

A Polynomial Time Nilpotence Test for Galois Groups and Related Results

V. Arvind¹ and Piyush P Kurur²

¹ Institute of Mathematical Sciences
C.I.T Campus, Chennai, India 600 113
`arvind@imsc.res.in`

² Department of Computer Science and Engineering,
Indian Institute of Technology, Kanpur,
Kanpur, UP 208016, India
`ppk@cse.iitk.ac.in` *

Abstract. We give a deterministic polynomial-time algorithm to check whether the Galois group $\text{Gal}(f)$ of an input polynomial $f(X) \in \mathbb{Q}[X]$ is nilpotent: the running time is polynomial in $\text{size}(f)$. Also, we generalize the Landau-Miller solvability test to an algorithm that tests if $\text{Gal}(f)$ is in Γ_d : this algorithm runs in time polynomial in $\text{size}(f)$ and n^d and, moreover, if $\text{Gal}(f) \in \Gamma_d$ it computes all the prime factors of $\#\text{Gal}(f)$.

1 Introduction

Computing the Galois group of a polynomial is a fundamental problem in algorithmic number theory. Asymptotically, the best known algorithm is due to Landau [3]: on input $f(X)$, it takes time polynomial in $\text{size}(f)$ and the order of its Galois group $\text{Gal}(f)$. If $f(X)$ has degree n then $\text{Gal}(f)$ can have $n!$ elements. Thus, Landau's algorithm takes time exponential in input size. It is a long standing open problem if there is an asymptotically faster algorithm for computing $\text{Gal}(f)$. Lenstra's survey [6] discusses this and related problems.

A different kind of problem is to test for a given $f(x)$ if $\text{Gal}(f)$ satisfies a specific property without explicitly computing it. Galois's seminal work showing $f(X)$ is solvable by radicals if and only if $\text{Gal}(f)$ is solvable is a classic example. Landau and Miller [4] gave a remarkable polynomial-time algorithm for testing solvability of the Galois group without computing the Galois group.

1.1 The results of this article

Our main result is a deterministic polynomial-time algorithm for testing if $\text{Gal}(f)$ is nilpotent. Although nilpotent groups are a proper subclass of solvable groups, the Landau-Miller solvability test does not give a nilpotence test. Basically, the Landau-Miller test is a method of testing that all composition factors of $\text{Gal}(f)$

* work done when the author was a PhD student at the Institute of Mathematical Sciences, Chennai.

are abelian, which tests solvability. Nilpotence however is a more “global” property, in the sense that it cannot be inferred by properties of the composition factors alone.

We note here that nilpotence testing of Galois groups has been addressed by other researchers with the goal of developing good practical algorithms. For example in [2] an algorithm for nilpotence testing is given which takes time polynomial in $\text{size}(f)$ and $\#\text{Gal}(f)$. However, ours is the first algorithm that is provably polynomial time, i.e. runs in time polynomial in $\text{size}(f)$, on all inputs.

Next, we show that the Landau-Miller solvability test can be extended to a polynomial-time algorithm for checking, given $f \in \mathbb{Q}[X]$, if $\text{Gal}(f)$ is in Γ_d for constant d . A group G is in Γ_d if there is a composition series $G = G_0 \triangleright \dots \triangleright G_t = \{1\}$ such that each nonabelian composition factor G_i/G_{i+1} is isomorphic to a subgroup of S_d . The class Γ_d often arises in permutation group algorithms (see e.g. [7]). Moreover, if $\text{Gal}(f) \in \Gamma_d$, the prime factors of $\#\text{Gal}(f)$ can be found in polynomial time.

1.2 Galois theory overview

We quickly recall some Galois theory (see, e.g. [5] for details). Let L and K be fields. If $L \supset K$, we say that L is an extension of K and denote it by L/K . If L/K then L is a vector space over K and by the *degree* of L/K , denoted by $[L : K]$, we mean its dimension. An extension L/K is *finite* if its degree $[L : K]$ is finite. If L/M and M/K are finite extensions then $[L : K] = [L : M].[M : K]$. The polynomial ring $K[X]$ is a unique factorisation domain: every polynomial can be uniquely (up to scalars) written as a product of irreducible polynomials. Let L/K be an extension. An $\alpha \in L$ is *algebraic* over K if $f(\alpha) = 0$ for some $f(X) \in K[X]$. For α algebraic over K , the *minimal polynomial* of α over K is the unique monic polynomial $\mu_\alpha[K](X)$ of least degree in $K[X]$ for which α is a root. We write $\mu_\alpha(X)$ for $\mu_\alpha[K](X)$ when K is understood. Elements $\alpha, \beta \in L$ are *conjugates* over K if they have the same minimal polynomial over K . The smallest subfield of L containing K and α is denoted by $K(\alpha)$.

The *splitting field* K_f of $f \in K[X]$ is the smallest extension of K containing all the roots of f . A finite extension L/K is *normal* if for all irreducible polynomials $f(X) \in K[X]$, either $f(X)$ splits or has no root in L . Any normal extension over K is the splitting field of some polynomial in $K[X]$. An extension L/K is *separable* if for all irreducible polynomials $f(X) \in K[X]$ there are no multiple roots in L . A normal and separable finite extension L/K is a *Galois extension*.

The *Galois group* $\text{Gal}(L/K)$ of L/K is the subgroup of automorphisms σ of L that leaves K fixed, i.e. $\sigma(\alpha) = \alpha$ for all $\alpha \in K$. The Galois group $\text{Gal}(f)$ of $f \in K[X]$ is $\text{Gal}(K_f/K)$. For a subgroup G of automorphisms of L , the *fixed field* L^G is the largest subfield of L fixed by G . We now state the fundamental theorem of Galois.

Theorem 1. [5, Theorem 1.1, Chapter VI] *Let L/K be a Galois extension with Galois group G . There is a one-to-one correspondence between subfields E of L*

containing K and subgroups H of G , given by $E \cong L^H$. The Galois group of $\text{Gal}(L/E)$ is H and E/K is a Galois extension if and only if H is a normal subgroup of G . If H is a normal subgroup of G and $E = L^H$ then $\text{Gal}(E/K)$ is isomorphic to the quotient group G/H .

