P. Spelier

The complexity of root-finding in orders

Bachelor thesis

June 30, 2018

Thesis supervisors: prof. em. H. W. Lenstra Jr.
dr. W. A. Kosters

Leiden University
Mathematical Institute

Contents

1 Introduction

An *order* is a commutative ring whose underlying additive group is isomorphic to $\mathbb{Z}^n$ for some integer $n \in \mathbb{Z}_{\geq 0}$ called the *rank* of that order, denoted by $\operatorname{rk} A$. An order is uniquely determined by how the standard basis vectors multiply; writing n for $\operatorname{rk} A$, we specify an order by listing structure constants $(a_{ijk})_{1 \leq i,j,k \leq n} \in \mathbb{Z}$ which describe the multiplication by $e_i \cdot e_j = \sum_{k=1}^{n} a_{ijk} e_k$ for a $\mathbb{Z}$-basis $e_1, \ldots, e_n$. This thesis treats problems about finding roots of polynomials in orders; specifically, we define the following problems.

Definition 1.1. Let $f \in \mathbb{Z}[X]$ be a polynomial. Then the problem Π_f is defined as: given as input an order A, determine whether $Z_A(f)$, the zero set of f in A, is non-empty.

Definition 1.2. Let A be an order. Then the problem Π_A is defined as: given as input a polynomial $f \in \mathbb{Z}[X]$, determine whether $Z_A(f)$, the zero set of f in A, is non-empty.

We use the terminology of polynomial, non-deterministic polynomial, and NP-complete problems to classify these problems; we refer to these classes as $\mathcal{P}, \mathcal{NP}, \mathcal{NPC}$, respectively. A short treatment of the subject can be found in Appendix A.

We say a polynomial in $\mathbb{Z}[X]$ is separable if it is separable over $\mathbb{Q}$, or equivalently if it has no double roots in $\overline{\mathbb{Q}}$. If f is non-separable or A is non-reduced, then we have little to no information about the respective problems Π_f, Π_A; one strays closely to undecidable problems, a taste of which is given in Section 6. There, we look at a more general problem.

Definition 1.3. Let $f \in \mathbb{Z}[X]$ be any polynomial. Then U_f is defined as: given as input two orders A, B, a ring homomorphism $g : A \to B$ and an element $b \in Z_B(f)$, determine whether $Z_A(f) \cap g^{-1}(b)$ is non-empty.

In that section we prove the following theorem.

Theorem 1.4. *The problem U_{X^2} is undecidable.*

For f separable or A reduced, we have more control: the two following theorems show that Π_f and Π_A are then decidable. In fact, the theorem about Π_A tells us exactly what happens for reduced A.

Theorem 1.5. *Let A be a reduced order. Then there is a polynomial time algorithm for Π_A.*

Theorem 1.6. *Let f be a separable polynomial. Then Π_f lies in $\mathcal{NP}$.*

This is proven in the beginning of Section 3. For the problems Π_f, we have the following conjecture.

Conjecture 1.7. *Let $f \in \mathbb{Z}[X]$ be separable. Then Π_f lies in $\mathcal{P}$ or in $\mathcal{NPC}$.*

Ideally, we would like a constructive proof of this statement: an algorithm that tells us for every separable polynomial f, whether Π_f admits a polynomial time algorithm or is NP-complete. In Section 3 some positive results are treated, and there are several NP-completeness theorems which work in specific cases. These NP-completeness theorems all have in common that they use a prime dividing the discriminant of $\Delta(f)$. We will use the following terminology.

Definition 1.8. Let R be any commutative ring, let $f \in R[X]$ be a polynomial, and let $a \in R$. Let $k \in \mathbb{Z}_{\geq 0}$. We say that a is a *k-fold zero* (double, triple, ...) of f in R if in $R[X]$ we have $(X-a)^k \mid f$. We say that a is a *zero of f of multiplicity k* in R if a is a k-fold zero but not a $(k+1)$-fold zero.

For the quadratic and cubic case we have proven the conjecture, culminating in the following two theorems.

Theorem 1.9. *For $f \in \mathbb{Z}[X]$ quadratic monic, we have $\Pi_f \in \mathcal{P}$ if $\Delta(f) = -4$ or $\Delta(f)$ is a square, and $\Pi_f \in \mathcal{NPC}$ otherwise.*

This statement is proven in Section 3.1.

Theorem 1.10. *For $f \in \mathbb{Z}[X]$ cubic monic, we have $\Pi_f \in \mathcal{P}$ if f is reducible, and $\Pi_f \in \mathcal{NPC}$ otherwise.*

In Section 3, enough general theorems and ad hoc lemmas are proven to classify all cubic polynomials but a small set: specifically, Proposition 3.33 tells us that for cubic monic irreducible f with discriminant not of the form $\pm 3^k$ we have $\Pi_f \in \mathcal{NPC}$. In Section 4 we treat the problem of finding all cubic monic polynomials with discriminant of the form $\pm 3^k$; in Theorem 4.1 we eventually find a minimal set S of polynomials such that for every cubic irreducible polynomial f that does not satisfy the conditions of Proposition 3.33 there exists $g \in S$ with $\Pi_f = \Pi_g$. Here equality of problems means that the sets of instances and the sets of yes-instances coincide.

Finally, we treat the remaining polynomials from S in Section 5 using ad hoc arguments, thereby completing the proof of Theorem 1.10.

An important tool in the NP-completeness proofs is a new family of algebraic problems.

Definition 1.11. Let R be a commutative ring that is finitely generated as a $\mathbb{Z}$-module, let G be a finite R-module (given by an addition table and how the generators of R act on G), and S a subset of G. Then define the problem $P^R_{G,S}$: with input $t \in \mathbb{Z}_{\geq 0}, x_* \in G^t$, the t-th Cartesian power of G, and H a submodule of G^t given by a list of generators, decide whether $(x_* + H) \cap S^t$ is non-empty. Write $P_{G,S}$ for $P^{\mathbb{Z}}_{G,S}$.

Definition 1.12. Let R be a commutative ring that is finitely generated as a $\mathbb{Z}$-module, let G be a finite R-module (given by an addition table and how the generators of R act on G), and S a subset of G. Then define the problem $\Pi^R_{G,S}$: with input $t \in \mathbb{Z}_{\geq 0}$ and H a submodule of G^t given by a list of generators, decide whether $H \cap S^t$ is non-empty. Write $\Pi_{G,S}$ for $\Pi^{\mathbb{Z}}_{G,S}$.

Remark 1.13. If R is not finitely generated as a $\mathbb{Z}$-module, we can replace it by its image in $\mathrm{End}(G)$.

Note these problems are certainly in $\mathcal{NP}$, as one can easily give an R-linear combination of the generators (and add x_* if necessary), and check that it lies in S^t. These problems are further studied in Section 2. For $R = \mathbb{Z}$, an R-module is just an abelian group; we prove two theorems that completely classify the problems $P_{G,S}, \Pi_{G,S}$, in the sense that for each problem we either have a polynomial time algorithm or a proof of NP-completeness.

Definition 1.14. With G an abelian group, $S \subset G$, we call S a *coset* if it is a coset of some subgroup of G.

Theorem 1.15. *If S is empty or a coset, then we have $P_{G,S} \in \mathcal{P}$. In all other cases, the problem is NP-complete.*

Theorem 1.16. *If S is empty or $\theta(S) := \bigcap_{a \in \mathbb{Z} | aS \subset S} aS$ is a coset, then we have $\Pi_{G,S} \in \mathcal{P}$. In all other cases, the problem is NP-complete.*

Remark 1.17. Note that if G is a group with order a prime power and S does not contain 0, then $\theta(S) = S$; if $0 \in S$, then $\theta(S) = \{0\}$, as will be proven in Lemma 2.14.

2 Group-theoretic NP-complete problems

In this section we will completely classify the group-theoretic problems $P_{G,S}$ and $\Pi_{G,S}$. Some of the lemmas we use to prove Theorems 1.15 and 1.16 we give for general $P^R_{G,S}$ (resp. $\Pi^R_{G,S}$) and some only for $P_{G,S}$ (resp. $\Pi_{G,S}$). Throughout this section, let R be a commutative ring, finitely generated as a $\mathbb{Z}$-module. All abelian groups and R-modules we consider in this section are finite.

2.1 Proof of Theorem 1.15

We first introduce four general lemmas that will help with the proof of Theorem 1.15.

Lemma 2.1. *The problem is translation invariant: we have $P^R_{G,S} \approx P^R_{G,S+g}$ for all $g \in G$ where $P_1 \approx P_2$ means that there is a reduction $P_1 \leq P_2$ and vice versa, as defined in Definition A.12.*

Proof. For the reduction $P^R_{G,S} \leq P^R_{G,S+g}$, we send an instance (t, x_*, H) to $(t, x_* + (g, \ldots, g), H)$. By symmetry, we also have $P^R_{G,S+g} \leq P^R_{G,S}$; by the definition of $\approx$, we are done. □

Lemma 2.2. *If G' is a submodule of G, then we have $P^R_{G',G'\cap S} \leq P^R_{G,S}$.*

Proof. Given an instance (t, x_*, H) of the first $P^R_{G',G'\cap S}$, we see it is also an instance of $P^R_{G,S}$, and as $H \cap S^t \subset G'^t \cap S^t = (G' \cap S)^t$, we see it is a yes-instance of the first problem exactly if it is a yes-instance of the second one. □

Lemma 2.3. *Let G' be a submodule of G and S' a subset of G, and define $S = S' + G'$. Then we have $P^R_{G/G',S'} \approx P^R_{G,S}$.*

Proof. For the reduction $P^R_{G/G',S'} \leq P^R_{G,S}$ we send an instance (t, x_*, H) to $(t, x_*, H + G'^t)$; this works exactly because of the property $S = S' + G'$. For the reduction $P^R_{G,S} \leq P^R_{G/G',S'}$, we pass everything through the map $G \to G/G'$. □

For the last lemma, we first introduce a definition.

Definition 2.4. Let G be an R-module. A *transformation* on G is a map $\varphi : G \to G$ of the form $x \mapsto c(x) + g$ with c an R-linear endomorphism of G and $g \in G$. For $S \subset G$, we write S_φ for $S \cap \varphi^{-1}(S)$.

Remark 2.5. Since R is commutative, multiplication by $r \in R$ is an R-linear endomorphism of G.

Lemma 2.6. *Let G be an R-module, S a subset of G and $\varphi : x \mapsto c(x) + g$ a transformation on G. Then $P^R_{G,S_\varphi} \leq P^R_{G,S}$.*

Proof. Let (t, x_*, H) be an instance of P_{G,S_φ}. Define $\Gamma = G^t \times G^t$ with π_1, π_2 the two projections, let $x'_* = (x_*, c(x_*) + g) \in \Gamma$ and $H' = \{(h, c(h)) \mid h \in H\} \subset \Gamma$. Note that H is naturally isomorphic to H' by $f : h \mapsto (h, c(h))$, as R is commutative. We then see that for $h \in H$ we have that $x'_* + f(h) \in S^{2t}$ if and only if $\pi_1(x'_* + f(h)), \pi_2(x'_* + f(h)) \in S^t$ if and only if $x_* + h \in S$ and $c(x_* + h) + g = \varphi(x_* + h) \in S$, which is equivalent to $x_* + h \in S_\varphi$. This shows that $(2t, x'_*, H')$ is a yes-instance of $P^R_{G,S}$ if and only if (t, x_*, H) is a yes-instance of P^R_{G,S_φ}. □

We will now prove the easy part of Theorem 1.15 with the following lemma.

Lemma 2.7. *If $S \subset G$ is empty or a coset of some subgroup of G, then $P^R_{G,S} \in \mathcal{P}$.*

Proof. As a submodule is in particular a subgroup, we have the inequality $P^R_{G,S} \leq P_{G,S}$, so it suffices to prove the lemma assuming that $R = \mathbb{Z}$. If S is empty, then the problem is easy — the answer is always no for $t > 0$ and yes for $t = 0$. If S is of the form $a + G'$ with $a \in G$ and G' a submodule of G, then by Lemma 2.3 the problem is equivalent to $P_{G/G',\{a\}}$. To solve $P_{G/G',\{a\}}$ in polynomial time, we only need to decide whether the single element $(a, \ldots, a) - x_*$ is in H: this is simply checking whether a linear system of equations over $\mathbb{Z}$ has a solution, which can be done in polynomial time as proven in §14 of [Len08]. Here we use that R is finitely generated as a $\mathbb{Z}$-module. Hence we indeed find that $P_{G/G',\{a\}}$ admits a polynomial time algorithm. □

We will prove the NP-complete part of Theorem 1.15 by induction on $|S|$. There are two base cases: $|S| = 2$ for any group G, and $|S| = |G| - 1$ for $G = \mathbb{F}_2^2$.

Lemma 2.8. *Let G be an R-module and S a subset of cardinality 2 which is not a coset of some subgroup. Then $P^R_{G,S}$ is NP-complete.*

Proof. Since S is not a coset of some subgroup, we can write $S = \{s, s + d\}$ with $s - d, s + 2d \notin S$. Because of Lemma 2.1 we can take $s = 0$. By Lemma 2.2 we are allowed to take $G = Rd$, and by renaming we can take $d = 1$, G some finite quotient of R and $S = \{0, 1\}$ with $-1, 2 \notin S$.

