not necessarily normal.

The first consequence of this is our inability to use [Hel18, Lemma 4.7], i.e. bounding the diameter of G by the product of the diameters of $N, G/N$ (a consequence of Schreier's lemma); on the other hand, the diameter of G/N is trivially bounded by the size of G/N itself, exactly because the small groups are small enough that we do not need anything more clever than that: therefore, we as well do not have any issue in using Schreier's lemma again (Lemma 6.2.5) and get a multiplication by the index $[G : N]$.

[3]We imagine a genealogical tree, with the ancestral root on top, rather than a real-life tree springing from the ground up. If ancient Berbers had conquered the world, maybe writing conventions and botany would have been in agreement today.

The second, and most dire, consequence is the fact that, as we cannot pass to the normal core of our subgroups (which on the contrary was possible in [Hel18, Lemma 4.2]), we cannot treat all orbits at the same time and reduce the subgroup tree to a subgroup chain: in this way we are forced to treat all the groups of the tree at once. The alternating groups can indeed be worked with horizontally quite well, thanks to the results on products of simple groups (Theorem 5.1.1, or [Hel18, Lemma 4.13]). A bound of the form $\prod_i m_i \leq n$ for the set of degrees m_i we need to consider is too strong to be within our reach: by Proposition 6.3.1(f) this holds on a single branch, but it is not sufficient if we are not passing to the normal core; as a consequence, the final bound in Theorem 6.3.6 is not polylogarithmic in $|G|$ as in [HS14], but it is still better than any e^{n^ε}, and more. The problem of treating the small indices horizontally is in that sense the only difficulty that lies in the way of producing a CFSG-free proof of a diameter bound for transitive groups.

Let us first introduce some notions that will define more clearly what we mean when we talk about a horizontal treatment of the tree.

Definition 6.3.3. *Let T be a tree as in Proposition 6.3.1.*

A horizontal cut of the tree is a set C of vertices and edges of T such that for any path from the root to a leaf there is a unique element of C lying on the path. If a horizontal cut is made only of vertices, we call it a horizontal section.

Two distinct horizontal cuts C_1, C_2 are non-crossing if, for every path from the root to a leaf, the vertex or edge of C_1 lying on the path always precedes or coincides with the vertex or edge of C_2 (or vice versa). Two horizontal cuts inside a set S of non-crossing cuts are consecutive if there are no other cuts in S lying between them.

A horizontal cut is a $(\mathcal{C}1)$-cut (respectively $(\mathcal{C}2)$-cut, $(\mathcal{C}3)$-cut) if it is not a horizontal section and all its edges are coloured "$(\mathcal{C}1)$" (respectively "$(\mathcal{C}2)$", "$(\mathcal{C}3)$").

Let us also define precisely what the gap in the argument for small indices is. We do so by formulating the following conjecture (as said after Lemma 6.2.5 we can adopt the convention that $\mathrm{diam}(H') = 1$ when $H' = \{e\}$, for ease of notation).

Conjecture 6.3.4. *Let $G \leq \mathrm{Sym}(n)$ be a transitive permutation subgroup, let $G_1, G_2, \dots, G_k$ be finite groups lying on a horizontal section of the tree built from G as in Proposition 6.3.1, and let G_i' be a subgroup of G_i for each $1 \leq i \leq k$. Let $H \leq G_1 \times \dots \times G_k$, and let $H' = H \cap (G_i' \times \dots \times G_k')$. Then, there are absolute constants $C_1, C_2 > 0$ such that*

$$\mathrm{diam}(H) \leq C_1 k^{C_2} \cdot \max\{[G_i : G_i'] | 1 \leq i \leq k\} \cdot \mathrm{diam}(H').$$

The dependence of the diameter of a group G on the product between $\mathrm{diam}(H)$ and $\mathrm{diam}(G, H)$ (see Lemma 6.2.5) on one hand, and the dependence of the product of diameters of simple groups on the maximum of the diameters of the factors (see Theorem 5.1.1) on the other, are the clear influences in the formulation of the conjecture above. The assumption is strong enough to be compatible with a proof of a diameter bound for transitive permutation subgroups that is as strong as in [HS14]; a result like Theorem 6.3.6, which provides a qualitatively weaker statement, can be proved even with a weaker version of Conjecture 6.3.4.

We remark that the condition that the groups G_i should be part of the same horizontal section inside the tree cannot be completely dropped. One can choose G_i to be the cyclic group generated by a p_i-cycle, where p_i is the i-th prime, G_i' to be the trivial subgroup, and H to be the whole product: in Theorem 5.1.1 we have bounded the diameter of $G_1 \times \ldots \times G_k$ by the product of the primes p_i, and as we said in §5.4 the bound is tight up to constant; since $p_k = (1 + o(1))k \log k$ and $\prod_{i=1}^{k} p_i = e^{(1+o(1))k \log k}$ by the prime number theorem, a bound like the one in Conjecture 6.3.4 for general G_i is false. From another perspective, the conjecture can be seen as limiting the possibilities for groups appearing in horizontal sections across all transitive groups G.

Before we move to the main theorem, where we use Conjecture 6.3.4 for our purposes, let us remark that the conjecture itself is true in the case $k = 1$, by Lemma 6.2.5 and the trivial bounds $\operatorname{diam}(H, H') \leq [H : H'] \leq [G_1 : G_1']$. In fact, for $k = 1$ we can easily prove even more and replace $[G_1 : G_1']$ by the tighter $\operatorname{diam}(G_1, G_1')$, thanks to the following result (which we also need in the course of the proof of the main theorem anyway).

Proposition 6.3.5. *Let G be a finite group and let $H \leq G$; let $G' \leq G$ and $H' = G' \cap H$. Then*

$$\operatorname{diam}(G', H') \leq \operatorname{diam}(G, H).$$

Proof. Let S' be a set of generators of G': we will prove that there is a set $S \supseteq S'$ of generators of G such that the Schreier graph $\operatorname{Sch}(\mathcal{L}(G', H'), S')$ is an induced subgraph of $\operatorname{Sch}(\mathcal{L}(G, H), S)$, so that in particular the diameter of the former is bounded from above by that of the latter.

First of all, we define an appropriate bijection φ between the set $\{g'H | g' \in G'\} \subseteq \mathcal{L}(G, H)$ and $\mathcal{L}(G', H')$, simply by $\varphi(g'H) = g'H'$. The map is well-defined: if $g_1'H = g_2'H$ then ${g_1'}^{-1} g_2' \in G' \cap H = H'$ and $g_1'H' = g_2'H'$; it is surjective because if $xH' \in \mathcal{L}(G', H')$ then in particular $xH' \subseteq G'$, which means that $x \in G'$, and it is injective because if $g_1'H' = g_2'H'$ then ${g_1'}^{-1} g_2' \in H' \leq H$ and $g_1'H = g_1' {g_1'}^{-1} g_2'H = g_2'H$. This bijection has also the property of respecting the edges of the graphs we are working with: for any $s' \in S'$ and any $g', g'' \in G'$, we have $s'(g'H) = g''H$ if and only if $s'(g'H') = g''H'$ (since $g''^{-1} s' g' \in G' \cap H = H'$); this means that the edges of $\operatorname{Sch}(\mathcal{L}(G, H), S)$ corresponding to elements of S' draw exactly the Schreier graph of $\mathcal{L}(G', H')$ on the vertices of the subset $\{g'H | g' \in G'\}$.

We have just to ensure that we can complete S' to a set of generators S of the whole G without introducing any new edges between the vertices of $\{g'H | g' \in G'\}$. That is however easy to do: it is sufficient to take a finite set $\{s_1, s_2, \ldots, s_k\}$ of new elements of G that do not belong to G' ensuring only that at every step $s_i \notin \langle G' \cup \{s_1, \ldots, s_{i-1}\}\rangle$, until we cannot do so anymore. The resulting set $S = S' \cup \{s_1, s_2, \ldots, s_k\}$ generates G, and since $s_i \notin G'$ we have $s_i g', s_i^{-1} g' \notin G'$ as well for any $g' \in G'$, so that an edge that starts from or ends into a vertex $g'H$ must have a coset gH with $g \notin G'$ as its other vertex. □

Now we move to the main theorem.

Theorem 6.3.6. *Assume that we can prove Conjecture 6.3.4 without using CFSG. Let n be large enough. Then, for any transitive permutation subgroup $G \leq \mathrm{Sym}(n)$, we can bound*

$$\mathrm{diam}(G) \leq e^{e^{\frac{1}{\log 2}(\log\log n)^2}}$$

without using CFSG.

