Primes via Zeros: interactive proofs for testing primality of natural classes of ideals

Abhibhav Garg [*] Rafael Oliveira [*] Nitin Saxena [†]

Abstract

Deciding whether a given system of polynomial equations is satisfiable (known as the Hilbert's nullstellensatz problem (HN)) is a central question in mathematics and computer science, with applications in several other fields of science. By the celebrated degree bounds for the nullstellensatz (Brownawell 1987, Kollar 1988, Jelonek 2005), algebraic algorithms that decide satisfiability of polynomial equations are known to be in PSPACE. On the other hand, the only hardness result for the HN problem is that it is NP-hard. Koiran's seminal work (Koiran, 1996) showed that, under the Generalized Riemann Hypothesis (GRH), this problem can be brought down to AM, significantly reducing the gap between complexity lower bounds and complexity upper bounds.

The above question can be thought of, in algebraic terms, as follows: given generators for an ideal $I := (f_1, \ldots, f_m) \subset \mathbb{C}[x_1, \ldots, x_n]$, decide whether $1 \in I$. Another central question in mathematics and computer science is the question of determining whether a given ideal I is prime, which geometrically corresponds to the zero set of I, denoted $Z(I)$, being irreducible. The case of principal ideals (i.e., $m = 1$) corresponds to the more familiar absolute irreducibility testing of polynomials, where the seminal work of (Kaltofen 1995) yields a randomized, polynomial time algorithm for this problem. However, when $m > 1$, the complexity of the primality testing problem seems much harder. The current best algorithms for this problem are only known to be in EXPSPACE.

Such drastic state of affairs has prompted research on the primality testing problem (and its more general variants, the primary decomposition problem, and the problem of counting the number of irreducible components) for natural classes of ideals. Notable classes of ideals are the class of radical ideals, complete intersections (and more generally Cohen-Macaulay ideals). For radical ideals, the current best upper bounds are given by (Bürgisser & Scheiblechner, 2009), putting the problem in PSPACE. For complete intersections, the primary decomposition algorithm of (Eisenbud, Huneke, Vasconcelos 1992) coupled with the degree bounds of (DFGS 1991), puts the ideal primality testing problem in EXP. In these situations, the only known complexity-theoretic lower bound for the ideal primality testing problem is that it is coNP-hard for the classes of radical ideals, and equidimensional Cohen-Macaulay ideals.

In this work, we significantly reduce the complexity-theoretic gap for the ideal primality testing problem for the important families of ideals I (namely, *radical ideals* and *equidimensional Cohen-Macaulay ideals*). For these classes of ideals, assuming the Generalized Riemann Hypothesis, we show that primality testing lies in $\Sigma_3^p \cap \Pi_3^p$. This significantly improves the upper bound for these classes, approaching their lower bound, as the primality testing problem is coNP-hard for these classes of ideals.

Another consequence of our results is that for equidimensional Cohen-Macaulay ideals, we get the first PSPACE algorithm for primality testing, exponentially improving the space and time complexity of prior known algorithms.

[*]University of Waterloo, Cheriton School of Computer Science. Email: {a65garg, rafael}@uwaterloo.ca
[†]IIT Kanpur, Department of CSE. Email: nitin@cse.iitk.ac.in

Contents

1 Introduction

Given a set of polynomials $f_1, \ldots, f_m \in \mathbb{C}[x_1, \ldots, x_n]$, a fundamental algorithmic task in computer algebra and algebraic geometry is the "factorization" of the ideal $I := (f_1, \ldots, f_m)$ into "irreducible" components, generalizing the task of factoring a given polynomial into its irreducible factors. Given the more complex structure of ideals in polynomial rings, the right generalization of the polynomial factoring problem is a primary decomposition of the ideal I, where one decomposes I into "primary ideals," which can be thought of, in the geometric sense, as the "irreducible components" of the zero set defined by the ideal I, accounted with their "multiplicities." This task was first undertaken in the pioneering works of Lasker and Hermann [Las05, Her26], and due to its paramount importance in computer algebra and algebraic geometry, it has been the subject of extensive research by algebraic geometers, computer algebraists and complexity theorists ever since, as can be seen in [Sei74, Sei78, GTZ88, EHV92, BM93] and references therein.

The above task naturally leads to the basic problem of deciding whether a given ideal is *prime*, which also has been extensively studied in theoretical computer science and algebraic geometry, as can be seen in [HS81, Kal95, BC04, BS07, Eis13] and references therein. In the Turing model, which is the model that we will focus on in this work, the ideal primality testing problem can be formally defined in the following way.

Problem 1.1 (Ideal primality testing)**.** *Given an algebraic circuit of size s computing polynomials with integer coefficients* $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$*, is the ideal* $(f_1, \ldots, f_m) \cdot \mathbb{C}[x_1, \ldots, x_n]$ *prime?*

The ideal primality testing problem is the direct generalization of the problem of testing whether a single polynomial is (absolutely) irreducible. More precisely, the absolute irreducibility testing problem for polynomials corresponds to Problem 1.1 with $m = 1$. A geometric version of the ideal primality testing problem was posed in [BC04, Problem 8.5]: what is the complexity to check whether a given algebraic set is irreducible over $\mathbb{C}$? Algebraically, the latter problem can be cast as: given polynomials $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$, is $\operatorname{rad}(f_1, \ldots, f_m)$ prime?

The $m = 1$ case of Problem 1.1, that is, the absolute irreducibility test of polynomials, can be done in randomized polynomial time, and by several different methods, due to the works [Kal85, BCGW93, Kal95, Gao03] (and references therein). However, when $m \geqslant 2$, Problem 1.1 seems to be notoriously more difficult, even for natural classes of ideals. Currently, the only algorithms to test whether a (general) ideal is prime are algorithms which compute a primary decomposition of the input ideal [Sei78, GTZ88, EHV92] (and references therein). Since these algorithms require the computation of Gröbner bases as a subroutine, the best complexity class which upper bounds the ideal primality testing problem is EXPSPACE, by applying the algorithm in [GTZ88] with the degree bounds for Gröbner basis obtained in [Dub90].

Due to the challenges posed by the general version of Problem 1.1 (and more generally for the primary decomposition problem), special cases of the problem(s) have been considered. In particular, natural classes of ideals have been studied, such as the cases where the ideal is a *complete intersection*[1], or when the given ideal is *radical*, or ideals with either constant dimension or constant codimension. When the input to Problem 1.1 is a complete intersection, the primary decomposition algorithm of [EHV92], combined with the Gröbner basis degree bounds of [DFGS91] yields an exponential time algorithm. When the input polynomials form a radical ideal, the works of [Chi86, Gri86] give an exponential time algorithm for computing the minimal primes of the given radical ideal, and [BS07, Theorem 4.1] gives a PSPACE-algorithm for counting the number of minimal primes of a given radical ideal. Thus, for the class of radical ideals, Problem 1.1 is in PSPACE.[2] For the case of radicals of ideals generated by constantly many polynomials, [BS10, Theorem 1.1] gives an algorithm in FRNC for computing the number of minimal primes, therefore solving the primality testing problem in RNC. Lastly, the dimension-dependent Gröbner basis degree bounds of [MT17], when combined with the primary decomposition algorithm of [GTZ88], yields a PSPACE-algorithm for primary decomposition of ideals of constant dimension, thereby also providing a PSPACE-algorithm for the ideal primality testing problem in this case. Apart from the above described algorithms, an important approach

[1] An ideal $I := (f_1, \ldots, f_m)$ is a complete intersection when the codimension of I equals the number of generators (in this case m). That is, each polynomial f_i "cuts" the dimension of the algebraic set by one, just as linearly independent linear forms would cut the dimension of the space by one. This is the algebraic generalization of the fact that a set of m linearly independent forms defines a space of dimension $n - m$.

[2] The works [Chi86, Gri86, BS07] work on the geometric setting of the question, as stated in [BC04, Problem 8.5]. When the input ideal is already promised to be radical, then the geometric setting and Problem 1.1 become the same problem.

to Problem 1.1 is to provide sufficient conditions/criteria on certain classes of ideals that will ensure primality. One prominent sufficient condition is given by the combination of Serre's criterion ([Eis13, Theorem 11.5], [Sta24, Tag 031O]) with the Jacobian criterion [Eis13, Theorem 16.19], which works for the important class of connected Cohen-Macaulay ideals[3] (see [Eis13, Theorem 18.15]).

Given the seemingly weak upper bounds for the ideal primality testing problem even for the natural classes of ideals above, one can ask if there are any complexity-theoretic lower bounds for the ideal primality testing problem. An easy argument shows that even for the natural class of zero-dimensional radical ideals, Problem 1.1 is coNP-hard. We prove this result in Proposition 2.7.