We reduce from 3-colorability. For a definition of this problem, see Definition A.1. Let C be our set of three colors, and let $C' = \{D \subset C \mid |D| = 2\}$ be the set of subsets of two different elements of C. Given a graph (V, E), we will construct a subgroup $H \subset \Gamma := G^{V \times C} \times G^{V \times C'} \times G^V \times G^{E \times C}$ and $x_* \in \Gamma$ such that $H + x_*$ has an element in $T := S^{V \times C} \times S^{V \times C'} \times S^V \times S^{E \times C}$ exactly if (V, E) is 3-colorable. Let $\pi_1, \pi_2, \pi_3, \pi_4$ denote the four projections from Γ on respectively $G^{V \times C}, G^{V \times C'}, G^V, G^{E \times C}$. We take H isomorphic to $G^{V \times C}$ with the isomorphism $H \to G^{V \times C}$ given by π_1, and the isomorphism the other way

given by

$$\begin{aligned}\varphi : G^{V\times C} &\to H\\ f &\mapsto (f, \rho(f), \sigma(f), \tau(f))\end{aligned}$$

where we define $\rho(f)(v, D) = \sum_{c\in D} f(v, c)$, $\sigma(f)(v) = \sum_{c\in C} f(v, c)$, $\tau(f)(e, c) = \sum_{v\in e} f(v, c)$, and $x_* = (0, 0, -1, 0)$. Note that $\pi_1(\varphi(f) + x_*)$ needs to be in $\{0, 1\}^{V\times C}$ for $\varphi(f) + x_*$ to be in T, and $\pi_1(\varphi(f) + x_*) \in \{0, 1\}^{V\times C}$ happens if and only if f itself is in $\{0, 1\}^{V\times C}$.

To prove this is truly a reduction, we interpret $\{0, 1\}^{V\times C}$ as assignments of subsets of C to the vertices V, using the bijection between $\mathrm{Fun}(V, \{0, 1\}^C)$ and $\mathrm{Fun}(V \times C, \{0, 1\}) = \{0, 1\}^{V\times C}$. A 3-coloring of (V, E) can then be equivalently redefined as such an assignment $f \in \{0, 1\}^{V\times C}$ with the property that for every vertex $v \in V$ we have $\sum_{c\in C} f(v, c) = 1$, i.e., each vertex gets a single color, and that for every $c \in C, \{i, j\} \in E$ we have $f(i)(c), f(j)(c)$ not both 1. It suffices to show that these colorings map under φ exactly to those $h \in \varphi(\{0, 1\}^{V\times C})$ with $h + x_* \in T$.

Let $f \in \{0, 1\}^{V\times C}$ be such a coloring with subsets of C, and let $h = \varphi(f)$ be the corresponding element of H. Then note that $\pi_2(h + x_*) \in \{0, 1\}^{V\times C'}$ if and only if for every two colors, at most one of them is used, or equivalently if for every vertex $v \in V$ it holds that $\sum_{c\in C} f(v, c)$ is at most 1. Also, $\pi_3(h+x_*)(v) = \sigma(f) - 1$, which is $-1 \notin S$ if $\sum_{c\in C} f(v, c) = 0$. Hence $(\pi_2(h + x_*), \pi_3(h + x_*)) \in \{0, 1\}^{V\times C'} \times \{0, 1\}^V$ if and only if for every vertex $v \in V$ we have $\sum_{c\in C} f(v, c) = 1$.

Finally, note that $\pi_4(h + x_*)(e, c)$ is in $\{0, 1\}$ exactly if the two endpoints of e do not both have color c. This completes the proof that the elements of $\{0, 1\}^{V\times C}$ that are 3-colorings correspond to $h \in H$ with $h + x_* \in T$, and hence the reduction is completed. □

Lemma 2.9. *Let G be an R-module of cardinality at least 3, and S a subset of cardinality $|G| - 1$. Then $P^R_{G,S}$ is NP-complete.*

Proof. By translating, we can assume $S = G \setminus \{0\}$. We will reduce from $|G|$-colorability.

Let (V, E) be an instance of $|G|$-colorability. Note that $G^V = \{(g_v)_{v\in V} \mid g_v \in G\}$ can be thought of as all ways of assigning elements of G to the vertices. Let f be the homomorphism from G^V to G^E defined by $(g_v)_{v\in V} \mapsto (g_u - g_v)_{(u,v)\in E}$, and note that an assignment in G^V is a $|G|$-coloring if and only if it is sent to an element of S^E. Then we can take H to be the submodule of G^E generated by the images of R-generators of G^V, of which we need at most $|G||V|$. This is a valid reduction as (V, E) will be $|G|$-colorable if and only if $H \cap S^E \neq \varnothing$; we can take x_* to be zero.

As we have $|G| \geq 3$, the $|G|$-colorability problem is NP-complete [GJ79] hence $P_{G,S}$ is NP-complete. □

2.1.1 Induction step for a special kind of group

First, we will do the induction step for a special family of finite groups: $G = \langle a, b\rangle$ with $a \neq b$ and S containing $0, a, b$, but not $a + b$.

Lemma 2.10. *Let G be a finite abelian group generated by elements $a \neq b$, and S a subset of G containing $0, a, b$ but not $a + b$. Assume Theorem 1.15 holds for all $\Pi_{G',S'}$ with $|G'| + |S'| < |G| + |S|$. Then $P_{G,S}$ is NP-complete.*

Proof. In this proof, we will heavily use Lemma 2.6. We restrict to bijective transformations of the form $x \mapsto cx + g$ with $c = \pm 1$. If φ is a transformation on G and P_{G,S_φ} is NP-complete, so is $P_{G,S}$. For the NP-completeness of the former, we only need $2 \leq |S_\varphi|$, $|S_\varphi| < |S|$ and S_φ not a coset, and then we are done by the induction hypothesis. We can also interpret this in another way: if two of the three conditions on S_φ hold, then either we are done immediately, or the third one does not hold, which gives us more information about S. If φ is bijective, then $|S_\varphi| \leq |S|$ with equality if and only if φ induces a bijection on S.

We will now prove the following claim: let $G' = \{g \in G \mid g, g + a, g + b \in S, g + a + b \notin S\}$. We already know that $0 \in G'$. We will prove that if $g \in G'$, then either $\Pi_{G,S}$ is NP-complete or $g - 2a \in G'$. For the proof, we can without loss of generality assume that $g = 0$.

We do a case distinction, based on whether $a - b$ is in S or not. First, we assume it is. Let $\varphi_1 : x \mapsto a + b - x$, and note that $a, b \in S_{\varphi_1}, 0 \notin S_{\varphi_1}$, so either we are done or we know that $a + \langle b - a\rangle \subset S$, which we now assume. Now let φ_2 be the transformation $x \mapsto a - b + x$. Note $0, a, b \in S_{\varphi_2}$, so $2 \leq |S|$ and S_{φ_2} is not a coset. This now tells us that either we are done, or $S = S_{\varphi_2}$, meaning $S = \Sigma + \langle b - a\rangle$. Writing $\Gamma = G/\langle b - a\rangle$ we see Σ is not a coset in Γ. Furthermore by Lemma 2.3 we know $P_{G,S} \approx P_{\Gamma,\Sigma}$. Since $b - a \neq 0$, we have that $|\Sigma| + |\Gamma|$ is strictly smaller than $|S| + |G|$, which means that by the induction hypothesis we know $P_{\Gamma,\Sigma}$ to be NP-complete. Hence $P_{G,S} \in \mathcal{NPC}$ as we wanted to show.

In the remaining case, we have $a - b \notin S$ and similarly we can assume that $b - a \notin S$ holds as well. Looking at $x \mapsto b - x$ or $x \mapsto a - x$ we see we can assume $\langle a\rangle, \langle b\rangle \subset S$. Now we look at $\varphi_3 : x \mapsto x - a - b$. If $-a - b \in S$, all conditions are met and we are done. So assume $-a - b \notin S$. Finally taking $\varphi_4 : x \mapsto a - b + x$, we can see that we must have $-a + \langle a - b\rangle \subset S$. We now have have $0 - 2a, a - 2a, b - 2a \in S, a + b - 2a \notin S$, hence $-2a \in G'$. This proves the claim.

Now G' is closed under $g \mapsto g - 2a$ and by symmetry also under $g \mapsto g - 2b$. As a, b are of finite order, we find $G' \subset 2G$ and hence $S = \{ka + \ell b \mid k, \ell \in \mathbb{Z}, k\ell \equiv 0 \bmod 2\}$ and $G \setminus S = a + b + \langle 2a, 2b\rangle$, meaning $G' = 2G$. Dividing out by G' and using Lemma 2.3 we see we $P_{G,S}$ is equivalent to $P_{G/G',\{0,a,b\}}$. Note that $|G/G'| = 4$; we know $0, a + b$ are different in G' as $x \in \langle 2a, 2b\rangle$ implies $x \in S$, and then a is non-zero as we have that $0 + b \in S$ but $a + b \notin S$, hence $a \notin \langle 2a, 2b\rangle$. So $G' = \mathbb{F}_2^2$ and $|S'| = |G'| - 1$. We have already proven this to be NP-complete, so we are done. □

2.1.2 Induction step for general case

Finally, we will prove Theorem 1.15 in the general case, by reducing to Lemma 2.10. For this, we first prove the following little lemma.

Lemma 2.11. *Let G be an abelian group, and S a subset of G. If S has at least three elements, and the following statement holds*

$$\forall s, a, b : (s, s+a, s+b \in S \wedge a \neq b) \Rightarrow s+a+b \in S,$$

then S is a coset.

Proof. Since the statement is translation invariant, assume $0 \in S$; we will prove that S is a subgroup. Let $\{0, x, y\}$ be a subset of S of size three. It suffices to prove that $2x, -x \in S$. Applying the property with $(s, a, b) = (0, x, y)$ we see $x+y \in S$. Then, with $(x+y, -x, -x-y)$ we see $-x \in S$ and with $(y, x, -y+x)$ we get $2x \in S$, concluding the proof. □

Now for any instance $P_{G,S}$ with S not a coset and with at least three elements, we can by contraposition of Lemma 2.11 find s, a, b with $s, s+a, s+b \in S$ and $a \neq b$ and $s+a+b \notin S$; by translating, we can assume $s = 0$. Then, we set $G' = \langle a, b \rangle$ and $S' = S \cap G'$. By Lemma 2.10, we know $P_{G',S'}$ is NP-complete, and then by Lemma 2.2 we find $P_{G,S}$ is NP-complete, as we wanted to show. This concludes the proof of Theorem 1.15.

2.2 Proof of Theorem 1.16

Theorem 1.16 follows immediately from Theorem 1.15 and a lemma that relates them. That lemma holds in the generality of R-modules, so we first will give a general definition of a special function θ.

Definition 2.12. Let G be an R-module, and S a subset of G. Then we write $\theta(S)$ for

$$\bigcap_{r \in R | rS \subset S} rS.$$

Lemma 2.13. *With G a finite R-module, $S \subset G$ we have the following equivalence of problems*

$$\Pi^R_{G,S} = \Pi^R_{G,\theta(S)} \approx \Pi^R_{R\theta(S),\theta(S)} \approx P^R_{R\theta(S),\theta(S)}$$

where $R\theta(S)$ means the R-module generated by $\theta(S)$.

Proof. We have to prove three equivalences, where the first is an equality. As defined in the appendix, two problems are the same if they have the same set of instances and the same set of yes-instances.

For the first one, note that if (t, H) is a yes-instance of $\Pi^R_{G,S}$ with certificate $h \in H \cap S^t$, then

$$\left(\prod_{r \in \mathrm{im}(R \to \mathrm{End}(G)) : rS \subset S} r \right) h$$

is in $\theta(S)^t$ as R is commutative, hence (t, H) is a yes-instance of $\Pi^R_{G,\theta(S)}$. The other way around, if (t, H) is a yes-instance of $\Pi^R_{G,\theta(S)}$, then it is a yes-instance of $\Pi^R_{G,S}$ since $\theta(S)$ is a subset of S, proving the first equality.

For the second one, write $S' = \theta(S), G' = RS'$ and note that $\Pi_{G,S'} \leq \Pi_{G',S'}$ by taking any instance H of the first problem and intersecting it with G'^t using the kernel algorithm from §14 of [Len08], since $H \cap S'^t = (H \cap G'^t) \cap S'^t$. And by Lemma 2.2, $\Pi_{G',S'} \leq \Pi_{G,S'}$ holds as well.

The real work happens in the third equivalence. Note $\Pi^R_{G',S'} \leq P^R_{G',S'}$ by taking $x_* = 0$. To show $P^R_{G',S'} \leq \Pi^R_{G',S'}$, let (t, x_*, H) be an instance of $P^R_{G',S'}$. We will construct an instance (t', H') of the first problem; we will take $t' = t + |S'|$ and $H' = H \times \{0\}^t + R(x_*, S')$.

We need to check that if (t, x_*, H) is a yes-instance, so is (t', H') and vice versa. The first implication is trivial; if $h \in H$ has $x_* + h \in S^t$, then $h' = y_* + (h, 0)$ is an element of H', and lies in S' on every coordinate, hence we see $h' \in H' \cap S'^{t'}$. For the other implication, let $h' = a_* y_* + (h, 0)$ be an element of $H' \cap S'^{t'}$. Looking at the last $|S'|$ coordinates, we see $a_* S' \subset S'$. But as $\theta(S') = \theta^2(S) = \theta(S) = S'$, we must have $a_* S' = S'$. Since a_* induces a bijection on S' and S' generates G' we see a_* is a unit in $\mathrm{End}(G')$, using the commutativity of R. Then some power of a_* is its inverse. Hence we can multiply h' with $a_*^{-1} \in \mathrm{End}(G')$, and since $a_*^{-1} S' = S'$ we see $y_* + a_*^{-1}(h, 0) \in S'^{t'}$. Restricting to the first t places, we see $x_* + h \in S'^t$, hence $(x_* + H) \cap S'^t \neq \varnothing$ as we wanted to show. □

Proof of Theorem 1.16. Theorem 1.16 now follows trivially from Theorem 1.15 and Lemma 2.13. □

As promised, we will also prove a short lemma about θ in a generalisation of the case that G has prime power cardinality.