As asserted at the beginning of this chapter, the bound above is worse than the ones reached using CFSG, namely (1.3.4) and (6.1.1), but it is a large improvement over the best known bounds that do not use CFSG, which are (1.3.1) for $G = \mathrm{Sym}(n), \mathrm{Alt}(n)$ and

$$\mathrm{diam}(G) \leq e^{4\sqrt{n}\log^2 n}$$

for any G primitive not giant, due to Babai [Bab82, Cor. 1.2]. For comparison, both bounds would correspond to having $\frac{1}{2}\log n + O(\log\log n)$ in the double exponential instead of $\frac{1}{\log 2}(\log\log n)^2$; the bound given in (1.3.5), due to Breuillard and Tointon and applying even to all G non-abelian simple groups, would have $\log n + \log\log n - O(1)$ (all the advantage of a small ε would just contribute to the size of the $O(1)$). On the other hand, the known bounds with CFSG would correspond to $(4 + o(1))\log\log n$, and Babai's conjecture to $\log\log n + O(1)$.

Proof. Let us draw the tree T associated with our G as described in Proposition 6.3.1, and call Ω the set of size n on which G acts. We are going to artificially lengthen it one step further: from every leaf $(\mathrm{Alt}(\Omega_i), \Omega_i)$, if $|\Omega_i| \geq 5$ we add one more $(\mathcal{C}3)$ edge to a new vertex $(\{e\}, \Omega_i)$, otherwise the same edge can be labelled as $(\mathcal{C}1)$; then, $(\mathcal{C}2)$ edges are added to split Ω_i into singletons. Now the leaves of the tree T are all of the form $(\{e\}, \{x\})$, and all properties of Proposition 6.3.1 are still respected.

In order to prove our bound, we are going to start from the root (G, Ω) and descend down the tree one horizontal section at the time, bounding every time the increase in diameter using Lemma 6.2.5, until we end at the leaves. To get the bound we desire, we will have to be careful in choosing how to descend along the various branches: we need to take advantage of the fact that many contemporaneous descents on multiple branches, either by alternating factors or by small factors, cost as much as only one of them by Theorem 5.1.1 and Conjecture 6.3.4 respectively. To do so, we will define appropriate horizontal cuts to work with.

Let us start with the $(\mathcal{C}3)$-cuts. By Proposition 6.3.1(c), all the vertices (G', Ω') are such that $G' \leq G_{\Omega'}|_{\Omega'}$. In the case of a $(\mathcal{C}3)$ edge $((G', \Omega'), (G'', \Omega'))$, the alternating group G'/G'' acts on a system of blocks in Ω' as $\mathrm{Alt}(m)$ acts on some $\binom{m}{k}$, and the blocks themselves are stabilized by G''; if $m \geq \frac{2}{3}n$ then the blocks are of size 1, G'' is the trivial subgroup and $k = 1$: therefore $G' = \mathrm{Alt}(\Omega')$ and, by Proposition 6.2.3, G must be a giant. For the following discussion on the tree of G, we will assume that our $G \leq \mathrm{Sym}(n)$ is transitive but not a giant, so that we are able to assume that every alternating group associated to a $(\mathcal{C}3)$ edge has degree $\leq \frac{2}{3}n$.

We construct a first $(\mathcal{C}3)$-cut C_1 in the following way. We start with the $(\mathcal{C}3)$ edge with the alternating group with the largest degree (or one of them arbitrarily

chosen, if more than one exist), and put it in C_1; then we discard all edges lying on any path from the root to a leaf passing from the edge we have chosen (in other words, all ancestors and descendants), we choose again the $(\mathcal{C}3)$ edge with the largest degree among all the remaining ones and we put it in C_1. We discard the edges lying on a path passing through the second edge we have chosen, and repeat the process until all the edges we have left (if any) are either $(\mathcal{C}1)$ or $(\mathcal{C}2)$: at this point, we arbitrarily choose vertices on the remaining paths one by one and put them in C_1, discarding every time all the edges lying on a path through the vertex we choose, until no edges at all are left. By construction, C_1 is a $(\mathcal{C}3)$-cut.

C_1 divides the tree T into two parts, the one closer to the root (a tree as well) and the one closer to the leaves (a forest); any vertex belonging to C_1 is defined to be in both parts, for the sake of simplicity (it will not matter in what follows, since by construction both the edge that precedes such a vertex and all the edges that follow it cannot be $(\mathcal{C}3)$). We repeat the construction of $(\mathcal{C}3)$-cuts as above, in both parts, and obtain two $(\mathcal{C}3)$-cuts C_2, C_3. Then we repeat the same construction on the four parts in which we have divided the original tree, and do so r times ($r \geq 1$ to be set later) obtaining in the end $(\mathcal{C}3)$-cuts $C_1, C_2, \ldots, C_{2^r-1}$: we call these the *thick cuts*.

If there are still some $(\mathcal{C}3)$ edges in T that have not been put in any thick cut, we will construct other $(\mathcal{C}3)$-cuts, which we call the *thin cuts*. For any of the 2^r parts in which T is divided by the thick cuts, we do the following: we take an arbitrary path from a root to a leaf (where the roots are now, quite naturally, the vertices that were the closest to the original root in T), choose the first $(\mathcal{C}3)$ edge we find and put it in the first $(\mathcal{C}3)$-cut (or we choose an arbitrary vertex, if no such edge exists), discard all the paths passing through our choice, take a second path and repeat until all the paths have been considered or discarded; after creating the first such cut (call it C), we discard completely its edges and all edges that precede C in the part of T we are examining, and start again as before with the construction of a second $(\mathcal{C}3)$-cut C'. We discard anything that precedes or belongs to C', and repeat until no $(\mathcal{C}3)$ edge is left in this part of T.

In this way, we have created a set of $(\mathcal{C}3)$-cuts, thick and thin, such that every $(\mathcal{C}3)$ edge sits in exactly one of them and such that any two cuts are non-crossing. More interestingly, if C is one of these cuts and $m(C)$ is the maximal degree among the alternating groups of all the $(\mathcal{C}3)$ edges of C, we can give bounds on $m(C)$ that will be useful to us.

By what we said before, we already have $m(C_1) \leq \frac{2}{3}n$; for the other cuts we can do better than that. Consider any $(\mathcal{C}3)$ edge $e_2 \in C_2$; by construction, there must be a path passing through it that contains a $(\mathcal{C}3)$ edge with a degree at least as large as the one of e_2, and this would be the unique edge e_1 belonging to both C_1 and that path: if all paths through e_2 intersected C_1 either in edges of smaller degree or in vertices, then e_2 itself would have belonged to C_1 in the first place. Hence, Proposition 6.3.1(f) implies that $m(C_2) \leq \sqrt{n}$ (and $m(C_3)$ as well).

For any $(\mathcal{C}3)$ edge $e_4 \in C_4$, by construction there must be a path with *two* edges with degrees at least as large as its degree. As before, there is a path with the degree of the edge e_1 lying in C_1 at least as large, and there is a path (built in the part of T defined by C_1 in which e_4 lies) with an edge e_2 in C_2 of degree

at least as large (say C_2 is the cut lying in the same part as e_4, otherwise we say the same for C_3); we have however to guarantee that the path is the same through both e_1 and e_2. If e_2 is closer than e_4 to the root of T, then the path through e_4 and e_1 passes through e_2 as well, and we are done; if e_4 is closer than e_2 to the root, then the path that we found for e_2 at the previous reasoning for C_2 (which has an edge $e_1' \in C_1$ of degree at least as large as e_2) passes through e_4 as well, and we are done again. Hence, Proposition 6.3.1(f) implies that $m(C_4) \leq \sqrt[3]{n}$ (and $m(C_5), m(C_6), m(C_7)$ as well).

We can work analogously by induction for all the thick C_i; start with C_{2^j}, and say that at every step $j' < j$ we have that $C_{2^{j'}}$ is the cut lying in the same part of T as C_{2^j} with respect to the subdivision of T yielded by the set of all the C_i with $i < 2^{j'}$ (we can rename C_i for $2^{j'} \leq i < 2^{j'+1}$ as we please, so there is no loss of generality here). If j' is maximal with respect to the property of having an edge $e_{2^{j'}}$ farther than e_{2^j} from the root, we take the path found for $e_{2^{j'}}$ in the case j' (which takes care of all edges for $j'' \leq j'$), and then all the j'' with $j' < j'' < j$ have edges lying on that same path as well, just by being closer than e_{2^j} to the root; as before, all these cuts really pass through edges and not vertices, or else e_{2^j} would have belonged instead to one of the $C_{2^{j'}}$ with $j' < j$.