In this work, assuming the Generalized Riemann Hypothesis (GRH), we substantially improve the complexity-theoretic gap between the upper and lower bounds for Problem 1.1 for two important classes of ideals: radical ideals and equidimensional Cohen-Macaulay ideals. Our main theorem is the following.

Theorem 1.2 (Interactive protocols for primality)**.** *Let C be an algebraic circuit of size s with integer constants that computes $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$. If $I := (f_1, \ldots, f_m)$ is either radical or equidimensional Cohen-Macaulay, and if the dimension of I is given, then the complexity of testing if I is prime lies in* coAM*, assuming GRH.*

Combining the above theorem with [Koi97, Theorem 4.1], we obtain that Problem 1.1 for radical ideals and for equidimensional Cohen-Macaulay ideals are in the third level of the Polynomial Hierarchy (PH).

Theorem 1.3 (Primality testing in PH)**.** *Let C be an algebraic circuit of size s with integer constants that computes $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$. If the ideal $I := (f_1, \ldots, f_m)$ is either radical or equidimensional Cohen-Macaulay then the complexity of testing if I is prime lies in $\Sigma_3^p \cap \Pi_3^p$, assuming GRH.*

In the case when the given ideal is radical, Theorem 1.3 makes substantial progress on [BC04, Problem 8.5], and improves on the previous best upper bound of [BS07] from PSPACE to $\Sigma_3^p \cap \Pi_3^p$. Note that the coNP lower bound from Proposition 2.7 also applies to radical ideals. Thus, our upper bound gets substantially close to the known lower bound for the problem. Two important distinctions must be made here: our protocols only work when the input ideal is radical, and we need to assume the GRH; whereas the algorithm from [BS07] computes the number of irreducible components for the radical of the ideal generated by $f_1, \ldots, f_m$, and their work is unconditional on the GRH.

In the case of equidimensional Cohen-Macaulay ideals, which contains as special cases the important classes of zero dimensional ideals and of complete intersection ideals, Theorem 1.3 yields a nearly tight gap on the complexity of Problem 1.1, since the coNP lower bound of Proposition 2.7 also applies to zero-dimensional radical ideals. Note that our result for zero-dimensional ideals improves upon the previous upper bound of PSPACE. If our input is a complete intersection, then the dimension is implicitly given to us, since the codimension of the ideal in this case is simply given by the number of input polynomials. Thus, in this case Theorem 1.2 tells us that Problem 1.1 is in coAM. This improves upon the previous best upper bound of exponential time.

The above results show that the most computationally expensive step in our protocols is to determine the dimension of the input ideal. Thus, any improved protocol to compute exactly the dimension of a given ideal would lead to improvements to our primality testing protocols.

As mentioned above, previous works solved the primality testing problem by either computing a primary decomposition of the input ideal, or by using de Rham cohomology to count the number of irreducible components of the associated algebraic set (for the case of radical ideals). Our approach to solve Problem 1.1 is different from the previous approaches. Our strategy is inspired by Koiran's approach [Koi96] to decide whether a system of polynomial equations has a solution, which we now describe.

Koiran, in his seminal work [Koi96], proposed a fundamentally different approach to the problem of deciding whether a system of polynomial equations with integer coefficients has a solution over $\mathbb{C}$ (known as the Hilbert's nullstellensatz problem). His approach to this problem was based on the following idea: let $\mathcal{F} := \{f_1 = 0, \ldots, f_m = 0\}$, where $f_i \in \mathbb{Z}[x_1, \ldots, x_n]$ with $\deg f_i \leqslant d$, be our system of polynomial equations. By Hilbert's nullstellensatz, $\mathcal{F}$ is unsatisfiable if, and only if, $1 \in (f_1, \ldots, f_m)$. In other words, $\mathcal{F}$ is unsatisfiable iff there are polynomials $g_1, \ldots, g_m \in \mathbb{C}[x_1, \ldots, x_n]$ such that $1 = f_1 g_1 + \cdots + f_m g_m$. Note that this last equation is simply a linear system of equations with integer entries, where the variables are

[3]The definition of Cohen-Macaulay ideals is somewhat technical, as can be seen in [Eis13, Chapter 18.2] and in [Sta24, Tag 00N7]. However, as discussed in [Eis13, Chapter 18.5], Cohen-Macaulay ideals form a rich class of ideals, and are central objects in commutative algebra. In particular, zero dimensional ideals and complete intersections are Cohen-Macaulay.

the coefficients of the g_i's. Thus if there is a solution for this system over $\mathbb{C}$, there must be a solution over $\mathbb{Q}$. Hence, by clearing denominators, there are polynomials $h_1, \ldots, h_m \in \mathbb{Z}[x_1, \ldots, x_n]$ and $a \in \mathbb{Z}_{>0}$ such that $a = f_1 h_1 + \cdots + f_m h_m$. In particular, if $\mathcal{F}$ is unsatisfiable, given any prime $p \in \mathbb{N}$ which does not divide a, we have $1 \equiv a^{-1} \cdot (f_1 \cdot h_1 + \cdots + f_m h_m) \bmod p$, which implies that $\mathcal{F}$, when seen as a system over $\mathbb{F}_p[x_1, \ldots, x_n]$, does not have any solutions over the field $\overline{\mathbb{F}}_p$. In particular, $\mathcal{F}$ has no solutions over $\mathbb{F}_p^n$. The main conceptual insight in [Koi96] is to ask (and answer) the question: if $\mathcal{F}$ is satisfiable, would we be able to find "simple" solutions (that is, solutions in $\mathbb{F}_p^n$, as opposed to solutions in $\overline{\mathbb{F}}_p^n$) for *enough* primes p?[4] His main theorem [Koi96, Theorem 1] proves (assuming the GRH) a *quantitative* version of the following fact: if $\mathcal{F}$ is unsatisfiable, then there are "very few" primes p such that $\mathcal{F}$ has a solution over $\mathbb{F}_p^n$; if $\mathcal{F}$ is satisfiable, then there are "enough" (small enough) primes p such that $\mathcal{F}$ has a solution over $\mathbb{F}_p^n$. Equipped with the above theorem, Koiran can now distinguish between the two cases by using a set lower bound protocol, hence putting the Hilbert's nullstellensatz problem in coAM.

It is important to emphasize that the main technical challenge with the above approach, as pointed out by Koiran (see [Koi96, Page 275]), is to prove such quantitative bounds on the number of "good primes," when one changes the base ring from $\mathbb{Z}$ to $\mathbb{F}_p$ (or geometrically, from $\mathbb{C}$ to $\overline{\mathbb{F}}_p$). The principle of changing the base field where an object of interest lives (in our case our algebraic set over $\mathbb{C}^n$) to extract some information of our geometric object based on properties of its image is called a *base change theorem*. Such results are well studied in algebraic geometry and number theory, as can be seen in [GD65] and [Poo08, Appendix C].[5] Often times, such base change theorems are non-constructive, which suffices for their mathematical purposes. However, for such a base change theorem to be useful for computation, as is the case in [Koi96], one needs to prove *effective versions* of such base change theorems. As usual in mathematics and computer science, making such theorems effective often requires a considerable amount of work.

In this work, we use the same high-level strategy developed by Koiran to give an AM protocol for proving that a given input ideal (from one of our classes of ideals) is not prime, assuming that we are given the dimension of our input ideal. We begin by noticing that in the classes of ideals that we study, there are 3 ways in which our ideal cannot be prime: the ideal can have two minimal primes of different dimensions, the ideal can have two distinct minimal primes of maximum dimension, or the ideal can have a single minimal prime with "multiplicity." For radical ideals, only the first and the second cases can happen, whereas for equidimensional Cohen-Macaulay ideals only the second and third cases can happen.

In the case where we have two minimal primes with different dimensions (which can only happen in the radical case) we can identify that the given ideal is not prime via the Jacobian. This allows us to reduce this case to an instance of the Hilbert's nullstellensatz problem, which can be handled by the protocol of [Koi96]. This is done in Section 5.1.

In the case where we have a single minimal prime with multiplicity, (which can only happen in the equidimensional Cohen-Macaulay case), we can use Serre's criterion, combined with the Jacobian criterion (as done in [Eis13, Theorem 18.15]) to identify the multiplicity via the codimension of the Jacobian ideal inside our ideal. Thus, in this case we can use the protocol to compute dimension lower bounds given by [Koi97]. This is done in Section 5.2.