1.3 Presenting algebraic numbers, number fields and Galois groups

The algorithms we describe take objects like algebraic numbers, number fields etc. as input. We define sizes of these objects. Integers are encoded in binary. A rational r is given by coprime integers a, b such that $r = a/b$. Thus, $\text{size}(r)$ is $\text{size}(a) + \text{size}(b)$. A polynomial $T(X) = a_0 + \dots + a_n X^n \in \mathbb{Q}[X]$ is given by a list of its coefficients. Thus, $\text{size}(T)$ is defined as $\sum \text{size}(a_i)$.

A *number field* is a finite extension of \mathbb{Q} . Let K/\mathbb{Q} be a number field of degree n . By the primitive element theorem [5, Theorem 4.6, Chapter V], there is an algebraic number $\eta \in K$ such that $K = \mathbb{Q}(\eta)$. Such an element is a *primitive element* of K/\mathbb{Q} and its minimal polynomial is a *primitive polynomial*. Let $\mu_\eta(X)$ be the minimal polynomial of η over \mathbb{Q} . Then the field K can be written as the quotient $K = \mathbb{Q}[X]/\mu_\eta(X)$. Thus K can be presented by giving a primitive polynomial for K/\mathbb{Q} . We can assume that η is an algebraic integer and hence its minimal polynomial $\mu_\eta(X)$ has integer coefficients [5, Proposition 1.1, Chapter VII]. When we say that an algorithm takes a number field K as input we mean that it takes a primitive polynomial $\mu_\eta(X)$ for K as input. Thus the input size for K , which we denote by $\text{size}(K)$, is defined to be $\text{size}(\mu_\eta)$.

Suppose $K = \mathbb{Q}(\eta)$ is presented by $\mu_\eta(X)$. Notice that each $\alpha \in K$ can be expressed as $\alpha = A_\alpha(\eta)$ for a unique polynomial $A_\alpha(X) \in \mathbb{Q}[X]$ of degree less than n . By $\text{size}(\alpha)$ we mean $\text{size}(A_\alpha(X))$. Note that the size of $\alpha \in K$ depends on the primitive element $\eta \in K$. Now, for a polynomial $f(X) = a_0 + \dots + a_m X^m$ in $K[X]$ we define $\text{size}(f)$ to be $\sum \text{size}(a_i)$.

Let $f(X) \in \mathbb{Q}[X]$ of degree n . For an algorithm purporting to compute $\text{Gal}(f)$, one possibility is that it outputs the complete multiplication table for $\text{Gal}(f)$. However, this could be exponential in $\text{size}(f)$ as $\text{Gal}(f)$ can be as large as $n!$. A succinct presentation of $\text{Gal}(f)$ is as a permutation group acting on the roots of f since elements of $\text{Gal}(f)$ permute the roots of f and are completely determined by their action on the roots of f . Thus, by numbering the roots of f , we can consider $\text{Gal}(f)$ as a subgroup of the symmetric group S_n (note here that $\text{Gal}(f)$ is determined only up to conjugacy as the numbering of the roots is arbitrary). Since any subgroup of S_n has a generator set of size $n - 1$ (see e.g. [8]), we can present $\text{Gal}(f)$ in size polynomial in n . Thus, by computing $\text{Gal}(f)$ we mean finding a small generator set for it as a subgroup of S_n . Determining $\text{Gal}(f)$ as a subgroup of S_n is a reasonable way of describing the output. Algorithmically, we can answer several natural questions about a subgroup G of S_n given by generator set in polynomial time. E.g. testing if G is solvable, finding a composition series for G etc. [8].

Previous complexity results As mentioned, the best known algorithm for computing the Galois group of a polynomial is due to Landau [3].

Theorem 2 (Landau). *There is a deterministic algorithm that takes as input a number field K , a polynomial $f(X) \in K[X]$ and a positive integer b in unary, and in time bounded by $\text{size}(f)$, $\text{size}(K)$ and b , decides if $\text{Gal}(K_f/K)$ has at most b elements, and if so computes $\text{Gal}(K_f/K)$ by finding the entire multiplication table of $\text{Gal}(K_f/K)$ (and hence also by giving the generating set of $\text{Gal}(K_f/K)$ as a permutation group on the roots of $f(X)$).*

The algorithm first computes a primitive element θ of K_f . Determining $\text{Gal}(f)$ amounts to finding the action of the automorphisms on θ . Subsequently, Landau and Miller [4] gave their polynomial-time solvability test.

Theorem 3 (Landau-Miller). *Given $f(X) \in \mathbb{Q}[X]$ there is a deterministic polynomial-time algorithm for testing if $\text{Gal}(f)$ is solvable.*

2 Preliminaries

We recall some permutation group theory from Wielandt's book [9]. Let Ω be a finite set. The *symmetric group* $\text{Sym}(\Omega)$ is the group of all permutations on Ω . By a *permutation group on Ω* we mean a subgroup of $\text{Sym}(\Omega)$. For $\alpha \in \Omega$ and $g \in \text{Sym}(\Omega)$, let α^g denote the image of α under the permutation g . For $A \subseteq \text{Sym}(\Omega)$, α^A denotes the set $\{\alpha^g : g \in A\}$. In particular, for $G \leq \text{Sym}(\Omega)$ the G -orbit containing α is α^G . The G -orbits form a partition of Ω . Given $G \leq \text{Sym}(\Omega)$ by a generating set A and $\alpha \in \Omega$, there is a polynomial-time algorithm to compute α^G [8].

