Theory of Computing Systems
https://doi.org/10.1007/s00224-022-10082-z

Equation Satisfiability in Solvable Groups

Paweł Idziak[1] · Piotr Kawałek[1] · Jacek Krzaczkowski[2] · Armin Weiß[3]

Accepted: 15 April 2022

Abstract

The study of the complexity of the equation satisfiability problem in finite groups had been initiated by Goldmann and Russell in (Inf. Comput. **178**(1), 253–262, 2002) where they showed that this problem is in P for nilpotent groups while it is NP-complete for non-solvable groups. Since then, several results have appeared showing that the problem can be solved in polynomial time in certain solvable groups G having a nilpotent normal subgroup H with nilpotent factor G/H. This paper shows that such a normal subgroup must exist in each finite group with equation satisfiability solvable in polynomial time, unless the Exponential Time Hypothesis fails.

Keywords Equations in groups · Solvable groups · Exponential time hypothesis · Fitting length

This article belongs to the Topical Collection: *Special Issue on International Colloquium on Automata, Languages and Programming (ICALP 2020)*
Guest Editors: Artur Czumaj and Anuj Dawar

The research of the first three authors was partially supported by Polish NCN Grant no 2014/14/A/ST6/00138. The fourth author has been funded by DFG project DI 435/7-1.

✉ Armin Weiß
armin.weiss@fmi.uni-stuttgart.de

Paweł Idziak
pawel.idziak@uj.edu.pl

Piotr Kawałek
piotr.kawalek@doctoral.uj.edu.pl

Jacek Krzaczkowski
krzacz@poczta.umcs.lublin.pl

[1] Faculty of Mathematics and Computer Science, Theoretical Computer Science Department, Jagiellonian University, Kraków, Poland

[2] Faculty of Mathematics, Physics and Computer Science, Institute of Computer Science, Maria Curie-Sklodowska University, Lublin, Poland

[3] Institut für Formale Methoden der Informatik (FMI), Universität Stuttgart, Universitätsstr. 38, 70569 Stuttgart, Germany

Published online: 04 July 2022

1 Introduction

The study of equations over algebraic structures has a long history in mathematics. Some of the first explicit decidability results in group theory are due to Makanin [25], who showed that equations over free groups are decidable. Subsequently several other decidability and undecidability results as well as complexity results on equations over infinite groups emerged (see [5, 9, 23, 29] for a random selection). Also the famous 10th Hilbert problem on Diophantine equations, that asks whether an equation of two polynomials over the ring of integers has a solution, was shown to be undecidable [26].

One can treat polynomials over a ring R to be terms over R with some variables already evaluated by elements of R. The same can be done with groups to define polynomials over a group G. Now the problem PolSat(G) takes as input an equation of the form $\mathbf{t}(x_1, \ldots, x_n) = \mathbf{s}(x_1, \ldots, x_n)$ (or equivalently $\mathbf{t}(x_1, \ldots, x_n) = 1$, by replacing $\mathbf{t} = \mathbf{s}$ by $\mathbf{ts}^{-1} = 1$), where $\mathbf{s}(\overline{x})$ and $\mathbf{t}(\overline{x})$ are polynomials over G, and asks whether this equation has a solution in G. Obviously working with terms $\mathbf{t}, \mathbf{s}$ rather than polynomials this problem trivializes by setting all the x_i's to 1. Likewise PolEqv(G) is the problem of deciding whether two polynomials $\mathbf{t}(\overline{x})$, $\mathbf{s}(\overline{x})$ are equal for all evaluations of the variables $\overline{x}$ in G.

While for infinite groups G the problems PolSat(G) and PolEqv(G) may be undecidable, they are solvable in exponential time in finite realms. In fact, PolSat(G) is in NP, whereas PolEqv(G) is in coNP. Actually the hardest possible groups that lead to NP-complete PolSat and coNP-complete PolEqv are all groups that are not solvable [10, 15]. On the other hand it is easy to see that both these problems can be solved in a linear time for all finite abelian groups.

Also in nilpotent groups both PolSat and PolEqv can be solved in polynomial time. While the running time of the first such algorithm for PolSat, due to Goldmann and Russell [10], is bounded by a polynomial of very high degree (as this bound was obtained by a Ramsey-type argument), the first algorithm for PolEqv (due to [3]) is much faster. For polynomials of length ℓ the running time for PolEqv(G) is bounded by $\mathcal{O}\left(\ell^{k+1}\right)$, where $k \leqslant \log |G|$ is the nilpotency class of the group G. Very recently two much faster algorithms for PolSat(G) have been described. One by [7] runs in $\mathcal{O}\left(\ell^{\frac{1}{2}|G|^2 \log |G|}\right)$ steps. The other one, provided in [21], runs even faster for all but finitely many nilpotent groups, i.e. in $\mathcal{O}\left(\ell^{|G|^2+1}\right)$ steps. The very same paper [21] concludes this race by providing randomized algorithms for PolSat and PolEqv working in linear time for all nilpotent groups.

However, the situation for solvable but non-nilpotent groups has been almost completely open. Due to [13] we know that PolSat and PolEqv for the symmetric group S_3 (and some others) can be done in polynomial time. More examples of such solvable but non-nilpotent groups are provided in [8, 12]. Actually already in 2004 Burris and Lawrence [3] conjectured that PolEqv for all solvable groups is in P. In 2011 Horváth renewed this conjecture and extended it to PolSat [11]. Actually these conjectures have been strongly supported also by recent results in [8], where many other examples of solvable non-nilpotent groups are shown to be tractable.

Up to recently, the smallest solvable non-nilpotent group with unknown complexity was the symmetric group S_4. One reason that prevented existing techniques for polynomial time algorithms to work for S_4 is that S_4 does not have a nilpotent normal subgroup with a nilpotent quotient. Somewhat surprisingly, in [18] the first three authors succeeded to show that neither PolSat(S_4) nor PolEqv(S_4) is in P as long as the Existential Time Hypothesis holds. Simultaneously, in [30] the fourth author proved super-polynomial lower bounds on PolSat and PolEqv for a broad class of finite solvable groups—again unless ETH fails. Both the lower bounds in [18] and [30] depended on the so-called Fitting length, which is defined as the length d of the shortest chain

$$1 = G_0 \leqslant G_1 \leqslant \ldots \leqslant G_d = G$$

of normal subgroups G_i of G with all the quotients G_{i+1}/G_i being nilpotent.

Indeed, the lower bounds in [30] apply to all finite solvable groups of Fitting length at least four and to certain groups of Fitting length three. However, this class of groups does not include S_4—although its Fitting length is three.

The present paper extends these results by showing super-polynomial lower bounds for the complexity of PolSat(G) and PolEqv(G)—again depending on the Fitting length. It strongly indicates that the mentioned conjectures by Burris and Lawrence and by Horváth fail by showing the following result.

Theorem 1 *If G is a finite solvable group of Fitting length $d \geqslant 3$, then both* PolSat(G) *and* PolEqv(G) *require at least $2^{\Omega(\log^{d-1} \ell)}$ steps unless ETH fails.*

The paper [2] contains all necessary pieces to provide for PolSat(G) an upper bound of the form $2^{\mathcal{O}(\log^r \ell)}$ with $r \geqslant 1$ depending on G whenever G is a finite solvable group. This upper bound relies on the *AND*-weakness conjecture saying that each CC^0 circuit for the n-input *AND* function has at least 2^{n^δ} gates. Thus, the *AND*-weakness conjecture implies that the lower bounds in Theorem 1 cannot be improved in an essential way.

Finally, we note that allowing to use definable polynomials as additional basic operations to build the input terms $\mathbf{t}, \mathbf{s}$ we may substantially shorten the size of the input. For example with the commutator $[x, y] = x^{-1}y^{-1}xy$ the expression $[\ldots[[x, y_1], y_2], \ldots, y_n]$ has linear size, while when presented in the pure group language it has exponential size. In this new setting PolSat (and PolEqv) have been shown [14, 22] to be NP-complete (or coNP-complete, respectively) for all non-nilpotent groups. Actually our proof of Theorem 1 shows this as well.

Moreover, the paper [17] shows (in a very broad context of an arbitrary algebra) that allowing such definable polynomials can be simulated by circuits over this algebra.

2 Preliminaries

Complexity and the Exponential Time Hypothesis We use standard notation from complexity theory as can be found in any textbook on complexity, e.g. [27].