Now, the only case we have left is when our ideal is equidimensional, but has more than one irreducible component (this can happen for both the radical and the equidimensional Cohen-Macaulay cases). This turns out to be the most challenging case for us. The base change theorems of [GD65] tell us that if our ideal is prime over $\mathbb{C}[x_1, \ldots, x_n]$, then, for many choices of primes p, our ideal over $\mathbb{F}_p[x_1, \ldots, x_n]$ will remain prime over $\overline{\mathbb{F}}_p[x_1, \ldots, x_n]$. Combining this fact with [Sei74], one also obtains that for many choices of primes p, if our ideal is equidimensional and not prime, then our ideal over $\mathbb{F}_p[x_1, \ldots, x_n]$ will remain equidimensional and not prime over $\overline{\mathbb{F}}_p[x_1, \ldots, x_n]$. Since we are only interested in distinguishing whether our ideal has more than one component of top dimension, by an effective version of Bertini's theorem (Corollary 2.26), we can assume that we are working with ideals of dimension 1. Thus, we have reduced our main problem to the task of distinguishing whether an algebraic curve (of potentially exponential degree) is absolutely irreducible or not.

While the above gives us hope to be able to distinguish this case by going modulo p (for a "good prime"

[4] We would like to emphasize that the part which takes a lot of work in this claim is the part that says that satisfiable systems will be satisfiable in $\mathbb{F}_p^n$. Checking that a system remains satisfiable over $\overline{\mathbb{F}}_p^n$ is also easy, as it can be solved by a similar argument as the one given for the unsatisfiability case.

[5] Sometimes base change theorems are also called *spreading out theorems*, as is done in the references given here.

p), there are two problems that we need to address: can we make this base change theorem *effective*? And in case we can make it effective, how can we witness the distinction between an absolutely irreducible curve (i.e. a prime ideal) and a reducible curve (i.e., non-prime ideal) in a "simple way"? Could we do the latter simply by looking at points in $\mathbb{F}_p^n$, as is done by Koiran? As it turns out, the answer to the last question is affirmative, due to effective versions of the celebrated Lang-Weil theorem (Theorem 2.18).

We have thus far reduced our primality testing problem to the problem of proving an effective base change theorem which preserves equidimensional components of given algebraic sets. However, some care needs to be taken here: the Lang-Weil theorem only guarantees enough solutions over $\mathbb{F}_p^n$ of algebraic varieties which are *defined* over $\mathbb{F}_p$ (that is, the polynomials defining the variety must have coefficients in $\mathbb{F}_p$). Thus, if we are to use the Lang-Weil theorem, we must ensure that for an equidimensional non-prime ideal, *at least two* of its minimal primes *will be* $\mathbb{F}_p$*-definable* for *enough* "good primes." At this point, one may ask:[6] why not just preserve every minimal prime? As discussed in [BM93, Pages 47 and 48], this is a very bad idea, since to do so would cause us to have to work over polynomials defined over extremely large extension fields, and therefore kill any attempt to get an effective base change theorem, since the coefficients become "too complex."

With the above in mind, we have now finally arrived at the precise effective base change theorem that we need to prove: given an equidimensional ideal $I := (f_1, \ldots, f_m)$ with $f_i \in \mathbb{Z}[x_1, \ldots, x_n]$ and $\dim I = 1$, we need to ensure that there are enough "good primes" such that *at least two* of its irreducible components are defined over $\mathbb{F}_p$. To prove such an effective base change theorem, we combine Kaltofen's factoring algorithm over the algebraic closure [Kal95] with results from algebraic number theory and elimination theory. As we have mentioned before, this is the crucial technical result that we need, and it is carried out in Sections 3 and 4, which culminate in our main technical theorems, Theorems 4.5 and 4.6.

We now present a summary of the above high-level discussion of our results, and in the following subsection we present a more complete overview of our proof, where we discuss in a more precise form all the issues that need to be overcome to prove correctness of our interactive protocols.

Summary of contributions.

1. In this work, we give AM protocols for non-primality for natural classes of ideals, assuming the Generalized Riemann Hypothesis. This significantly tightens the complexity gap of these problems. More precisely, we have:

   - For the class of radical ideals, we prove that Problem 1.1 can be solved in coAM, whereas Proposition 2.7 shows that Problem 1.1 is coNP-hard for this class of problems.
   - For the class of equidimensional Cohen-Macaulay ideals (which includes the class of dimension zero ideals, and the class of complete intersection ideals), we prove that Problem 1.1 can be solved in coAM, whereas Proposition 2.7 shows that Problem 1.1 is also coNP-hard for this class of ideals.

2. On the technical side, our main contribution is to prove effective base change theorems for geometric irreducibility (and reducibility) of algebraic sets with equations defined over the integers. This is the content of Theorems 4.5 and 4.6.

   - One important remark is that our base change theorems, together with the effective Chebotarev density theorem (Theorem 2.19), yields nearly optimal density bounds for the good primes preserving (ir)reducibility, as discussed in Remark 5.3. This result is interesting on its own right.
   - The proof of our base change theorem involves the combination of (effective) techniques from algebraic geometry, elimination theory, efficient algorithms for factorization over the algebraic closure of a field (and its corresponding absolute irreducibility test), as well as techniques from algebraic number theory (to control the bit complexity of the algebraic numbers involved in the operations that we need to consider).

[6] As mentioned in [BM93, pages 47 and 48], "one who has never done any calculations."

1.1 Proof overview

We now give a more detailed overview of the proof of our main technical theorem. Recall that we are given polynomials $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$, defining the ideal $I := (f_1, \ldots, f_m)$ over $\mathbb{C}[x_1, \ldots, x_n]$. Let $V := Z(I) \subseteq \mathbb{C}^n$ be the zero set of the ideal I, and $r := \dim V$ be the dimension of the ideal. Since we will be interested in the bit-complexity of the polynomials and (algebraic) integers involved, we denote the *logarithmic height* of an integer $a \in \mathbb{Z}$ by $\mathrm{ht}(a) := \lceil \log(|a| + 1) \rceil$ (i.e., its bit complexity) and extend this definition to denote the logarithmic height of a polynomial by the maximum logarithmic height of any of its coefficients.

As stated earlier, for the types of ideals we consider, there are only a few ways in which the ideal I can fail to be prime. If I is radical, then it can fail to be prime because V consists of at least two components, either both of dimension r, or one of dimension r and one of dimension less than r. If I is equidimensional CM, it can fail to be prime either because V consists of two components of dimension r, or because V consists of a single component of dimension r, but this component occurs with "multiplicity."

Handling the case where V has a component of dimension less than r (which can happen only in the radical case), or the case where V has a single component of dimension r but with multiplicity (which can happen only in the CM case) is straightforward, as we have discussed in the previous section. Thus, we focus on the case when I is equidimensional, where deciding primality reduces to deciding if V has only one component of dimension r. Also, as mentioned before, by an effective version of Bertini's theorem, we can further assume that $r = 1$.

As discussed above, our approach involves studying how the ideal I and the zeroset V behave when the coefficients are changed from $\mathbb{Z}$ to $\mathbb{F}_p$ for some prime p. Let I_p be the ideal generated by $f_1 \pmod p, \cdots, f_m \pmod p$, and $V_p := Z(I_p) \subseteq \overline{\mathbb{F}}_p^n$. We need to show that three properties are preserved. The first is that for all but a small number of primes, $\dim V = \dim V_p$. The second is that if V is irreducible, then for all but a small number of primes, V_p is irreducible. The last is that if V is reducible, then for sufficiently many primes, V_p is also reducible, and crucially, V_p has at least two distinct components that are $\mathbb{F}_p$-definable.

Before we sketch our proof techniques for the above facts, we show that we cannot hope to get much stronger statements than the ones we claim. While the examples we will give here may appear simple, it is easy to generalize their main idea to produce examples with more complex behavior.

For the first property: dimension. Consider for example $I = (x_1, x_1 + ax_2)$ for some $a \in \mathbb{N}$. We have $\dim V = n - 2$. However, for any prime $p \mid a$, we have $I_p = (x_1)$ therefore $\dim V_p = n - 1$. Thus, we cannot hope that $\dim V_p$ remains unchanged for every base change from $\mathbb{Z}$ to $\mathbb{F}_p$.

For the second property: irreducibility. In this case, let $I = (x_1(x_1 - 1) - ax_2)$. Here, we have that V is irreducible, but for any prime $p|a$, the zeroset V_p is reducible.

And now for the third property: reducibility. Let's assume $n = 2$. Let $g \in \mathbb{Z}[x_1]$ be a monic polynomial of degree d, with Galois group S_d (the full symmetric group on d elements). Note that a random polynomial satisfies this property. Let f be the polynomial obtained by homogenizing g with respect to x_2, and let $I = (f)$. Then $\dim V = 1$, and V has d irreducible factors. The smallest extension of $\mathbb{Q}$ that contains all the factors of f is exactly the splitting field of g, which has degree $d!$, which is exponential in d. However, there are extensions of degree just d^2 that contain two irreducible factors, namely extensions generated by any pair of roots of g.

Now that we are aware of the pitfalls along the way, we can give a more precise idea of the quantitative bounds that we obtain. As is evident by the examples, the size of the set of bad primes for each of the statements will depend on the logarithmic height of the given polynomials. We will show that the size of the set of bad primes for the first two properties is polynomial in the degrees and logarithmic heights of $f_1, \ldots, f_m$, and exponential in the number of variables. As the logarithmic heights are themselves exponential in s (the size of the circuit), the bounds we obtain are exponential in the input size.