Lemma 2.14. *Let G be an R-module such that $R/\operatorname{Ann}_R(G)$ is local. Let S be a subset of G. Then $\theta(S)$ equals $\{0\}$ if $0 \in S$ and S otherwise.*

Proof. Obviously, if $0 \in S$ then for every integer a we have $0 \in aS$ so $\{0\} \subset \theta(S)$, and $0S \subset S$ hence $\theta(S) \subset \{0\}$, proving the first part. For the second part, every element of $\mathrm{im}(R \to \mathrm{End}(G))$ is either invertible or nilpotent. If $r \in \mathrm{im}(R \to \mathrm{End}(G))$ is nilpotent, and $0 \notin S$, then $rS \not\subset S$; if it would, then $r^k S \subset S$ for every $k \in \mathbb{Z}_{>0}$, contradiction with nilpotency of r and $0 \notin S$. That means that if $rS \subset S$ then on G, we have that r induces an automorphism, and $rS \subset S$ then implies by cardinality that $rS = S$. Hence in this case $\theta(S) = S$, as we set out to prove. □

Remark 2.15. Some important examples of when the conditions are satisfied, are the case where R itself is local, and the case where $R = \mathbb{Z}$ and G has prime power cardinality.

3 General results on NP-completeness of Π_f

In this section we prove some general results on when Π_f is NP-complete. First we will give an algorithm that shows that for reduced orders A we have $\Pi_A \in \mathcal{P}$, which after a slight modification also proves that $\Pi_f \in \mathcal{NP}$ for separable $f \in \mathbb{Z}[X]$ (i.e., those with no double roots over $\mathbb{Q}$). The real work is in the proofs of NP-completeness; we will give a short explanation about the problem in general, including an explanation of when Π_f, Π_g are equal, some polynomial algorithms and a lemma that allows us to restrict to monic polynomials.

Algorithm 3.1. *We take as input A an order, $f \in \mathbb{Z}[X]$ a polynomial such that either f is separable or A is reduced. The algorithm returns whether f has a zero in A.*

1. *If f is separable, replace A by A_{sep}, the subring consisting of elements of A that are the zero of some separable polynomial in $\mathbb{Z}[X]$, using Algorithm 4.2 of [LS17].*
2. *Apply Algorithm 7.2 of [LS18] to $E := A \otimes_{\mathbb{Z}} \mathbb{Q}$ to find irreducible $g_1, \ldots, g_s$ with $E \cong \prod_{i=1}^{s} K_i$ where $K_i = \mathbb{Q}[X]/(g_i)$, together with an isomorphism $\varphi : \prod_{i=1}^{s} K_i \to E$.*
3. *Use the LLL algorithm [Len84] to find $Z_{K_i}(f)$ for every K_i.*
4. *For every $(\alpha_i)_{i=1}^{s} \in \prod_{i=1}^{s} Z_{K_i}(f)$, use the isomorphism $\prod_{i=1}^{s} K_i \to E$ to compute $\varphi((\alpha_i)_{i=1}^{s})$ with respect to the $\mathbb{Z}$-basis $e_1, \ldots, e_{\operatorname{rk} A}$ of A, and test whether all coefficients are integral. If all coefficients are integral, then f has a zero in A; the answer is yes.*
5. *If no zeroes of f in A were found in the previous step, the answer is no.*

Proposition 3.2. *The time complexity is*

$$O\left(p\left((1+\operatorname{rk} A)(1+\deg f)\log\left(2+\sum_{1\le i,j,k\le n}|a_{ijk}|\right)\right)(1+\deg f)^{|\operatorname{Spec}(A\otimes_{\mathbb{Z}}\mathbb{Q})|}\right)$$

where $p(m) = O(m^{\ell})$ for some fixed integer ℓ.

Proof. The adding of 1 or 2 at several positions in the time complexity is done to correctly handle the degenerate cases $\deg f = 0, \operatorname{rk} A = 0$. The standard operations as multiplication, addition, all take polynomial time in $(1+\operatorname{rk} A)(1+\deg f)\log\left(2+\sum_{1\le i,j,k\le n}|a_{ijk}|\right)$. Note that the s we have found in the second step equals $|\operatorname{Spec}(A \otimes_{\mathbb{Z}} \mathbb{Q})|$, and that in every field of characteristic zero f has at most $\deg f$ zeroes, so we check at most $(\deg f)^{|\operatorname{Spec}(A\otimes_{\mathbb{Z}}\mathbb{Q})|}$ candidates. Then it takes time $O(n^2)$ to apply φ, and time $O(n)$ to compute whether that zero of f in E indeed lies in A. □

Remark 3.3. A special case is where A is a domain, where Algorithm 3.1 always runs in polynomial time.

Proof of Theorem 1.5. Using Algorithm 3.1 together with Proposition 3.2 this theorem is now trivial; in fact, Algorithm 3.1 works in polynomial time even if we only fix $|\operatorname{Spec}(A \otimes_{\mathbb{Z}} \mathbb{Q})|$. □

Proof of Theorem 1.6. We use the definition of $\mathcal{NP}$ as given in Definition A.8. To prove that for f separable, Π_f lies in $\mathcal{NP}$, we need for each yes-instance A a certificate $c(A) \in A$ such that there is an algorithm that given $A, c \in A$ outputs "yes" if $c = c(A)$ and "no" if A is a no-instance, in polynomial time in the size of the input. Note that by encoding A in $\mathbb{Z}_{>0}$, this is indeed equivalent to Definition A.8. Our algorithm is very short.

1. Calculate whether $f(c) = 0$.

If we take $c(A) \in Z_A(f)$, then Algorithm 3.1 also shows that the size of $c(A)$ is polynomial in the size of the input, hence our algorithm works in polynomial time. Hence the problem Π_f lies in $\mathcal{NP}$. □

Remark 3.4. This does not necessarily work for non-separable polynomials, as then we cannot guarantee that if $Z_A(f)$ is non-empty, it contains a small element.

Lemma 3.5. *Let $f \in \mathbb{Z}[X]$ be non-zero, and let f_{mon} be its largest degree monic divisor. Then $\Pi_f = \Pi_{f_{\mathrm{mon}}}$.*

Proof. It suffices to show for any order A that $Z_A(f) \neq \varnothing$ holds if and only if $Z_A(f_{\mathrm{mon}}) \neq \varnothing$ holds. Obviously, if f_{mon} has a zero in A, then so does f. If f has a zero α in A, then, as A is an order, α is the zero of some monic polynomial g. If we then use again that A is torsion free, we see that α is a zero of the monic polynomial $\gcd(g, f)$. Any monic polynomial that divides f also divides f_{mon}, so f_{mon} has a zero in A. This in fact proves the stronger statement $Z_A(f) = Z_A(f_{\mathrm{mon}})$. □

Definition 3.6. Let $f, g \in \mathbb{Z}[X]$ be two polynomials. Then we say that f and g are *equivalent*, notation $f \sim g$, if and only if there exist ring homomorphisms $\varphi : \mathbb{Z}[X]/(f) \to \mathbb{Z}[X]/(g), \psi : \mathbb{Z}[X]/(g) \to \mathbb{Z}[X]/(f)$.

Example 3.7. For any $f \in \mathbb{Z}[X]$, we have $f \sim f(\pm X + k)$ with $k \in \mathbb{Z}$.

Example 3.8. For $n \in \mathbb{Z}_{\geq 1}$, write $n = 2^r s$ with s odd. Then $X^n + 1 \sim X^{2^r} + 1$.

The following lemma motivates this definition.

Lemma 3.9. *Let $f, g \in \mathbb{Z}[X]$ be two monic polynomials. Then $\Pi_f = \Pi_g$ holds if and only if f and g are equivalent.*

Proof. First note that the functor **Rings** $\to$ **Sets** : $R \mapsto Z_R(f)$ is represented by $\mathbb{Z}[X]/(f)$, i.e., there is a functorial bijection between $Z_R(f)$ and $\mathrm{Hom}_{\mathbf{Rings}}(\mathbb{Z}[X]/(f), R)$. It is given by sending a zero α of f in R to $\mathbb{Z}[X]/(f) \to R, g + (f) \mapsto g(\alpha)$, and the other way around by sending $\varphi : \mathbb{Z}[X]/(f) \to R$ to $\varphi(X)$.

First we will show the implication $\Pi_f = \Pi_g \Rightarrow f \sim g$. Note that $\mathbb{Z}[X]/(f)$ is an order, and in fact a yes-instance of Π_f so also of Π_g, which means that $\mathrm{Hom}(\mathbb{Z}[X]/(g), \mathbb{Z}[X]/(f))$ is non-empty. By symmetry, there is also a map the other way, hence we have $f \sim g$.

Now we will prove the other implication. Assume $f \sim g$ holds, and let A be any order. Consider the composition map $\mathrm{Hom}(\mathbb{Z}[X]/(g), \mathbb{Z}[X]/(f)) \times \mathrm{Hom}(\mathbb{Z}[X]/(f), A) \to \mathrm{Hom}(\mathbb{Z}[X]/(g), A)$. As $\mathrm{Hom}(\mathbb{Z}[X]/(g), \mathbb{Z}[X]/(f))$ is non-empty, we now have that if A is a yes-instance of Π_f, it is a yes-instance of Π_g and by symmetry also vice versa. □

Now we will treat the few polynomial cases known so far. We start with a rather trivial lemma. We refer to Definition A.2 for the definition of trivial problems; note that they always lie in $\mathcal{P}$.

Lemma 3.10. *Let $f \in \mathbb{Z}[X]$ be a polynomial with $Z_{\mathbb{Z}}(f) \neq \varnothing$. Then Π_f is trivial.*

Proof. As $\mathbb{Z}$ is an initial object in the category of rings, for any order A we have a morphism $\mathbb{Z} \to A$ and hence a morphism $Z_{\mathbb{Z}}(f) \to Z_A(f)$, hence $Z_A(f) \neq \varnothing$. □

There is one family of polynomials for which a non-trivial polynomial time algorithm is known, as proven in the following theorem.

Theorem 3.11. *Let $n \in \mathbb{Z}_{\geq 1}$. Then for $f = X^n + 1$ we have $\Pi_f \in \mathcal{P}$.*

Proof. A zero of f is necessarily a root of unity. By Theorem 1.2 of [LS17], we can find a set of generators S for $\mu(A)$, the group of roots of unity. Then asking whether f has a root in A is asking whether in $\mu(A)$ the element -1 is an n-th power, i.e., if -1 is in the subgroup generated by $\{s^n \mid s \in S\}$. Theorem 1.3 of the mentioned article allows us to compute this in polynomial time, hence this gives a polynomial time algorithm for Π_f. □

Remark 3.12. Note that we have found a polynomial time algorithm for the n-th cyclotomic polynomial, where n is a power of two. Strangely enough, Theorem 3.15 will tell us that for $X^2 + X + 1$, the third cyclotomic polynomial, the problem is NP-complete as $(X + 1)^2 + (X + 1) + 1 \equiv X^2 \bmod 3$.

Now we will prove two general theorems that can be used to classify problems Π_f as NP-complete. First we will state a general lemma that we will use multiple times to prove NP-completeness; although it cannot be applied in every proof in Section 3.1 and Section 3.2, the general idea will be used in all proofs.

Lemma 3.13. *Let $f \in \mathbb{Z}[X]$ be a polynomial, A an order, $\psi : A \to B$ a surjective ring homomorphism with B finite, R a subring of B, and G an R-module inside B. Assume that $G \cap R = 0$ and the multiplication on B restricted to $G \times G$ is the zero map. Let $a \in R$ such that $\psi(Z_A(f)) = a + S$ with $S \subset G$. Then $\Pi^R_{G,S} \leq \Pi_f$.*

Proof. Let (t, H) be an instance of $\Pi^R_{G,S}$. Note that R has a unique R-linear ring homomorphism into B^t, the diagonal map. We write this as an inclusion; in that way, we have $R[H] \subset B^t$. By the condition that multiplication on G is the zero map and $G \cap R = 0$ we have that $R[H]$ is as an R-module isomorphic to $R \oplus H$. Now we see that $R[H] \cap (a + S^t)$ is in bijection with $H \cap S^t$, by the map $x \mapsto x - a$. Let $A_H \subset A^t$ be the inverse image of $R[H]$ with respect to the map $A^t \to B^t$; as $R[H]$ is a ring, so is A_H. We see that we end up with a surjective map $Z_{A_H}(f) \to H \cap S^t$. A surjective map has the property that the domain is empty if and only if the codomain is empty, hence $H \cap S^t \neq \varnothing$ if and only if $Z_{A_H}(f) \neq \varnothing$. So we produce A_H as an instance of Π_f, completing the reduction. □

Remark 3.14. Note that if $R = \mathbb{Z} \cdot 1 \subset B$, then an R-module is just an abelian group G that satisfies $|R|G = 0$, with no further structure. Hence then $\Pi^R_{G,S}$ equals $\Pi_{G,S}$.

Theorem 3.15. *Let f be a monic irreducible polynomial over $\mathbb{Z}$ of degree $n > 1$, and $p \nmid n$ a prime such that $f \equiv X^n \bmod p$. Then Π_f is NP-complete.*

Proof. We will use Lemma 3.13.

Let $\alpha_1, \ldots, \alpha_n$ be the zeroes of f in $\overline{\mathbb{Q}}$, and let A be the order $\mathbb{Z}[\alpha_1, \ldots, \alpha_n]$. Let I be the A-ideal generated by $\alpha_1, \ldots, \alpha_n$.

Now let $B := A/(pA + I^2)$, let $R = \mathbb{F}_p \subset B$, let $\overline{\alpha_i}$ be the image of α_i in B, let $G = \langle \overline{\alpha_1}, \ldots, \overline{\alpha_n} \rangle$, and $S = \{\overline{\alpha_1}, \ldots, \overline{\alpha_n}\}$. We will prove that $\Pi^{\mathbb{F}_p}_{G,S} = \Pi_{G,S}$ is NP-complete.

As G is a group with order a power of p, by Theorem 1.16 and Lemma 2.14 it suffices to check that $0 \notin S$ and that S is not a coset.