The Exponential Time Hypothesis (ETH) is the conjecture that there is some $\delta > 0$ such that every algorithm for 3SAT needs time $\Omega(2^{\delta n})$ in the worst case where n is the number of variables of the given 3SAT instance. By the Sparsification Lemma [20, Thm. 1] this is equivalent to the existence of some $\epsilon > 0$ such that every algorithm for 3SAT needs time $\Omega(2^{\epsilon(m+n)})$ in the worst case where m is the number of clauses of the given 3SAT instance (see also [4, Thm. 14.4]). In particular, under ETH there is no algorithm for 3SAT running in time $2^{o(n+m)}$.

Another classical NP-complete problem is C-COLORING for $C \geqslant 3$. Given an undirected graph $\Gamma = (V, E)$ the question is whether there is a valid C-coloring of Γ, i.e. a map $\chi : V \to \{1, \ldots, C\}$ satisfying $\chi(u) \neq \chi(v)$ whenever $\{u, v\} \in E$. Moreover, by [4, Thm. 14.6], 3-COLORING cannot be solved in time $2^{o(|V|+|E|)}$ unless ETH fails. Since 3-COLORING can be reduced to C-COLORING for fixed $C \geqslant 3$ by introducing only a linear number of additional edges and a constant number of vertices, it follows for every $C \geqslant 3$ that also C-COLORING cannot be solved in time $2^{o(|V|+|E|)}$ unless ETH fails.

Groups and Commutators Throughout, we only consider finite groups G. We follow the notation of [28]. For groups G and H we write $H \leqslant G$ if H is a subgroup of G, or $H < G$ if H is a proper subgroup of G. Similarly we write $H \trianglelefteq G$ (or $H \lhd G$) if H is a normal subgroup of G (or a proper normal subgroup). For a subset $X \subseteq G$ we write $\langle X \rangle$ for the subgroup generated by X, and $\langle\langle X \rangle\rangle = \langle x^g \mid x \in X, g \in G \rangle$ for the normal subgroup generated by X.

We write $[x, y] = x^{-1}y^{-1}xy$ for the commutator and $x^y = y^{-1}xy$ for the conjugation. Moreover, we write $[x_1, \ldots, x_n] = [[x_1, \ldots, x_{n-1}], x_n]$ for $n \geqslant 3$.

We will be also using commutator of (normal) subgroups (or even subsets) $X, Y, X_1, \ldots, X_k \subseteq G$ defined by $[X, Y] = \langle [x, y] \mid x \in X, y \in Y \rangle$ and $[X_1, \ldots, X_k] = [[X_1, \ldots, X_{k-1}], X_k]$. Note here that the commutator $[H, K]$ is a normal subgroup of G whenever H and K are. Finally, we put $[x, {}_k\, y] = [x, \underbrace{y, \ldots, y}_{k \text{ times}}]$ and $[X, {}_k\, Y] = [X, \underbrace{Y, \ldots, Y}_{k \text{ times}}]$.

We will also need the concept of a centralizer of a subset X in G, which is defined as $C_G(X) = \{ g \in G \mid [g, h] = 1 \text{ for all } h \in X \}$. If N is a normal subgroup, then $C_G(N)$ is a normal subgroup as well.

Below we collect some basic facts about commutators of elements and subgroups.

(2.1) For $g, x, y, z, x_1, \ldots x_n, y_1, \ldots y_n \in G$ and normal subgroups K_1, K_2, M, N of a group G and we have

(i) $[xy, z] = [x, z]^y [y, z]$ and $[x, yz] = [x, z][x, y]^z$.

(ii) $[K_1, K_2] = [K_2, K_1] \leqslant K_1 \cap K_2$ and $[K_1 K_2, N] = [K_1, N][K_2, N]$.

(iii) If $x \equiv y \mod N$ and $g \in M$, then for all $k \in \mathbb{N}$ we have

$$[x, {}_k\, g] \equiv [y, {}_k\, g] \mod [N, {}_k\, M].$$

(iv) If $g \in M$ and $x_i \equiv y_i \mod N$, then

$$[g, x_1, \ldots, x_n] \equiv [g, y_1, \ldots, y_n] \mod [M, N].$$

(v) For all $f \in C_G(N)$, $g \in G$, $h \in N$ and $k \in \mathbb{N}$ we have

$$[hf, {}_k\, g] = [h, {}_k\, g]\,[f, {}_k\, g]\,.$$

Proof (i) is a straightforward standard calculation (see also [28, 5.1.5]):

$$\begin{aligned}[x, z]^y[y, z] &= y^{-1}(x^{-1}z^{-1}xz)y\ y^{-1}z^{-1}yz \\ &= y^{-1}x^{-1}z^{-1}xyz = (xy)^{-1}z^{-1}(xy)z = [xy, z]\end{aligned}$$

The first part of (ii) is clear from the definition, while the second one follows immediately from (i). To see (iii) and (iv), let $g \in M$, $x, y \in G$ and $h \in N$ with $hx = y$ to see that

$$\begin{aligned}[hx, g] &= [h, g]^x[x, g] \in [N, M][x, g] \qquad \text{and} \\ [g, hx] &= [g, x][g, h]^x \in [g, x][M, N].\end{aligned}$$

Then our statements follow by induction.

Finally, for (v), let $f \in C_G(N) = \{\, g \in G \mid [f, h] = 1 \text{ for all } h \in N \,\}$ and $g \in G$, $h \in N$. Then we have

$$[hf, g] = [h, g]^f[f, g] = [h, g][f, g].$$

Since $C_G(N)$ is a normal subgroup, also $[f, g] \in C_G(N)$ so that we can then induct on k. □

Since G is finite, for all $x, y \in G$, there are $i < j$ such that $[x, {}_i\, y] = \big[x, {}_j\, y\big]$. Writing $k = j - i$, we get $[x, {}_i\, y] = \big[x, {}_{i+k}\, y\big]$ for all sufficiently large i's. For each choice of x and y we might get a different value for k; yet, by taking a common multiple of all the k's, we obtain some $\omega \in \mathbb{N}$ such that for all $x, y \in G$ and all $i \geqslant \omega$ we have $[x, {}_i\, y] = \big[x, {}_{i+\omega}\, y\big]$.

Since for normal subgroups M, N of G we have

$$M \geqslant [M, {}_1\, N] \geqslant [M, {}_2\, N] \geqslant \ldots \geqslant [M, {}_i\, N] \geqslant \big[M, {}_{i+1}\, N\big] \geqslant \ldots,$$

the finiteness of G ensures us that there is some $k_0 \in \mathbb{N}$ such that $\big[M, {}_{k_0}\, N\big] = [M, {}_k\, N]$ for all $k \geqslant k_0$ and all normal subgroups M, N of G. We can assume that $\omega \geqslant k_0$. It is clear that $\omega = |G|!$ is large enough, but typically much smaller values suffice. Thus, we have:

(2.2) For $x, y \in G$, $M, N \trianglelefteq G$ and $i, j \geqslant \omega$ we have

- $[x, {}_i\, y] = \big[x, {}_{i+\omega}\, y\big]$,
- $[M, {}_i\, N] = \big[M, {}_j\, N\big]$.

We fix $\omega = \omega(G)$ throughout. Be aware that it depends on the specific group G.

Nilpotency and Fitting Series The k-th term of the lower central series is $\gamma_k(G) = [G, {}_k\, G]$. The *nilpotent residual* of G is defined as $\bigcap_{k \geqslant 0} \gamma_k(G) = \gamma_\omega(G)$ where ω is as above (i.e. $\gamma_\omega(G) = \gamma_i(G)$ for every $i \geqslant \omega$). Recall that a finite group G is nilpotent if and only if $\gamma_\omega(G) = 1$.

The *Fitting* subgroup $\mathrm{Fit}(G)$ is the union of all nilpotent normal subgroups. Let G be a finite solvable group. It is well-known that $\mathrm{Fit}(G)$ itself is a nilpotent normal

subgroup (see e.g. [16, Satz 4.2]). We will need the following characterization of the Fitting subgroup due to Baer ([1, Satz L'], which is also an immediate consequence of [28, 12.3.8]).

(2.3) $\mathrm{Fit}(G) = \{\, g \in G \mid [h, {}_{\omega}\, g] = 1 \text{ for all } h \in G \,\}$.