With the above in mind, in order for the set size lower bound protocols to work, when V is reducible, the number of good primes (that is, the number of primes p for which there are at least two $\mathbb{F}_p$ definable components) has to also be exponential, within a bound A of primes up to which we can check roots. In Remark 5.3 we show that the best one could hope to do in terms of the density of good primes is inverse exponential. Our main theorem will show that we can indeed achieve this inverse exponential density. This will show that there are sufficiently many good primes.

There are three main ideas/tools we use to prove the above properties. The first is the fact that ideal

membership can be written as a linear system, if bounds are known for the degrees of the witnesses. Here, by a witness to the membership of $g \in I$ we mean polynomials $h_1, \ldots, h_m$ such that $g = \sum f_i h_i$. In general these bounds can be doubly exponential on the input size [Her26]. However in the special cases of radical ideals and ideals of constant dimension, single exponential bounds are known ([Jel05], [MR13]). This allows us to show that certain memberships (and non-memberships) in ideals continue to hold when going modulo p, for all but a small number of primes. The second tool is an effective Bertini-Noether theorem [Kal95]. The Bertini-Noether theorem states that if a polynomial $f \in \mathbb{Z}[x_1, \ldots, x_n]$ is absolutely irreducible, then for all but finitely many primes p it remains absolutely irreducible in $\mathbb{F}_p$. The effective version of this theorem states gives an upper bound on the number of bad primes for f that depends on the logarithmic height of f. The third main tool we use is the factoring algorithm of [Kal95]. This algorithm gives a bound on the logarithmic heights of the coefficients of factors of polynomials. In this work, we combine the above ideals with a number of standard tools from commutative algebra and algebraic number theory, which can be found in Sections 2 and 3.

We now return to the three statements we want to prove.

Dimension of V_p. We want to show that $\dim V_p = 1$ for all but a small number of primes p (recall that we are assuming $\dim V = 1$). Since $\dim V = 1$, we have $I \cap \mathbb{Z}[x_i, x_j] \neq (0)$ for all pairs $i \neq j$. This is because $I \cap \mathbb{Z}[x_i, x_j]$ corresponds to projection to coordinates i, j, and such projections also have dimension at least 1. Further, if $\dim V = 1$, then $I \cap \mathbb{Z}[x_i] = (0)$ for some i, say $i = 1$. This is because projection to every coordinate cannot be finite, since V is not finite. These two properties characterise the fact that $\dim V = 1$.

By the effective Nullstellensatz, if $I \cap \mathbb{Z}[x_i, x_j] \neq (0)$, then it contains a polynomial g_{ij} with $\deg g_{ij} \leqslant d^n$, and such that $g_{ij} = \sum f_k h_{ijk}$ with $h_{ijk} \in \mathbb{Q}[x_1, \ldots, x_n]$ having $\deg h_{ijk} \leqslant d^n$. Moreover, by the effective nullstellensatz, a similar linear-algebraic condition can be derived (with the same degree bounds) which established that there is no non-zero polynomial $g_1 \in I \cap \mathbb{Z}[x_1]$. Since the degrees of the potential h_{ijk} are bounded, we can write the condition for existence of g_{ij}, g_1 as a linear system, where the unknowns are the coefficients of h_{ijk}. Using standard height bounds, we can deduce that for all but exponentially many primes, the existence and non-existence of solutions is preserved mod p, since these facts only depend on the vanishing and non vanishing of certain minors of the matrix of this linear system. Therefore, for all but exponentially many primes we have $\dim V_p = 1$. This result is formally stated and proved in Lemma 4.1.

Irreducibility of V_p. We want to show that if V is irreducible, then V_p is irreducible for all but a small number of primes p. Suppose $n = 2$, and $\dim V = 1$. Then V is irreducible if and only if $g \in \operatorname{rad}(I)$ for some absolutely irreducible polynomial g. Further, any such g must divide every generator of I. This easily allows us to get a bound on the coefficients of g, and by the effective Nullstellensatz, we can deduce that $g^e = \sum f_i h_i$ with $\deg g^e, \deg h_i \leqslant d^n$. By Cramer's rule, we can upper bound the common denominator of the polynomials h_i. For any prime p that does not divide the common denominator, we have $g^e \in I_p$. Further, the effective Bertini-Noether theorem, for all but a few primes p, the polynomial $g \pmod p$ remains absolutely irreducible. Finally, for all but a few primes, $\dim V_p = \dim V$. If p is any prime that passes all the above conditions, then V_p is irreducible.

The general case, where n is arbitrary, is slightly more involved. It is no longer true that V is irreducible if and only if $g \in \operatorname{rad}(I)$ for some absolutely irreducible polynomial g. For example, V could consist of two curves that both lie on the same irreducible hypersurface. To overcome this, we project to a random two dimensional linear subspace, call this projection π. Since V is irreducible, $\pi(V)$ is irreducible, and $\pi(V) = Z(g)$ for an absolutely irreducible polynomial g. A technical (and somewhat involved) argument gives us bounds on the coefficients of g in this more general setting (Lemma 3.14). Proceeding as above, we can show that for all but a few primes p, we have $\dim V_p = 1$ and $\pi(V_p)$ is irreducible.

However, the above is not enough to conclude that V_p is irreducible, since V_p could have been reducible, and all of its components could have been mapped to $\pi(V_p)$ under π. To fix this, instead of just considering a single random projection π, we consider a large number of random projections $\pi_1, \ldots, \pi_k$, and repeat the above argument. If V_p has two components, $\pi_i(V_p)$ can only be irreducible for a small fraction of these projections, and we can use this to bound all primes p such that V_p is not irreducible. This result is formally stated and proved in Theorem 4.5.

Reducibility of V_p and $\mathbb{F}_p$-definability of some of its components. We want to show that if V is reducible, then V_p is reducible, and has at least two $\mathbb{F}_p$-definable components for sufficiently many primes p. Again, let us start by assuming $n = 2$. If V has k irreducible components, then $V = Z(g)$ for a polynomial $g \in \mathrm{rad}(I)$ with exactly k absolutely irreducible factors. As before, g must divide the generators of I, which allows us to get a bound on the coefficients of g. By the effective Nullstellensatz, we can deduce that $g^e = \sum f_i h_i$ with $\deg g^e, \deg h_i \leqslant d^n$, and we can deduce that for all but a few primes p, we have $g^e \in I_p$. However, this does not imply that V_p is reducible: it might be the case that $V_p = Z(g')$ for some g' that is an absolutely irreducible factor of $g \pmod p$. Therefore, the strategy we used to show that irreducibility is preserved does not work here.

To prove that the above cannot happen (that is, that $V_p = Z(g')$), we try and find defining equations for some components of V. In the case when $n = 2$, these are exactly the factors of g. Suppose $g = \prod_{i=1}^{k} g_i$. Each g_i has coefficients in some algebraic extension of $\mathbb{Q}$. The smallest extension of $\mathbb{Q}$ that contains every g_i might be of prohibitively high degree, but from the factoring algorithm of Kaltofen (Lemma 3.15) we can deduce that there is a small enough extension of $\mathbb{Q}$ that contains g_1, g_2. Now we consider the ideals $I_1 := I_1 + g_1$ and $I_2 := I + g_2$. These are ideals of dimension 1 that are irreducible. Therefore, we can use our base change and irreducibility arguments to deduce that $I_{1,p}$ and $I_{2,p}$ are irreducible curves, which will be components of I_p. Of course we have to be careful here: g_1 and g_2 have coefficients in $\mathbb{Q}(\alpha)$ for some algebraic number α, therefore it is not even clear at first what it means to go modulo p. Here, we need to invoke some standard results from algebraic number theory. If G is the monic equation of α, then for any p such that $G \pmod p$ has a root, the operation of going mod p is well defined on $\mathbb{Z}[\alpha][x_1, \ldots, x_n]$. This allows us to not only make sense of the mod p operation and invoke our previous result, but it also guarantees that the components of V_p that we find this way are $\mathbb{F}_p$ definable. It is important to note here that there are a number of technical challenges in extending our earlier statements to this general algebraic number theoretic setting. These challenges are dealt with in Section 3.

Now suppose n is arbitrary. We again find equations that define the components of V. As before we do this by random projection to a two dimensional affine subspace, and then we apply (some of) our arguments from the case of $n = 2$. It turns out that a single extra polynomial g_1 is not enough to define a component of V. However, we show any component of V can be defined by adding at most two equations to I. As before, we give bounds on the coefficients of these equations, and then invoke our earlier arguments to deduce that V_p is reducible, and for sufficiently many primes it has at least two components that are $\mathbb{F}_p$ definable. This result is formally stated and proved in Theorem 4.6.