For $\Delta \subseteq \Omega$ and $g \in \text{Sym}(\Omega)$, Δ^g denotes $\{\alpha^g : \alpha \in \Delta\}$. The setwise stabilizer of Δ , i.e. $\{g \in G : \Delta^g = \Delta\}$, is denoted by G_Δ . If Δ is the singleton set $\{\alpha\}$ we write G_α instead of $G_{\{\alpha\}}$. For any Δ by $G|_\Delta$ we mean G_Δ restricted to Δ . An often used result is the orbit-stabilizer formula stated below [9, Theorem 3.2].

Theorem 4 (Orbit-stabilizer formula). *Let G be a permutation group on $\text{Sym}(\Omega)$ and let α be any element of Ω then the order of the group G is given by $\#G = \#G_\alpha \cdot \#\alpha^G$.*

A permutation group G on Ω is *transitive* if there is a single G -orbit. Suppose $G \leq \text{Sym}(\Omega)$ is transitive. Then a non-empty subset Δ of Ω is a G -block if for all $g \in G$ either $\Delta^g = \Delta$ or $\Delta^g \cap \Delta = \emptyset$. For every G , Ω is a block and each singleton $\{\alpha\}$ is a block. These are the *trivial blocks* of G . A transitive group G is *primitive* if it has only trivial blocks and it is *imprimitive* if it has nontrivial blocks. A G -block Δ is a *maximal subblock* of a G -block Σ if $\Delta \subset \Sigma$ and there is no G -block Υ such that $\Delta \subset \Upsilon \subset \Sigma$. Let Δ and Σ be two G -blocks. A chain $\Delta = \Delta_0 \subset \dots \subset \Delta_t = \Sigma$ is a *maximal chain* of G -blocks between Δ and Σ if for all i , Δ_i is a maximal subblock of Δ_{i+1} .

For a G -block Δ and $g \in G$, Δ^g is also a G -block such that $\#\Delta = \#\Delta^g$. Let Δ and Σ be two G -blocks such that $\Delta \subseteq \Sigma$. The Δ -block system of Σ , is the collection

$$\mathcal{B}(\Sigma/\Delta) = \{\Delta^g : g \in G \text{ and } \Delta^g \subseteq \Sigma\}.$$

The set $\mathcal{B}(\Sigma/\Delta)$ is a partition of Σ . It follows that $\#\Delta$ divides $\#\Sigma$ and by index of Δ in Σ , which we denote by $[\Sigma : \Delta]$, we mean $\#\mathcal{B}(\Sigma/\Delta) = \frac{\#\Sigma}{\#\Delta}$. We will use $\mathcal{B}(\Delta)$ to denote $\mathcal{B}(\Omega/\Delta)$. We state the connection between blocks and subgroups [9, Theorem 7.5].

Theorem 5 (Galois correspondence of blocks). *Let $G \leq \text{Sym}(\Omega)$ be transitive and $\alpha \in \Omega$. For $G \geq H \geq G_\alpha$ the orbit $\Delta = \alpha^H$ is a G -block and $G_\Delta = H$. The correspondence $\alpha^H = \Delta \iff G_\Delta = H$ is a one-to-one correspondence between G -blocks Δ containing α and subgroups H of G containing G_α . Furthermore for G -blocks $\Delta \subseteq \Sigma$ we have $[G_\Sigma : G_\Delta] = [\Sigma : \Delta]$.*

Let $G \leq \text{Sym}(\Omega)$ be transitive and Δ and Σ be two G -blocks such that $\Delta \subseteq \Sigma$. Let $G(\Sigma/\Delta)$ denote the group $\{g \in G : \Upsilon^g = \Upsilon \text{ for all } \Upsilon \in \mathcal{B}(\Sigma/\Delta)\}$. We write G^Δ for the group $G(\Omega/\Delta)$. For any $g \in G_\Sigma$, since g setwise stabilises Σ , g permutes the elements of $\mathcal{B}(\Sigma/\Delta)$. Hence for any $\Upsilon \in \mathcal{B}(\Sigma/\Delta)$ we have $\Upsilon^{g^{-1}G(\Sigma/\Delta)g} = \Upsilon$. Thus, $G(\Sigma/\Delta)$ is a normal subgroup of G_Σ . In particular, G^Δ is a normal subgroup of G .

Remark. The following two lemmata are quite standard in permutation group theory. For the reader's convenience we have included short proofs. The following lemma lists important properties of G^Δ .

Lemma 1.

1. For a G -block $\Delta \subseteq \Sigma$, $G(\Sigma/\Delta)$ is the largest normal subgroup of G_Σ contained in G_Δ .
2. Let Σ be a G -block then $G^\Sigma \hookrightarrow \prod_{\Upsilon \in \mathcal{B}(\Sigma)} G|_\Upsilon$.
3. Let Δ be a G -subblock of Σ then $\frac{G_\Sigma}{G(\Sigma/\Delta)}$ is a faithful permutation group on $\mathcal{B}(\Sigma/\Delta)$ and is primitive when Δ is a maximal subblock.
4. The quotient group G^Σ/G^Δ can be embedded as a subgroup of $\left(\frac{G_\Sigma}{G(\Sigma/\Delta)}\right)^l$ for some l .

Proof. Let $N \subseteq G_\Delta$ be a normal subgroup of G_Σ . Since $\Delta^N = \Delta$, and since G_Σ acts transitively on $\mathcal{B}(\Sigma/\Delta)$, for any $\Upsilon \in \mathcal{B}(\Sigma/\Delta)$ there is a $g \in G_\Sigma$ such that $\Upsilon = \Delta^g$. Therefore, $\Upsilon^N = \Delta^{gN} = \Delta^{Ng} = \Upsilon$ for each $\Upsilon \in \mathcal{B}(\Sigma/\Delta)$. Thus $N \subseteq G(\Sigma/\Delta)$. Since $G(\Sigma/\Delta) \trianglelefteq G_\Sigma$ we have proved part 1.