Now we define the *upper Fitting series*

$$1 = \mathcal{U}_0(G) \lhd \mathcal{U}_1(G) \lhd \cdots \lhd \mathcal{U}_k(G) = G$$

by $\mathcal{U}_{i+1}(G)/\mathcal{U}_i(G) = \mathrm{Fit}(G/\mathcal{U}_i(G))$. If the group is clear, we simply write $\mathcal{U}_i$ for $\mathcal{U}_i(G)$. The number of factors k is called the *Fitting length* of G (denoted by $\mathrm{FitLen}(G)$).

The following fact can be derived by a straightforward induction from the characterization of $\mathrm{Fit}(G)$ as largest nilpotent normal subgroup.

(2.4) For $H \leqslant G$ and $g \in G$ we have

- $\mathcal{U}_i(H) = \mathcal{U}_i \cap H$, for all i,
- $\mathrm{FitLen}(H) \leqslant i$ if and only if $H \leqslant \mathcal{U}_i$,
- $\mathrm{FitLen}\,\langle\!\langle g \rangle\!\rangle = i$ if and only if $g \in \mathcal{U}_i \setminus \mathcal{U}_{i-1}$.

Example 1 The symmetric group on four elements S_4 has Fitting length 3 with $1 \leqslant C_2 \times C_2 \leqslant A_4 \leqslant S_4$ being the upper (and also lower) Fitting series.

Example 2 If $G_1, \ldots, G_k$ are nilpotent groups, then the Fitting length of the wreath product $G_1 \wr \cdots \wr G_k$ is at most k (for a definition of wreath products, we refer to any standard textbook like [28]). The Fitting length is exactly k if and only if there are primes $p_1, \ldots, p_k$ with $p_i \mid |G_i|$ and $p_i \neq p_{i+1}$ for all i.

More generally, every group of Fitting length k is a divisor (a quotient of a subgroup) of such a wreath product of k nilpotent groups. As we do not rely on these characterizations, we leave the proofs to the reader.

Equations in Groups A *term* (in the language of groups) is a word over an alphabet $\mathcal{X} \cup \mathcal{X}^{-1}$ where $\mathcal{X}$ is a set of variables. A *polynomial* over a group G is a term where some of the variables are replaced by constants—i.e., a word over the alphabet $G \cup \mathcal{X} \cup \mathcal{X}^{-1}$. Since we are dealing with finite groups only, a symbol $X^{-1} \in \mathcal{X}^{-1}$ for $X \in \mathcal{X}$ can be considered as an abbreviation for $X^{|G|-1}$. We write $\mathbf{s}(x_1, \ldots, x_n)$ or short $\mathbf{s}(\overline{x})$ for a polynomial (resp. term) $\mathbf{s}$ with variables from $\{\, x_1, \ldots, x_n \,\}$. There is a natural composition of terms and polynomials: if $\mathbf{r}(x_1, \ldots, x_n), \mathbf{s}_1, \ldots, \mathbf{s}_n$ are polynomials (resp. terms), we write $\mathbf{r}(\mathbf{s}_1, \ldots, \mathbf{s}_n)$ for the polynomials (resp. terms) obtained by substituting each occurrence of a variable x_i by the polynomial (resp. term) $\mathbf{s}_i$.

A tuple $(g_1, \ldots, g_n) \in G^n$ is a *satisfying* assignment for $\mathbf{s}$ if $\mathbf{s}(g_1, \ldots, g_n) = 1$ in G. The problems $\textsc{PolSat}(G)$ and $\textsc{PolEqv}(G)$ are as follows: for both of them the input is a polynomial $\mathbf{s}(x_1, \ldots, x_n)$. For $\textsc{PolSat}(G)$ the question is whether there *exists* a satisfying assignment, for $\textsc{PolEqv}(G)$ the question is whether *all* assignments are satisfying. Note here that these problems have many other names. For example in in [10, 30], PolSat is denoted by EQN-SAT and PolEqv by EQN-ID.

Inducible Subgroups According to [10], we call a subset $S \subseteq G$ *inducible* if $S = \{\, \mathbf{s}(g_1, \ldots, g_n) \mid g_1, \ldots, g_n \in G \,\}$ for some polynomial $\mathbf{s}(x_1, \ldots, x_n)$ of G.

The importance of inducible subgroups lies in the observation that one can restrict variables in equations to inducible subgroups (simply by replacing each variable by the polynomial defining the inducible subgroup). This immediately gives the following lemma.

Lemma 1 ([10, Lemma 8], [14, Lemma 9, 10]) *If H is an inducible subgroup of G, then*

- PolSat(H) *is polynomial time many-one reducible to* PolSat(G),
- PolEqv(H) *is polynomial time many-one reducible to* PolEqv(G).

We will use this lemma to restrict our consideration for an appropriate subgroup of the form $\gamma_k(G)$. We will see that such subgroups are inducible.

3 Proof of Theorem 1

The proof of the theorem is based on coding (by group polynomials) functions that imitate the behaviour of conjunctions. Unfortunately, the lengths of such n-ary conjunction-like group polynomials are not bounded by any polynomial in n and, therefore, they cannot be used to show NP-completeness of PolSat.[1] However, the group polynomials we are going to produce have length bounded by $2^{\mathcal{O}(n^{\frac{1}{d-1}})}$ where $d = \mathrm{FitLen}(G)$. Given such relatively short conjunction-like group polynomials we reduce graph coloring or 3SAT, depending on whether $|G/H| \geqslant 3$ for a carefully chosen large subgroup H of G. In any case such reduction, together with the ETH, would give the lower bound $2^{\Omega(\log^{d-1} \ell)}$ for PolSat(G).

To see how to produce such relatively short conjunction-like polynomials, we start with the upper Fitting series of G

$$1 = \mathcal{U}_0 \lhd \mathcal{U}_1 \lhd \cdots \lhd \mathcal{U}_d = G$$

to go downwards along this series and consecutively carefully choose $h_\alpha \in \mathcal{U}_\alpha \setminus \mathcal{U}_{\alpha-1}$ on each level $\alpha = d, d-1, \ldots, 1$ of this sequence. Then we get two different cosets $\mathcal{U}_{\alpha-1}$ and $h_\alpha \cdot \mathcal{U}_{\alpha-1}$ which are supposed to simulate false and true values, respectively.

The conjunction-like polynomials are based on the terms $\tilde{\mathbf{q}}^{(k)}(z, x_1, \ldots, x_k)$ and $\mathbf{q}^{(k)}(z, x_1, \ldots, x_k, w)$ for $k \geqslant 0$ defined by

$$\begin{aligned}
\tilde{\mathbf{q}}^{(0)}(z) &= z, \\
\tilde{\mathbf{q}}^{(k)}(z, x_1, \ldots, x_k) &= \left[\tilde{\mathbf{q}}^{(k-1)}(z, x_1, \ldots, x_{k-1}),\, \omega\, x_k\right], && \text{for } k \geqslant 1, \text{ and} \\
\mathbf{q}^{(k)}(z, x_1, \ldots, x_k, w) &= \tilde{\mathbf{q}}^{(k+1)}(z, x_1, \ldots, x_k, w), && \text{for } k \geqslant 0.
\end{aligned}$$

[1] In fact, the mentioned *AND*-weakness conjecture prevents the existence of such short – polynomial size – "conjunction-like" expressions. On the other hand, our construction also shows that the strongest version of the *AND*-weakness conjecture–a $2^{\Omega(n)}$ lower bound – does not hold.

Note that our definition of the $\mathbf{q}^{(k)}$'s immediately yields

$$\mathbf{q}^{(k+1)}(z, x_1, \ldots, x_k, w, w) = \mathbf{q}^{(k)}(z, x_1, \ldots, x_k, w) \tag{1}$$

The conjunction-like behaviour of the $\mathbf{q}^{(k)}$'s on the $\mathcal{U}_\alpha$-cosets is precisely described in the following lemma.