1.1.1 Putting it all together

We now show how all the results mentioned above are put together to obtain our main protocol. For brevity, we only consider the case when I is radical and $\dim V = 1$, as the other cases are very similar. We now sketch an AM protocol that shows that I is not prime.

In the radical case, I can fail to be prime for 2 reasons: V may have a component of strictly smaller dimension (hence in this case V will have at least one isolated point as a solution), or V is equidimensional with at least two components.

Case 1: V has an isolated point as a solution. Consider the Jacobian matrix $\mathcal{J}$, whose entries are $\mathcal{J}_{ij} = \partial_i f_j$. If α is an isolated point of V, then the rank of $\mathcal{J}(\alpha)$ is n. Merlin first sends Arthur the index of n columns of $\mathcal{J}$ that are linearly independent. If this submatrix is $\mathcal{J}'$, then Arthur simply has to check if the system $f_1 = 0, \ldots, f_m = 0, \det \mathcal{J}' \neq 0$ is satisfiable. This can be rewritten as $f_1 = 0, \ldots, f_m = 0, y \det \mathcal{J}' - 1 = 0$ where y is a new variable. Now, the satisfiability protocol of [Koi96] can be used to decide I is not prime, as in this case the protocol of [Koi96] will return that the system that we constructed will have a solution (which must correspond to an isolated point).

On the other hand, if V only has components of dimension 1, then at every point on V, the rank of $\mathcal{J}$ is at most $n-1$. Therefore, no matter what choice of columns Merlin picks, the above protocol will fail to accept with high probability.

Case 2: V has two components of dimension 1. For every prime p such that V_p has at least two irreducible components, the Lang-Weil theorem says that the number of $\mathbb{F}_p$ points in V_p is at least $2p$ (up to an error

term). If V has two components, then we have shown that there are enough primes with this property. On the other hand, if V is irreducible, for all except a finite number of primes, V_p is irreducible, and therefore the Lang-Weil theorem says that the number of $\mathbb{F}_p$ points in V_p is at most p (up to an error term). Therefore, Merlin can prove to Arthur that V is reducible, by performing the set lower bound protocol on the set of primes for which the number of $\mathbb{F}_p$ points in V_p is at least $2p$. Merlin also has to convince Arthur that there are in fact $2p$ points in V_p, since the primes p are too large for Arthur to check this himself. For this, the set lower bound protocol is used again, on the set of $\mathbb{F}_p$ points of V_p.

1.2 Acknowledgements

N.S. is grateful to the University of Waterloo for an invite to the distinguished lecture series in Jun'24. The discussions done in Waterloo's conducive environment led to this work. In addition, N.S. thanks the funding support from DST-SERB (JCB/2022/57) and N.Rama Rao Chair. A.G. would like to thank TIFR for very generously hosting him during the later part of this project. Part of this work took place while R.O. and N.S. participated in the workshop on "Algebraic and Analytic Methods in Computational Complexity" in Dagstuhl (Sep'24). The authors would like to thank Schloss Dagstuhl for providing excellent hospitality and conditions to conduct research discussions. The authors would like to thank Peter Bürgisser and Pascal Koiran for helpful conversations about this project during the workshop at Dagstuhl.

2 Preliminaries

In this section we establish the notation that will be used throughout the paper, along with the background results that we will need in the later sections.

2.1 Notation

Let $R := \mathbb{Z}[x_1, \ldots, x_n]$ denote the n-variate polynomial ring with integer coefficients, and let $S := \overline{\mathbb{Q}}[x_1, \ldots, x_n]$ denote the extension of scalars to the algebraic closure of $\mathbb{Q}$. All the ideals we discuss are ideals of S, unless stated otherwise. The same holds for any property of the ideals we discuss, for example the property of an ideal being prime, or Cohen-Macaulay, among others. We use $\mathbb{A}^n$ to denote the affine space in n-dimensions, with underlying field $\overline{\mathbb{Q}}$. We use $Z(I)$ to denote the zeroset of an ideal I.

In the course of our proofs, we will have to deal with polynomials with coefficients in finite extensions of $\mathbb{Q}$. Given a monic irreducible polynomial $q \in \mathbb{Z}[z]$, and a root $\alpha \in \overline{\mathbb{Q}}$ of q, we use A to denote the ring $\mathbb{Z}[\alpha][x_1, \ldots, x_n]$. Note that $\mathbb{Z}[\alpha]$ is not necessarily the ring of integers of $\mathbb{Q}(\alpha)$. Elements of A will be represented as polynomials in $R[z]$, with z-degree less than $\deg q$. This representation of elements of A is unique. The ring A depends on q, however q will always be clear from context, and therefore we suppress it from notation. In this setting, by $\deg f$ we mean the degree of f in the x-variables, unless stated otherwise.

The logarithmic height of an integer c, denoted by $\mathrm{ht}(c)$, is the bit-complexity of c. This notion of logarithmic heights extends naturally to polynomials. Given a polynomial $f \in R$, the logarithmic height of f, also denoted by $\mathrm{ht}(f)$, is the maximum logarithmic height of the coefficients of f. If $f \in A$, then $\mathrm{ht}(f)$ is defined to be the logarithmic height of f treated when written as a polynomial in $R[z]$ with z-degree less than $\deg q$.

Given polynomials $f_1, \ldots, f_m, g$ with $\deg f_i, \deg g \leqslant d$, the condition $g \in (f_1, \ldots, f_m)$ is equivalent to the existence of polynomials $h_1, \ldots, h_m$ such that $g = \sum f_i h_i$. If a bound on the degrees of h_i is known a priori (which is usually the case), then the above condition can be written as a linear system where the coefficients of $h_1, \ldots, h_m$ are the unknowns. If this bound on the degrees of h_i is D, we use $M_D(f_1, \ldots, f_m)$ to denote the matrix corresponding to this linear system. The entries of this matrix are coefficients of $f_1, \ldots, f_m$. When $f_1, \ldots, f_m$ are clear from context we denote this matrix by just M_D. Observe that the total number of columns of M_D, that is, the number of unknowns in the system is at most $m \cdot \binom{D+n}{n}$. The total number of equations is at most $\binom{D+d+n}{n}$. Both these estimates easily follow from counting the number of monomials of given degrees. When reasoning using this linear system, we will slightly overload notation, and use the same symbols to refer to both polynomials and their coefficient vectors. For example, we write

the condition $g \in (f_1, \ldots, f_m)$ as $g = M_D v$ where v is a vector of unknowns. We further say that $h_1, \ldots, h_m$ are a solution to the system.

2.2 Results from complexity theory

We begin by describing an AM protocol that lower bounds the size of sets that have an efficient membership test. The protocol is due to [GS86], and the following statement is from [AB09, Section 8.4.1].

Lemma 2.1. *Suppose $S \subset \{0,1\}^n$ is a set such that the problem of membership in S is in NP. Suppose further that a number K is known, and S is guaranteed to either satisfy $|S| \geqslant 2K$ or $|S| < K$. Then the problem of deciding if $|S| \geqslant 2K$ is in* AM.

We sketch the protocol here, since it will make it easier to explain the generalisation we require. However, we omit proofs. We also assume for convenience that K is a power of 2.

Algorithm 1: Goldwasser-Sipser set lower bound protocol

Input : Boolean formula $\phi_S(x, y)$ such that $x \in S$ if and only if there exists y satisfying $\phi_S(x, y)$, and an integer K.
Arthur : Let $k := \log_2(2K)$, and pick a random hash function $h : \{0,1\}^n \to \{0,1\}^{k+1}$ from a pairwise independent hash function collection, and pick $u \in \{0,1\}^{k+1}$ uniformly at random. Send h, u to Merlin.
Merlin : Find x, y such that $h(x) = u$ and $\phi_S(x, y)$ is true. Send x, y to Arthur.
Arthur : Accept if and only if $h(x) = u$ and $\phi_S(x, y)$ is true.

We now state a slight generalisation of this protocol, to the case when membership in S itself can be verified by an AM protocol.

Corollary 2.2. *Suppose for each $x \in \{0,1\}^n$ there exists a subset $S_x \subset \{0,1\}^{m(x)}$, and an integer $K(x)$, such that $m(x), K(x)$ are polynomially bounded functions of x. Suppose further that there is a uniform algorithm that runs in polynomial time that given input x, returns the number $K(x)$ and also returns a boolean formula ϕ_x, such that $z \in S_x$ if and only if there exists y such that $\phi_x(z, y)$ is true. Suppose further than an integer K is known.*

If $S \subset \{0,1\}^n$ is the set of elements x such that $|S_x| \geqslant 2K(x)$, and if S is promised to either satisfy $|S| \geqslant 2K$ or $|S| \leqslant K$, then the problem of deciding if $|S| \geqslant 2K$ is in AM.