Part 2 directly follows from the definition of G^Σ . Part 3 follows from the fact that $g, h \in G_\Sigma$ have the same action on $\mathcal{B}(\Sigma/\Delta)$ precisely when $gG(\Sigma/\Delta) = hG(\Sigma/\Delta)$. The nontrivial $\frac{G_\Sigma}{G(\Sigma/\Delta)}$ -blocks of $\mathcal{B}(\Sigma/\Delta)$ are in 1-1 correspondence with the G -blocks properly between Δ and Σ . Thus, $\frac{G_\Sigma}{G(\Sigma/\Delta)}$ is primitive if and only if Δ is a maximal subblock of Σ .

For Part 4 notice that we have the group isomorphism

$$\frac{G|_\Upsilon}{G(\Upsilon/\Delta)|_\Upsilon} \cong \frac{G_\Upsilon}{G(\Upsilon/\Delta)},$$

for each $\Upsilon \in \mathcal{B}(\Sigma)$. As $G^\Delta = G^\Sigma \cap \prod G(\Upsilon/\Delta_\Upsilon)|_\Upsilon$ we have

$$G^\Sigma/G^\Delta \hookrightarrow \prod_{r \in \mathcal{B}(\Sigma)} \frac{G|_r}{G(\Upsilon/\Delta_\Upsilon)|_r} = \prod_{r \in \mathcal{B}(\Sigma)} \frac{G_r}{G(\Upsilon/\Delta_\Upsilon)}.$$

Let $g \in G$ such that $\Delta^g = \Delta_\Upsilon$. Then, $G_\Upsilon = g^{-1}G_\Sigma g$ and $G(\Upsilon/\Delta_\Upsilon) = g^{-1}G(\Sigma/\Delta)g$. Thus, $\frac{G_\Sigma}{G(\Sigma/\Delta)}$ and $\frac{G_\Upsilon}{G(\Upsilon/\Delta_\Upsilon)}$ are isomorphic, which implies that G^Σ/G^Δ is isomorphic to a subgroup of $\left(\frac{G_\Sigma}{G(\Sigma/\Delta)}\right)^l$ for some l .

Lemma 2. *Let $G \leq \text{Sym}(\Omega)$ be transitive and $N \triangleleft G$. Let $\alpha \in \Omega$. Then the N -orbit α^N is a G -block and the collection of N -orbits is an α^N -block system of Ω under G action. If $N \neq \{1\}$ then $\|\alpha^N\| > 1$. Furthermore, if $G_\alpha \leq N \neq G$ then the α^N -block system is nontrivial implying that G is not primitive.*

Proof. Let $\alpha \in \Omega$ and $g \in G$. Then $(\alpha^N)^g = \alpha^{Ng} = \alpha^{gN} = (\alpha^g)^N$. Thus $(\alpha^N)^g$ and α^N are N -orbits, and hence are identical or disjoint. Thus, α^N is a G -block and the N -orbits form a block system. Clearly, if $\alpha^N = \{\alpha\}$ then $N = \{1\}$. Finally, by the Orbit-Stabilizer formula $\#G = \#\Omega \cdot \#G_\alpha$ and $\#N = \#\alpha^N \cdot \#G_\alpha$. Thus, if $\{1\} \neq N \neq G$ then α^N is a proper G -block.

3 Nilpotence testing for Galois groups

First we recall crucial properties of nilpotent transitive permutation groups. These are standard group theoretic facts that we assemble together and, for the sake of completeness, provide proof sketches where necessary. We start with a characterization of finite nilpotent groups. Let G be a finite group and p_1, \dots, p_k be the prime factors of $\#G$. For each i , let G_{p_i} be a p_i -Sylow subgroup of G . Then G is *nilpotent* if and only if G is the (internal) direct product $G_{p_1} \times \dots \times G_{p_k}$. Consequently, G_{p_i} is the unique p_i -Sylow subgroup of G for each i and hence $G_{p_i} \triangleleft G$.

Lemma 3. *Let $G \leq \text{Sym}(\Omega)$ be transitive and nilpotent, and p be any prime. Then*

- (1) *The prime p divides $\#G$ if and only if p divides $\#\Omega$.*
- (2) *If $p \mid \#G$ and $\alpha \in \Omega$ then there is a block Σ_p^α containing α such that $\#\Sigma_p^\alpha$ is the highest power of p that divides $\#\Omega$.*
- (3) *Let Δ be any G -block containing α such that $\#\Delta = p^l$ and suppose p divides $\#G$. Then $\Delta \subseteq \Sigma_p^\alpha$. Also, for $q \neq p$, the q -Sylow subgroup of G_Δ is given by $G_q \cap G_\Delta = G_q \cap G_\alpha$.*

Proof. Part (1): As G is transitive, $\#\Omega$ divides $\#G$. Hence, each prime factor of $\#\Omega$ divides $\#G$. Let p be a prime factor of $\#G$. For $\alpha \in \Omega$, let $\Sigma_p^\alpha = \alpha^{G_p}$. Since G_p is transitive on Σ_p^α , it follows from the Orbit-Stabilizer formula that $\#\Sigma_p^\alpha$ divides $\#G_p$. Hence $\#\Sigma_p^\alpha$ is p^l for some l . Since $G_p \triangleleft G$, by Lemma 2 it

follows that its orbit Σ_p^α is G -block which contains more than one element of Ω . Hence $\#\Sigma_p^\alpha = p^l$ for some $l > 0$. Since p divides the cardinality of a G -block Σ_p^α , p divides $\#\Omega$.

Part (2): From the Galois correspondence of G -blocks (Theorem 5) we have $[\Omega : \Sigma_p^\alpha] = [G : G_{\Sigma_p^\alpha}]$. Notice that p is not a factor of $[G : G_p]$ as G_p is the p -Sylow subgroup of G . Since $G_p \triangleleft G_{\Sigma_p^\alpha}$ it follows that p is not a factor of $[G : G_{\Sigma_p^\alpha}]$. Hence p is not a factor of $[\Omega : \Sigma_p^\alpha]$.