Lemma 2 *For any level $1 \leqslant \alpha \leqslant d-1$ and $h_{\alpha+1} \in \mathcal{U}_{\alpha+1} \setminus \mathcal{U}_\alpha$ there is some $h_\alpha \in \mathcal{U}_\alpha \setminus \mathcal{U}_{\alpha-1}$ such that for each $k \in \mathbb{N}$ we have*

$$\mathbf{q}^{(k)}(h_\alpha, x_1, \ldots, x_k, h_{\alpha+1}) \in \begin{cases} h_\alpha \cdot \mathcal{U}_{\alpha-1}, & \text{if } x_i \in h_{\alpha+1} \cdot \mathcal{U}_\alpha \text{ for all } i, \\ \mathcal{U}_{\alpha-1}, & \text{if } x_i \in \mathcal{U}_\alpha \text{ for some } i. \end{cases}$$

Proof In this proof we may, without loss of generality, factor out our group G by $\mathcal{U}_{\alpha-1}$, or equivalently assume that $\alpha = 1$. This means that $\mathcal{U}_\alpha = \mathrm{Fit}(G)$ and so, by Baer's theorem (2.3), there is some $a \in G$ with $\left[a, {}_\omega\, h_{\alpha+1}\right] \neq 1$. Let $\beta \in \mathbb{N}$ be maximal such that $\left[a, {}_\omega\, h_{\alpha+1}\right] \in \gamma_\beta(\mathcal{U}_\alpha) \setminus \{1\}$ for some $a \in G$. Now, we simply put $h_\alpha = \left[a, {}_\omega\, h_{\alpha+1}\right]$, to observe that $h_\alpha = \left[h_\alpha, {}_\omega\, h_{\alpha+1}\right]$ and $\langle\langle h_\alpha \rangle\rangle \leqslant \gamma_\beta(\mathcal{U}_\alpha)$. The last inclusion gives that for all $x_1, \ldots, x_{k+1} \in G$ we have $\mathbf{q}^{(k)}(h_\alpha, x_1, \ldots, x_{k+1}) \in \gamma_\beta(\mathcal{U}_\alpha)$.

Suppose now that one of the x_i's is in $\mathcal{U}_\alpha$. Then $\tilde{\mathbf{q}}^{(i)}(h_\alpha, x_1, \ldots, x_i) = \left[\tilde{\mathbf{q}}^{(i-1)}(h_\alpha, x_1, \ldots, x_{i-1}), {}_\omega\, x_i\right] \in [\mathcal{U}_\alpha, {}_\omega\, \mathcal{U}_\alpha] = \gamma_\omega(\mathcal{U}_\alpha) = \{1\}$. Hence, also $\mathbf{q}^{(k)}(h_\alpha, x_1, \ldots, x_k, h_{\alpha+1}) = 1$.

On the other hand, if all the x_i's are in the coset $h_{\alpha+1}\mathcal{U}_\alpha$, then, by (2.1.iv), we have $\mathbf{q}^{(k)}(h_\alpha, x_1, \ldots, x_k, h_{\alpha+1}) \equiv \mathbf{q}^{(k)}(h_\alpha, h_{\alpha+1}, \ldots, h_{\alpha+1}, h_{\alpha+1}) = h_\alpha$ modulo $[\langle\langle h_\alpha \rangle\rangle, \mathcal{U}_\alpha] \leqslant [\gamma_\beta(\mathcal{U}_\alpha), \mathcal{U}_\alpha] \leqslant \gamma_{\beta+1}(\mathcal{U}_\alpha)$. Hence, $\mathbf{q}^{(k)}(h_\alpha, \overline{x}, h_{\alpha+1}) = h_\alpha f$ for some $f \in \gamma_{\beta+1}(\mathcal{U}_\alpha)$. Thus, all we have to show is that $f \in \mathcal{U}_{\alpha-1}$, or – in our setting – that $f = 1$. To do this we induct on $j \geqslant \beta + 1$ to show that $f \in \gamma_j(\mathcal{U}_\alpha)$ for all j's.

Starting with $f \in \gamma_j(\mathcal{U}_\alpha) \leqslant \gamma_{\beta+1}(\mathcal{U}_\alpha)$, we also have $\left[f, {}_\omega\, h_{\alpha+1}\right] \in \gamma_{\beta+1}(\mathcal{U}_\alpha)$. But now, maximality of β ensures us that

$$\left[f, {}_\omega\, h_{\alpha+1}\right] = 1. \tag{2}$$

Obviously $[f, g] \in \gamma_{j+1}(\mathcal{U}_\alpha)$ whenever $f \in \gamma_j(\mathcal{U}_\alpha)$ and $g \in \mathcal{U}_\alpha$. This simply means that $f \in C_{G/\gamma_{j+1}(\mathcal{U}_\alpha)}(\mathcal{U}_\alpha/\gamma_{j+1}(\mathcal{U}_\alpha))$, and by (2.1.v) we obtain

$$\left[h_\alpha f, {}_\omega\, h_{\alpha+1}\right] \equiv \left[h_\alpha, {}_\omega\, h_{\alpha+1}\right]\left[f, {}_\omega\, h_{\alpha+1}\right] \mod \gamma_{j+1}(\mathcal{U}_\alpha). \tag{3}$$

Summing up we get

$$\begin{aligned} h_\alpha f &= \left[h_\alpha f, {}_\omega\, h_{\alpha+1}\right] && \text{(by (1))} \\ &\equiv \left[h_\alpha, {}_\omega\, h_{\alpha+1}\right]\left[f, {}_\omega\, h_{\alpha+1}\right] \mod \gamma_{j+1}(\mathcal{U}_\alpha) && \text{(by (3))} \\ &= h_\alpha \cdot 1, && \text{(by (2))} \end{aligned}$$

so that $f \in \gamma_{j+1}(\mathcal{U}_\alpha)$.

Going along the j's we arrive to the conclusion that $f \in \gamma_\omega(\mathcal{U}_\alpha) = \{1\}$, as promised. □

Now, picking $h_d \in G \setminus \mathcal{U}_{d-1}$, the consecutive use of Lemma 2 supplies us with elements $h_{d-1}, \ldots, h_1$ that allow us to define conjunction-like polynomials

$\mathbf{q}^{(k)}_{\alpha+1}(x_1,\ldots,x_k) = \mathbf{q}^{(k)}(h_\alpha, x_1,\ldots,x_k,h_{\alpha+1})$. Note here that, since the terms $\mathbf{q}^{(k)}$ use iterated commutators ($\omega\cdot(k+2)$ times), their sizes are exponential in k. However, to get a conjunction on $n = k^{d-1}$ elements we first split these elements into k^{d-2} groups, each having k elements. If there were only two cosets of G of $\mathcal{U}_{d-1}$, then applying to each such k element group the polynomial $\mathbf{q}^{(k)}_d$ everything would be sent into $\mathcal{U}_{d-2} \cup h_{d-1}\cdot\mathcal{U}_{d-2}$. Now, we group the obtained k^{d-2} values into k^{d-3} groups, each of size k and apply $\mathbf{q}^{(k)}_{d-1}$ to each such group. Repeating this procedure we finally arrive into $\mathcal{U}_1$ ensuring that the appropriate composition of the $\mathbf{q}^{(k)}_\alpha$'s returns either the value 1 or h_1. One can easily notice that the size (i.e., length as a word) of such composed polynomial is $2^{\mathcal{O}(k)} = 2^{\mathcal{O}(n^{\frac{1}{d-1}})}$.

Unfortunately, the behaviour of the $\mathbf{q}^{(k)}_d$'s and the entire long compositions can be controlled only on two cosets of $\mathcal{U}_{d-1}$. This requires $|G/\mathcal{U}_{d-1}| = 2$— which very seldom is the case. Thus, the very top level requires a very careful treatment. First, we replace the group G with a smaller subgroup G_0 of the same Fitting length but such that G_0 is abelian over its $\mathcal{U}_{d-1}$. Then we find a normal subgroup $\mathcal{U}_{d-1} \leqslant H \lhd G_0$ so that we will be able to control the behaviour of the $\mathbf{q}^{(k)}$'s on all cosets of H in G_0. The first step towards realizing this idea is described in the next observation.

Lemma 3 *In each finite solvable group G there is a subgroup G_0 satisfying:*

- *G_0 is inducible,*
- *$\mathrm{FitLen}(G_0) = \mathrm{FitLen}(G) = d$, and*
- *$G_0/\mathcal{U}_{d-1}(G_0)$ is abelian.*

Proof We simply set $G_0 = \gamma_m(G)$ where m is maximal with $\gamma_m(G) \not\leqslant \mathcal{U}_{d-1}(G)$. This secures $\mathrm{FitLen}(G_0) = d$. To see that all groups $\gamma_j(G)$ in the lower central series

$$G = \gamma_0(G) \geqslant \gamma_1(G) \geqslant \cdots \geqslant \gamma_\omega(G)$$

are inducible, we induct on j and argue like in [10, Lemma 5]. Let $\gamma_j(G)$ be the image of the polynomial $\mathbf{p}(\overline{x})$. Every element in $\gamma_{j+1}(G) = [\gamma_j(G), G]$ is a product of at most $|G|$ elements of the form $[z, y]$, where z ranges over $\gamma_j(G)$ and y over entire G. Thus, introducing new sequences of pairwise different variables $\overline{x}^1,\ldots,\overline{x}^{|G|}$ we can produce $\gamma_{j+1}(G)$ as the image of the polynomial $\prod_{i=1}^{|G|}[\mathbf{p}(\overline{x}^i), y_i]$.