Moreover, we repeat the same reasoning for any thin cut, treating it as if it were a thick cut at the step $r+1$ (disregarding all the other thin cuts). Therefore, we have in the end the following bounds:

$$\begin{aligned} m(C_1) &\leq \frac{2}{3}n, \\ m(C_i) &\leq n^{\frac{1}{j+1}} && \text{for all } 1 \leq j < r, 2^j \leq i < 2^{j+1}, \\ m(C) &\leq n^{\frac{1}{r+1}} && \text{for all } C \text{ thin.} \end{aligned} \tag{6.3.1}$$

Finally, Proposition 6.3.1(f) implies a bound on the number of thin cuts as well. If every path has the $(\mathcal{C}3)$ edges satisfy such a relation, the number of $(\mathcal{C}3)$ edges themselves on the path is bounded by $\frac{\log n}{\log 5}$; in the worst case, every two thick cuts C_i, C_j have at least one path whose $(\mathcal{C}3)$ edges all lie between them, apart from the two $(\mathcal{C}3)$ edges that belong already to C_i, C_j. On the other hand, the number of thin cuts between two thick cuts is by construction the same as the maximal number of $(\mathcal{C}3)$ edges on a single path between them: thus, there are at most $\frac{2^r \log n}{\log 5}$ thin cuts.

After that, we move to the $(\mathcal{C}1)$-cuts. Take any part of T between two consecutive $(\mathcal{C}3)$-cuts: we construct $(\mathcal{C}1)$-cuts on it in the same way as we constructed thin cuts before, i.e. constructing the first $(\mathcal{C}1)$-cut by taking every time the first $(\mathcal{C}1)$ edge and discarding all the paths passing through it, then the second $(\mathcal{C}1)$-cut by doing the same with the edges left out from the first one, and repeating until all $(\mathcal{C}1)$ edges have been taken. We can again bound the number of total $(\mathcal{C}1)$-cuts: by Proposition 6.3.1(e), the number of $(\mathcal{C}1)$ edges on a path is bounded by $O(\log^2 n)$, so that reasoning as before there will be at most $O(2^r \log^3 n)$ $(\mathcal{C}1)$-cuts in the whole tree.

Finally, we move to the $(\mathcal{C}2)$ edges: by how we constructed them, a $(\mathcal{C}2)$ edge cannot be followed by another $(\mathcal{C}2)$ edge, so for every two consecutive cuts among

the ($\mathcal{C}1$)-cuts and ($\mathcal{C}3$)-cuts already defined we simply take the unique ($\mathcal{C}2$)-cut that we are allowed to have between them (if any).

At this point, we have defined on T a set of horizontal cuts that are pairwise non-crossing and such that every edge is contained in a unique cut. What we do now is start from the root and descend the tree, bounding the diameter one cut at a time by some factor. Between any two consecutive horizontal cuts among those we have defined, there is a unique horizontal section: at every step we suppose that we have already bounded $\mathrm{diam}(G)$ by some factor times $\mathrm{diam}(H)$, where H is a subgroup of the product of all the G_i in a given horizontal section, and we prove that we can move to the next section at the cost of a new factor; at the end, we will then bound all the factors we have collected. The base case, obviously, is the section made of the sole root, with the tautological bound $\mathrm{diam}(G) \leq 1 \cdot \mathrm{diam}(G)$.

Say that at the horizontal section $\{(G_i, \Omega_i)\}_{i \in I}$ we have already shown the bound $\mathrm{diam}(G) \leq C \cdot \mathrm{diam}(H)$, for some $H \leq \prod_i G_i$ and some $C > 0$; call $\{(G'_i, \Omega'_i)\}_{i \in I'}$ the next horizontal section, and observe that $|I|, |I'| \leq n$ since the Ω_i and the Ω'_i both form partitions of Ω. If the next horizontal cut is a ($\mathcal{C}1$)-cut, we have $I = I'$ and each G'_i is a subgroup of G_i with $[G_i : G'_i] \leq n^{C_3 \log^5 n}$ for some $C_3 > 0$ by Proposition 6.3.1(e); calling $H' = H \cap \prod_i G'_i$ and using Conjecture 6.3.4,

$$\mathrm{diam}(H) \leq C_1 n^{C_2} \cdot n^{C_3 \log^5 n} \cdot \mathrm{diam}(H'), \tag{6.3.2}$$

so that we have a bound in terms of the next horizontal section with the extra factor $C_1 n^{C_2 + C_3 \log^5 n}$ besides C. If the next cut is a ($\mathcal{C}3$)-cut (call it K), then $I = I'$ again and G_i / G'_i is either an alternating group or the trivial group; we call $H' = H \cap \prod_i G'_i$ as before, and we use Lemma 6.2.5, Proposition 6.3.5 and [Hel18, Lemma 4.13][4] to get

$$\begin{aligned} \mathrm{diam}(H) &\leq 4\mathrm{diam}(H, H')\mathrm{diam}(H') \\ &\leq 4\mathrm{diam}\left(\prod_i (G_i / G'_i)\right) \mathrm{diam}(H') \\ &< 4 \cdot \frac{196}{243} \cdot 5n^4 \mathrm{diam}(\mathrm{Alt}(m(K))) \cdot \mathrm{diam}(H'), \end{aligned} \tag{6.3.3}$$

thus giving a new bound with an extra factor of $17n^4\mathrm{diam}(\mathrm{Alt}(m(K)))$, say. If the next cut is a ($\mathcal{C}2$)-cut, then for every (G_i, Ω_i) there is a subset $\{(G'_{ij}, \Omega'_{ij})\}_{j \in J(i)}$ of the next section with $G_i|_{\Omega'_{ij}} = G'_{ij}$ for all j and $\bigcup_j \Omega'_{ij} = \Omega$: in that case $G_i \leq \prod_j G'_{ij}$ in the obvious way, and we only need to reembed H appropriately so as to make it into a subgroup of $\prod_{i,j} G'_{ij}$; we have passed to the next horizontal section without changing the bound, since H and its diameter have remained the same.

Combining (6.3.2) and (6.3.3) with the bound on the number of thick, thin and ($\mathcal{C}1$)-cuts, and recalling that all the leaves are trivial (so that the subgroup

[4] This is the version of Theorem 5.1.1 holding for alternating groups only. We cannot use Theorem 5.1.1 itself because its case-by-case subdivision and Ore's conjecture depend on CFSG, while Helfgott's version for the alternating group uses the much older result in [Mil99]. The proof of Theorem 5.1.1 would work the same way though, so we can insert its constants inside (6.3.3) as well.

H at the last step must be $\{e\}$, and $\mathrm{diam}(H) = 1$ by our notational convention of Lemma 6.2.5 and Conjecture 6.3.4), we obtain

$$\begin{aligned}\mathrm{diam}(G) &\leq (C_1 n^{C_2 + C_3 \log^5 n})^{O(2^r \log^3 n)} \cdot (17n^4)^{2^r \log n + 2^r - 1} \cdot \prod_K \mathrm{diam}(\mathrm{Alt}(m(K))) \\ &= n^{C_4 2^r \log^8 n} \cdot \prod_K \mathrm{diam}(\mathrm{Alt}(m(K))) \end{aligned} \tag{6.3.4}$$

for any $G \leq \mathrm{Sym}(n)$ transitive not giant, where C_4 is some absolute constant and the product is on all $(\mathcal{C}3)$-cuts K. This will play the same role as (1.3.3), which is [Hel18, Prop. 4.15]: as said before, the essential weakening is that we lost the stronger bound on the indices of the alternating groups in the product.

From here, we proceed along the lines of [Hel18, §6]: we will not go over the details, except for the calculations that differ from the original route. Assume as inductive hypothesis that we have proved Theorem 6.3.6 for all $n' \leq e^{-\frac{1}{10}} n$ and all $G' \leq \mathrm{Sym}(n')$ transitive. Let $G \leq \mathrm{Sym}(n)$ transitive: if G is not a giant we have (6.3.4), while if $G = \mathrm{Sym}(n), \mathrm{Alt}(n)$ we have

$$\mathrm{diam}(G) \leq e^{C(\log n)^3 (\log\log n)^2} \mathrm{diam}(G') \tag{6.3.5}$$

for some $C > 0$, where either $G' = \mathrm{Sym}(n'), \mathrm{Alt}(n')$ with $n' \leq e^{-\frac{1}{10}} n$ or $G' \leq \mathrm{Sym}(n)$ is transitive not giant; this is a consequence of Theorem 6.1.1, which does not use CFSG. In the first case we are done by induction, since

$$C(\log n)^3 (\log\log n)^2 + e^{\frac{1}{\log 2}(\log\log n')^2} < e^{\frac{1}{\log 2}(\log\log n)^2}$$

for n large; in the second case, we use (6.3.4) on G' and absorb the factor on the RHS of (6.3.5), so that we obtain the same bound as in (6.3.4) even for G giant, with $C_4 + 1$ instead of C_4 (as long as n is large enough) and where K are $(\mathcal{C}3)$-cuts on the tree of a different transitive group G' (with the same degree, though). Thus, we only have to see whether the bound in (6.3.4) is enough to imply the statement of the theorem.