Part (3): notice that G_Δ is a nilpotent group with the unique normal q -Sylow subgroup $G_q \cap G_\Delta$. Thus, $G_\Delta = \prod_q (G_q \cap G_\Delta)$. By Theorem 5 we have

$$\#\Delta = [G_\Delta : G_\alpha] = \prod_q [G_q \cap G_\Delta : G_q \cap G_\alpha]. \quad (1)$$

Since $G_q \cap G_\Delta$ is a q -group, p divides $[G_q \cap G_\Delta : G_q \cap G_\alpha]$ if and only if $q = p$. However, in Equation 1, $\#\Delta$ is a power of p . This forces $[G_q \cap G_\Delta : G_q \cap G_\alpha] = 1$ for all $q \neq p$. Thus $G_q \cap G_\Delta = G_q \cap G_\alpha$ for $q \neq p$. Therefore, G_Δ is the product group $G_p \cap G_\Delta \times \prod_{q \neq p} G_q \cap G_\alpha$. Since $G_{\Sigma_p^\alpha}$ contains both G_p and G_α we have $G_{\Sigma_p^\alpha} \geq G_\Delta$. Thus, Δ is a G -subblock of Σ_p^α .

We recall a result about permutation p -groups (see e.g. Luks [7, Lemma 1.1]).

Lemma 4. *Let $G \leq \text{Sym}(\Omega)$ be a transitive p -group and Δ be a maximal G -block. Then $[\Omega : \Delta] = p$ and $G_\Delta = G(\Omega/\Delta) = G^\Delta$ is a normal group of index p in G .*

The next lemma is an easy consequence of Lemma 4 and it states a useful property of permutation p -groups.

Lemma 5. *Let $H \leq \text{Sym}(\Omega)$ be a transitive p -group and $\alpha \in \Omega$. Let $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_t = \Omega$ be any maximal chain of H -blocks. Then*

1. $[\Delta_{i+1} : \Delta_i] = p$ for all $0 \leq i < t$.
2. $H(\Delta_{i+1}/\Delta_i) = H_{\Delta_i}$. Hence, $H_{\Delta_i} \triangleleft H_{\Delta_{i+1}}$ and the quotient $H_{\Delta_{i+1}}/H_{\Delta_i}$ is cyclic of order p .

We continue with the notation of Lemma 3. In the next lemma we show that the block structure of transitive nilpotent permutation group G is similar to the block structure p -groups.

Lemma 6. *Let G be a nilpotent transitive permutation group on Ω and let p be a prime factor of $\#G$. Let Δ be any subset of Σ_p^α . Then Δ is a G -block if and only if Δ is a G_p block (in its transitive action on Σ_p^α).*

Proof. Let H denote the p -Sylow subgroup G_p . Let \widehat{H} denote the product $\prod_{q \neq p} G_q$ of all other Sylow subgroups of G . Then $G = H \times \widehat{H}$. Recall that Σ_p^α is the H -orbit of α .

Firstly any G -block $\Delta \subseteq \Sigma_p^\alpha$ is an H -block. To prove the converse consider any H -block $\Sigma \subseteq \Sigma_p^\alpha$. Consider the group $G' = H_\Sigma \times (\widehat{H} \cap G_\alpha)$. Firstly notice that the group G' is a subgroup of $G_{\Sigma_p^\alpha}$. Also since G_α is nilpotent, we have $G_\alpha = H_\alpha \times (\widehat{H} \cap G_\alpha)$. Furthermore since Σ is a H -block, we have $H_\Sigma \geq H_\alpha$. Therefore $G' \geq G_\alpha$ and by the Galois correspondence of blocks (Theorem 5), $\Sigma = \alpha^{G'}$ is a G -block and $G_\Sigma = G'$.

We give a characterisation of nilpotent transitive permutation groups by properties of maximal chains of G -blocks between $\{\alpha\}$ and Σ_p^α which is crucial for our polynomial-time nilpotence test. This characterization is probably well-known to group theorists. However, as we haven't seen it anywhere, we include a proof.

Theorem 6. *Let $G \leq \text{Sym}(\Omega)$ be a transitive permutation group satisfying properties (1) and (2) of Lemma 3 (which are necessary conditions for nilpotence of G). Fix an $\alpha \in \Omega$. The following statements are equivalent.*

- (1) G is nilpotent.
- (2) For each prime factor p of $\#G$, every maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m = \Sigma_p^\alpha$ has the property that $[\Delta_{i+1} : \Delta_i] = p$, G_{Δ_i} is a normal subgroup of $G_{\Delta_{i+1}}$, and p does not divide the order of G/G^{Δ_m} .
- (3) For each prime p dividing $\#G$, there is a maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m = \Sigma_p^\alpha$ with the property that $[\Delta_{i+1} : \Delta_i] = p$, G_{Δ_i} is a normal subgroup of $G_{\Delta_{i+1}}$, and p does not divide the order of G/G^{Δ_m} .

Proof. Clearly (2) implies (3). It suffices to show that (3) implies (1) and (1) implies (2).

To see that (3) implies (1) it is enough to show that each Sylow subgroup of G is normal. To this end, let p be a prime factor of $\#G$ and let $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m = \Sigma_p^\alpha$ be a maximal chain of G -blocks having the properties mentioned in (3).

Firstly, since $G(\Delta_{i+1}/\Delta_i)$ is the largest normal subgroup of $G_{\Delta_{i+1}}$ that is contained in G_{Δ_i} (part 1 of Lemma 1), (3) implies that $G_{\Delta_i} = G(\Delta_{i+1}/\Delta_i)$. Furthermore it follows from Lemma 1 that there is a positive integer l_i for each i such that the quotient group $G^{\Delta_{i+1}}/G^{\Delta_i}$ is embeddable in an l_i -fold product of copies of $\frac{G^{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)} = G_{\Delta_{i+1}}/G_{\Delta_i}$. Since $[G_{\Delta_{i+1}} : G_{\Delta_i}] = p$ it follows that $G^{\Delta_{i+1}}/G^{\Delta_i}$ is a p -group for each i . As $\#G^{\Delta_m} = \prod_{i=0}^{m-1} [G^{\Delta_{i+1}} : G^{\Delta_i}]$, G^{Δ_m} is also a p -group. Since $G^{\Delta_m} \triangleleft G$ and p does not divide $[G : G^{\Delta_m}]$ it follows that G^{Δ_m} is a normal p -Sylow subgroup of G . The nilpotence of G follows as this holds for all prime factors of $\#G$.