By the condition on $f \bmod p$, we see that $J := I + pA$ is nilpotent in A/pA. As $A = \mathbb{Z}[I]$ we have $A/pA = \mathbb{F}_p[J]$. Since $n > 1$, we have $\operatorname{rk} A \geq 2$ hence $|A/pA| \geq p^2$, which implies that $J \neq \{0\}$. As J is nilpotent, that implies that $J^2 \subsetneq J$. So at least one of $\overline{\alpha_1}, \ldots, \overline{\alpha_n}$ is non-zero, and by the transitivity of the Galois action, all of them are non-zero.

We have to prove that S is not a coset. Since a coset has p-power cardinality, it suffices to prove that $|S| > 1$ and $|S| \mid n$. Assume $|S| = 1$. Then in B, we have $\overline{\alpha_1} = \cdots = \overline{\alpha_n}$; as $\alpha_1 + \cdots + \alpha_n \in p\mathbb{Z}$ we have $n\overline{\alpha_1} = 0$. We know n is a unit in $\mathbb{F}_p$, hence $\overline{\alpha_1} = 0$, contradiction. The fact $|S| \mid n$ follows immediately from the action of the Galois group. As said, we find that S is not a coset.

Now note that $G \cdot G = 0$ and $G \cap \mathbb{F}_p = 0$, so with $a = 0$ all of the conditions of Lemma 3.13 are satisfied. Together with the NP-completeness of $\Pi_{G,S}$, this implies that Π_f is NP-complete. □

Proposition 3.16. *Let f be a monic irreducible polynomial over $\mathbb{Z}$ of degree $n > 1$ and $p \mid \Delta(f)$ an odd prime. Let A be an order with α_1, α_2 two distinct zeroes of f in A. Further, let $a \in \overline{\mathbb{F}_p}$, let $\mathbb{F}_q = \mathbb{F}_p(a)$ and let $\psi : \mathbb{Z}[\alpha_i, \alpha_j] \to \mathbb{F}_q[X]/(X^2) = \mathbb{F}_q[\varepsilon]$ be a ring homomorphism such that $\psi(\alpha_1) = a + \varepsilon$ and $\psi(\alpha_2) = a - \varepsilon$. Finally assume, that we have that $Z_A(f) = \{\alpha_1, \alpha_2\}$ or we have both that $Z_{\mathbb{Z}[\alpha_1]}(f) = \{\alpha_1\}$ and that all zeroes in $Z_A(f) \setminus \{\alpha_1, \alpha_2\}$ get sent under $(\mathbb{F}_q[\varepsilon] \to \mathbb{F}_q) \circ \psi$ to something different from a. Then Π_f is NP-complete.*

Proof. In the case that $Z_A(f) = \{\alpha_1, \alpha_2\}$, we can directly use Lemma 3.13. Let $B = \mathbb{F}_q[\varepsilon]$, let $R = \mathbb{F}_q$. Now let $G = \varepsilon\mathbb{F}_q, S = \{\pm\varepsilon\}$. As $G \cdot G = 0$ and $G \cap R = 0$, all of the conditions of Lemma 3.13 hold, hence $\Pi^R_{G,S} \leq \Pi_f$. By Lemma 2.8 we see that $\Pi^R_{G,S}$ is NP-complete as $p > 2$, hence so is Π_f.

If we do not have $Z_A(f) = \{\alpha_1, \alpha_2\}$, we have to slightly change the proof as we cannot use Lemma 3.13 directly. We still reduce from the NP-complete problem $\Pi^{\mathbb{F}_q}_{\mathbb{F}_q,\{\pm1\}}$. Let (t, H) be an instance of $\Pi^{\mathbb{F}_q}_{\mathbb{F}_q,\{\pm1\}}$. Let $B = \mathbb{F}_q[\varepsilon]$, let $R = \mathbb{F}_q$, and let $R_H = R[\varepsilon\mathbb{F}_q \times \varepsilon H] \subset B^{t+1}$. Let $\varphi : \mathbb{Z}[\alpha_1] \times A^t \to B^{t+1}$ be ψ on every coordinate, and let $A_H = \varphi^{-1}(R_H)$. We want to find which elements in R_H are the image under φ of a zero of f. Such an element is of the form $x + \varepsilon(y, h)$ with $x \in R \subset B^t, y \in \mathbb{F}_q$ and $h \in H$. Since $Z_{\mathbb{Z}[\alpha_1]}(f) = \{\alpha_1\}$, on the first coordinate we must get $a + \varepsilon$, meaning $x = a$ and $y = 1$. On the last t coordinates, we then have $a + \varepsilon h$. By assumption any zero $\alpha \in Z_A(f) \setminus \{\alpha_1, \alpha_2\}$ is sent under ψ to $b + c\varepsilon$ with $b \neq a$. Hence the fact that $a + \varepsilon(y, h)$ is the image under φ of some zero of f in A_H, is equivalent to it lying in the image under φ of some element in $\{\alpha_1\} \times \{\alpha_1, \alpha_2\}^t$. So we see that $Z_{A_H}(f)$ is non-empty if and only if $H \cap S^t$ is non-empty. □

For the proof of the second general theorem we first state a definition, following Section 1.6 of [Gio13].

Definition 3.17. Let R be a commutative ring, and fix $f \in R[X]$ monic of degree n. Then we define $A_0 = R, f_0 = f$ and recursively for $0 \leq i < n$ we define $A_{i+1} = A_i[x_{i+1}]/(f_i(x_{i+1})), \alpha_{i+1} = \overline{x_{i+1}} \in A_{i+1}$ and $f_{i+1}(X) = \frac{f_i(X)}{X - \alpha_{i+1}}$ as element of $A_{i+1}[X]$.

We will only use this definition in the case $R = \mathbb{Z}$. Note that if the Galois group of f over $\mathbb{Q}$ is S_n, then A_i is isomorphic to $\mathbb{Z}[\alpha_1, \ldots, \alpha_i]$ where $\alpha_1, \ldots, \alpha_n$ are the zeroes of f in $\overline{\mathbb{Q}}$. For any other Galois group, this is never the case for $i = n$, as the rank of A_n is $n!$ while the rank of $\mathbb{Z}[\alpha_1, \ldots, \alpha_n]$ is $|\operatorname{Gal}(f)| < n!$. Furthermore, note that in $A_i[X]$ we have $\prod_{j=1}^{i}(X - \alpha_i) \mid f$, and by Theorem 1.6.7 of [Gio13], the ring A_i is universal with this property, i.e., in the case $R = \mathbb{Z}$ we have that A_i represents the functor $S \mapsto \{(\alpha_1, \ldots, \alpha_i) \in S^i \mid \prod_{j=1}^{i}(X - \alpha_j)$ divides f in $S[X]\}$.

Theorem 3.18. *Let f be a monic irreducible polynomial over $\mathbb{Z}$ of degree $n \geq 2$ with either Galois group acting triply transitively on $Z_{\overline{\mathbb{Q}}}(f)$ or $n = 2$. Let p be*

an odd prime factor of $\Delta(f)$. Assume that either $n \in \{2\} \cup \mathbb{Z}_{\geq 4}$ or we have both that n equals 3 and that f has a zero of multiplicity 2 modulo p. Then Π_f is NP-complete.

Proof. This proof works by showing that the conditions of Proposition 3.16 hold, where we choose $a \in \overline{\mathbb{F}_p}$ to be a zero of f of multiplicity at least 2 (or exactly 2 if $n = 3$). Write $\mathbb{F}_q$ for $\mathbb{F}_p(a)$.

First, we construct the map ψ as needed. Our A will be A_2. Let α_1, α_2 be the roots x_1, x_2 of f in A_2. As $\mathrm{Gal}(f)$ acts triply transitively on $Z_{\overline{\mathbb{Q}}}(f)$ or $n = 2$, we have that A_2 is isomorphic to $\mathbb{Z}[\alpha, \beta]$ where α, β are any two zeroes of f in $\overline{\mathbb{Q}}$. We use the universal property of A_2 to construct $\psi : A_2 \to \mathbb{F}_q[\varepsilon]$ with $\psi(\alpha_1) = a + \varepsilon, \psi(\alpha_1) = a - \varepsilon$; the universal property implies that this map exists, since $(X - (a + \varepsilon))(X - (a - \varepsilon)) = (X - a)^2$, which divides f modulo p exactly because a is a double zero of f.

Now, let $n = 2$ or $n > 3$. Then the conditions tell us that A_2 only contains α_1, α_2 and no other roots of f, which means the conditions for Proposition 3.16 hold. If $n = 3$, then as f_1 is irreducible over $\mathbb{Z}[\alpha_1]$ we have that $Z_{\mathbb{Z}[\alpha_1]}(f) = \{\alpha_1\}$, and as a is a root of multiplicity 2, we can also apply Proposition 3.16. □

Remark 3.19. It immediately follows that for $n = 2, n > 3$ the only polynomials with Galois group S_n for which we have not proven NP-completeness yet, are those with discriminant $\pm 2^k$. For $n = 3$, Theorem 1.10 tells us that for any polynomial f with Galois group S_3 the problem Π_f is NP-complete. Since for a fixed degree n, a monic polynomial has Galois group S_n with probability 1, and discriminant not of the form $\pm 2^k$ with probability 1, we see that if the degree $n \geq 2$ is fixed, then the problem Π_f is almost surely NP-complete.

3.1 Quadratic polynomials

In this section we will fully treat the quadratic case. Let $f \in \mathbb{Z}[X]$ be a quadratic monic polynomial. If f is reducible, then by Lemma 3.10 the problem is in $\mathcal{P}$. If f is irreducible and there is an odd prime dividing $\Delta(f)$, then we can use Theorem 3.15 or Theorem 3.18 to find that Π_f is NP-complete. The only case that remains is f irreducible, $\Delta(f) = \pm 2^k$ with $k \in \mathbb{Z}_{\geq 0}$. Since an irreducible polynomial of degree ≥ 2 has a prime factor in its discriminant by Minkowski's theorem, we have $k \geq 1$. Hence f has a double root modulo 2, so the coefficient of X is even. By translating, we may assume $f = X^2 - a$, with discriminant $4a$.

Hence the only polynomials that remain are of the form $X^2 - a$ with $|a|$ a power of 2 and a not a square. For $X^2 + 1$, we have given a polynomial time algorithm in Theorem 3.11. In all other cases, we have $2 \mid a$ and we can use the following theorem.

Theorem 3.20. *Let $f = X^2 - a$ with $2 \mid a$ and a not a square. Then Π_f is NP-complete.*

Proof. We will reduce from $\Pi_{\mathbb{Z}/8\mathbb{Z},\{\pm 1\}}$, which is NP-complete by Theorem 1.16. Let $A = \mathbb{Z}[\sqrt{a}]$ and $B = \mathbb{Z}/8\mathbb{Z}[\sqrt{a}], R = \mathbb{Z}/8\mathbb{Z} \subset B$; let $S \subset R$ be $\{\pm\sqrt{a}\}$

and $G = \sqrt{a}\mathbb{Z}/8\mathbb{Z} \subset B$. Now let (t, H) be an instance of $\Pi_{\mathbb{Z}/8\mathbb{Z},\{\pm 1\}}$, let $C = \{(x_1, \ldots, x_t) \in B^t \mid x_1 \equiv \cdots \equiv x_t \bmod 2B\}$. As $\sqrt{a} \equiv -\sqrt{a} \bmod 2B$, we have $S^t \subset C$. Letting $H' = H\sqrt{a} \cap C$, we see that (t, H) is a yes-instance if and only if $H' \cap S^t$ is non-empty. Note that $H' = \langle (1 + 2B^t) \cap H' \rangle \cup (2B^t \cap H')$; if $H' \subset 2B^t$, the answer is trivially no. Otherwise, we see $H' = \langle (1 + 2B^t) \cap H' \rangle$ hence $H' \cdot H' \cdot H'$ is generated by elements of the form $\prod_{i=1}^{3}(\sqrt{a} + \sqrt{a}x_i)$ with $x_1, x_2, x_3 \in 2B^t$ and $\sqrt{a}+\sqrt{a}x_1, \sqrt{a}+\sqrt{a}x_2, \sqrt{a}+\sqrt{a}x_3 \in H'$. Using that $4a = 0$ in B, this product equals $a\sqrt{a}(1 + x_1 + x_2 + x_3)$. Now using that $ax_3 = -ax_3$, we see this equals $a(\sqrt{a}+\sqrt{a}x_1+\sqrt{a}+\sqrt{a}x_2-(\sqrt{a}+\sqrt{a}x_3)) \in H'$. This implies that $H' \cdot H' \cdot H'$ is a subset of H', which means that $\mathbb{Z}/8\mathbb{Z}[H']$ is as an additive group $\mathbb{Z}/8\mathbb{Z} + H' + H' \cdot H'$, with $\mathbb{Z}/8\mathbb{Z}[H'] \cap S^t = H' \cap S^t$. Then we define A_H to be the inverse image under the natural map $A^t \to B^t$ of $\mathbb{Z}/8\mathbb{Z}[H']$, and we see that A_H contains a zero of f exactly if $H' \cap S^t = 0$, completing the reduction. □

This concludes the proof of Theorem 1.9.

3.2 Cubic polynomials

In this part we will prove NP-completeness for many monic cubic polynomials. Note that reducible monic cubic polynomials have a zero in $\mathbb{Z}$, and hence Π_f is trivial according to Lemma 3.10. Therefore we will consider only irreducible monic polynomials. To concisely state our many lemmas, we first state a definition, using the terminology of Definition 3.17.

Definition 3.21. Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic. We define the $\mathbb{Z}$-*rank* of f, written $\mathrm{rk}_{\mathbb{Z}}(f)$, to be the rank of the smallest A_i which contains three zeroes of f.

Note that we have $\mathrm{rk}_{\mathbb{Z}}(f) = 6$ if f_1 is irreducible over A_1, and $\mathrm{rk}_{\mathbb{Z}}(f) = 3$ otherwise.