We give a four round protocol that decides $|S| \geqslant 2K$. The result then follows from the fact that AM = AM[c] for all constants $c \geqslant 2$. The proof of correctness of the protocol is the exact same as the proof of correctness of Lemma 2.1, therefore we omit the proof.

Algorithm 2: Goldwasser-Sipser set lower bound protocol with membership in AM

Input : An integer K, and an algorithm that on input x outputs the integers $K(x)$ and the circuit ϕ_x described in Corollary 2.2
Arthur : Let $k := \log_2(2K)$, and pick a random hash function $h : \{0,1\}^n \to \{0,1\}^{k+1}$ from a pairwise independent hash function collection, and pick $u \in \{0,1\}^{k+1}$ uniformly at random. Send h, u to Merlin.
Merlin : Find x such that $h(x) = u$ and $|S_x| \geqslant 2K(x)$. Send x to Arthur.
Arthur : Reject if $h(x) \neq u$. If $h(x) = u$, let $k_x := \log_2(2K(x))$, and pick a random hash function $h_x : \{0,1\}^{m(x)} \to \{0,1\}^{k_x+1}$ from a pairwise independent hash function collection, and pick $u_x \in \{0,1\}^{k_x+1}$ uniformly at random. Send h_x, u_x to Merlin.
Merlin : Find y, z such that $h_x(z) = u_x$ and such that $\phi_x(z, y)$ is true. Send z, y to Arthur.
Arthur : Accept if and only if $h_x(z) = u_x$ and $\phi_x(z, y)$ is true.

Remark 2.3. The choice of the constant 2 in the above protocol is arbitrary, and 2 can be replaced by a slightly smaller constant, say 1.9.

We now state the main theorems of Koiran, giving interactive protocols for the problem of deciding whether a system of polynomial equations has a solution, and for the problem of estimating the dimension of an algebraic set. We first state the main theorem of [Koi96], on the Hilbert nullstellensatz problem.

Theorem 2.4. *Assume GRH. Given* $f_1, \ldots, f_m \in R$*, there is an* AM *protocol that decides if* $Z(f_1, \ldots, f_m) \neq \emptyset$*.*

The above theorem was further generalised by Koiran in [Koi97, Theorem 4.1], where now one wants to decide a lower bound on the dimension of the given algebraic set.

Theorem 2.5. *Assume GRH. Given* $f_1, \ldots, f_m \in R$*, and an integer* r*, there is an* AM *protocol that decides if* $\dim Z(f_1, \ldots, f_m) \geqslant r$*.*

Remark 2.6. We make two remarks on the proofs of Theorem 2.4 and Theorem 2.5 that will be crucial to the way we invoke these results. The first of these is regarding the representation of input polynomials. The AM protocols for the above problems involve evaluating the polynomials $f_1, \ldots, f_m \pmod p$ at points in $\mathbb{F}_p$, where p is a prime of bit complexity polynomial in the input. Therefore, the polynomials $f_1, \ldots, f_m$ are allowed to be given either as white box circuits of polynomial size (here the size includes the bit complexity of the constants in the circuit), or more generally as black-boxes that allow mod p queries.

The second remark is regarding the parameters. Suppose $\deg f_i \leqslant d$ and $\mathrm{ht}(f_i) \leqslant h$. The length of the messages and the computation done by Arthur in the protocols in Theorem 2.4 and Theorem 2.5 is polynomial in $\log\left(h \cdot 2^{(n \log \sigma)^c}\right)$, for a universal constant c, where $\sigma := dm + 2$. Therefore, if we create a system of polynomials $g_1, \ldots, g_m$ with $\deg g_i \leqslant d^n$ and $\mathrm{ht}\left(h \cdot 2^{(n \log \sigma)^c}\right)$, then the protocols in Theorem 2.4 and Theorem 2.5 applied to $g_1, \ldots, g_m$ still run in time poly(n, m, d, h) as long as we ensure that $g_1, \ldots, g_m$ have circuits of size poly(n, m, d, h). This will be crucial in our proofs.

Lastly, we show that the ideal primality testing problem is coNP-hard, even for the special case of equidimensional Cohen-Macaulay ideals.

Proposition 2.7 (coNP-hardness of ideal primality testing)**.** *Given a 3CNF* Φ *on* n *variables, there exists a polynomial time algorithm whose output is a circuit* C_Φ *computing polynomials* $f_1, \ldots, f_m \in R$*, along with their partial derivatives* $\partial_i f_j$*, such that the ideal* $I := (f_1, \ldots, f_m)S \subset S$ *has the following properties:*

- I *is a radical, equidimensional Cohen-Macaulay ideal*
- I *is prime if and only if* Φ *is unsatisfiable.*

Proof. Let $x_1^2 - x_1, x_2^2 - x_2, \ldots, x_n^2 - x_n, g_1, \ldots, g_r$ be the arithmetization of Φ. The instance Φ is satisfiable if and only if the arithmetized system has a solution.

The reduction algorithm proceeds as follows. First check if $(0, \ldots, 0)$ is a solution to the above system. If it is, then the algorithm returns the system $f_1 = x_1(x_1 - 1)$, clearly there is a constant-sized circuit that computes $f_1, \partial_i f_1$.

Suppose $(0, \ldots, 0)$ is not a solution to the above system. Consider the set of polynomials $\{x_i^2 - x_i\}_{i \in [n]} \cup \{x_j g_i\}_{i \in [r], j \in [n]}$. Since each g_i has a small circuit, there is a small circuit that computes all the above polynomials and all their derivatives. Further, the set of zeroes of the above system is exactly those points on the boolean hypercube that correspond to a satisfying assignment of ϕ, and also the point $(0, \ldots, 0)$.

In each of the above cases, by [Sei74, Lemma 92], the ideal constructed is radical. In the first case the ideal is a complete intersection therefore CM. In the second case the ideal is zero dimensional, therefore CM. If Φ is satisfiable, then in either case the system created has at least two components and is therefore not prime. If Φ is not satisfiable, the system created has unique solution $(0, \ldots, 0)$, thus the ideal is prime. □

2.3 Algebraic circuits

The inputs to our algorithms are given as algebraic circuits. Algebraic circuits are natural models for computing polynomials. They consist of a directed acyclic graph, with nodes marked with $+$ and $\times$. The source nodes are marked either with variables x_i or with the constant 1. Each internal node computes the natural polynomial: nodes marked with $+$ compute the sums of their inputs, and nodes marked with $\times$ compute

the product. The edges are marked with constants, if (u, v) is an edge with constant α, then the input to v corresponding to u is αf_u, where f_u is the polynomial computed at the node u. We refer the reader to the excellent survey [SY10] for more background on algebraic circuits.

Circuits are studied both as uniform and as non-uniform models of computation. Further, they are studied both in the setting where all constants are considered to have size 1, and also in the setting where the logarithmic height of the constants is part of the size of the circuit. Since we are using circuits as inputs to a problem in the Turing machine model of computation, we are naturally in the setting where the size of circuit includes the logarithmic heights of the constants of the circuit. We assume that the constants in the circuit are all integers.

Size of a circuit: The size of an algebraic circuit is the sum of the logarithmic heights of the constants on the edges of circuits. We assume that every edge has a constant (potentially 1), therefore the size of a circuit is a lower bound on the number of edges of the circuit.

Evaluating a circuit modulo p: Given a circuit C, a prime p, and a point $\alpha \in \mathbb{F}_p^n$, there is a polynomial time algorithm that computes the evaluations at α of all the mod p reductions of all polynomials computed by C. The algorithm recursively computes the evaluation $\pmod p$ of the inputs to a given node in the circuit, and then performs arithmetic in $\mathbb{F}_p$ to compute the evaluation of the gate itself.

Structural results: We need the following technical result [Bür13, Lemma 4.16] that bounds the logarithmic heights and degrees of polynomials computed by circuits of size s. The result follows by induction on the circuit, and using the fact that every coefficient in the circuit has logarithmic height at most s.

Lemma 2.8. *If C is a circuit of size s computing polynomials $f_1, \ldots, f_m \in R$ then $\mathrm{ht}(f_i) \leqslant 2^{2s}$ and $\deg f_i \leqslant 2^s$.*

We also need the following result from [BS83] that proves that given a circuit C, there is a circuit C' of similar size that also computes the partial derivatives of the polynomials computed by C.

Lemma 2.9. *If C is a circuit of size s computing polynomials $f_1, \ldots, f_m \in R$ there is a circuit C' of size $5sm$ that computes $f_1, \ldots, f_m$ and all the partial derivatives $\partial_i f_j$.*

2.4 Results from linear algebra

Cramer's rule is a well known explicit formula for the solution of a linear system of the form $Ax = b$, when A is a $n \times n$ matrix with $\det(A) \neq 0$. Under these conditions, the unique solution to the above system is given by $x_i = \frac{\det A_i}{\det A}$, where A_i is the matrix obtained by replacing the i^{th} column of A with the vector b. The following easy consequence shows that a similar formula exists for under and overdetermined systems, provided a solution is promised to exist. Note that the lemma applies for matrices with coefficients in any domain.