Recall (6.3.1): we can use the inductive hypothesis on each of the diameters in the product of (6.3.4) since $e^{-\frac{1}{10}} n$ is larger than all the $m(K)$, and therefore we have

$$\begin{aligned}\log \mathrm{diam}(G) \leq (C_4 + 1) 2^r \log^9 n + e^{\frac{1}{\log 2}\left(\log\log\left(\frac{2}{3}n\right)\right)^2} + \sum_{j=1}^{r-1} 2^j e^{\frac{1}{\log 2}\left(\log\log\left(n^{1/(j+1)}\right)\right)^2} \\ + \frac{2^r \log n}{\log 5} e^{\frac{1}{\log 2}\left(\log\log\left(n^{1/(r+1)}\right)\right)^2}.\end{aligned}$$

The largest term on the RHS is the second one, which we can bound from above for n large as

$$\begin{aligned}e^{\frac{1}{\log 2}\left(\log\log\left(\frac{2}{3}n\right)\right)^2} &= e^{\frac{1}{\log 2}\log^2\left(\log n - \log\frac{3}{2}\right)} \leq e^{\frac{1}{\log 2}\left(\log\log n - \frac{\log 3/2}{\log n}\right)^2} \\ &\leq e^{\frac{1}{\log 2}(\log\log n)^2} e^{-\frac{\log 3/2}{\log 2}\frac{\log\log n}{\log n}}\end{aligned}$$

$$\leq e^{\frac{1}{\log 2}(\log\log n)^2} - \frac{\log 3/2}{2\log 2}\frac{\log\log n}{\log n}e^{\frac{1}{\log 2}(\log\log n)^2}.$$

The last term depends on our choice of r. We choose $r = 3$, and for n large we get

$$\begin{aligned}\frac{2^r \log n}{\log 5} e^{\frac{1}{\log 2}\left(\log\log\left(n^{1/(r+1)}\right)\right)^2} &= \frac{8\log n}{\log 5} e^{\frac{1}{\log 2}(\log\log n - \log 4)^2} \\ &\leq \frac{8\log n}{\log 5} e^{\frac{1}{\log 2}(\log\log n)^2} e^{-\frac{\log 4}{\log 2}\log\log n} \\ &\leq \frac{8/\log 5}{\log n} e^{\frac{1}{\log 2}(\log\log n)^2}.\end{aligned}$$

The first term is bounded by a constant times $\log^9 n$. Finally, for the sum we can obtain for n large

$$\sum_{j=1}^{r-1} 2^j e^{\frac{1}{\log 2}\left(\log\log\left(n^{1/(j+1)}\right)\right)^2} \leq 6e^{\frac{1}{\log 2}(\log\log n - \log 2)^2} \leq \frac{6}{\log n} e^{\frac{1}{\log 2}(\log\log n)^2}.$$

Combining all of the bounds, we obtain the result. □

6.4 Concluding remarks

It is easy to see that, as long as no deeper analysis is conducted on what possibilities arise for a tree like the one in Proposition 6.3.1, one cannot even prove Theorem 6.3.6 with a small improvement as putting $(\log\log n)^{2-\varepsilon}$ in the double exponent.

For instance, there could be a permutation subgroup $G \leq \mathrm{Sym}(n)$ and two disjoint sets $\Omega_1, \Omega_2 \subseteq [n]$ with $|\Omega_1| = |\Omega_2| = \frac{n}{10}$ such that in the tree relative to G the two vertices $(G_1, \Omega_1), (G_2, \Omega_2)$ appear with $G_1 = \mathrm{Alt}\left(\frac{n}{10}\right)$ and $G_2 = \mathrm{Alt}\left(\sqrt{\frac{n}{10}}\right) \wr \mathrm{Sym}(2)$. In that situation, after (G_1, Ω_1) there is forcibly a $(\mathcal{C}3)$ edge where the quotient is the whole $\mathrm{Alt}\left(\frac{n}{10}\right)$, while after (G_2, Ω_2) there is a $(\mathcal{C}3)$ edge with quotient $\mathrm{Alt}\left(\sqrt{\frac{n}{10}}\right)$, followed by $(\mathcal{C}2)$ edges and then by other $(\mathcal{C}3)$ edges again with quotients $\mathrm{Alt}\left(\sqrt{\frac{n}{10}}\right)$: if n is large enough, all other routes in the proof of Theorem 3.2.1 in fact cannot occur. In that case, it is not possible to give diameter bounds without having at least to treat, in two separate instances, both $\mathrm{diam}\left(\mathrm{Alt}\left(\frac{n}{10}\right)\right)$ and $\mathrm{diam}\left(\mathrm{Alt}\left(\sqrt{\frac{n}{10}}\right)\right)$; hence, if we assume that we have bounds with $2-\varepsilon$ instead of 2 for those two factors, the recursion process does not work since for any $C_1, C_2 > 0$ we have

$$\begin{aligned}\mathrm{diam}(G) &\geq C_1 e^{C_2(\log\log\frac{n}{10})^{2-\varepsilon}} + C_1 e^{C_2(\log\log\sqrt{\frac{n}{10}})^{2-\varepsilon}} \\ &> C_1 e^{C_2(\log\log n)^{2-\varepsilon}} (e^{-C_2'\frac{(\log\log n)^{1-\varepsilon}}{\log n}} + e^{-C_2''(\log\log n)^{1-\varepsilon}}) \\ &> C_1 e^{C_2(\log\log n)^{2-\varepsilon}} \left(\frac{1}{(\log n)^{\frac{C_2'}{(\log\log n)^{\varepsilon}}}} + 1 - C_2''\frac{(\log\log n)^{1-\varepsilon}}{\log n}\right),\end{aligned}$$

and the last expression in parenthesis is > 1 for any $\varepsilon > 0$, provided that we choose n large enough.

The author, as a matter of fact, believes that such vertices cannot occur in the tree for any G: after all, better bounds that use CFSG exist, at least for $\mathrm{Alt}(n)$. An analysis of which possibilities are excluded from the trees in Proposition 6.3.1 is therefore in line with both proving Conjecture 6.3.4 and improving the overall bound in Theorem 6.3.6. In fact, that would be the most likely route towards proving the conjecture: a first step involving a description of which groups can (or cannot) appear in such trees, and a second step that proves the conjecture only for those ones that may actually show up. One could even weaken one of the two steps, investigating a larger class of groups or showing a weaker bound for them, and Theorem 6.3.6 would still work, albeit with less strong bounds (but not necessarily so: we have already observed that there is margin for weakening the conjecture without affecting the final result). Conjecture 6.3.4 in this sense shows a discrete deal of flexibility, whether the reader deems it to be a virtue or a defect.

Bibliography

[AB09] M. J. Atallah and M. Blanton. *Algorithms and Theory of Computation Handbook: General Concepts and Techniques*. Chapman & Hall/CRC, Boca Raton (USA), second edition, 2009.

[Alo86] N. Alon. Eigenvalues, geometric expanders, sorting in rounds and Ramsey theory. *Combinatorica*, 6(3):207–219, 1986.

[AM85] N. Alon and V. D. Milman. λ_1, isoperimetric inequalities for graphs, and supeconcentrators. *J. Combin. Theory Ser. B*, 38:73–88, 1985.

[AS85] M. Aschbacher and L. Scott. Maximal subgroups of finite groups. *J. Algebra*, 92:44–80, 1985.

[AS04a] M. Aschbacher and S. D. Smith. *The Classification of Quasithin Groups: I. Structure of Strongly Quasithin $\mathcal{K}$-groups*, volume 111 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 2004.

[AS04b] M. Aschbacher and S. D. Smith. *The Classification of Quasithin Groups: II. Main Theorems: The Classification of Simple QTKE-groups*, volume 112 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 2004.

[AS16] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley Publishing, fourth edition, 2016.

[Bab79] L. Babai. Monte-Carlo algorithms in graph isomorphism testing. Technical Report 79–10, Dép. Math. et Stat., Université de Montréal, 1979.

[Bab80] L. Babai. Isomorphism testing and symmetry of graphs. *Ann. Discrete Math.*, 8:101–109, 1980.

[Bab81] L. Babai. On the order of uniprimitive permutation groups. *Ann. of Math. (2)*, 113:553–568, 1981.

[Bab82] L. Babai. On the order of doubly transitive permutation groups. *Invent. Math.*, 65:473–484, 1982.

[Bab83] L. Babai. *Permutation Groups, Coherent Configurations and Graph Isomorphism*. PhD thesis, Hungarian Academy of Sciences, Budapest (Hungary), 1983. In Hungarian.