If the $\mathbb{Z}$-rank of some cubic monic irreducible polynomial f is 6 and $\mathrm{Gal}(f) = \mathrm{A}_3$, then one can check that $A_2 \otimes_{\mathbb{Z}} \mathbb{Q}$ contains 9 zeroes of f; the following important lemma controls the number of zeroes in A_2.

Lemma 3.22. *Let f be monic irreducible cubic, with $\mathbb{Z}$-rank 6. Then f has exactly three zeroes in A_2.*

Proof. If $\mathrm{Gal}(f)$ is S_3, then the statement is trivial, as then A_2 is isomorphic to $\mathbb{Z}[Z_{\overline{\mathbb{Q}}}(f)]$. From now on, assume $\mathrm{Gal}(f) = \mathrm{A}_3$, with $\mathrm{Gal}(\mathbb{Q}(\alpha)/\mathbb{Q}) = \langle\sigma\rangle$. Note f has at least three zeroes $\alpha_1, \alpha_2, \alpha_3$ in A_2, obtained by the construction of A_2. Let $\alpha, \beta = \sigma(\gamma), \gamma = \sigma(\beta)$ be the zeroes of f in $\overline{\mathbb{Q}}$, where we pick our algebraic closure of $\mathbb{Q}$ such that $\alpha = \alpha_1$. Note that $A_2 \otimes_{\mathbb{Z}} \mathbb{Q}$ is naturally isomorphic to $\mathbb{Q}(\alpha) \times \mathbb{Q}(\alpha)$, with $\alpha_2 \otimes 1$ being sent to (β, γ). Under this isomorphism, $A_1 = \mathbb{Z}[\alpha]$ is sent to $\mathbb{Z}[\alpha] \subset \mathbb{Q}(\alpha) \times \mathbb{Q}(\alpha)$ by the diagonal. All in all we have the injections

in the following diagram

$$\begin{array}{ccc} \mathbb{Z}[\alpha,\alpha_2] & \longrightarrow & \mathbb{Q}(\alpha)\times\mathbb{Q}(\alpha) \\ \uparrow & & \uparrow \\ \mathbb{Z}[\alpha] & \longrightarrow & \mathbb{Q}(\alpha) \end{array}$$

Now we define an equivalence relation on the nine zeroes of f in $\mathbb{Q}(\alpha)\times\mathbb{Q}(\alpha)$ by $x \sim y$ if the fields they generate inside $\mathbb{Q}(\alpha)\times\mathbb{Q}(\alpha)$ are the same. The nine zeroes fall into three equivalence classes: the corresponding fields are $K_i = \{(x,y) \in \mathbb{Q}(\alpha)\times\mathbb{Q}(\alpha) \mid y = \sigma^i(x)\}$ for $i = 0,1,2$, each isomorphic to $\mathbb{Q}(\alpha)$. Specifically, we see that $\alpha_2 \not\sim \alpha$, and hence $\alpha_3 \not\sim \alpha, \alpha_2$ by the symmetry. We claim that of each equivalence class, only one zero lies in $\mathbb{Z}[\alpha,\alpha_2]$. By the Galois action, we only need to prove it for the equivalence class containing α, consisting of $(\alpha,\alpha),(\beta,\beta),(\gamma,\gamma)$. Note that $\mathbb{Z}[\alpha,\alpha_2]$ has basis $1,\alpha_2$ as $\mathbb{Z}[\alpha]$-module and $\mathbb{Q}(\alpha)\times\mathbb{Q}(\alpha)$ has basis $1,\alpha_2$ as $\mathbb{Q}(\alpha)$-module. So $\mathbb{Z}[\alpha,\alpha_2]\cap\mathbb{Q}(\alpha) = \mathbb{Z}[\alpha]$, and hence β,γ are not in $\mathbb{Z}[\alpha,\alpha_2]$ as they are not in $\mathbb{Z}[\alpha]$. This proves that each equivalence class contains only one zero in $\mathbb{Z}[\alpha,\alpha_2]$, so $\mathbb{Z}[\alpha,\alpha_2] = A_2$ has exactly three zeroes of f. □

Lemma 3.23. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic of $\mathbb{Z}$-rank 6. If there is a prime p with $p \neq 2$ such that f has a zero of multiplicity 2 modulo p, then Π_f is NP-complete.*

Proof. If $\mathrm{Gal}(f) = \mathrm{S}_3$, then this is a special case of Theorem 3.18. For $\mathrm{Gal}(f) = \mathrm{A}_3$, we can prove NP-completeness directly from Proposition 3.16 and Lemma 3.22. Write $f \equiv (X-a)^2(X-b) \bmod p$ with $a \not\equiv b \bmod p$. As in the proof of Theorem 3.18 we take $A = A_2$ and use that f has exactly three zeroes $\alpha_1,\alpha_2,\alpha_3$ in A. Then we construct by the universal property of A a ring homomorphism $\psi : A \to \mathbb{F}_3[\varepsilon]$ with $\psi(\alpha_1) = a+X$ and $\psi(\alpha_2) = a-X$. By $\alpha_1+\alpha_2+\alpha_3 = 2a+b$ we have $\psi(\alpha_3) = b$. Also, note that $Z_{\mathbb{Z}[\alpha_1]}(f) = \{\alpha_1\}$ since f has $\mathbb{Z}$-rank 6. □

Lemma 3.24. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic. If there is a prime p with $p \neq 3$ and $f \equiv (X-a)^3 \bmod p$ for some $a \in \mathbb{F}_p$, then Π_f is NP-complete.*

Proof. Special case of Theorem 3.15. □

Lemma 3.25. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic of $\mathbb{Z}$-rank 3. Then there is no prime p such that f has a zero of multiplicity 2 modulo p.*

Proof. Let α be one of the zeroes of f in $\overline{\mathbb{Q}}$. Note that since $\mathbb{Z}[\alpha]$ already contains the other two zeroes of f, the Galois group of f acts on $\mathbb{Z}[\alpha]$. Let p be a rational prime. We will show that $\mathrm{Gal}(f)$ acts transitively on the primes above p; if it would not, there would be $\mathfrak{p} \mid p, \mathfrak{q} \mid p$ with $\mathfrak{p},\mathfrak{q}$ in different $\mathrm{Gal}(f)$-orbits. Then using the Chinese remainder theorem, there is an element x such that $x \in \mathfrak{p}$, but not in any $\sigma^{-1}(\mathfrak{p})$ for $\sigma \in \mathrm{Gal}(f)$. Then $N_{K/\mathbb{Q}}(x) = \prod_{\sigma\in\mathrm{Gal}(f)} \sigma(x)$ is contained

in $\mathfrak{p}$ but not in $\mathfrak{q}$; but $N_{K/\mathbb{Q}}(x) \in \mathbb{Z}$ and both $\mathfrak{p}, \mathfrak{q}$ have intersection (p) with $\mathbb{Z}$, so that is impossible.

Then, the ramification indices of the primes over p must be equal. We conclude the proof by observing that 2 does not divide 3. □

Remark 3.26. This lemma becomes false if one replaces the rank condition by the condition $\mathrm{Gal}(f) = \mathrm{A}_3$. For example take X^3+6X^2-X-5 with discriminant 65^2; modulo 5 this factors as $X(X+3)^2$.

Proposition 3.27. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic. If $\Delta(f) \neq \pm 2^k 3^\ell$ with $k, \ell \in \mathbb{Z}_{\geq 0}$, then Π_f is NP-complete. Further, if $f \equiv (X-a)^3 \bmod 2$ or $f \equiv (X-a)(X-b)^2 \bmod 3$ with $b \not\equiv a \bmod 3$, then Π_f is NP-complete as well.*

Proof. If $\Delta(f)$ contains a prime factor $p > 3$, then f has a zero of multiplicity 3 or 2 modulo p. In the first case, Π_f is NP-complete by Lemma 3.24. In the second case, by contraposition of Lemma 3.25 we have $\mathrm{rk}_{\mathbb{Z}}(f) = 6$ and hence by Lemma 3.23 the problem Π_f is NP-complete.

Furthermore, if $f \equiv (X-a)^3 \bmod 2$ or $f \equiv (X-a)(X-b)^2 \bmod 3$ with $b \neq a$ then we can again use respectively Lemma 3.24 or Lemma 3.25 followed by Lemma 3.23. □

Lemma 3.28. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic such that f has a zero of multiplicity 2 modulo 2 and one of multiplicity 3 modulo 3. Then Π_f is NP-complete.*

Proof. Note that by Lemma 3.25 the $\mathbb{Z}$-rank of f is 6. Let A be the A_2 corresponding to f, and let $\alpha_1, \alpha_2, \alpha_3$ be the three zeroes of f in A_2 given by Lemma 3.22. As in the proof of Theorem 3.18 using that f has a zero a of order 2 modulo 2, we find that for any $t \in \mathbb{Z}_{\geq 0}$ there is a homomorphism $\varphi : A_1 \times A_2^t \to \mathbb{F}_2$ obtained from applying on each coordinate the morphism $\psi : A_2 \to \mathbb{F}_2$ where one sends both α_1 and α_2 to a. Then one can check that under ψ the zero α_3 is sent to $a+1$. We let $A' \subset A_1 \times A_2^t$ be the inverse image of $\mathbb{F}_2 \subset \mathbb{F}_2^{t+1}$; the zeroes of f in A' are exactly $\{\alpha_1\} \times \{\alpha_1, \alpha_2\}^t$. Let a' be a zero of f of multiplicity 3 modulo 3 and let ψ' be the ring homomorphism $A_2 \to \mathbb{F}_3[\varepsilon]$ given by $\psi'(\alpha_1) = a + \varepsilon, \psi'(\alpha_2) = a - \varepsilon$. We reduce from $\Pi_{\mathbb{F}_3, \{\pm 1\}}$; letting (t, H) be an instance of $\Pi_{\mathbb{F}_3, \{\pm 1\}}$ and $R_H = \mathbb{F}_3[\varepsilon \mathbb{F}_3 \times \varepsilon H]$ we see that the inverse image of R_H with respect to $A' \to \mathbb{F}_3[\varepsilon]^t$ contains a zero of f if and only if $H \cap \{\pm 1\}^t$ is non-empty, completing the reduction. As $\Pi_{\mathbb{F}_3, \{\pm 1\}}$ is NP-complete, this completes the proof. □

Lemma 3.29. *There are no irreducible cubic polynomials with discriminant $\pm 2^k$ with $k \in \mathbb{Z}_{\geq 0}$.*

Proof. Let f be such a polynomial — we will derive a contradiction. Let $\mathbb{Z}[\alpha] = \mathbb{Z}[X]/(f)$ with $\alpha = \overline{X}$, let $K = \mathbb{Q}(\alpha)$ and let Δ be the discriminant of K (and note that Δ is also a power of 2, up to sign). We now have the following inclusion

of fields:

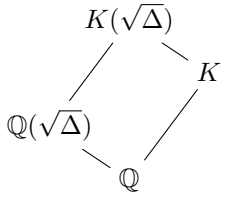

Using that the Minkowski bound is at least 1 (see for example Corollary 5.10 of [Ste17]), we find Δ is in absolute value at least 13. However, the discriminant of $\mathbb{Q}(\sqrt{\Delta})$ is one of $1, -4, \pm 8$ (it is 1 exactly if $\mathrm{Gal}(f) = \mathrm{A}_3$). From this we will derive a contradiction, using the discriminant of $K(\sqrt{\Delta})$ in between. We do this by looking at the splitting behavior of (2).

Since $2 \mid \Delta$, the prime (2) ramifies over $K/\mathbb{Q}$. We see that in $\mathcal{O}_K$ either $(2) = \mathfrak{p}^3$ or $(2) = \mathfrak{p}^2\mathfrak{q}$ with $\mathfrak{p} \neq \mathfrak{q}$. In the first case, since then $\mathfrak{p}$ ramifies tamely, we have $2^2 \| \Delta$, so $\Delta = \pm 4$, contradiction with the upper bound $|\Delta| \geq 13$ we found earlier. The other case is a bit more complex. Note that in this case f has Galois group S_3 as $K/\mathbb{Q}$ is clearly not Galois; in a Galois extension, every ramification index of a prime over 2 is equal. That $K/\mathbb{Q}$ is not Galois implies that the discriminant of $\mathbb{Q}(\sqrt{\Delta})$ is not 1, so it is divisible by 2. Hence (2) factors as $\mathfrak{r}^2$ in $\mathbb{Q}(\sqrt{\Delta})$. Since $K(\sqrt{\Delta})/\mathbb{Q}$ is a Galois extension, we see that in $K(\sqrt{\Delta})$ we have $(2) = (\mathfrak{tuv})^2$ with $\mathfrak{tuv} = \mathfrak{r}$ and $\mathfrak{tu} = \mathfrak{p}$ and $\mathfrak{v}^2 = \mathfrak{q}$. We see that $K(\sqrt{\Delta})/\mathbb{Q}(\sqrt{\Delta})$ is unramified, and hence $\Delta_{K(\sqrt{\Delta})} = \Delta^3_{\mathbb{Q}(\sqrt{\Delta})}$. Note we also have $\Delta_{K(\sqrt{\Delta})} \geq \Delta^2$. Now we make another small case distinction: if $\Delta_{\mathbb{Q}(\sqrt{\Delta})} = -4$, we find $|\Delta| \leq 8$, contradiction. If $\Delta_{\mathbb{Q}(\sqrt{\Delta})} = \pm 8$, we find $|\Delta| \leq 22$, but Δ is a power of two with an odd number of factors 2 and it is in absolute value at least 13, and we again arrive at contradiction.

We conclude that there is no cubic number field with discriminant $\pm 2^k$, so also no irreducible cubic polynomial with such a discriminant. □

We again summarise the results in a proposition.