Next, we show that (1) implies (2). Suppose G is nilpotent. Let p be a prime factor of $\#G$ and $\alpha \in \Omega$. Let H be the p -Sylow subgroup G_p of G and let $\widehat{H} = \prod_{q \neq p} G_q$ be the product of all its other Sylow subgroups. Let $\{\alpha\} = \Delta_0 \subset \Delta_1 \subset \dots \subset \Delta_m = \Sigma_p^\alpha$ be any maximal chain of G -blocks between α and Σ_p^α . It follows from Lemma 6 that the chain $\{\Delta\}_{0 \leq i \leq m}$ is a maximal chain of G_p -blocks. By Lemma 5 we have $[\Delta_{i+1} : \Delta_i] = p$, $H_{\Delta_i} \triangleleft H_{\Delta_{i+1}}$, and $H_{\Delta_{i+1}}/H_{\Delta_i}$ is cyclic

of order p . The group $G_{\Delta_i} = H_{\Delta_i} \times \widehat{H}_{\Delta_i}$ and $G_{\Delta_{i+1}} = H_{\Delta_{i+1}} \times \widehat{H}_{\Delta_{i+1}}$. Also since \widehat{H}_{Δ_i} is the product of q -Sylow subgroups of H_{Δ_i} where q varies over all prime factors of $\#G$ different from p , it follows from Lemma 3 that $\widehat{H}_{\Delta_i} = \widehat{H}_\alpha$. Therefore $G_{\Delta_i} \triangleleft G_{\Delta_{i+1}}$ and quotient group $G_{\Delta_{i+1}}/G_{\Delta_i} \cong H_{\Delta_{i+1}}/H_{\Delta_i}$. The group G/G^{Δ_m} acts faithfully on $\mathcal{B}(\Omega/\Delta_m)$ and is transitive under this action. Since $p \nmid [\Omega : \Delta_m]$, p cannot divide the order of G/G^{Δ_m} (Lemma 3).

The following lemma is crucial for the nilpotence testing algorithm. If G is nilpotent then, for each prime factor p of $\#G$, the lemma implies that no matter how the maximal chain of blocks Δ_i of Theorem 6 is constructed, it must terminate in Σ_p^α .

Lemma 7. *Let G be a transitive nilpotent permutation group on Ω . Let p be any prime dividing $\#G$. Let Δ be any G -block such that $\#\Delta = p^l$ for some integer $l \geq 0$. Let m be the highest power of p that divides $\#\Omega$. If $l < m$ then we have*

1. *There exists a G -block Σ such that Δ is a maximal G -subblock of Σ and $[\Sigma : \Delta] = p$.*
2. *For all G -blocks Σ such that Δ is a maximal G -subblock of Σ and $[\Sigma : \Delta] = p$, G_Δ is a normal subgroup of G_Σ .*

Proof. Since $\#\Delta$ is p^l it follows that Δ is a G -subblock of Σ_p^α (Lemma 3). It follows from Lemma 6 that Δ is a G_p -block on the transitive action of G_p on Σ_p^α . Furthermore if $l < m$ there is a G_p -block Σ (and hence by Lemma 6 a G -block) such that $\Sigma_p^\alpha \supseteq \Sigma \supset \Delta$ and $[\Sigma : \Delta] = p$. This proves part 1.

Let $\alpha \in \Delta$. It follows from Lemma 3 that for $q \neq p$ the q -Sylow subgroup of G_Σ and G_Δ are both $G_q \cap G_\alpha$. Let \widehat{G}_p be $\prod_{q \neq p} G_q$. The groups G_Σ and G_Δ are $(G_p \cap G_\Sigma) \times (\widehat{G}_p \cap G_\alpha)$ and $(G_p \cap G_\Delta) \times (\widehat{G}_p \cap G_\alpha)$ respectively. Moreover, $G_p \cap G_\Sigma$ and $G_p \cap G_\Delta$ are p -groups with index $[G_p \cap G_\Sigma : G_p \cap G_\Delta] = [G_\Sigma : G_\Delta] = [\Sigma : \Delta] = p$. Therefore, $G_p \cap G_\Delta$ is normal in $G_p \cap G_\Sigma$. Thus, $G_\Delta = (G_p \cap G_\Delta) \times (\widehat{G}_p \cap G_\alpha)$ is normal in $G_\Sigma = (G_p \cap G_\Sigma) \times (\widehat{G}_p \cap G_\alpha)$ and $\frac{G_\Sigma}{G_\Delta} = \frac{G_p \cap G_\Sigma}{G_p \cap G_\Delta}$ is isomorphic to \mathbb{Z}_p .

3.1 The nilpotence test

Given $f(X) \in \mathbb{Q}[X]$ our goal is to test if $\text{Gal}(f)$ is nilpotent. We can assume that $f(X)$ is irreducible. For, otherwise we can compute the irreducible factors of $f(X)$ over \mathbb{Q} using the LLL algorithm, and perform the nilpotence test on each distinct irreducible factor. This suffices because nilpotent groups are closed under products and subgroups. Let G be $\text{Gal}(f)$. We consider G as a subgroup of $\text{Sym}(\Omega)$, where Ω is the set of roots of $f(X)$. Since f is irreducible, G is transitive on Ω .

For any G -block Δ , let \mathbb{Q}_Δ be the fixed field of the splitting field \mathbb{Q}_f under the automorphisms of G_Δ . Let Δ be a G -block containing α . Since $G_\Delta \geq G_\alpha$, \mathbb{Q}_Δ is a subfield of $\mathbb{Q}_{\{\alpha\}} = \mathbb{Q}(\alpha)$.