[Bab06] L. Babai. On the diameter of Eulerian orientations of graphs. In *SODA '06 - Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 822–831, New York (USA), 2006. Association for Computing Machinery (ACM).

[Bab16a] L. Babai. Graph isomorphism in quasipolynomial time. `arXiv:1512.03547v2`, 2016.

[Bab16b] L. Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In D. Wichs and Y. Mansour, editors, *STOC '16 - Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 684–697, Cambridge (USA), 2016. Association for Computing Machinery (ACM).

[Bab19] L. Babai. Canonical form for graphs in quasipolynomial time: preliminary report. In M. Charikar and E. Cohen, editors, *STOC '19 - Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1237–1246, Cambridge (USA), 2019. Association for Computing Machinery (ACM).

[Bau06] B. Baumslag. A simple way of proving the Jordan-Hölder-Schreier theorem. *Amer. Math. Monthly*, 113(10):933–935, 2006.

[BBB+99] A. Blokhuis, S. Ball, A. E. Brouwer, L. Storme, and T. Szőnyi. On the number of slopes of the graph of a function defined on a finite field. *J. Combin. Theory Ser. A*, 86:187–196, 1999.

[BBS04] L. Babai, R. Beals, and A. Seress. On the diameter of the symmetric group: polynomial bounds. In *SODA '04 - Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1108–1112. Society for Industrial and Applied Mathematics, 2004.

[BES80] L. Babai, P. Erdős, and S. M. Selkow. Random graph isomorphism. *SIAM J. Comput.*, 9(3):628–635, 1980.

[BF55] R. Brauer and K. A. Fowler. On groups of even order. *Ann. of Math. (2)*, 62(3):565–583, 1955.

[BG82] L. Babai and C. D. Godsil. On the automorphism groups of almost all Cayley graphs. *European J. Combin.*, 3:9–15, 1982.

[BG08] J. Bourgain and A. Gamburd. Uniform expansion bounds for Cayley graphs of $\mathrm{SL}_2(\mathbb{F}_p)$. *Ann. of Math. (2)*, 167:625–642, 2008.

[BG10] J. Bourgain and A. Gamburd. Strong uniform expansion in $\mathrm{SL}(2, p)$. *Geom. Funct. Anal.*, 20:1201–1209, 2010.

[BGH+14] J. Bamberg, N. Gill, T. P. Hayes, H. A. Helfgott, Á. Seress, and P. Spiga. Bounds on the diameter of Cayley graphs of the symmetric group. *J. Algebraic Combin.*, 40:1–22, 2014.

[BGT10] E. Breuillard, B. Green, and T. Tao. Linear approximate groups. *Electron. Res. Announc. Math. Sci.*, 17:57–67, 2010.

[BGT11] E. Breuillard, B. Green, and T. Tao. Approximate subgroups of linear groups. *Geom. Funct. Anal.*, 21(4):774–819, 2011.

[BGT12] E. Breuillard, B. Green, and T. Tao. The structure of approximate groups. *Publ. math. Inst. Hautes Études Sci.*, 116:115–221, 2012.

[BGT13] E. Breuillard, B. Green, and T. Tao. Small doubling in groups. In L. Lovász, I. Z. Ruzsa, and V. T. Sós, editors, *Erdős Centennial*, pages 129–151. Springer, Berlin (Germany), 2013.

[BH05] L. Babai and T. P. Hayes. Near-independence of permutations and an almost sure polynomial bound on the diameter of the symmetric group. In *SODA '05 - Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1057–1066. Association for Computing Machinery (ACM), 2005.

[Bir34] G. Birkhoff. Transfinite subgroup series. *Bull. Amer. Math. Soc.*, 40(12):847–850, 1934.

[BK79] L. Babai and L. Kučera. Canonical labelling of graphs in linear average time. In *FOCS '79 - Proceedings of the 20th Annual IEEE Symposium on Foundations of Computer Science*, pages 39–46. Institute of Electrical and Electronics Engineers (IEEE), 1979.

[BKL80] L. Babai, P. Klingsberg, and E. M. Luks. Canonical labelling for vertex coloured graphs. Unpublished, 1980.

[BKL83] L. Babai, W. M. Kantor, and E. M. Luks. Computational complexity and the classification of finite simple groups. In *FOCS '83 - Proceedings of the 24th Annual IEEE Symposium on Foundations of Computer Science*, pages 162–171. Institute of Electrical and Electronics Engineers (IEEE), 1983.

[BKT04] J. Bourgain, N. Katz, and T. Tao. A sum-product estimate in finite fields, and applications. *Geom. Funct. Anal.*, 14:27–57, 2004.

[BL83] L. Babai and E. M. Luks. Canonical labeling of graphs. In *STOC '83 - Proceedings of the 15th Annual ACM Symposium on Theory of Computing*, pages 171–183. Association for Computing Machinery (ACM), 1983.

[BLS87] L. Babai, E. M. Luks, and Á. Seress. Permutation groups in NC. In *STOC '87 - Proceedings of the 19th Annual ACM Symposium on Theory of Computing*, pages 409–420. Association for Computing Machinery (ACM), 1987.

[Bra79] R. Brauer. Blocks of characters and structure of finite groups. *Bull. Amer. Math. Soc. (N.S.)*, 1(1):21–38, 1979.

[BRD15] J. Button and C. M. Roney-Dougal. An explicit upper bound for the Helfgott delta in $\mathrm{SL}(2, p)$. *J. Algebra*, 421:493–511, 2015.

[Bre14] E. Breuillard. A brief introduction to approximate groups. In E. Breuillard and H. Oh, editors, *Thin Groups and Superstrong Approximation*, volume 61 of *MSRI Publications*, pages 23–50. Cambridge University Press, New York (USA), 2014.

[BS88] L. Babai and Á. Seress. On the diameter of Cayley graphs of the symmetric group. *J. Combin. Theory Ser. A*, 49(1):175–179, 1988.

[BS92] L. Babai and Á. Seress. On the diameter of permutation groups. *European J. Combin.*, 13(4):231–243, 1992.

[BT16] E. Breuillard and M. C. H. Tointon. Nilprogressions and groups with moderate growth. *Adv. Math.*, 289:1008–1055, 2016.

[BY17] A. Biswas and Y. Yang. A diameter bound for finite simple groups of large rank. *J. Lond. Math. Soc. (2)*, 95(2):455–474, 2017.

[Cam81] P. J. Cameron. Finite permutation groups and finite simple groups. *Bull. Lond. Math. Soc.*, 13:1–22, 1981.

[Cam99] P. J. Cameron. *Permutation Groups*. Cambridge University Press, Cambridge (UK), 1999.

[Cau13] A. L. Cauchy. Recherches sur les nombres. *J. École Polytech.*, 9:99–123, 1813. In French.

[Cay78] A. Cayley. Desiderata and suggestions. No. 2. - The theory of groups: graphical representation. *Amer. J. Math.*, 1(2):174–176, 1878.

[CCN+85] J. H. Conway, R. T. Curtis, S. P. Norton, R. A. Parker, and R. A. Wilson. *ATLAS of Finite Groups: Maximal Subgroups and Ordinary Characters for Simple Groups*. Clarendon Press, Oxford (UK), 1985.

[CFI92] J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.

[Che70] J. Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. In R. C. Gunning, editor, *Problems in Analysis: A Symposium in Honor of Salomon Bochner (PMS-31)*, pages 195–200. Princeton University Press, 1970.

[CLRS01] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, Cambridge (USA), second edition, 2001.

[Col80] M. J. Collins. *Finite Simple Groups II.* Academic Press, London (UK), 1980.

[CPSS83] P. J. Cameron, C. E. Praeger, J. Saxl, and G. M. Seitz. On the Sims conjecture and distance transitive graphs. *Bull. Lond. Math. Soc.*, 15:499–506, 1983.

[CS10] E. Croot and O. Sisask. A probabilistic technique for finding almost-periods of convolutions. *Geom. Funct. Anal.*, 20(6):1367–1396, 2010.

[Dav35] H. Davenport. On the addition of residue classes. *J. London Math. Soc.*, 10:30–32, 1935.

[DF03] D. S. Dummit and R. M. Foote. *Abstract Algebra.* John Wiley & Sons, third edition, 2003.

[Din11] O. Dinai. Growth in SL_2 over finite fields. *J. Group Theory*, 14:273–297, 2011.

[Dix69] J. D. Dixon. The probability of generating the symmetric group. *Math. Z.*, 110:199–205, 1969.

[DM96] J. D. Dixon and B. Mortimer. *Permutation Groups.* Springer-Verlag, New York (USA), 1996.