Proposition 3.30. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic. If $\Delta(f)$ has a prime factor other than 3 or f does not have a triple zero modulo 3, then Π_f is NP-complete.*

Proof. If $\Delta(f)$ has a prime factor bigger than 3, the problem is already NP-complete by Proposition 3.27; from now on, assume that it does not have such a prime factor. If $\Delta(f)$ is divisible by 2, then by contraposition of Lemma 3.29 it is also divisible by 3, and unless f has a zero of multiplicity 2 modulo 2 and a zero of multiplicity 3 modulo 3, the problem is NP-complete by Proposition 3.27; if we are in that case, we can use Lemma 3.28 to prove NP-completeness. This proves the first part of the statement.

If $|\Delta(f)|$ is a power of 3, then it is divisible by 3 by the Minkowski bound $|\Delta(f)| \geq 13$. If it does not have a triple zero modulo 3, it must have a zero of

multiplicity 2, which means that the problem is NP-complete by Proposition 3.27. □

We finish this section with two lemmas that tell us what happens if the polynomial has a triple zero modulo a power of 3, for the $\mathbb{Z}$-rank 6 and 3 cases separately.

Lemma 3.31. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic with $\mathbb{Z}$-rank 6 with a triple root modulo 9. Then Π_f is NP-complete.*

Proof. Assume by translation that $f \equiv X^3 \bmod 9$. Let $B = (\mathbb{Z}/9\mathbb{Z})[\omega, \varepsilon]$, where $1 + \omega + \omega^2 = 0$ and $\varepsilon^2 = 0$. Let $R = \mathbb{Z}/9\mathbb{Z}$ and let $a = 0$. Letting $\alpha_1, \alpha_2, \alpha_3$ be the three zeroes of f in A_2, we use the universal property of A_2 to give a map $A_2 \to R$ sending α_1 to ε and α_2 to $\omega\varepsilon$. Using that $9 \mid \alpha_1 + \alpha_2 + \alpha_3$ we see $\alpha_3 \mapsto \omega^2\varepsilon$. Now we can reduce from the NP-complete problem $\Pi_{\mathbb{Z}/9\mathbb{Z}[\omega],\{1,\omega,\omega^2\}}$ by Lemma 3.13. □

Lemma 3.32. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic with $\mathbb{Z}$-rank 3. Then f does not have a triple root modulo 27.*

Proof. We will argue by contradiction. Let f be as in the conditions, and assume by translating that f has 0 as a triple root modulo 27. Let α, β, γ be the zeroes of f in $\overline{\mathbb{Q}}$. We define $R := \mathbb{Z}[\alpha]/(27) \cong (\mathbb{Z}/27\mathbb{Z})[\eta]$ where $\eta^3 = 0$. As f splits as $(X-\alpha)(X-\beta)(X-\gamma)$ in $\mathbb{Z}[\alpha]$, we find that X^3 totally splits over R with one of the factors being $X - \eta$. It can be seen that if X^3 factors over R as $(X-\eta)(X-a)(X-b)$ then $(X-a)(X-b) = X^2+\eta X+\eta^2$. Since $X^2+\eta X+\eta^2$ splits over R, the discriminant $-3\eta^2$ is a square of R. Let $x \in R$ be such that $-3\eta^2 = x^2$. Let $\mathfrak{m} = (3, \eta)$ be the maximal ideal in R. We see $-3\eta^2 \in \mathfrak{m}^3 \setminus \mathfrak{m}^4$, hence $x \in \mathfrak{m} \setminus \mathfrak{m}^2$. But if $3u + v\eta$ is an element of $\mathfrak{m} \setminus \mathfrak{m}^2$, then u or v is a unit, and in both cases the square is in $\mathfrak{m}^2 \setminus \mathfrak{m}^3$ as $9, 3\eta, \eta^2$ form a basis of the $\mathbb{F}_3$ vector space $\mathfrak{m}^2/\mathfrak{m}^3$. This means that $-3\eta^2$ is not a square, contradiction, so no such f exists. □

These lemmas all together give us the following proposition.

Proposition 3.33. *Let f be monic irreducible cubic. Then Π_f is NP-complete if at least one of the following conditions holds:*

- *$\Delta(f)$ has a prime factor other than 3;*
- *f does not have a triple zero modulo 3;*
- *$\mathrm{rk}_{\mathbb{Z}}(f) = 6$ and f has a triple zero modulo 9;*
- *f has a triple zero modulo 27.*

Proof. This proposition consists of four statements; the first two are given by Proposition 3.30, the third by Lemma 3.31, and the fourth by the contraposition of Lemma 3.32 followed by Lemma 3.31. □

4 Cubic polynomials with discriminant $\pm 3^\ell$

In the previous section we have proven that for any cubic monic irreducible polynomial $f \in \mathbb{Z}[X]$ whose discriminant has a prime factor that is not 3, the problem Π_f is NP-complete. This motivates the following theorem; the exact conditions of the theorem complement Proposition 3.33, in the sense that if f cubic monic irreducible does not satisfy these conditions, then it holds that $\Pi_f \in \mathcal{NPC}$ by Proposition 3.33. We refer to Definitions 3.21 and 3.6 for the definitions of $\mathbb{Z}$-rank and equivalence of polynomials respectively.

Theorem 4.1. *Let $f \in \mathbb{Z}[X]$ be monic irreducible cubic, with discriminant of the form $\pm 3^k$ with $k \in \mathbb{Z}_{\geq 0}$. Assume that f has a zero of multiplicity 3 modulo 3, and not a triple zero modulo 27. Also, assume that if the $\mathbb{Z}$-rank is 6, then f does not have a triple zero modulo 9. Then f is equivalent to one of the polynomials in Table 4.2.*

For the proof of Theorem 4.1, we first state two definitions and a lemma about integral points on a family of elliptic curves.

Definition 4.2. Let S be a finite set of primes, and $\overline{S}$ the multiplicative subset of $\mathbb{Z}$ generated by S. Then $\mathbb{Z}_S$, the ring of *S-integers*, is defined as $\{\frac{a}{b} \mid a \in \mathbb{Z}, b \in \overline{S}\}$.

Throughout the rest of the section, we take $S = \{3\}$.

Definition 4.3. Let $a \in \mathbb{Z} \setminus \{0\}$. Then C_a is the elliptic curve given by the equation $y^2 = x^3 + a$.

Lemma 4.4. *Let $\ell \in \mathbb{Z}_{\geq 0}$, and let ℓ' be the unique number with $0 \leq \ell' < 6$ and $\ell \equiv \ell' \bmod 6$. Let $a = 3^k$ with $\ell' + 6k = \ell$. Then there is a bijection $C_{\pm 2^4 3^{\ell'}}(\mathbb{Z}_S) \to C_{\pm 2^4 3^\ell}(\mathbb{Z}_S)$, given by $(x, y) \mapsto (a^2 x, a^3 y)$.*

Proof. It only remains to check that $(x, y) \in C_{\pm 2^4 3^{\ell'}}$ if and only if $(a^2 x, a^3 y) \in C_{\pm 2^4 3^\ell}$. This follows because the equation $y^2 = x^3 \pm 2^4 3^{\ell'}$ holds if and only if $a^6 y^2 = a^6 x^3 \pm a^6 2^4 3^{\ell'}$ does. □

Proof of Theorem 4.1. We will write down a family of elliptic curve equations in the coefficients of the polynomials. By introducing S-integers, we can restrict to finitely many of these equations, and use the software package Sage [Sag18] to find the resulting polynomials.

Let f be a cubic irreducible polynomial with discriminant $\pm 3^\ell$ with $\ell \geq 1$, satisfying the conditions of the theorem. Since f has a triple zero modulo 3, the coefficient corresponding to X^2 is divisible by 3. Hence we can put f into the form $X^3 + pX + q$ with $p, q \in \mathbb{Z}$ by translation.

Now $\Delta(f)$ has the simple formula $-4p^3 - 27q^2$. Setting $-4p^3 - 27q^2 = \pm 3^\ell$, we find a family of elliptic-curve-like diophantine equations. Multiplying such an

equation by $2^4 3^3$, and substituting $x = -2^2 3p, y = 2^2 3^3 q$, we find the equation

$$C_{\mp 2^4 3^{\ell'}} : y^2 = x^3 \mp 2^4 3^{\ell'}$$

with $\ell' = \ell + 3$. It suffices to find all integral points (x, y) on one of these curves. Lemma 4.4 lets us do even more: we can parametrise all points on $\bigcup_{\ell \geq 0, s=\pm 1} C_{s2^4 3^\ell}(\mathbb{Z}_S)$ by $\bigcup_{6 > \ell \geq 0, s=\pm 1} C_{s2^4 3^\ell}(\mathbb{Z}_S) \times \mathbb{Z}_{\geq 0}$. Now theorem 4.3 of [Sil09] tells us there are only finitely many S-integral point on the curves $C_{\pm 2^4 3^\ell}$ with $0 \leq \ell < 6$. The author used Sage to explicitly find these points. The list of parametrised corresponding polynomials up to the transformation $f(X) \mapsto -f(-X)$ can be seen in Table 4.1. The reducible polynomials are those with Galois group of cardinality 1 or 2. Next to the irreducible polynomials are the values of $t \in \mathbb{Z}_{\geq 0}$ such that the polynomial has integral coordinates. Now we observe that all polynomials have a triple root modulo 27 for $t \geq 2$, and that $X^3 + 9$ and $X^3 - 54X + 153$ both have a triple zero modulo 9 and $\mathbb{Z}$-rank 6. This almost give the final Table 4.2; it only remains to observe that $X^3 - 3X + 1 \sim X^3 - 21X + 37$, as for $f \in \{X^3 - 3X + 1, X^3 - 21X + 37\}$ we have that $\mathbb{Z}[X]/(f)$ is the ring of integers of $\mathbb{Q}(\zeta_9 + \zeta_9^{-1})$ where ζ_9 is a primitive ninth root of unity; $\zeta_9 + \zeta_9^{-1}$ is a zero of $X^3 - 3X + 1$, and $3(\zeta_9 + \zeta_9^{-1})^2 + \zeta_9 + \zeta_9^{-1} - 6$ is a zero of $X^3 - 21X + 37$.

Furthermore, note that the three polynomials in Table 4.2 are pairwise non-equivalent, as all of the discriminants are different. □

Polynomial	Cardinality of Galois group	All $t \in \mathbb{Z}_{\geq 0}$ for which the polynomial is integral
$X^3 - \frac{1}{3} \cdot 3^{2\,t} X + \frac{1}{27} \cdot 3^{3\,t}$	3	≥ 1
$X^3 - \frac{73}{108} \cdot 3^{2\,t} X + \frac{595}{2916} \cdot 3^{3\,t}$	1	
$X^3 - \frac{7}{3} \cdot 3^{2\,t} X + \frac{37}{27} \cdot 3^{3\,t}$	3	≥ 1
$X^3 - 3^{2\,t} X + \frac{1}{3} \cdot 3^{3\,t}$	3	≥ 1
$X^3 - \frac{193}{12} \cdot 3^{2\,t} X + \frac{2681}{108} \cdot 3^{3\,t}$	2	
$X^3 + \frac{1}{27} \cdot 3^{3\,t}$	2	≥ 1
$X^3 - \frac{1}{12} \cdot 3^{2\,t} X + \frac{7}{108} \cdot 3^{3\,t}$	6	
$X^3 + \frac{1}{9} \cdot 3^{3\,t}$	6	≥ 1
$X^3 + \frac{2}{3} \cdot 3^{2\,t} X + \frac{7}{27} \cdot 3^{3\,t}$	2	≥ 1
$X^3 + \frac{1}{3} \cdot 3^{3\,t}$	6	≥ 1
$X^3 - \frac{3}{4} \cdot 3^{2\,t} X + \frac{5}{12} \cdot 3^{3\,t}$	6	
$X^3 - 6 \cdot 3^{2\,t} X + \frac{17}{3} \cdot 3^{3\,t}$	6	≥ 1

Table 4.1: A list containing all monic cubic polynomials in $\mathbb{Z}[X]$, up to the substition $f(X) \mapsto -f(-X)$, that have a triple zero modulo 3 and discriminant of the form $\pm 3^k$, together with the Galois group.

Polynomial	Discriminant	Factorisation of discriminant
$X^3 - 3$	-243	-3^5
$X^3 - 3X + 1$	81	3^4
$X^3 - 9X + 9$	729	3^6

Table 4.2: A minimal set S of polynomials such that every monic cubic irreducible polynomial that satisfies the conditions of Theorem 4.1 is equivalent to a polynomial in S.

5 NP-completeness for difficult cubic polynomials

In this section we prove NP-completeness for the problems Π_f with f in Table 4.2, at the end concluding the proof of Theorem 1.10.

Lemma 5.1. *Let $f = X^3 - 3$. Then Π_f is NP-complete.*

Proof. Let α, β, γ be the three zeroes of f in $\overline{\mathbb{Q}}$. Let B be the finite ring $\mathbb{Z}/9\mathbb{Z}[\pi] := \mathbb{Z}/9\mathbb{Z}[X]/(X^6+3)$. Note B has a $\mathbb{Z}/3\mathbb{Z}$-grading $B = B_0 \oplus B_1 \oplus B_2$ with $B_i = \pi^i\mathbb{Z}/9\mathbb{Z} \oplus \pi^{i+3}\mathbb{Z}/9\mathbb{Z}$. We denote $\zeta := -\frac{1}{2} + \frac{1}{2}\pi^3$; observe that $\zeta^2 + \zeta + 1 = 0$. In this ring, $X^3 - 3$ has a factorisation as $(X + \pi^2)(X + \zeta\pi^2)(x + \zeta^2\pi^2)$. As $\mathrm{Gal}(f) = \mathrm{S}_3$, the order $A := \mathbb{Z}[\alpha, \beta, \gamma]$ is naturally isomorphic to to A_2. Then by the universal property of A_2, we have a morphism $\psi : A \to B$ given by $\psi(\alpha) = -\pi^2, \psi(\beta) = -\zeta\pi^2, \psi(\gamma) = -\zeta^2\pi^2$. Letting $S = \{\psi(\alpha), \psi(\beta), \psi(\gamma)\}$, we note $S \subset B_2$. We define $G = B_2$. Note that S is not a coset in G and does not contain zero. By Theorem 1.16 and Lemma 2.14, this means $\Pi_{G,S}$ is NP-complete.