Lemma 2.10. *Suppose A is an $n \times m$ matrix, and suppose the linear system $Ax = b$ is guaranteed to have a solution. Then there exists a solution where each x_i is either 0 or of the form $\frac{\det M_i}{\det N}$, where $\det M_i, \det N$ are minors of the augmented matrix $A|b$.*

Proof. Let $r := \operatorname{rank} A$, so $r \leqslant m, n$. The fact that $Ax = b$ has a solution is equivalent to the fact that the augmented matrix $[A|b]$ also has rank r. After rearranging the columns, we can assume that the first r columns of A are linearly independent. Since b lies in the column span of A, it also lies in the column span of the first r columns of A. Equivalently, if we set the variables $x_{m-r+1}, \ldots, x_m$ to 0, the resulting system still has a solution. It suffices to show the result for the resulting system, therefore we can assume that A has full column rank, so $r = m$.

After further rearranging rows, we can assume that the first r rows of A are linearly independent, and the remaining rows are linear combinations of the first r rows. The augmented matrix $[A|b]$ also has rank r, therefore the last $n - r$ rows of $[A|b]$ are linear combinations of the first r rows. If we write $A'x = b'$ for the linear system consisting of the first r rows of $Ax = b$, then any solution to $A'x = b'$ is also a solution to $Ax = b$. Applying Cramer's rule to $A'x = b'$ gives us the desired result. □

2.5 Results from algebraic geometry

In this subsection, we collect some useful basic results from algebraic geometry which we will need in the later sections. We begin by stating a characterisation of the dimension of an algebraic set.

Lemma 2.11. *Suppose $I \subset S$ is an ideal, and V is the zeroset of I. Then V has dimension r if and only if the following two conditions hold.*

- *There is some subset $U \subset [n]$ of size r such that the elimination ideal I_U is empty.*
- *For every subset $U' \subset [n]$ of size $r+1$, the elimination ideal $I_{U'}$ is nonempty.*

We now state some standard facts about tangent spaces, which can be found in [SR94, Chapter 2].

Let I be a radical ideal generated by $f_1, \ldots, f_m$. Let $\mathcal{J}$ be the Jacobian of I, which is defined to be the matrix with $\mathcal{J}_{ij} = \partial f_i / \partial x_j$. At any point $x \in V$, the tangent space $T_x(V)$ is isomorphic to the kernel of $\mathcal{J}(x)$, where $\mathcal{J}(x)$ is the Jacobian entrywise evaluated at the point x [SR94, Chapter 2, Section 1, Theorem 2.1]. At every nonsingular point x, the dimension of the tangent space is exactly equal to dimension of the component of V passing through x [SR94, Chapter 2, Section 1, Theorem 2.3]. At every singular point x, the dimension of the tangent space is greater than the dimension of the components of V passing through x. The singular locus of V is a proper algebraic subset of V, and does not contain any irreducible component of V.

We deduce that there is a point $x_0 \in V$ such that $\dim T_{x_0}(V) < r$ if and only if V has an irreducible component of dimension less than r. Equivalently, there is a point $x_0 \in V$ such that $\operatorname{rank} \mathcal{J}(x_0) > n - r$ if and only if V has an irreducible component of dimension less than r.

The next two statements can be found in [Hei83, Theorems 1 and 2], and are referred to as Bézout's inequality.

Theorem 2.12 (Bézout's inequality)**.** *Let $\mathbb{K}$ be an algebraically closed field and $X, Y \subseteq \mathbb{A}^n_{\mathbb{K}}$ be constructible sets. Then, we have*

$$\deg(X \cap Y) \leqslant \deg X \cdot \deg Y.$$

Theorem 2.13 (Degree bounds for constructible sets)**.** *Let $\mathbb{K}$ be an algebraically closed field and $f_1, \ldots, f_m \in \mathbb{K}[x_1, \ldots, x_n]$ be polynomials of degree at most d. Then, we have*

$$\deg V(f_1, \ldots, f_m) \leqslant d^{\min\{m,n\}}.$$

The following lemma gives effective bounds on hyperplane sections that reduce the dimension of projective varieties, a proof can be found in [GSS19, Lemma 19]

Lemma 2.14. *Suppose $V \subset \mathbb{P}^n$ is a projective zeroset of dimension r and degree D. If ℓ is a linear form where each coefficient is chosen uniformly and independently from a set $B \subset \mathbb{N}$, then $\dim V \cap Z(\ell) = r - 1$ with probability at least $1 - D/|B|$. If $\ell_1, \ldots, \ell_{r+1}$ are linear forms chosen the same way then $V \cap Z(\ell_1, \ldots, \ell_{r+1}) = \emptyset$ with probability at least $1 - (r+1)D/|B|$.*

Suppose $W \subset \mathbb{A}^n$ is an affine zeroset of dimension r and degree D. If ℓ is a linear polynomial where each coefficient is chosen uniformly and independently from a set $B \subset \mathbb{N}$, then $\dim W \cap Z(\ell) = r - 1$ with probability at least $1 - 2D/|B|$. If $\ell_1, \ldots, \ell_{r+1}$ are linear polynomials chosen the same way then $W \cap Z(\ell_1, \ldots, \ell_{r+1}) = \emptyset$ with probability at least $1 - 2(r+1)D/|B|$.

The following is a sufficient condition for a linear map to be a Noether normalising map. This statement, and its proof can be found in [SR94, Chapter 1, Section 5, Theorem 1.15].

Lemma 2.15. *Suppose $V \subset \mathbb{P}^n$ is a projective variety disjointed from a k-dimensional linear subspace $E \subset \mathbb{P}^n$. Then the projection $\pi : X \to \mathbb{P}^{n-k-1}$ with center E defines a finite map $X \to \pi(X)$.*

Combining the two lemmas above gives us the following effective Noether Normalisation theorem.

Lemma 2.16 (Effective Noether Normalisation)**.** *Suppose $V \subset \mathbb{P}^n$ is a projective zeroset of dimension r and degree D. If $\ell_1, \ldots, \ell_{r+1}$ are linear forms where each coefficient is chosen uniformly and independently from a set $B \subset \mathbb{N}$, then the map $\pi : V \to \mathbb{P}^r$ with coordinate functions ℓ_i is a well defined finite map with probability at least $1 - (r+1)D/|B|$.*

Suppose $W \subset \mathbb{A}^n$ is an affine zeroset of dimension r and degree D. If $\ell_1, \ldots, \ell_r$ are linear polynomials where each coefficient is chosen uniformly and independently from a set $B \subset \mathbb{N}$, then the map $\pi : W \to \mathbb{A}^r$ with coordinate functions ℓ_i is a well defined finite map with probability at least $1 - 2(r+1)D/|B|$.

Proof. For the projective case, by Lemma 2.14, the zeroset $Z(\ell_1, \ldots, \ell_{r+1})$ is a linear space disjointed from V with probability at least $1-(r+1)D/|B|$. If this holds, then π is a finite and well defined map by Lemma 2.15. The affine case follows by applying the projective case to the closure of W_p, and picking $\ell_{r+1} = x_0$. □

2.6 Results from number theory

We state some algebraic number theory facts that we will need. We do not provide any proofs since these facts are standard, [Mil20, Chapter 2, 3] is an excellent exposition of all of these results.

Let $q \in \mathbb{Z}[z]$ be a monic irreducible polynomial with $\deg q = e$, and let α be a root of q. Let $\mathbb{F} := \mathbb{Q}(\alpha)$ be the algebraic extension of $\mathbb{Q}$ generated by α, we have $\mathbb{F} \cong \mathbb{Q}[z]/(q)$. Let $O_{\mathbb{F}}$ be the ring of integers of $\mathbb{F}$, that is, the set of elements $\beta \in \mathbb{F}$ such that β satisfies a monic equation with coefficients in $\mathbb{Z}$. We have $\mathbb{Z} \subset O_{\mathbb{F}}$, and $O_{\mathbb{F}}$ is integrally closed in $\mathbb{F}$. The discriminant of q, denoted $\operatorname{disc}_z(q)$ is defined to be $\operatorname{res}_z(q, \partial q)$.

Usually the ring of integers $O_{\mathbb{F}}$ is different from $\mathbb{Z}[\alpha]$. We have $\mathbb{Z}[\alpha] \subset O_{\mathbb{F}}$, but this inequality might be strict. However, it holds that $O_{\mathbb{F}} \subset (1/\operatorname{disc}_z(q))\mathbb{Z}[\alpha]$. Therefore we can represent any algebraic integer as an element of $\mathbb{Q}[z]$, where the coefficients have common denominator $\operatorname{disc}_z(q)$.