[Don18] D. Dona. On short expressions for cosets of permutation subgroups. `arXiv:1805.12031`, 2018.

[Don19a] D. Dona. The diameter of products of finite simple groups. `arXiv:1902.06932`, 2019.

[Don19b] D. Dona. Number of directions determined by a set in $\mathbb{F}_q^2$ and growth in $\mathrm{Aff}(\mathbb{F}_q)$. `arXiv:1910.06752`, 2019.

[Don19c] D. Dona. The Weisfeiler-Leman algorithm and the diameter of Schreier graphs. *Groups Geom. Dyn.*, 13(4):1235–1253, 2019.

[DSV16] E. Dobson, P. Spiga, and G. Verret. Cayley graphs on abelian groups. *Combinatorica*, 36(4):371–393, 2016.

[Dvi73] Y. Dvir. Covering properties of permutation groups. In Z. Arad and M. Herzog, editors, *Products of Conjugacy Classes in Groups*, pages 197–221. Springer-Verlag, Berlin (Germany), 1973.

[EJ20] S. Eberhard and U. Jezernik. Babai's conjecture for high-rank classical groups with random generators. `arXiv:2005.09990`, 2020.

[EP99] S. Evdokimov and I. Ponomarenko. On highly closed cellular algebras and highly closed isomorphisms. *Electron. J. Combin.*, 6, 1999. Article no. R18.

[EP09] S. Evdokimov and I. Ponomarenko. Permutation group approach to association schemes. *European J. Combin.*, 30:1456–1476, 2009.

[ES83] P. Erdős and E. Szemerédi. On sums and products of integers. In P. Erdős, L. Alpár, G. Halász, and A. Sárközy, editors, *Studies in Pure Mathematics: To the Memory of Paul Turán*, pages 213–218. Birkhäuser, Basel (Switzerland), 1983.

[FHL80] M. Furst, J. Hopcroft, and E. Luks. Polynomial-time algorithms for permutation groups. In *FOCS '80 - Proceedings of the 21th Annual IEEE Symposium on Foundations of Computer Science*, pages 36–41. Institute of Electrical and Electronics Engineers (IEEE), 1980.

[FJR+98] J. Friedman, A. Joux, Y. Roichman, J. Stern, and J.-P. Tillich. The action of a few permutations on r–tuples is quickly transitive. *Random Structures Algorithms*, 12(4):335–350, 1998.

[Fre73] G. A. Freiman. *Foundations of a Structural Theory of Set Addition*, volume 37 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence (USA), 1973.

[FSS83] M. Fürer, W. Schnyder, and E. Specker. Normal forms for trivalent graphs and graphs of bounded valence. In *STOC '83 - Proceedings of the 15th Annual ACM Symposium on Theory of Computing*, pages 161–170. Association for Computing Machinery (ACM), 1983.

[FST13] S. L. Fancsali, P. Sziklai, and M. Takáts. The number of directions determined by less than q points. *J. Algebraic Combin.*, 37:27–37, 2013.

[FT63] W. Feit and J. G. Thompson. Solvability of groups of odd order. *Pacific J. Math.*, 13:775–1029, 1963.

[GAA+13] G. Gonthier, A. Asperti, J. Avigad, Y. Bertot, C. Cohen, F. Garillot, S. Le Roux, A. Mahboubi, R. O'Connor, S. Ould Biha, I. Pasca, L. Rideau, A. Solovyev, E. Tassi, and L. Théry. A machine-checked proof of the odd order theorem. In *Interactive Theorem Proving*, pages 163–179, Berlin, 2013. Springer.

[Gal46a] É. Galois. Lettre de Galois à M. Auguste Chevalier. *J. Math. Pures Appl. (1)*, 11:408–416, 1846. In French; letter from 1832.

[Gal46b] É. Galois. Mémoire sur les conditions de résolubilité des équations par radicaux. *J. Math. Pures Appl. (1)*, 11:417–433, 1846. In French; manuscript from 1831.

[GH11] N. Gill and H. A. Helfgott. Growth of small generating sets in $\mathrm{SL}_n(\mathbb{Z}/p\mathbb{Z})$. *Int. Math. Res. Not. IMRN*, 2011(18):4226–4251, 2011.

[GH14] N. Gill and H. A. Helfgott. Growth in solvable subgroups of $\mathrm{GL}_r(\mathbb{Z}/p\mathbb{Z})$. *Math. Ann.*, 360:157–208, 2014.

[GHR15] N. Gill, H. A. Helfgott, and M. Rudnev. On growth in an abstract plane. *Proc. Amer. Math. Soc.*, 143(8):3593–3602, 2015.

[GHS+09] A. Gamburd, S. Hoory, M. Shahshahani, A. Shalev, and B. Virág. On the girth of random Cayley graphs. *Random Structures Algorithms*, 35:100–117, 2009.

[GK07] A. A. Glibichuk and S. V. Konyagin. Additive properties of product sets in fields of prime order. In *Additive Combinatorics*, volume 43 of *CRM Proceedings and Lecture Notes*, pages 279–286. American Mathematical Society, 2007.

[GLS94] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups*, volume 40.1 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 1994.

[GLS96] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 2. Part I, Chapter G: General Group Theory*, volume 40.2 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 1996.

[GLS98] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 3. Part I, Chapter A: Almost Simple $\mathcal{K}$-Groups*, volume 40.3 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 1998.

[GLS99] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 4. Part II, Chapters 1-4: Uniqueness Theorems*, volume 40.4 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 1999.

[GLS02] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 5. Part III, Chapters 1-6: The Generic Case, Stages 1-3a*, volume 40.5 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 2002.

[GLS05] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 6. Part IV: The Special Odd Case*, volume 40.6 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 2005.

[GLS18a] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 7. Part III, Chapters 7-11: The Generic Case, Stages 3b and 4a*, volume 40.7 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 2018.

[GLS18b] D. Gorenstein, R. Lyons, and R. Solomon. *The Classification of the Finite Simple Groups, Number 8. Part III, Chapters 12-17: The Generic Case, Completed*, volume 40.8 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence (USA), 2018.

[GMW86] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *FOCS '86 - Proceedings of the 27th Annual IEEE Symposium on Foundations of Computer Science*, pages 174–187. Institute of Electrical and Electronics Engineers (IEEE), 1986.

[GNS18] M. Grohe, D. Neuen, and P. Schweitzer. A faster isomorphism test for graphs of small degree. In *FOCS '18 - Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, pages 89–100. Institute of Electrical and Electronics Engineers (IEEE), 2018.

[God81] C. D. Godsil. GRRs for nonsolvable groups. In *Algebraic methods in graph theory*, volume I, II (Szeged, 1978) of *Colloq. Math. Soc. János Bolyai*, pages 221–239. Amsterdam, 1981.

[Gor82] D. Gorenstein. *Finite Simple Groups: An Introduction to their Classification*. Plenum Press, New York (USA), 1982.

[GR07] B. Green and I. Z. Ruzsa. Freiman's theorem in an arbitrary abelian group. *J. Lond. Math. Soc. (2)*, 75:163–175, 2007.

[Gro17] M. Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, volume 47 of *Lecture Notes in Logic*. Cambridge University Press, Ithaca (USA), 2017.

[Hal20] Z. Halasi. Diameter of Cayley graphs of $\mathrm{SL}(n,p)$ with generating sets containing a transvection. `arXiv:2002.10443`, 2020.

[HBD17] H. A. Helfgott, J. Bajpai, and D. Dona. Graph isomorphisms in quasi-polynomial time. `arXiv:1710.04574`, 2017.

[Hel08] H. A. Helfgott. Growth and generation in $\mathrm{SL}_2(\mathbb{Z}/p\mathbb{Z})$. *Ann. of Math. (2)*, 167:601–623, 2008.

[Hel11] H. A. Helfgott. Growth in $\mathrm{SL}_3(\mathbb{Z}/p\mathbb{Z})$. *J. Eur. Math. Soc. (JEMS)*, 13:761–851, 2011.

[Hel15] H. A. Helfgott. Growth in groups: ideas and perspectives. *Bull. Amer. Math. Soc. (N.S.)*, 52(3):357–413, 2015.

[Hel18] H. A. Helfgott. Growth in linear algebraic groups and permutation groups: towards a unified perspective. `arXiv:1804.03049`, 2018.

[Hel19a] H. A. Helfgott. Growth and expansion in algebraic groups over finite fields. To appear in the proceedings of the Arizona Winter School 2016, `arXiv:1902.06308v3`, 2019.

[Hel19b] H. A. Helfgott. Isomorphismes de graphes en temps quasi-polynomial [d'après Babai et Luks, Weisfeiler-Leman, ...] (Exp. no. 1125). In *Séminaire Bourbaki, Vol. 2016/2017, Exposés 1120-1135*, volume 407 of *Astérisque*, pages 135–182, 2019.