We describe the main idea. By Theorem 6, G is nilpotent if and only if for all primes p that divide the order of G , there is a maximal chain of G -blocks

$\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m$ satisfying conditions of part (3) of Theorem 6. We show these conditions can be verified in polynomial time once the tower of fields $\mathbb{Q}(\alpha) = \mathbb{Q}_{\Delta_0} \supset \dots \supset \mathbb{Q}_{\Delta_m}$ are known. Thus, for testing nilpotence of G we will first need a polynomial-time algorithm that computes \mathbb{Q}_{Δ_i} . The following theorem is essentially due to Landau and Miller [4] restated in a form suitable for our application.

Theorem 7. *Let $f(X) \in \mathbb{Q}[X]$ be irreducible, $G = \text{Gal}(f)$ be its Galois group and Ω be the set of roots of f . Let $\Delta \subseteq \Omega$ be any G -block and $\alpha \in \Delta$. There is an algorithm that given a primitive polynomial $\mu_\Delta(X) \in \mathbb{Q}[X]$ of \mathbb{Q}_Δ , runs in time polynomial in $\text{size}(f)$ and $\text{size}(\mu_\Delta)$ and computes a primitive polynomial $\mu_\Sigma(X) \in \mathbb{Q}[X]$ of \mathbb{Q}_Σ for all G -blocks Σ such that Δ is a maximal block of Σ . Moreover $\text{size}(\mu_\Sigma)$ is at most a polynomial in $\text{size}(f)$ and is independent of $\text{size}(\mu_\Delta)$.*

We now give the algorithm for testing nilpotence.

Input: A polynomial $f(X) \in \mathbb{Q}[X]$ of degree n
Output: “Accept” if $\text{Gal}(f)$ is nilpotent; “Reject” otherwise
Verify that $f(X)$ is solvable;

- 1 Compute the set P of all the prime factors of $\#\text{Gal}(f)$;
Let $G \leq \text{Sym}(\Omega)$ denote the Galois group of f , where Ω is the set of roots of f .
- 2 **for every** $p \in P$ **do**
 - if** p does not divide n **then**
 - | **print** *Reject*
 - end**
 - Let m be the highest power of p dividing n .
 - 3 Attempt to compute the tower $\mathbb{Q}_{\Delta_m} \subset \dots \subset \mathbb{Q}_{\Delta_0}$ for a maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m$ such that $[\mathbb{Q}_{\Delta_{i+1}} : \mathbb{Q}_{\Delta_i}] = p$.
 - 4 **if** Step 3 fails **or** $\mathbb{Q}_{\Delta_{i+1}}$ is not normal over \mathbb{Q}_{Δ_i} **then**
 - | **print** *Reject*
 - end**
 - Let $\mu_{\Delta_m}(X)$ be the primitive polynomial for \mathbb{Q}_{Δ_m} .
 - 5 **if** p divides $\#\text{Gal}(\mu_{\Delta_m})$ **then**
 - | **print** *Reject*
 - end**
- end**
- print** *Accept*

Algorithm 1: Nilpotence test

We prove that Algorithm 1 runs in polynomial time. For the steps 1 and 5 note that for polynomials f with solvable Galois groups, as a byproduct of the Landau-Miller test [4], the prime factors of $\#\text{Gal}(f)$ can be found in polynomial time (see also Theorem 11). We explain how step 3 can be done in polynomial time using Theorem 7. We construct \mathbb{Q}_{Δ_i} inductively starting with $\mathbb{Q}_{\Delta_0} = \mathbb{Q}(\alpha)$.

Assume we have computed \mathbb{Q}_{Δ_i} . Using Theorem 7 we compute \mathbb{Q}_{Σ} for each G -block Σ containing Δ_i as a maximal G -subblock. Among them choose a \mathbb{Q}_{Σ} for which $[\Sigma : \Delta_i] = p$ and let $\mathbb{Q}_{\Delta_{i+1}}$ be \mathbb{Q}_{Σ} . The inductive construction of $\mathbb{Q}_{\Delta_{i+1}}$ from \mathbb{Q}_{Δ_i} can be done in time bounded by a polynomial in size (f) . Putting it together we have the following proposition.

Proposition 1. *Algorithm 1 runs in time polynomial in size (f) .*

We now argue its correctness. Part (1) of Theorem 6 implies that if G is nilpotent then Algorithm 1 accepts. Conversely, suppose the algorithm accepts. Then for each prime p dividing $\#G$ we have a maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m$ such that $\mathbb{Q}_{\Delta_i}/\mathbb{Q}_{\Delta_{i+1}}$ are normal extensions for each $0 \leq i < m$ (this we verify in step 4 of Algorithm 1). Recall that \mathbb{Q}_{Δ_i} is the fixed field of \mathbb{Q}_f w.r.t. G_{Δ_i} . Hence by checking $\mathbb{Q}_{\Delta_i}/\mathbb{Q}_{\Delta_{i+1}}$ is a normal extension we have verified that $G_{\Delta_i} \triangleleft G_{\Delta_{i+1}}$. Also, the splitting field of the primitive polynomial $\mu_{\Delta_m}(X)$ is the normal closure of \mathbb{Q}_{Δ_m} over \mathbb{Q} . It follows from Lemma 1 and Theorem 1 that $\text{Gal}(\mu_{\Delta_m})$ is G^{Δ_m} . Hence, by checking p does not divide $\#\text{Gal}(\mu_{\Delta})$ we have verified that p does not divide $\#G/G^{\Delta_m}$. Thus, we have verified that the maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_m$ satisfies the conditions of Part(3) of Theorem 6 implying that G is nilpotent. Putting it all together we have the following theorem.

Theorem 8. *There is a polynomial-time algorithm that takes $f \in \mathbb{Q}[X]$ as input and tests if $\text{Gal}(f)$ is nilpotent.*

4 Generalizing the Landau-Miller solvability test

In this section we show that the Landau-Miller solvability test can be adapted to test if the Galois group of $f(X) \in \mathbb{Q}[X]$ is in Γ_d for constant d . Note that for $d < 5$, Γ_d is the class of solvable groups and hence our result is a generalization of the result of Landau-Miller [4]. We first recall a well-known bound on the size of primitive permutation groups in Γ_d .