We will give a reduction $\Pi_{G,S} \leq \Pi_f$. Let (t, H) be an instance of $\Pi_{G,S}$. Let C be the subring of B^t given by $C = \{(x_1, \ldots, x_t) \in B^t \mid x_1 \equiv \cdots \equiv x_n \bmod \zeta - 1\}$. Note $S^t \subset C$, as $\psi(\alpha) \equiv \psi(\beta) \equiv \psi(\gamma) \bmod \zeta - 1$. Let H' be H intersected with C. We may assume H' is not contained in $(\zeta - 1)B^t$; if it is, clearly $H' \cap S^t = \varnothing$. Note that in $B_2/(\zeta-1)B_2$ we have $\pi^2(\zeta-1) = 3\pi^2 + 5\pi^5 = 0$ and $\pi^5(\zeta - 1) = -3\pi^2 + 3\pi^5 = 0$. We deduce that $B_2/(\zeta - 1)B_2 = \{0, \pi^2, 2\pi^2\}$. Writing $H' = \langle H' \cap (\pi^2 + (\zeta - 1)\pi^2 B_0^t)\rangle \cup (H' \cap (\zeta - 1)\pi^2 B_0^t)$, we see $H' = \langle H' \cap (\pi^2 + (\zeta - 1)\pi^2 B_0^t)\rangle$. That means that $H' \cdot H' \cdot H' \cdot H'$ is generated by elements of the form $\prod_{i=1}^4 (\pi^2 + (\zeta - 1)\pi^2 x_i)$ with $x_i \in B_0, \pi^2 + (\zeta - 1)\pi^2 x_i \in H'$ for $i = 1, \ldots, 4$. Using that $\pi^6 = -3$ and that the exponent of B equals 9, we see that the product equals $-3\pi^2 \left(1 + \sum_{i=1}^4 x_i(\zeta - 1)\right)$, which equals $-3\sum_{i=1}^4 (\pi^2 + (\zeta - 1)\pi^2 x_i) \in H'$. Hence $H' \cdot H' \cdot H' \cdot H' \subset H'$, meaning that $\mathbb{Z}/9\mathbb{Z}[H']$ is equal to $\mathbb{Z}/9\mathbb{Z} + H' + H' \cdot H' + H' \cdot H' \cdot H'$. Using the grading of B, the intersection $\mathbb{Z}/9\mathbb{Z}[H'] \cap S^t$ hence equals H' itself. Now we can define A_H to be the inverse image under $A^t \to B^t$ of $\mathbb{Z}/9\mathbb{Z}[H']$, and we see $Z_f(A_H)$ is non-empty exactly if $H \cap S^t$ is non-empty. This completes the reduction. □

Lemma 5.2. *Let $f = (X - 1)^3 - 3(X - 1) + 1 = X^3 - 3X^2 + 3$. Then Π_f is NP-complete.*

Proof. Let R be the ring $\mathbb{F}_3[\varepsilon] = \mathbb{F}_3[X]/(X^2)$, and let G be a free R-module of rank 1, with generator m. Let $S = \{0, m, -m - \varepsilon m\} \subset G$. Note that $P^R_{G,S}$ is NP-complete, as by Lemma 2.6 with $\varphi : x \mapsto m - x$ we have $P^R_{G,\{0,m\}} \leq P^R_{G,S}$, and by Lemma 2.8 we know $P^R_{G,\{0,m\}}$ is NP-complete. Now we define a new problem P: the input is $t \in \mathbb{Z}_{>0}$, a submodule H of G^t and $x_* \in G^t$ with $(m, \ldots, m) - \varepsilon x_* \in H$; the output is whether $(x_* + H) \cap S^t$ is non-empty. Obviously, $P \leq P^R_{G,S}$. We will also prove $P^R_{G,S} \leq P$. Let (t, H, x_*) be an instance

of $P^R_{G,S}$. For ease of notation, from now on we write $\underline{x}$ for a vector consisting of all x'es. Let $t' = t+1, H' = H \times \{0\} + R \cdot (\underline{m} - \varepsilon(x_*, 0))$ and $x'_* = (x_*, 0)$. Then (t', H', x'_*) is an instance of P. If (t, H, x_*) is a yes-instance of $P^R_{G,S}$ with $h \in H$ such that $h + x_* \in S^t$, then $(h, 0) \in H'$ and $(h, 0) + (x_*, 0) \in S^{t'}$, so (t', H', x'_*) is a yes-instance of P. Conversely, if (t', H', x'_*) is a yes-instance of P, then there is an $h' \in H'$ with $h' + (x_*, 0) \in S^{t'}$. By looking at the last coordinate, we see that h' can be written as $(h, 0) + v(\underline{m} - \varepsilon(x_*, 0))$ with $h \in H, v \in R$. Note that $\sigma : G \to G, x \mapsto (1+\varepsilon)(x - m)$ gives a bijection on S which cycles S. If we also denote σ for the map $G^{t'} \to G^{t'}$ that applies σ coordinatewise, we see that $\sigma(x'_* + H') = x'_* + (1+\varepsilon)H' - (1+\varepsilon)(\underline{m} - \varepsilon x'_*)$ which implies that $\sigma(x'_* + H')$ lies in $x'_* + H'$. Then clearly σ also acts on $(x'_* + H') \cap S^{t'}$. Therefore without loss of generality $(h, 0) + v(\underline{m} - \varepsilon(x_*, 0)) + (x_*, 0)$ is zero on the last coordinate, meaning $v = 0$. Then $(h, 0) + (x_*, 0) \in S^{t'}$ hence $h + x_* \in S^t$, so we find (t, H, x_*) is a yes-instance of $P^R_{G,S}$. We have now proven that $P \approx P^R_{G,S}$, so we have $P \in \mathcal{NPC}$.

We will now reduce from P to Π_f. Let α be a zero of f in $\overline{\mathbb{Q}}$, and note that with $A := \mathbb{Z}[\alpha]$ we have $Z_A(f) = \{\alpha, \alpha^2 - 2\alpha, \alpha - \alpha^2 + 3\}$. Let $\mathfrak{p}$ be the prime ideal in A over 3 generated by α; then (3) factorises as $\mathfrak{p}^3$. Define $B = A/\mathfrak{p}^4$. Note that as $-3 = \alpha^2(\alpha - 3)$ we have $-3 = \alpha^3$ in B. Finally, note that $Z_A(f)$ can be written as $\alpha + \{0, \alpha^2 - 3\alpha, -\alpha^2 + 3\}$ where $\{0, \alpha^2 - 3\alpha, -\alpha^2 + 3\}$ is a subset of $\mathfrak{p}^2$, and modulo $\mathfrak{p}^4$ it is equal to $\{0, \alpha^2, -\alpha^2 - \alpha^3\}$. This means the image of $Z_A(f)$ in B is $\alpha + \{0, \alpha^2, -\alpha^2 - \alpha^3\}$.

We take m, the generator of G, to be $\alpha^2 \in B$, with $\varepsilon \in R$ acting on G as multiplication by α. Let (t, H, x_*) be an instance of P. Now define $R_H \subset B^t$ as $\mathbb{Z}/9\mathbb{Z} + \mathbb{Z}/3\mathbb{Z}(\underline{\alpha} + x_*) + H$. Note that this is in fact a ring; the only non-trivial requirement is that $(\underline{\alpha} + x_*)^2 \in R_H$, but $(\underline{\alpha} + x_*)^2 = \underline{\alpha^2} + 2\alpha x_* = \underline{m} - \varepsilon x_*$, which is an element of H by definition of the problem P. Also, note that $\underline{3} \in R_H$ is $\underline{-\alpha^3}$, and $\underline{\alpha^3} = \varepsilon(\underline{m} - \varepsilon x_*) \in H$. This tells us that $R_H \cap G^t = H$. Now we see that $(\underline{\alpha} + S^t) \cap R_H = (\underline{\alpha} + x_*) + (-x_* + S^t) \cap R_H$ is in bijection with to $(-x_* + S^t) \cap R_H = (-x_* + S^t) \cap H$. Let A_H be the inverse image under $A^t \to B^t$ of R_H. As $\alpha + S$ is the image of $Z_A(f)$ in B, we see $Z_f(A_H)$ is non-empty exactly if $(x_* + H) \cap S^t$ is non-empty. This completes the reduction. □

Lemma 5.3. *Let $f = X^3 - 9X + 9$. Then Π_f is NP-complete.*

Proof. Let α be a zero of f in $\overline{\mathbb{Q}}$, and note that with $A = \mathbb{Z}[\alpha]$ we have $Z_A(f) = \{\alpha, \alpha + \alpha^2 - 6, -2\alpha - \alpha^2 + 6\}$. Let $B = A/(9, 3\alpha^2), R = \mathbb{Z}/9\mathbb{Z} \subset B$. We see B is a ring of cardinality 3^5, generated as an additive group by $1, \alpha, \alpha^2$ of order $9, 9, 3$ respectively. To prove NP-completeness, let $G = \langle \alpha^2 + 3 \rangle \subset B$. This is a group of cardinality 3. Let $S = \{\pm(\alpha^2 + 3)\} \subset G$. Note that S is not a coset and does not contain 0, so $\Pi_{G,S}$ is NP-complete. We will reduce from this problem to Π_f. Let (t, H) be an instance of $\Pi_{G,S}$. Write T for the image of $Z_A(f)$ in B.

Let $t' = 2t + 1$. Let $H' = \{(x, -x, 0) \mid x \in H\} \subset B^{t'}$. Let $x_* = (\alpha - (\alpha^2 + 3), \ldots, \alpha - (\alpha^2 + 3), \alpha)$, and let $R_H = R[H', x_*]$. Note that as an additive group, this is generated by $\mathbb{Z}/9\mathbb{Z}, H', x_*, x_*^2, H'x_*$. We see x_* and x_*^2 have order 9 and 3 respectively.

Claim: $R_H \cap T^{t'}$ is non-empty if and only if $H \cap S^t$ is non-empty. We prove this by examining an element $x \in R_H \cap T^{t'}$. Let $\sigma : B \to B, x \mapsto x^2 + x + 3$ and $\tau : B \to B, x \mapsto -2x - (x^2 + 3)$ be two maps, and note that by virtue of the Galois group still acting on T, we have that σ, τ induce a transitive permutations on T, and $\sigma|_T = \tau|_T^{-1}$. Denoting σ, τ for the two maps $B^{t'} \to B^{t'}$ that coordinatewise perform σ respectively τ it is clear that $\sigma(R_H), \tau(R_H)$ are subsets of R_H, as R_H is a ring. This tells us that $\sigma(x), \tau(x)$ also lie in $R_H \cap T^{t'}$. Letting $\pi : B^{t'} \to B$ denote the projection onto the last coordinate, we see that this means that $R_H \cap T^{t'}$ is non-empty if and only if $R_H \cap T^{t'} \cap \pi^{-1}(\alpha)$ is non-empty.

As $\pi(H') = 0$, we have that $\pi(R_H) = \pi(\mathbb{Z}/9\mathbb{Z} + \mathbb{Z}/9\mathbb{Z}x_* + \mathbb{Z}/3\mathbb{Z}x_*^2)$. As $\pi(1) = 1, \pi(x_*) = \alpha, \pi(x_*^2) = \alpha^2$ we see that $R_H \cap \pi^{-1}(\alpha) = x_* + H' + H'x_*$. Finally, we will prove that $(x_* + H' + H'x_*) \cap T^{t'}$ is non-empty if and only if $H \cap S^t$ is non-empty. Note that $x_* + H' + H'x_*$ contains an element of $T^{t'}$ if and only if there are $h_1, h_2 \in H'$ with $x_* + h_1 + h_2 x_* \in T^{t'}$. Writing $h_1 = (x, -x, 0)$ and $h_2 = (y, -y, 0)$ with $x, y \in H$, this is equivalent to $(x, -x) + (y, -y)(\alpha - (\alpha^2 + 3)) \in \{\pm(\alpha^2 + 3), -3\alpha\}^{2t}$. As $(\alpha^3 + 3)G = 0$, we can write $(x, -x) + (y, -y)(\alpha - (\alpha^2 + 3) = (x, -x) + \alpha(y, -y)$ with $\alpha G = \langle 3\alpha \rangle$. Then we see that $(y, -y)$ is an element of $\{0, -(\alpha^2 + 3)\}^{2t}$, implying $y = 0$. So we find $R_H \cap T^{t'} \neq \varnothing \Leftrightarrow \exists x \in H : (x, -x) \in \{\pm(\alpha^2 + 3)\}^{2t}$. This is clearly equivalent to $H \cap S^t \neq 0$, proving the claim.

That means that we have constructed a subring $R_H \subset B^{t'}$ such that $R_H \cap T^t \neq \varnothing \Leftrightarrow H \cap S^t \neq \varnothing$ holds. Letting A_H be the inverse image of R_H under the natural map $A^{t'} \to B^{t'}$, we have completed the reduction. □

Proof of Theorem 1.10. Let $f \in \mathbb{Z}[X]$ be cubic, monic. By Lemma 3.10 we have $\Pi_f \in \mathcal{P}$ if f is reducible. Assume that f is irreducible. If it satisfies the conditions of Proposition 3.33, then Π_f is NP-complete by that proposition. Otherwise, it satisfies the conditions of Theorem 4.1, and hence is equivalent to one of the three polynomials in Table 4.2. For those three polynomials NP-completeness has been proven in this section. This completes the proof. □

6 An undecidability result

In this section we prove the undecidability of U_{X^2}. We first state the negative result of Hilberts tenth problem, proven by Davis, Putnam, Robinson and Matiyasevich. One can find a proof in [Mat93].

Definition 6.1. *Hilberts tenth problem* is: given an $n \in \mathbb{Z}_{\geq 1}$ and a polynomial p in $\mathbb{Z}[X_1, \ldots, X_n]$, determine whether p has a zero in $\mathbb{Z}^n$.