Finally, $G_0/\mathcal{U}_{d-1}(G_0)$ is abelian as we have $[G_0, G_0] = [\gamma_m(G), \gamma_m(G)] \leqslant [\gamma_m(G), G] = \gamma_{m+1}(G) \leqslant \mathcal{U}_{d-1}$, where the last inclusion is a consequence of the maximality of m. □

From now on we simply change notation and replace our starting group G by G_0, or in other words we assume that $G/\mathcal{U}_{d-1}(G)$ is abelian. Now, to construct (and control) the promised normal subgroup H first we pick $K \trianglelefteqslant G$ among the minimal (with respect to inclusion) normal subgroups satisfying:

- $[K, G] = K$ and
- $\mathrm{FitLen}(K) = d - 1$.

Since $\gamma_\omega(G)$ satisfies both above conditions, such K indeed exists.

(3.1) K is indecomposable, i.e. if $K = K_1K_2$ for some $K_1, K_2 \trianglelefteq G$ then $K = K_1$ or $K = K_2$.

Proof Suppose that (K_1, K_2) is a minimal pair (coordinatewise) with $K = K_1K_2$. Since $K = [K, G] = [K_1K_2, G] = [K_1, G][K_2, G]$ and $[K_i, G] \leqslant K_i$, we immediately get $[K_i, G] = K_i$ for both $i = 1, 2$. Now if $K_i < K$, then minimality of K gives $\mathrm{FitLen}(K_i) \leqslant d - 2$. If this happens for both K_1 and K_2, then $d - 1 = \mathrm{FitLen}(K) = \mathrm{FitLen}(K_1K_2) = \max\{\mathrm{FitLen}(K_1), \mathrm{FitLen}(K_2)\} \leqslant d - 2$, a contradiction. □

By (3.1) we know that there exists the unique $K_0 \trianglelefteq G$ with $K_0 < K$ and such that there is no normal subgroup of G that lies strictly between K_0 and K.

Note that, if $a \in K \setminus K_0$, we cannot have $\langle\langle a \rangle\rangle \leqslant K_0$. This gives

(3.2) For all $a \in K \setminus K_0$ we have $\langle\langle a \rangle\rangle = K$.

The other consequence of the fact that the solvable group G has no normal subgroups strictly between K_0 and K is the following.

(3.3) K/K_0 is abelian.

We will also need:

(3.4) $[K_0, {}_\omega G] \leqslant \mathcal{U}_{d-2}(K)$.

Proof By our choice of ω, we have $[[K_0, {}_\omega G], G] = [K_0, {}_\omega G]$. Since $[K_0, {}_\omega G] \leqslant K_0$ is strictly contained in K and K was chosen to be minimal with $[K, G] = K$ and $\mathrm{FitLen}(K) = d - 1$, we must have $\mathrm{FitLen}([K_0, {}_\omega G]) \leqslant d - 2$. □

Now we are ready to define the normal subgroup H of G. We simply put H to be the centralizer in G of K modulo K_0, i.e the largest normal subgroup with $[H, K] \leqslant K_0$. Then obviously $H = \{ g \in G \mid [K, g] \leqslant K_0 \}$.

(3.5) $\mathcal{U}_{d-1} \leqslant H < G$. In particular, G/H is abelian.

Proof To see that $H < G$ suppose otherwise, i.e. $[K, G] \leqslant K_0$. This, however, contradicts our choice of K to satisfy $[K, G] = K$.

The first inclusion is simply equivalent to $[K, \mathcal{U}_{d-1}] \leqslant K_0$. Indeed, since $\mathrm{FitLen}(K) = d - 1$, we have $\left[K, {}_\omega \mathcal{U}_{d-1}\right] \leqslant \gamma_\omega(\mathcal{U}_{d-1}) \leqslant \mathcal{U}_{d-2}$ and, thus, $[K, \mathcal{U}_{d-1}] < K$. Since we assumed $G/\mathcal{U}_{d-1}$ to be abelian, the second part of the statement follows. □

Directly from our definitions, we know that $[x, y] \in K_0$ whenever $x \in K$ and $y \in H$. But the reason for our careful choice of K and then H was to have a precise control over the behaviour of $[x, y]$ for y in other cosets of H (and x still in K.)

Thus, for any $g \in G$ we define a map $\varphi_g : K \to K/K_0$ by $\varphi_g(x) = [x, g] \cdot K_0$. Since by (3.3) is K/K_0 is abelian, using (2.1.i), one can easily check that φ_g is a group homomorphism for all $g \in G$. Also we have $\varphi_g(K_0) \leqslant K_0$, i.e. the kernel

of this homomorphism contains K_0 so that φ_g actually induces a homomorphism $K/K_0 \to K/K_0$. We also write φ_g for this induced homomorphism.

(3.6) If $g \in G \setminus H$, then $\varphi_g : K/K_0 \to K/K_0$ is an isomorphism.

Proof We start with showing that for $g \in G$

$$\varphi_g(x^b) = \varphi_g(x)^b \tag{4}$$

whenever $x \in K$ and $b \in G$. Indeed, by (3.3), we can write $bg = hgb$ for some $h \in H$. Then we have

$$\begin{aligned}
\varphi_g(x^b) &= [x^b, g] \cdot K_0 \\
&= (x^b)^{-1}\, g^{-1} b^{-1} x b g \cdot K_0 \\
&= (x^b)^{-1}\, b^{-1} g^{-1} h^{-1} x h g b \cdot K_0 \\
&= (x^b)^{-1}\, b^{-1} g^{-1} x g b \cdot K_0 \qquad \text{(since } h \in H\text{)} \\
&= (x^{-1}\, g^{-1} x g)^b \cdot K_0 \\
&= \varphi_g(x)^b.
\end{aligned}$$

To see that the kernel of the original φ_g is K_0, pick $a \in K \setminus K_0$, so that, by (3.2), every element $x \in K$ can be represented as $x = a^{g_1} \cdots a^{g_n}$ for some $g_1, \ldots, g_n \in G$. Now, if $\varphi_g(a) = K_0$, then (4) gives $\varphi_g(x) = K_0$ for all $x \in K$. This would however put g into the centralizer H, contrary to our assumption. □

Note that (4) means that φ_g is not only a group homomorphism but actually a homomorphism of G-modules. Here K/K_0 is a G-module under the action of G on K/K_0 via conjugation. In terms of modules the proof of (3.6) is stated even easier: The kernel of φ_g has to be a submodule of K/K_0. However, by (3.2) K/K_0 is generated, as a G-module, by any of its non-trivial elements.

Remark 1 Notice, that for (3.6), we need G/H to be abelian. Indeed, in general, if N is a minimal (and, thus, indecomposable) normal subgroup with $[N, G] = N$, the map $N \to N$ defined by $x \mapsto [x, g]$ is *not* necessarily bijective for all $g \notin C_G(N)$. For instance take the semidirect product $(C_3 \times C_3) \rtimes D_4$ where $D_4 = \langle a, b \mid a^2 = b^2 = (ab)^4 = 1 \rangle$ is the dihedral group of order 8 and a acts by exchanging the two components of $C_3 \times C_3$ and b by inverting the second one. Then, $N = C_3 \times C_3$ is an indecomposable normal subgroup and $[N, G] = N$ but $a \notin C_G(N)$ and $[(1, 1), a] = [(2, 2), a] = 1$, so $x \mapsto [x, g]$ is not bijective on N (here we use an additive notation for $C_3 = \{0, 1, 2\}$).

We summarize our observations in the following claim.

(3.7) For all $x \in K$ we have

$$\tilde{\mathbf{q}}^{(1)}(x, y) \in \begin{cases} xK_0, & \text{if } y \notin H, \\ K_0, & \text{if } y \in H. \end{cases}$$

Proof Note first that ω was chosen to satisfy $[x, {}_{\omega}\, y] = [x, {}_{2\omega}\, y]$. Moreover, for a fixed $g \in G$ the unary polynomial $\tilde{\mathbf{q}}^{(1)}(x, g)$ acts on K as the composition φ_g^{ω} of φ_g with itself ω times. Now, if $g \notin H$, then (3.6) yields that φ_g^{ω} is the identity on the quotient K/K_0. Moreover, φ_g^{ω} is constant K_0 for $g \in H$. □

With claim (3.7) we are ready to construct polynomials that will allow us to code coloring or 3SAT at the very top level.