[Hig60] G. Higman. Enumerating p-groups. I: Inequalities. *Proc. Lond. Math. Soc. (3)*, 10(1):24–30, 1960.

[HMPQ19] Z. Halasi, A. Maróti, L. Pyber, and Y. Qiao. An improved diameter bound for finite simple groups of Lie type. *Bull. Lond. Math. Soc.*, 51:645–657, 2019.

[Höl89] O. Hölder. Zurückführung einer beliebigen algebraischen Gleichung auf eine Kette von Gleichungen. *Math. Ann.*, 34:26–56, 1889. In German.

[Höl92] O. Hölder. Die einfachen Gruppen im ersten und zweiten Hundert der Ordnungszahlen. *Math. Ann.*, 40:55–88, 1892. In German.

[Hru12] E. Hrushovski. Stable group theory and approximate subgroups. *J. Amer. Math. Soc.*, 25(1):189–243, 2012.

[HS14] H. A. Helfgott and Á. Seress. On the diameter of permutation groups. *Ann. of Math. (2)*, 179:611–658, 2014.

[HSZ15] H. A. Helfgott, Á. Seress, and A. Zuk. Random generators of the symmetric group: Diameter, mixing time and spectral gap. *J. Algebra*, 421:349–368, 2015.

[HT71] J. Hopcroft and R. Tarjan. A V^2 algorithm for determining isomorphism of planar graphs. *Inform. Process. Lett.*, 1:32–34, 1971.

[Ivi85] A. Ivić. *The Riemann Zeta-function*. John Wiley & Sons, New York (USA), 1985.

[JK81] G. James and A. Kerber. *The representation theory of the symmetric group*. Addison-Wesley, Reading (USA), 1981.

[Jor70] C. Jordan. *Traité des substitutions et des équations algébriques*. Gauthier-Villars, Paris (France), 1870. In French.

[Kas07a] M. Kassabov. Symmetric groups and expander graphs. *Invent. Math.*, 170:327–354, 2007.

[Kas07b] M. Kassabov. Universal lattices and unbounded rank expanders. *Invent. Math.*, 170:297–326, 2007.

[KMS84] D. Kornhauser, G. Miller, and P. Spirakis. Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In *FOCS '84 - Proceedings of the 25th Annual IEEE Symposium on Foundations of Computer Science*, pages 241–250, Singer Island (USA), 1984. Institute of Electrical and Electronics Engineers (IEEE).

[Kne53] M. Kneser. Abschätzung der asymptotischen Dichte von Summenmengen. *Math. Z.*, 58:459–484, 1953. In German.

[Knu76] D. E. Knuth. Big Omicron and big Omega and big Theta. *ACM SIGACT News*, 8(2):18–24, 1976.

[KZA17] M. Klin and M. Ziv-Av. A non-Schurian coherent configuration on 14 points exists. *Des. Codes Cryptogr.*, 84:203–221, 2017.

[Lie84] M. W. Liebeck. On minimal degrees and base sizes of primitive permutation subgroups. *Arch. Math. (Basel)*, 43(1):11–15, 1984.

[LL98] R. Lawther and M. W. Liebeck. On the diameter of a Cayley graph of a simple group of Lie type based on a conjugacy class. *J. Combin. Theory Ser. A*, 83:118–137, 1998.

[LM88] E. M. Luks and P. McKenzie. Parallel algorithms for solvable permutation groups. *J. Comput. System Sci.*, 37:39–62, 1988.

[LO81] H. J. Landau and A. M. Odlyzko. Bounds for eigenvalues of certain stochastic matrices. *Linear Algebra Appl.*, 38:5–15, 1981.

[Lor18] O. Lorscheid. $\mathbb{F}_1$ for everyone. *Jahresber. Deutsch. Math.-Verein.*, 120:83–116, 2018.

[LOST10] M. W. Liebeck, E. A. O'Brien, A. Shalev, and P. H. Tiep. The Ore conjecture. *J. Eur. Math. Soc. (JEMS)*, 12:939–1008, 2010.

[LPS88] M. W. Liebeck, C. Praeger, and J. Saxl. On the O'Nan-Scott theorem for finite primitive permutation groups. *J. Aust. Math. Soc. Ser. A*, 44:389–396, 1988.

[LS01] M. W. Liebeck and A. Shalev. Diameters of finite simple groups: sharp bounds and applications. *Ann. of Math. (2)*, 154:383–406, 2001.

[LS19] M. W. Liebeck and A. Shalev. Girth, words and diameter. *Bull. Lond. Math. Soc.*, 51:539–546, 2019.

[Lub94] A. Lubotzky. *Discrete Groups, Expanding Graphs and Invariant Measures*, volume 125 of *Progress in Mathematics*. Birkhäuser Verlag, Basel (Switzerland), 1994.

[Luk82] E. M. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *J. Comput. System Sci.*, 25:42–65, 1982.

[Mar02] A. Maróti. On the orders of primitive groups. *J. Algebra*, 258(2):631–640, 2002.

[Mil99] G. A. Miller. On the commutators of a given group. *Bull. Amer. Math. Soc.*, 6(3):105–109, 1899.

[MR96] J.-P. Massias and G. Robin. Bornes effectives pour certaines fonctions concernant les nombres premiers. *J. Théor. Nombres Bordeaux*, 8:215–242, 1996. In French.

[Mur17] B. Murphy. Upper and lower bounds for rich lines in grids. `arXiv:1709.10438`, 2017.

[MW20] B. Murphy and J. Wheeler. Growth in some finite three-dimensional matrix groups. `arXiv:2005.05077`, 2020.

[NP11] N. Nikolov and L. Pyber. Product decompositions of quasirandom groups and a Jordan type theorem. *J. Eur. Math. Soc. (JEMS)*, 13:1063–1077, 2011.

[Ore51] O. Ore. Some remarks on commutators. *Proc. Amer. Math. Soc.*, 2:307–314, 1951.

[Plü70] H. Plünnecke. Eine zahlentheoretische Anwendung der Graphentheorie. *J. Reine Angew. Math.*, 243:171–183, 1970. In German.

[PPSS12] C. E. Praeger, L. Pyber, P. Spiga, and E. Szabó. Graphs with automorphism groups admitting composition factors of bounded rank. *Proc. Amer. Math. Soc.*, 140(7):2307–2318, 2012.

[PS14] L. Pyber and E. Szabó. Growth in linear groups. In E. Breuillard and H. Oh, editors, *Thin Groups and Superstrong Approximation*, volume 61 of *MSRI Publications*, pages 253–268. Cambridge University Press, New York (USA), 2014.

[PS16] L. Pyber and E. Szabó. Growth in finite simple groups of Lie type. *J. Amer. Math. Soc.*, 29(1):95–146, 2016.

[PSV17] P. Potočnik, P. Spiga, and G. Verret. Asymptotic enumeration of vertex-transitive graphs of fixed valency. *J. Combin. Theory Ser. B*, 122:221–240, 2017.

[Pyb93] L. Pyber. On the orders of doubly transitive permutation groups, elementary estimates. *J. Combin. Theory Ser. A*, 62:361–366, 1993.

[Pyb16] L. Pyber. A CFSG-free analysis of Babai's quasipolynomial GI algorithm. `arXiv:1605.08266`, 2016.

[Raz14] A. A. Razborov. A product theorem in free groups. *Ann. of Math. (2)*, 179:405–429, 2014.

[Rob55] H. Robbins. A remark on Stirling's formula. *Amer. Math. Monthly*, 62(1):26–29, 1955.

[RS18] M. Rudnev and I. D. Shkredov. On growth rate in $\mathrm{SL}_2(\mathbb{F}_p)$, the affine group and sum-product type implications. `arXiv:1812.01671`, 2018.

[RS20] M. Rudnev and S. Stevens. An update on the sum-product problem. `arXiv:2005.11145`, 2020.

[RT85] I. Z. Ruzsa and S. Turyányi. A note on additive bases of integers. *Publ. Math. Debrecen*, 32(1-2):101–104, 1985.

[Ruz79] I. Z. Ruzsa. On the cardinality of $A+A$ and $A-A$. In *Combinatorics (Coll. Math. Soc. Bolyai, Keszthely 1976)*, volume 18, pages 933–938. Akadémiai Kiadó, 1979.

[Ruz94] I. Z. Ruzsa. Generalized arithmetical progressions and sumsets. *Acta Math. Hungar.*, 65(4):379–388, 1994.

[Ruz96] I. Z. Ruzsa. Sums of finite sets. In D. V. Chudnovsky, G. V. Chudnovsky, and M. B. Nathanson, editors, *Number Theory: New York Seminar 1991-1995*, pages 281–293. Springer-Verlag, New York (USA), 1996.