Theorem 6.2. *Hilberts tenth problem is undecidable.*

We first prove the following theorem, which resembles Theorem 6.2 but is more useful in our context.

Theorem 6.3. *There is no algorithm that decides, given an $m \in \mathbb{Z}_{\geq 1}$ and a finite set of polynomials S in $\mathbb{Z}[X_1, \ldots, X_m]$ of degree at most 2, whether the polynomials in S have a common zero in $\mathbb{Z}^m$, that is whether the set $Z_{\mathbb{Z}^m}(S)$ of common zeroes of S in $\mathbb{Z}^m$ is non-empty.*

Proof. We reduce from Theorem 6.2. Let (n, p) be the input for Hilberts tenth problem. We briefly sketch the reduction, producing (m, S) such that the polynomials in S have a common zero exactly if p has a zero.

1. Let $m := n, S := \{p\}$.
2. While S contains a polynomial q containing a monomial $c\prod_{i=1}^{k} X_{n_i}, c \neq 0$ where $n_i \in \{1, \ldots, m\}$ for $i = 1, \ldots, k$ of degree k strictly bigger than 2, make $m := m+1, S := S \cup \{X_m - X_{n_1}X_{n_2}\}$ and in q replace the monomial $c\prod_{i=1}^{k} X_{n_i}$ with $cX_m \prod_{i=3}^{k} X_{n_i}$, lowering the degree of that monomial.

Note that the zero set of S is conserved in each step, and that step 2 always terminates. Hence this proves the theorem. □

Proof of Theorem 1.4. We reduce from Theorem 6.3. Let $n \in \mathbb{Z}_{\geq 1}$ and $S = \{p_1, \ldots, p_m\}$ a subset of $\mathbb{Z}[X_1, \ldots, X_n]$ consisting of polynomials of degree at most 2 be given. We will construct input (A, B, g, b) for U_{X^2} that is a yes-instance if and only if the polynomials in S have a common zero in $\mathbb{Z}^n$.

Embed $\mathbb{Z}[X_1, \ldots, X_n]$ in $\mathbb{Z}[X_0, \ldots, X_n]$. We now multiply every monomial in one of the polynomials of S by a power of X_0 such that the polynomial is homogeneous of degree 2; call the resulting homogeneous polynomials $q_1, \ldots, q_m$. For $0 \leq i, j \leq n, 1 \leq k \leq m$ let c_{ijk} be $1 + \delta_{ij}$ times the coefficient of the monomial X_iX_j in q_k. Note that for $1 \leq k \leq m$ we have $\sum_{0 \leq i,j \leq n} c_{ijk}X_iX_j = 2q_k(X_0, \ldots, X_n)$ and $q_k(1, X_1, \ldots, X_n) = p_k(X_1, \ldots, X_n)$.

Let $1, v_0, v_1, \ldots, v_n, w_1, \ldots, w_m$ be formal variables and define $V = \bigoplus_{i=1}^{n} v_i\mathbb{Z}$ and $W = \bigoplus_{i=1}^{k} w_i\mathbb{Z}$. We then choose A additively equal to $1 \cdot \mathbb{Z} \oplus V \oplus W$. We define a multiplication on A by making multiplication by 1 the identity, multiplication on $V \times W, W \times W$ the zero map, and giving a bilinear symmetric map $\varphi : V \times V \to W$. We see that A is automatically commutative, and

the multiplication is associative. We define φ on the basis $v_0, v_1, \ldots, v_n$; we let $\varphi(v_i, v_j) = \sum_{k=1}^{m} c_{ijk} w_k$.

Finally, let B be the quotient group $\mathbb{Z} \oplus v_0\mathbb{Z}$ of A with $g : A \to B$ the projection and multiplication defined by $g(xy) = g(x)g(y)$, and let $b = v_0$. Note that $b^2 = 0$ as $g(v_0^2) \in g(W) = \{0\}$.

It remains to prove that $Z_A(X^2) \cap g^{-1}(b)$ is non-empty if and only if $Z_{\mathbb{Z}^n}(S)$ is non-empty. Let $x = \sum_{i=1}^{n} a_i v_i + \sum_{i=1}^{n} b_i w_i$ with $a_0 = 1$ be an element of $g^{-1}(b)$. We will show that $x^2 = 0$ if and only if $(a_1, \ldots, a_n)$ is contained in $Z_{\mathbb{Z}^n}(S)$. By definition of the multiplication of A, we see x^2 equals $\sum_{i=1}^{m} a_i a_j c_{ijk} w_k$, which is zero if and only if $2q_k(a_0, \ldots, a_n) = 2p_k(a_1, \ldots, a_n)$ is zero for every k. As 2 is not a zero divisor, this concludes the reduction, so U_{X^2} is undecidable. $\square$

A $\mathcal{P}$ and $\mathcal{NP}$

In this appendix we will briefly treat some of the P and NP theory and terminology. We will do so mostly informally, and in no way completely. For more information, the reader is directed to a standard work like [GJ79].

A.1 Problems

We start with a non-formal definition of n-colorability, to serve as a running example throughout the appendix.

Definition A.1. Let $n \in \mathbb{Z}_{\geq 1}$. Suppose (V, E) is a graph with V denoting the set of vertices and E the set of edges of the graph. We interpret $\{1, ..., n\}^V$ as the set consisting of assignments $f : V \to \{1, ..., n\}$. Such an assignment f is called an *n-coloring* (or a *coloring* if n is clear from the context) of (V, E) if for every two vertices i, j with $\{i, j\} \in E$ we have $f(i) \neq f(j)$; adjacent vertices are not allowed to have the same color. Now define n-COL: given a graph (V, E), determine whether it has an n-coloring.

Now we will give the formal definition of a decision problem; we only treat decision problems.

Definition A.2. A *decision problem*, briefly a problem, is defined to be a pair of sets (I, Y) with $Y \subset I \subset \mathbb{Z}_{>0}$, where we call I the set of *instances* of the problem, and Y the set of *yes-instances*. A problem is called *trivial* if $Y = I$ or $Y = \varnothing$.

Example A.3. We can formulate the primality problem PRIMES as (I, Y) with $I = \mathbb{Z}_{>0}$ and Y consisting of all primes.

Example A.4. For a problem as n-COL, the input is not a number — it is a graph. This can be dealt with by encoding graphs as natural numbers in an algorithmically nice way. Then I would consist of the encodings of graphs, and Y of the encodings of n-colorable graphs. In the present thesis we will simply gloss over this, thinking of the input as actually being a graph instead of the encoding of one.

A.2 $\mathcal{P}$

We will not define algorithms formally; we treat an algorithm as a step-by-step calculation, usually specified in a programming language or in a natural language. The most common formal definition is that of a Turing machine; for information on the subject, one can read [HU69].

We can now say when an algorithm solves a problem in polynomial time.

Definition A.5. An algorithm *solves a problem* (I, Y) *in polynomial time* if it runs in polynomial time in the length of the input $i \in I$, and ends in an

accepting state if and only if $i \in Y$, and otherwise in a rejecting state; informally, outputting "yes" and "no" respectively. The length of the input i is defined to be $\log(i)$. Contrary to [GJ79], we do not put any restrictions on what the algorithm does given input in $\mathbb{Z}_{>0} \setminus I$.

Note that an encoding usually has the property that the length does not change up to a polynomial; for example, for a graph we might as well say that the algorithm has a graph (V, E) as input, and should run in polynomial time in $|V| + |E|$.

Now we can define the class of problems $\mathcal{P}$.

Definition A.6.

$$\mathcal{P} = \{\text{problems for which there exists a polynomial time algorithm}\}.$$

Example A.7. We have that PRIMES, 2-COL $\in \mathcal{P}$. For the first one, see the famous article [AKS04]. Note that a polynomial algorithm for 2-COL is obtained by just trying to colour the graph with two colors, at each step coloring either an uncolored neighbor of a colored vertex or if no such vertex exists, an arbitrary new vertex.

A.3 $\mathcal{NP}$

We will now give a definition of $\mathcal{NP}$. Traditionally, this works by defining non-deterministic algorithms, and saying what it means for a non-deterministic algorithm to solve a problem: it should always classify a no-instance as a no-instance, and it should sometimes output "yes" for a yes-instance. For example, the following non-deterministic algorithm solves 3-COL: guess a 3-coloring, and check whether it is correct.

We present an equivalent way to describe $\mathcal{NP}$ without using random number generators.

Definition A.8. We say a problem (I, Y) is in $\mathcal{NP}$ if there exists a set theoretic function $c : Y \to \mathbb{Z}_{>0}$ with $\log c(y)$ polynomially bounded in $\log y$ and a function $t : I \times \mathbb{Z}_{>0} \to \{\text{yes, no}\}$ that can be calculated in polynomial time in the length of the input such that $t|_{(I \setminus Y) \times \mathbb{Z}_{>0}}$ is the constant "no" function and for all $y \in Y$ we have that $t(y, c(y))$ equals "yes". This $c(y)$ is also called a *certificate* for y.

Example A.9. We see that $\mathcal{P} \subset \mathcal{NP}$, by letting $t(i, m)$ equal "yes" if and only if $i \in Y$; we can pick c the constant 1 function. Also, n-COL $\in \mathcal{NP}$, by letting c send an encoding of a n-colorable graph to an encoding of one of its n-colorings.

A.4 Reductions

Next we will introduce the important notion of polynomial time reductions.

Definition A.10. Let $\Pi_1 = (I_1, Y_1), \Pi_2 = (I_2, Y_2)$ be two problems. We say Π_1 is *reducible* to Π_2, notation $\Pi_1 \leq \Pi_2$, if there exists a function $f : I_1 \to I_2$, called a *reduction*, that can be calculated by an algorithm in polynomial time, such that $Y_1 = f^{-1}(Y_2)$ or equivalently for every $i \in I_1$ we have $i \in Y_1 \Leftrightarrow f(i) \in Y_2$.

Informally, this means we can model an instance of Π_1 as an instance of Π_2, such that the "yes"-ness does not change. We can also interpret this as stating Π_2 is at least as difficult as Π_1; if we have an oracle that solves Π_2 and a reduction $\Pi_1 \leq \Pi_2$, then we can solve Π_1 in polynomial time. To give an idea of how a reduction works, we will prove the following lemma.

Lemma A.11. n-COL $\leq (n+1)$-COL.

Proof. Let $G = (V, E)$ be an instance of n-COL. Now define $V' = V \sqcup \{v_*\}$, and $E' = E \sqcup \{\{v, v_*\} \mid v \in V\}$. One easily checks that (V, E) is n-colorable if and only if (V', E') is $(n+1)$-colorable, which shows that this is indeed a reduction. □

As suggested by the notation, $\leq$ is indeed a transitive relation on the set of problems. However, it is not anti-symmetric. With that in mind, we can define the following equivalence relation.

Definition A.12. Let Π_1, Π_2 be two problems. We say Π_1 is *equivalent* to Π_2, notation $\Pi_1 \approx \Pi_2$, if $\Pi_1 \leq \Pi_2$ and $\Pi_2 \leq \Pi_1$.

Note that the non-trivial problems in $\mathcal{P}$ form one of the equivalence classes of this relation.

A.5 $\mathcal{NPC}$

With the notation of the last two subsections, we can define NP-complete problems.

Definition A.13.

$$\mathcal{NPC} = \{\Pi \in \mathcal{NP} \mid \forall R \in \mathcal{NP} : R \leq \Pi\}.$$

Informally, this consists of the hardest problems in $\mathcal{NP}$. It is a priori not clear that $\mathcal{NPC}$ is even non-empty. As famously proven by Stephen Cook in [Coo71], we have that SAT $\in \mathcal{NPC}$ (for a definition of SAT, see that article). Then, using a reduction SAT $\leq$ 3-COL (see [GJS76]) it follows that 3-COL $\in \mathcal{NPC}$, and by Lemma A.11 n-colorability is NP-complete for all $n \geq 3$. Another important observation is that $\mathcal{NPC}$ also forms an equivalence class of $\approx$, conjecturally disjoint from $\mathcal{P}$.

References

[AKS04] M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Ann. of Math. (2)*, 160(2):781–793, 2004.

[Coo71] S. A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158, New York, NY, USA, 1971. ACM.

[Gio13] A. Gioia. *On the Galois closure of commutative algebras*. PhD thesis, Leiden University, 2013.

[GJ79] M. R. Garey and D. S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.

[GJS76] M. R. Garey, D. S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoret. Comput. Sci.*, 1(3):237–267, 1976.

[HU69] J. E. Hopcroft and J. D. Ullman. *Formal languages and their relation to automata*. Addison-Wesley Publishing Co., Reading, Mass., 1969.

[Len84] A. K. Lenstra. Factoring multivariate integral polynomials. *Theoret. Comput. Sci.*, 34(1-2):207–213, 1984.

[Len08] H. W. Lenstra, Jr. Lattices. In *Algorithmic number theory: lattices, number fields, curves and cryptography*, volume 44 of *Math. Sci. Res. Inst. Publ.*, pages 127–181. Cambridge Univ. Press, Cambridge, 2008.

[LS17] H. W. Lenstra, Jr. and A. Silverberg. Roots of unity in orders. *Found. Comput. Math.*, 17(3):851–877, 2017.

[LS18] H. W. Lenstra, Jr. and A. Silverberg. Algorithms for commutative algebras over the rational numbers. *Found. Comput. Math.*, 18(1):159–180, 2018.

[Mat93] Y. V. Matiyasevich. *Hilbert's tenth problem*. Foundations of Computing Series. MIT Press, Cambridge, MA, 1993.

[Sag18] Sage Developers. *SageMath, the Sage Mathematics Software System (Version 8.2)*, 2018. `http://www.sagemath.org`.

[Sil09] J. H. Silverman. *The arithmetic of elliptic curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, Dordrecht, second edition, 2009.

[Ste17] P. Stevenhagen. Number rings, 2017.