Lemma 4 *There is $h \in K \setminus \mathcal{U}_{d-2}$ and families of polynomials*

$$\mathbf{r}^{(k)}(y_1, \ldots, y_k) \qquad \text{and}$$
$$\mathbf{s}^{(k)}(y_{1,1}, y_{1,2}, y_{1,3} \ldots, y_{k,1}, y_{k,2}, y_{k,3})$$

of length $2^{\mathcal{O}(k)}$ such that

$$\mathbf{r}^{(k)}(\overline{y}) \in \begin{cases} h \cdot \mathcal{U}_{d-2}, & \text{if } y_i \notin H \text{ for all } i, \\ \mathcal{U}_{d-2}, & \text{if } y_i \in H \text{ for some } i, \end{cases} \tag{5}$$

and

$$\mathbf{s}^{(k)}(\overline{y}) \in \begin{cases} h \cdot \mathcal{U}_{d-2}, & \text{if for all } i \text{ there is some } j \text{ with } y_{i,j} \in H, \\ \mathcal{U}_{d-2}, & \text{if } y_{i,1}, y_{i,2}, y_{i,3} \notin H \text{for some } i. \end{cases} \tag{6}$$

Proof First, we use (3.7) and induct on k in order to see that for all $a \in K \setminus K_0$ we have

$$\tilde{\mathbf{q}}^{(k)}(a, y_1, \ldots, y_k) \in \begin{cases} aK_0, & \text{if } y_i \notin H \text{ for all } i, \\ K_0, & \text{if } y_i \in H \text{ for some } i. \end{cases}$$

Now we fix some arbitrary $a \in K \setminus K_0$ and $g \in G \setminus H$. Then obviously also $h = [a, {}_{\omega}\, g] \in K \setminus K_0$. Actually $h \notin \mathcal{U}_{d-2}$, as otherwise $h \in \mathcal{U}_{d-2} \cap K \leqslant K_0$.

Now, by (3.4) we know that $M := [K_0, {}_{\omega}\, G] \leqslant \mathcal{U}_{d-2}(K)$. By (2.1.iii) it follows that

$$\mathbf{q}^{(k)}(a, y_1, \ldots, y_k, g) \in \begin{cases} hM, & \text{if } y_i \notin H \text{ for all } i, \\ M, & \text{if } y_i \in H \text{ for some } i. \end{cases}$$

Thus, $\mathbf{r}^{(k)}(y_1, \ldots, y_k) = \mathbf{q}^{(k)}(a, y_1, \ldots, y_k, g)$ satisfies (5). Clearly, its length is in $2^{\mathcal{O}(k)}$.

To construct the polynomials $\mathbf{s}^{(k)}$, we first define

$$\mathbf{p}(x, y_1, y_2, y_3) = x \cdot \tilde{\mathbf{q}}^{(3)}(x, y_1, y_2, y_3)^{-1}.$$

Then for all $x \in K$, by (3.7), we have

$$\mathbf{p}(x, y_1, y_2, y_3) \in \begin{cases} K_0, & \text{if } y_j \notin H \text{ for all } j, \\ xK_0, & \text{if } y_j \in H \text{ for some } j. \end{cases}$$

Now, with a, g, h and M as above, we proceed as with the $\mathbf{r}^{(k)}$'s to define

$$\tilde{\mathbf{s}}^{(k)}(\overline{y}) = \mathbf{p}(\cdots \mathbf{p}(a, y_{1,1}, y_{1,2}, y_{1,3}), \ldots, y_{k,1}, y_{k,2}, y_{k,3})$$

and

$$\mathbf{s}^{(k)}(\overline{y}) = \left[\tilde{\mathbf{s}}^{(k)}(\overline{y}), {}_{\omega}\, g\right].$$

As previously, (6) follows from (2.1.iii). □

Our next claim summarizes Lemma 2 and 4.

Lemma 5 *For $1 \leqslant \alpha \leqslant d-1$ there are elements $h_\alpha \neq 1$ and families of polynomials*

$$\mathbf{r}_\alpha^{(m)}(y_1, \ldots, y_m) \qquad \text{and}$$
$$\mathbf{s}_\alpha^{(m)}(y_{1,1}, y_{1,2}, y_{1,3} \ldots, y_{m,1}, y_{m,2}, y_{m,3})$$

of length $2^{\mathcal{O}(m^{\frac{1}{d-\alpha}})}$ such that

$$\mathbf{r}_\alpha^{(m)}(\overline{y}) \in \begin{cases} h_\alpha \cdot \mathcal{U}_{\alpha-1}, & \text{if } y_i \notin H \text{ for all } i, \\ \mathcal{U}_{\alpha-1}, & \text{if } y_i \in H \text{ for some } i, \end{cases}$$

and

$$\mathbf{s}_\alpha^{(m)}(\overline{y}) \in \begin{cases} h_\alpha \cdot \mathcal{U}_{\alpha-1}, & \text{if for all } i \text{ there is some } j \text{ with } y_{i,j} \in H, \\ \mathcal{U}_{\alpha-1}, & \text{if } y_{i,1}, y_{i,2}, y_{i,3} \notin H \text{ for some } i. \end{cases}$$

Proof We induct downwards on $\alpha = d-1, \ldots, 2, 1$. To start with we refer to 4 to set $h_{d-1} = h$ while $\mathbf{r}_{d-1}^{(m)}(\overline{y}) = \mathbf{r}^{(m)}(\overline{y})$ and $\mathbf{s}_{d-1}^{(m)}(\overline{y}) = \mathbf{s}^{(m)}(\overline{y})$.

Now let $\alpha < d-1$ and set $k = \lceil \sqrt[d-\alpha]{m} \rceil$ and $\ell = \lceil \frac{m}{k} \rceil$. By possibly duplicating some of the variables we may assume that $m = k\ell$.

To define $\mathbf{r}_\alpha^{(m)}(\overline{y}) = \mathbf{r}_\alpha^{(m)}(y_1, \ldots, y_m)$ we first refer to Lemma 2 to get h_α from $h_{\alpha+1}$ and then we set

$$\mathbf{r}_\alpha^{(m)}(\overline{y}) = \mathbf{q}^{(k)}\Big(h_\alpha, \mathbf{r}_{\alpha+1}^{(\ell)}(y_1, \ldots, y_\ell), \ldots, \mathbf{r}_{\alpha+1}^{(\ell)}(y_{m-\ell+1}, \ldots, y_m), h_{\alpha+1}\Big),$$

where the polynomial $\mathbf{r}_{\alpha+1}^{(\ell)}$ is supplied by the induction hypothesis. From Lemma 2 it should be clear that $\mathbf{r}_\alpha^{(m)}$ satisfies the condition claimed for it.

Also its length can be bounded inductively. Substituting to the polynomial $\mathbf{q}^{(k)}(h_\alpha, x_1, \ldots, x_k, h_{\alpha+1})$ of length $2^{\mathcal{O}(k)}$ (by Lemma 2) the $k = m^{\frac{1}{d-\alpha}}$ copies of the polynomial $\mathbf{r}_{\alpha+1}^{(\ell)}$ of length $2^{\mathcal{O}\left(\ell^{\frac{1}{d-\alpha-1}}\right)}$ and using $\ell = m^{\frac{d-\alpha-1}{d-\alpha}}$ we arrive at the following bound for the length of $\mathbf{r}_\alpha^{(m)}$

$$\begin{aligned} 2^{\mathcal{O}(k)} \cdot 2^{\mathcal{O}(\ell^{\frac{1}{d-\alpha-1}})} &= 2^{\mathcal{O}\left(m^{\frac{1}{d-\alpha}} + \left(m^{\frac{d-\alpha-1}{d-\alpha}}\right)^{\frac{1}{d-\alpha-1}}\right)} \\ &= 2^{\mathcal{O}\left(m^{\frac{1}{d-\alpha}} + m^{\frac{d-\alpha-1}{d-\alpha} \cdot \frac{1}{d-\alpha-1}}\right)} \\ &= 2^{\mathcal{O}\left(m^{\frac{1}{d-\alpha}}\right)}. \end{aligned}$$

In a very similar way we produce $\mathbf{s}_\alpha^{(m)}(\overline{y})$ from the $\mathbf{s}_{\alpha+1}^{(\ell)}$'s by simply putting

$$\begin{aligned} \mathbf{s}_\alpha^{(m)}(\overline{y}) = \mathbf{q}^{(k)}\big(&h_\alpha, \mathbf{s}_{\alpha+1}^{(\ell)}(y_{1,1}, y_{1,2}, y_{1,3}, \ldots, y_{\ell,1}, y_{\ell,2}, y_{\ell,3}), \ldots \\ &\ldots, \mathbf{s}_{\alpha+1}^{(\ell)}(y_{m-\ell+1,1}, y_{m-\ell+1,2}, y_{m-\ell+1,3}, \ldots, y_{m,1}, y_{m,2}, y_{m,3}), h_{\alpha+1}\big). \end{aligned}$$

□

Now we are ready to conclude our proof of Theorem 1. Recall that due to Lemma 3 we are working in the group G in which $G/\mathcal{U}_{d-1}G$ is abelian. We are going to reduce 3SAT or C-COLORING to POLSAT(G) and POLEQV(G) depending on whether $C = |G/H| > 2$ or not. In either case the reduction from C-COLORING to POLSAT(G) and POLEQV(G) works; however, the case $C = 2$ has to be treated in a different way since 2-COLORING is decidable in polynomial time.