Theorem 9 ([1]). *Let $G \leq S_n$ be a primitive permutation group in Γ_d for a constant d . Then $\#G \leq n^{O(d)}$.*

Theorem 10. *For constant $d > 0$, there is an algorithm that takes as input $f(X) \in \mathbb{Q}[X]$ and in time polynomial in size (f) and $n^{O(d)}$ decides whether $\text{Gal}(f)$ is in Γ_d .*

Proof. We sketch the proof. Assume without loss of generality that $f(X)$ is irreducible. Let $G = \text{Gal}(f)$ as a subgroup of $\text{Sym}(\Omega)$, where Ω is the set of roots of f . Let $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_t = \Omega$ be any maximal chain of G -blocks. The series $\{1\} = G^{\Delta_0} \triangleleft \dots \triangleleft G^{\Delta_t} = G$ gives a normal series for G . By closure properties of Γ_d , $G \in \Gamma_d$ iff $\frac{G^{\Delta_{i+1}}}{G^{\Delta_i}} \in \Gamma_d$ for each i . If G is in Γ_d so are $G_{\Delta_{i+1}}$ and $G(\Delta_{i+1}/\Delta_i)$ and hence their quotient $\frac{G_{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)}$. On the other hand since

$\frac{G^{\Delta_{i+1}}}{G^{\Delta_i}}$ is isomorphic to a subgroup of $\left(\frac{G^{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)}\right)^l$ for some l (Lemma 1), $\frac{G^{\Delta_{i+1}}}{G^{\Delta_i}} \in \Gamma_d$ if $\frac{G^{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)} \in \Gamma_d$. Hence $G \in \Gamma_d$ iff $\frac{G^{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)}$ is in Γ_d for each i . We give a polynomial-time algorithm to verify the above fact for some maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_t = \Omega$.

First, by Theorem 7 we compute $K_i = \mathbb{Q}_{\Delta_i}$ for a maximal chain of G -blocks $\{\alpha\} = \Delta_0 \subset \dots \subset \Delta_t = \Omega$. Let L_i be the fixed field of \mathbb{Q}_f with respect to the automorphisms of $G(\Delta_{i+1}/\Delta_i)$ then L_{i+1} is the normal closure of K_i over K_{i+1} . This follows because $G(\Delta_{i+1}/\Delta_i)$ is the largest proper normal subgroup of $G_{\Delta_{i+1}} = \text{Gal}(\mathbb{Q}_f/\mathbb{Q}_{\Delta_{i+1}})$. Hence $\text{Gal}(L_{i+1}/K_{i+1})$ is $\frac{G^{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)}$, and it suffices to check that each $\text{Gal}(L_i/K_i)$ is in Γ_d .

The group $\frac{G^{\Delta_{i+1}}}{G(\Delta_{i+1}/\Delta_i)}$ acts faithfully and primitively on $\Omega' = \mathcal{B}(\Delta_{i+1}/\Delta_i)$, by Lemma 1 and since Δ_i is a maximal subblock of Δ_{i+1} . If $G \in \Gamma_d$ then $[L_{i+1} : K_{i+1}] = \#\text{Gal}(L_{i+1}/K_{i+1}) \leq n^{O(d)}$ and degrees $[L_i : \mathbb{Q}]$ are all less than $n^{O(d)}$. We can use Theorem 2 to compute $\text{Gal}(L_i/K_i)$ as a multiplication table in time polynomial in size(f) and n^d for each i . We then verify that $\text{Gal}(L_i/K_i) \in \Gamma_d$ by computing a composition series for it and checking that each composition factor is in Γ_d . At any stage in the computation of $\text{Gal}(L_i/K_i)$ if the sizes of the fields becomes too large, i.e. larger than the bound of Theorem 9 we abort the computation and decide that $\text{Gal}(f)$ is not in Γ_d . Clearly, these steps can be done in polynomial time.

It follows from the proof of Theorem 10 that a prime p divides $\#\text{Gal}(f)$ if and only if it divides $[L_i : K_i]$ for some $1 \leq i \leq t$.

Theorem 11. *Given $f(X) \in \mathbb{Q}[X]$ with Galois group in Γ_d there is an algorithm running in time polynomial in size(f) and n^d that computes all the prime factors of $\#\text{Gal}(f)$.*

References

1. Lazlo Babai, Peter J. Cameron, and P. P. Palfy. On the order of primitive groups with restricted nonabelian composition factors. *Journal of Algebra*, 79:161–168, 1982.
2. Pilar Fernandez-Ferreiros and M. Angeles Gomez-Molleda. Deciding the nilpotency of the Galois group by computing elements in the centre. *Mathematics of Computation*, 73(248), 2003.
3. Susan Landau. Polynomial time algorithms for Galois groups. In John Fitch, editor, *EUROSAM 84 Proceedings of International Symposium on Symbolic and Algebraic Computation*, volume 174 of *Lecture Notes in Computer Sciences*, pages 225–236. Springer, July 1984.
4. Susan Landau and Gary L. Miller. Solvability by radicals is in polynomial time. *Journal of Computer and System Sciences*, 30:179–208, 1985.
5. Serge Lang. *Algebra*. Addison-Wesley Publishing Company, Inc, third edition, 1999.
6. Hendrik W. Lenstra Jr. Algorithms in algebraic number theory. *Bulletin of the American Mathematical Society*, 26(2):211–244, April 1992.

7. Eugene M. Luks. Isomorphism of graphs of bounded valence can be tested in polynomial time. *Journal of Computer and System Sciences*, 25(1):42–65, 1982.
8. Eugene M. Luks. Permutation groups and polynomial time computations. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 11:139–175, 1993.
9. Helmut Wielandt. *Finite Permutation Groups*. Academic Press, New York, 1964.