[Ruz99] I. Z. Ruzsa. An analog of Freiman's theorem in groups. *Astérisque*, 258:323–326, 1999.

[San12] T. Sanders. On the Bogolyubov-Ruzsa lemma. *Anal. PDE*, 5(3):627–655, 2012.

[San13] T. Sanders. The structure theory of set addition revisited. *Bull. Amer. Mat. Soc. (N.S.)*, 50(1):93–127, 2013.

[Sch27] O. Schreier. Die Untergruppen der freien Gruppen. *Abh. Math. Semin. Univ. Hambg.*, 5:161–183, 1927. In German.

[Sch88] U. Schöning. Graph isomorphism is in the low hierarchy. *J. Comput. System Sci.*, 37:312–323, 1988.

[Sco80] L. L. Scott. Representations in characteristic p. *Proc. Sympos. Pure Math.*, 37:319–331, 1980.

[Sel65] A. Selberg. On the estimation of Fourier coefficients of modular forms. In *Proceedings of Symposia in Pure Mathematics, Volume VIII*, pages 1–15, Providence (USA), 1965. American Mathematical Society.

[Ser03] Á Seress. *Permutation Group Algorithms*. Cambridge University Press, Cambridge (UK), 2003.

[Shr59] S. S. Shrikhande. The uniqueness of the L_2 association scheme. *Annals of Mathematical Statistics*, 30(3):781–798, 1959.

[Sim67] C. C. Sims. Graphs and finite permutation groups. *Math. Z.*, 95:76–86, 1967.

[Sol01] R. Solomon. A brief history of the classification of the finite simple groups. *Bull. Amer. Math. Soc. (N.S.)*, 38(3):315–352, 2001.

[Spi12] P. Spiga. Two local conditions on the vertex stabiliser of arc-transitive graphs and their effect on the Sylow subgroups. *J. Group Theory*, 15:23–35, 2012.

[SW15] X. Sun and J. Wilmes. Faster canonical forms for primitive coherent configurations: extended abstract. In R. Servedio and R. Rubinfeld, editors, *STOC '15 - Proceedings of the 47th Annual ACM SIGACT Symposium on Theory of Computing*, pages 693–702, Cambridge (USA), 2015. Association for Computing Machinery (ACM).

[SW16] X. Sun and J. Wilmes. Faster canonical forms for primitive coherent configurations. `arXiv:1510.02195v2`, 2016.

[Sző96] T. Szőnyi. On the number of directions determined by a set of points in an affine Galois plane. *J. Combin. Theory Ser. A*, 74:141–146, 1996.

[Sző99] T. Szőnyi. Around Rédei's theorem. *Discrete Math.*, 208/209:557–575, 1999.

[Tan11] Y. S. Tan. On the diameter of Cayley graphs of finite groups. University of Chicago VIGRE REU, 2011.

[Tao08] T. Tao. Product set estimates for non-commutative groups. *Combinatorica*, 28(5):547–594, 2008.

[Tao10] T. Tao. Freiman's theorem for solvable groups. *Contrib. Discrete Math.*, 5(2):137–184, 2010.

[Toi14] M. C. H. Tointon. Freiman's theorem in an arbitrary nilpotent group. *Proc. Lond. Math. Soc. (3)*, 109:318–352, 2014.

[TV06] T. Tao and V. H. Vu. *Additive Combinatorics*, volume 105 of *Cambridge studies in advanced mathematics*. Cambridge University Press, Cambridge (UK), 2006.

[Vin11] L. A. Vinh. The Szemerédi-Trotter type theorem and the sum-product estimate in finite fields. *European J. Combin.*, 32:1177–1181, 2011.

[Wei66] L. Weinberg. A simple and efficient algorithm for determining isomorphism of planar triply connected graphs. *IEEE Trans. Circuit Theory*, 13(2):142–148, 1966.

[Wie34] H. Wielandt. Abschätzungen für den Grad einer Permutationsgruppe von vorgeschriebenem Transitivitätsgrad. *Schriften Math. Seminars Inst. Angew. Math. Univ. Berlin*, 2:151–174, 1934. In German.

[Wil09] R. A. Wilson. *The Finite Simple Groups*, volume 251 of *Graduate Texts in Mathematics*. Springer, London (UK), 2009.

[WL68] B. Weisfeiler and A. Leman. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsiya*, 9:12–16, 1968. In Russian.

[Zem70] V. N. Zemlyachenko. Canonical numbering of trees. *Proc. Seminar on Comb. Anal. at Moscow State. Univ.*, 1970. In Russian.

[Zis89] I. Zisser. The covering numbers of the sporadic simple groups. *Israel J. Math.*, 67(2):217–224, 1989.

[ZKT85] V. N. Zemlyachenko, N. M. Korneenko, and R. I. Tyshkevich. Graph isomorphism problem. *J. Soviet Math.*, 29(4):1426–1481, 1985. Translated from Russian from *Zapiski Nauchnykh Seminarov LOMI* 118:83–158, 1982.

Explicit dissertatio,
lege feliciter.

Curriculum vitae

Address: Mathematisches Institut, Georg-August-Universität Göttingen, Bunsenstraße 3-5, 37073 Göttingen, Germany.
E-mail: daniele.dona@mathematik.uni-goettingen.de

Education

- September 2016 - present: PhD at Georg-August-Universität Göttingen (Göttingen, Germany).

  Supervisors: Harald Andrés Helfgott, Valentin Blomer.

- September 2013 - October 2015: Master's degree under the ALGANT Master programme, edition 2013-2015. First year at Concordia University (Montréal, Canada), second year at Université de Bordeaux (Bordeaux, France).

  Erasmus Mundus Scholarship for the academic years 2013-2014 and 2014-2015.

- September 2010 - July 2013: Bachelor's degree at Università degli Studi di Torino (Torino, Italy).

  INdAM Scholarship for the academic years 2010-2011, 2011-2012 and 2012-2013.

Publications and preprints

- *Number of directions determined by a set in* $\mathbb{F}_q^2$ *and growth in* $\mathrm{Aff}(\mathbb{F}_q)$, arXiv:1910.06752 (submitted).
- *Explicit* L^2 *bounds for the Riemann* ζ *function*, with Harald A. Helfgott and Sebastian Zuniga Alterman, arXiv:1906.01097.
- *The diameter of products of finite simple groups*, arXiv:1902.06932 (submitted).
- *On short expressions for cosets of permutation subgroups*, arXiv:1805.12031 (submitted).
- *Notes on well-distributed minimal sub-BIBDs for* $\lambda = 1$, arXiv:1803.10545 (expository notes).
- *Graph isomorphisms in quasi-polynomial time*, with Harald Andrés Helfgott and Jitendra Bajpai, arXiv:1710.04574 (translation).
- *The Weisfeiler-Leman algorithm and the diameter of Schreier graphs*, Groups Geom. Dyn., 13(4):1235–1253, 2019.

Talks

- October 2019: "Benjamini-Schramm convergence and invariant random subgroups", *Summer School on L^2-Torsion and Symmetric Spaces*, Georg-August-Universität Göttingen, Göttingen, Germany.
- May 2019: "Two applications of Cameron's classification of permutation groups", *2019 AAN Symposium*, Benasque, Spain.

Other conferences, seminars, workshops, schools attended

- March 2020: *Mariaspring retreat for Helfgott's research group*, Georg-August-Universität Göttingen, Mariaspring, Germany.
- November 2018: *Solving problems in Florence*, Université de Paris, Firenze, Italy.
- August 2018: *Mahler measures and special values of L-functions*, Københavns Universitet, København, Denmark.
- July 2018: *Escuela AGRA III*, Academia Nacional de Ciencias, Córdoba, Argentina.
- January 2018: *Workshop Model Theory and Combinatorics*, Institut Henri Poincaré, Paris, France.
- January 2018: *Advanced school: Entropies and Soficity*, Université de Lyon, Lyon, France.
- September 2015: *Summer School on Perfectoid Spaces*, Università degli Studi di Padova, Bressanone, Italy.
- 2015: *Séminaire Théorie des Nombres*, Université de Bordeaux, Bordeaux, France.
- November 2013: *Montréal-Toronto Workshop in Number Theory*, Fields Institute, Toronto, Canada.
- 2013-2014: *Québec-Vermont Number Theory Seminar*, Concordia University and McGill University, Montréal, Canada.

Other

- Organizer of the *Junior Mathematics Colloquium*, Georg-August-Universität Göttingen, October 2018 - January 2019.
- English language: level C2. IELTS certificate: 8.5, December 2015.
- French language: level B1. Alliance Française (Bordeaux Aquitaine), August 2014.