In our reduction the formula Φ from 3SAT (or a graph Γ from C-COLORING) is transformed to a polynomial $\mathbf{s}_\Phi$ (or $\mathbf{r}_\Gamma$) and a group element h_1 so that the following will hold:

(A) the length of $\mathbf{s}_\Phi$ (resp. $\mathbf{r}_\Gamma$) is in $2^{\mathcal{O}(\sqrt[d-1]{m})}$ where m is the number of clauses (resp. the number of edges),
(B) $\mathbf{s}_\Phi$ (resp. $\mathbf{r}_\Gamma$) can be computed in time $2^{\mathcal{O}(\sqrt[d-1]{m})}$ (i.e., polynomial in the length of $\mathbf{s}_\Phi$ (resp. $\mathbf{r}_\Gamma$)),
(C) if Φ is satisfiable (resp. Γ has a valid C-coloring), then $\mathbf{s}_\Phi = h_1$ (resp. $\mathbf{r}_\Gamma = h_1$) is satisfiable, and,
(D) if Φ is *not* satisfiable (resp. Γ does *not* have a valid C-coloring), then $\mathbf{s}_\Phi = 1$ (resp. $\mathbf{r}_\Gamma = 1$) holds under *all* evaluations.

The latter two points imply that $\mathbf{s}_\Phi = h_1$ (resp. $\mathbf{r}_\Gamma = h_1$) is satisfiable if and only if Φ is satisfiable (resp. Γ has a valid C-coloring) and $\mathbf{s}_\Phi = 1$ (resp. $\mathbf{r}_\Gamma = 1$) holds identically in G if and only if Φ is *not* satisfiable (resp. Γ does *not* have a valid C-coloring).

Now, if ℓ denotes the input length for POLSAT or POLEQV (i.e. the size of $\mathbf{s}_\Phi$ or $\mathbf{r}_\Gamma$), then an algorithm for POLSAT or POLEQV working in $2^{o(\log^{d-1}\ell)}$-time would solve 3SAT (resp. C-COLORING) in time

$$2^{\mathcal{O}(\sqrt[d-1]{m})} + 2^{o(\log^{d-1}(2^{\sqrt[d-1]{m}}))} = 2^{o(m)},$$

contradicting ETH.

We start with describing the reduction from C-COLORING to POLSAT(G) and POLEQV(G) where $C = |G/H|$. The quotient $|G/H|$ serves as the set of colors. For a graph $\Gamma = (V, E)$ with $E \subseteq \binom{V}{2}$, $|V| = n$ and $|E| = m$, we use variables x_v for $v \in V$. For an edge $\{u, v\} \in E$ the value of $x_u x_v^{-1}$ (modulo H) decides whether the vertices u, v have the same color. To control whether the coloring of Γ is proper we define the polynomial $\mathbf{r}_\Gamma$ by putting

$$\mathbf{r}_\Gamma((x_v)_{v\in V}) = \mathbf{r}_1^{(m)}\Big((x_u x_v^{-1})_{\{u,v\}\in E}\Big)$$

where $\mathbf{r}_1^{(m)}$ and h_1 are supplied by Lemma 5—and, thus, meet the length bound (A). Point (B) is clear from the definition of the polynomial. Notice that the edges can be fed into $\mathbf{r}_1^{(m)}$ in any order without affecting the final value of polynomials. Every evaluation of the variables x_v by elements of G defines a coloring $\chi : V \to G/H$ in a natural way. If this coloring is valid (i.e. $\chi(u) \not\equiv \chi(v) \mod H$ for every edge $\{u, v\} \in E$), then all the expressions $\chi(u)\chi(v)^{-1}$ are not in H and Lemma 5 ensures us that $\mathbf{r}_\Gamma((x_v)_{v\in V}) = h_1$. This shows (C).

Conversely, by Lemma 5, for every evaluation of the x_v's by elements of G that does not satisfy the equation $\mathbf{r}_\Gamma((x_v)_{v\in V}) = 1$, we have $x_u x_v^{-1} \notin H$ for all edges $\{u, v\}$. This obviously yields a valid coloring of Γ—hence, it proves (D).

As 2-COLORING is solvable in polynomial time in the case $|G/H| = 2$, we interpret 3SAT and use the two cosets of H in G as the true/false boolean values. We start with the formula

$$\Phi = (A_{1,1} \vee A_{1,2} \vee A_{1,3}) \wedge \cdots \wedge (A_{m,1} \vee A_{m,2} \vee A_{m,3}),$$

where each literal $A_{i,j}$ is either one of the boolean variables $X_1, \ldots, X_n$ or its negation. First, we transform the literals $A_{i,j}$ into the expressions $x_{i,j}$ that are supposed to range over G by picking $g \in G \setminus H$ and then setting

$$x_{i,j} = \begin{cases} gx_k, & \text{if } A_{i,j} = X_k, \\ x_k, & \text{if } A_{i,j} = \neg X_k. \end{cases}$$

Finally, we set

$$\mathbf{s}_\Phi(x_1, \ldots, x_n) = \mathbf{s}_1^{(m)}\big(x_{1,1}, x_{1,2}, x_{1,3}, \ldots, x_{m,1}, x_{m,2}, x_{m,3}\big)$$

where again $\mathbf{s}_1^{(m)}$ is supplied by Lemma 5.

Now, given an assignment to the boolean variables $X_1, \ldots, X_n$, we obtain an assignment for $x_1, \ldots, x_n$ by setting $x_i = g$ if X_i is *true* and $x_i = 1$ if X_i is *false*. It can be easily checked using Lemma 5 that the original assignment was satisfying for Φ if and only if $\mathbf{s}_\Phi(\overline{x}) = h_1$ is satisfied (notice that $g^2 \in H$). This shows (C). On the other hand, if $\mathbf{s}_\Phi(\overline{x}) \neq 1$, then, by Lemma 5, for all i there is some j with $x_{i,j} \in H$. Hence, if we assign *true* to X_k if and only if $x_k \notin H$, we obtain a satisfying assignment for Φ—proving (D).

Notice that also in the case $|G/H| \geqslant 3$ it would be possible to describe a reduction of 3SAT to POLSAT(G). However, in order to encode negations of literals, we need to restrict the variables to only two possible values (modulo H). This can be done using the polynomials we constructed for the reduction from C-COLORING (which we can use to "forbid" any undesired value). Nevertheless, the total construction would be more complicated than the two individual reductions we described above.

4 Conclusion

With Theorem 1 in mind, one could suspect that finite solvable groups of Fitting length 2 have polynomial time algorithms for POLSAT. As we have already mentioned, the very recent paper [8] shows that POLSAT is in P for many such groups, in particular, for all semidirect products $G_p \rtimes A$, where G_p is a p-group and A is abelian. This, however, does not cover e.g. the dihedral group D_{15}. In fact, in [19] POLSAT(D_{15}) is shown not to be in P, unless ETH fails. On the other hand, POLEQV$(D_{15}) \in$ P. Actually from [8] we know that POLEQV$(G) \in$ P for each semidirect product $G = N \rtimes A$ where N is nilpotent and A is abelian. In fact, D_{15} is the first known example of a group with polynomial time POLEQV and non-polynomial (under ETH) POLSAT. The converse situation cannot happen as, for a group G, POLSAT$(G) \in$ P implies POLEQV$(G) \in$ P. Indeed, to confirm that

$\mathbf{t}(\overline{x}) = 1$ holds for all possible values of the $\overline{x}$'s, we check that for no $g \in G \setminus \{1\}$ the equation $\mathbf{t}(\overline{x}) = g$ has a solution.

We conclude our paper with two obvious questions.

Problem 1 Characterize finite solvable groups (of Fitting length 2) with PolSat decidable in polynomial time.

Problem 2 Characterize finite solvable groups (of Fitting length 2) with PolEqv decidable in polynomial time.

Finally, we want to point out the consequences of our main result to another problem: For a finitely generated (but possibly infinite) group with a finite set of generators Σ the power word problem is as follows: The input is a tuple $(p_1, x_1, p_2, x_2, \ldots, p_n, x_n)$ where the p_i are words over Σ and the x_i are integers encoded in binary. The question is whether $p_1^{x_1} \cdots p_n^{x_n}$ evaluates to the identity of the group. The complexity of the power word problem in a wreath product $G \wr \mathbb{Z}$ where G is a finite group has a similar behaviour as PolEqv: if G is nilpotent, the power word problem of $G \wr \mathbb{Z}$ is in polynomial time [6] (actually even in TC^0) and, if G is non-solvable, it is coNP-complete [24]. Indeed, in [6] a surprising connection to PolEqv has been pointed out: if G is a finite group, then PolEqv(G) can be reduced in polynomial time to the power word problem of the wreath product $G \wr \mathbb{Z}$. In particular, Theorem 1 implies that the power word problem of $G \wr \mathbb{Z}$ where G is a finite solvable group of Fitting length at least three is not in P assuming ETH.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

1. Baer, R.: Engelsche elemente Noetherscher Gruppen. Math. Ann. **133**, 256–270 (1957). https://doi.org/10.1007/BF02547953
2. Barrington, D.A.M., McKenzie, P., Moore, C., Tesson, P., Thérien, D.: Equation satisfiability and program satisfiability for finite monoids. In: Mathematical Foundations of Computer Science 2000, 25th International Symposium, MFCS 2000, Proceedings, Lecture Notes in Computer Science, vol. 1893, pp. 172–181. Springer (2000). https://doi.org/10.1007/3-540-44612-5_13
3. Burris, S., Lawrence, J.: Results on the equivalence problem for finite groups. Algebra Universalis **52**(4), 495–500 (2005). https://doi.org/10.1007/s00012-004-1895-8
4. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer. https://doi.org/10.1007/978-3-319-21275-3 (2015)

5. Diekert, V., Elder, M.: Solutions of twisted word equations, EDT0l languages, and context-free groups. In: ICALP 2017, Proceedings, LIPIcs, vol. 80, pp. 96:1–96:14, Dagstuhl (2017). https://doi.org/10.4230/LIPIcs.ICALP.2017.96. http://drops.dagstuhl.de/opus/volltexte/2017/7397
6. Figelius, M., Ganardi, M., Lohrey, M., Zetzsche, G.: The complexity of knapsack problems in wreath products. In: 47th International colloquium on automata, languages, and programming, ICALP 2020, July 8–11, 2020, Saarbrücken, Germany (Virtual Conference), pp. 126:1–126:18 (2020). https://doi.org/10.4230/LIPIcs.ICALP.2020.126
7. Földvári, A.: The complexity of the equation solvability problem over nilpotent groups. J. Algebra **495**, 289–303 (2018). https://doi.org/10.1016/j.jalgebra.2017.10.002
8. Földvári, A., Horváth, G.: The complexity of the equation solvability and equivalence problems over finite groups. Int. J. Algebra Comput. **30**(03), 607–623 (2020). https://doi.org/10.1142/S0218196720500137
9. Garreta, A., Miasnikov, A., Ovchinnikov, D.: Diophantine problems in solvable groups. Bull. Math. Sci. https://doi.org/10.1142/S1664360720500058 (2020)
10. Goldmann, M., Russell, A.: The complexity of solving equations over finite groups. Inf. Comput. **178**(1), 253–262 (2002). https://doi.org/10.1006/inco.2002.3173
11. Horváth, G.: The complexity of the equivalence and equation solvability problems over nilpotent rings and groups. Algebra Universalis **66**(4), 391–403 (2011). https://doi.org/10.1007/s00012-011-0163-y
12. Horváth, G.: The complexity of the equivalence and equation solvability problems over meta-Abelian groups. J. Algebra **433**, 208–230 (2015). https://doi.org/10.1016/j.jalgebra.2015.03.015
13. Horváth, G., Szabó, C.A.: The complexity of checking identities over finite groups. IJAC **16**(5), 931–940 (2006). https://doi.org/10.1142/S0218196706003256
14. Horváth, G., Szabó, C.: The extended equivalence and equation solvability problems for groups. Discrete. Math. Theor. Comput. Sci. **13**(4), 23–32 (2011)
15. Horváth, G., Mérai, L., Szabó, C., Lawrence, J.: The complexity of the equivalence problem for non-solvable groups. Bull. Lond. Math. Soc. **39**(3), 433–438 (2007). https://doi.org/10.1112/blms/bdm030
16. Huppert, B.: Endliche Gruppen. I. Die Grundlehren der Mathematischen Wissenschaften Band, vol. 134. Springer, Berlin-New York (1967)
17. Idziak, P.M., Krzaczkowski, J.: Satisfiability in multi-valued circuits. In: Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '18, pp. 550-558. Association for Computing Machinery, New York (2018). https://doi.org/10.1145/3209108.3209173. Full version in SIAM Journal on Computing, 53(2022), 337–378 https://doi.org/10.1137/18M1220194
18. Idziak, P.M., Kawałek, P.., Krzaczkowski, J.: Intermediate problems in modular circuits satisfiability. In: Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '20, pp. 578–590. Association for Computing Machinery, New York (2020). https://doi.org/10.1145/3373718.3394780
19. Idziak, P.M., Kawalek, P., Krzaczkowski, J.: Complexity of modular circuits. In 37th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS) (LICS '22), August 2–5, 2022, Haifa, Israel. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3531130.3533350 (2021)
20. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity. J. Comput. Syst. Sci. **63**(4), 512–530 (2001). https://doi.org/10.1006/jcss.2001.1774
21. Kawałek, P., Krzaczkowski, J.: Even faster algorithms for CSAT over supernilpotent algebras. In: 45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020), Leibniz International Proceedings in Informatics (LIPIcs), vol. 170, pp. 55:1–55:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl (2020). https://doi.org/10.4230/LIPIcs.MFCS.2020.55
22. Kompatscher, M.: Notes on extended equation solvability and identity checking for groups. Acta Math. Hungar. **159**(1), 246–256 (2019). https://doi.org/10.1007/s10474-019-00924-7
23. Lohrey, M., Sénizergues, G.: Theories of HNN-extensions and amalgamated products. In: ICALP 2006, Proceedings, pp. 504–515 (2006). https://doi.org/10.1007/11787006_43
24. Lohrey, M., Weiß, A.: The power word problem. In: 44h International symposium on mathematical foundations of computer science, MFCS 2019, proceedings, LIPIcs, vol. 138, pp. 43:1–43:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik (2019). http://www.dagstuhl.de/dagpub/978-3-95977-117-7
25. Makanin, G.S.: Decidability of the universal and positive theories of a free group. Izv. Akad. Nauk SSSR Ser. Mat. **48**, 735–749 (1984). In Russian; English translation in: Math. USSR Izvestija, 25, 75–88, 1985

26. Matijasevic, Y.V.: Enumerable sets are diophantine. Soviet. Math. Dokl. **11**, 354–358 (1970). https://ci.nii.ac.jp/naid/10009422455/en/
27. Papadimitriou, C.H.: Computational Complexity. Addison Wesley, Reading (1994)
28. Robinson, D.J.S. A Course in the Theory of Groups Graduate Texts in Mathematics, 2nd edn, vol. 80. Springer, New York (1996). https://doi.org/10.1007/978-1-4419-8594-1
29. Roman'kov, V.: Equations in free metabelian groups. Siberian Math. J. **20**. https://doi.org/10.1007/BF00969959 (1979)
30. Weiß, A.: Hardness of equations over finite solvable groups under the exponential time hypothesis. In: 47th International Colloquium on Automata, Languages, and Programming (ICALP 2020), Leibniz International Proceedings in Informatics (LIPIcs), vol. 168, pp. 102:1–102:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl (2020). https://doi.org/10.4230/LIPIcs.ICALP.2020.102