The ring $O_{\mathbb{F}}$ is a Dedekind domain, and every nonzero prime ideal of $O_{\mathbb{F}}$ is maximal. For any prime $\mathfrak{p} \subset O_{\mathbb{F}}$, there is a unique prime $p \in \mathbb{N}$ such that $\mathfrak{p} \cap \mathbb{Z} = (p)$. If this happens, then we say $\mathfrak{p}$ lies above p. While $O_{\mathbb{F}}$ is not factorial in general, it admits unique factorisation of ideals. Every ideal $I \subset O_{\mathbb{F}}$ can be the written uniquely as a product of primes, $I = \prod \mathfrak{p}_i^{e_i}$. In particular, the ideal (p) for a prime $p \in \mathbb{N}$ itself can be factored as $(p) = \prod \mathfrak{p}_i^{e_i}$. The set of primes ideals of $O_{\mathbb{F}}$ occurring in this factorisation of (p) are exactly the prime ideals of $O_{\mathbb{F}}$ that lie above p. When $p \nmid \operatorname{disc}_z(q)$, $e_i = 1$ for all i, if this holds we say that p is unramified.

For any unramified p, and $\mathfrak{p}$ lying above p, the quotient map $O_{\mathbb{F}} \to O_{\mathbb{F}}/\mathfrak{p}$ is well structured. Since $\mathfrak{p}$ is maximal, and since $p \in \mathfrak{p}$, the quotient $O_{\mathbb{F}}/\mathfrak{p}$ is a field of characteristic p, and is in fact a finite extension of $\mathbb{F}_p$. The prime $\mathfrak{p}$ corresponds to an irreducible factor q_1 of $q \pmod p$, and $O_{\mathbb{F}}/\mathfrak{p} \cong \mathbb{F}_p[z]/q_1(z)$. Composing with the inclusion map $\mathbb{Z}[\alpha] \to O_{\mathbb{F}}$ gives us a map $\mathbb{Z}[\alpha] \to \mathbb{F}_p[z]/q_1(z)$. We usually apply this to the case when q_1 is a linear polynomial, and therefore $\mathbb{F}_p[z]/q_1(z) = \mathbb{F}_p$. This map therefore extends the usual map $\mathbb{Z} \to \mathbb{F}_p$.

We now state a corollary of Gauss's lemma. The following formulation is from [WR76, Lemma 7.1], we continue to refer to it as Gauss's lemma.

Lemma 2.17 (Gauss's lemma)**.** *Suppose $\delta \in O_{\mathbb{F}}$ and $f \in (1/\delta)O_{\mathbb{F}}$ is a monic polynomial with factorisation $f = gh$ in $\mathbb{F}$. If g, h are monic then $g, h \in (1/\delta)O_{\mathbb{F}}$.*

We will need three more technical results from number theory. The first of these is an effective Lang-Weil bound. The Lang-Weil bounds are upper and lower bounds on the number of $\mathbb{F}_p$ points on irreducible varieties that are $\mathbb{F}_p$ definable. Classically these bounds have error terms, and effective versions of such bounds give bounds on the coefficients of the error terms. The following is a version of such a bound from [CM06, Theorem 7.1].

Theorem 2.18 (Effective Lang-Weil bound)**.** *Let $V \subset \mathbb{F}_p^n$ be an absolutely irreducible affine variety defined over $\mathbb{F}_p$ of dimension r and degree D. Suppose also that $p > 2(r+1)D^2$. Then*

$$|V(\mathbb{F}_p) - p^r| < (D-1)(D-2)p^{r-1/2} + 5D^{13/3}p^{r-1}.$$

The second number theoretic statement we need is a lower bound on the number of primes p for which $q \pmod p$ has a root in $\mathbb{F}_p$. Observe that the Chebotarev density theorem states that this fraction of primes is the same as the fraction of the Galois group of q that has at least one fixed point, and this fraction is at least $1/e$. We require an effective version of this statement. The following statement is from [Koi96, Corollary 1]. The proof itself is based on a bound by [AO83] which in turn is based on an effective version of the Chebotarev density theorem [LO77]. We note that this last result is conditional, and assumes the GRH, and this is also what makes our result conditional.

Theorem 2.19. *Assuming the GRH holds, there exists an absolute constant c such that*

$$\pi_q(x) \geqslant \frac{1}{e}\left(\pi(x) - \log \operatorname{disc}_z(q) - cx^{1/2}\log\left(\operatorname{disc}_z(q)\, x^e\right)\right),$$

where $\pi(x)$ is the prime counting function and $\pi_q(x)$ is the number of primes $p \leqslant x$ such that $q \pmod p$ has a root in $\mathbb{F}_p$.

Finally, we need an unconditional lower bound on the prime counting function $\pi(x)$. One such estimate is provided in [Dus10, Theorem 6.9].

Theorem 2.20. *For $x \geqslant 600$ we have $\pi(x) > x/\ln x$.*

2.7 Degree bounds for polynomial ideals

We now recollect some important complexity results for polynomial ideals. We begin with Jelonek's effective Nullstellensatz, from [Jel05, Theorem 1.1].

Theorem 2.21 (Effective Nullstellensatz)**.** *Let $\mathbb{K}$ be an algebraically closed field, and $X \subset \mathbb{K}^m$ be an affine n-dimensional variety of degree D. Let $f_1, \ldots, f_m \in \mathbb{K}[X]$ be non-constant polynomials without common zeros, where $d_i := \deg f_i$, and $d_1 \geqslant \cdots \geqslant d_m$. Lastly, let*

$$N(d_1, \ldots, d_m; n) := \begin{cases} \prod_{i=1}^{m} d_i & \text{if } n > 1 \text{ and } n \geqslant m \\ d_m \cdot \prod_{i=1}^{n-1} d_i & \text{if } m > n > 1 \\ d_1 & \text{if } n = 1 \end{cases}$$

There exist polynomials $g_i \in \mathbb{K}[X]$ with

$$\deg(f_i g_i) \leqslant \begin{cases} D \cdot N(d_1, \ldots, d_m; n), & \text{if } m \leqslant n \\ 2D \cdot N(d_1, \ldots, d_m, n) - 1, & \text{if } m > n \end{cases}$$

such that $1 = \sum_{i=1}^{m} f_i g_i$.

We now state some useful bounds on the complexity of representation of the reduced Gröbner basis of any ideal. The bounds that we state below are dependent on the Krull dimension of the given ideal, and this will be crucial in our applications in the later sections. These bounds are from [MR13, Theorem 4].

Theorem 2.22 (Gröbner basis complexity)**.** *Let $\mathbb{K}$ be an infinite field and $I = (f_1, \ldots, f_m) \subseteq \mathbb{K}[x_1, \ldots, x_n]$ be an ideal of dimension r, where $d_j := \deg(f_j)$ and $d_1 \geqslant \cdots \geqslant d_m$. Let*

$$B := 2 \cdot \left(\frac{1}{2} \left((d_1 \cdots d_{n-r})^{2(n-r)} + d_1 \right) \right)^{2^r}.$$

Given any admissible monomial ordering, if $G = \{g_1, \ldots, g_t\}$ is the reduced Gröbner basis of I, then we have:

$$\deg(g_i) \leqslant B \quad \forall i \in [t].$$

Moreover, for each $i \in [t]$, there are polynomials $h_{i1}, \ldots, h_{im} \in \mathbb{K}[x_1, \ldots, x_n]$ satisfying $\deg(f_j \cdot h_{ij}) \leqslant B$ such that

$$g_i = \sum_{j=1}^{m} f_j h_{ij}.$$

The above bounds, when combined with the division algorithm, allow us to derive the following upper bound on the dimension of the linear system (over the base field) needed to check membership in polynomial ideals.

Corollary 2.23 (Representation degree)**.** *Let $\mathbb{K}$ be an infinite field and $I = (f_1, \ldots, f_m) \subseteq \mathbb{K}[x_1, \ldots, x_n]$ be an ideal of dimension r, where $d_j := \deg(f_j)$ and $d_1 \geqslant \cdots \geqslant d_m$. Let*

$$B := 2 \cdot \left(\frac{1}{2} \left((d_1 \cdots d_{n-r})^{2(n-r)} + d_1 \right) \right)^{2^r}.$$

Given any polynomial $f \in \mathbb{K}[x_1, \ldots, x_n]$ with $\deg(f) = D$, we have that $f \in I$ if, and only if, there exist $g_1, \ldots, g_m \in \mathbb{K}[x_1, \ldots, x_n]$ such that $\deg(f_j g_j) \leqslant \max(B, D)$ and $f = \sum_{j=1}^{m} f_j g_j$.

2.8 An Effective Bertini Theorem

Bertini's second theorem states that the intersection of an irreducible variety of dimension $r \geqslant 2$ with a random hyperplane of dimension $n - r + 1$ is an irreducible curve. This statement allows us to reduce the problem of irreducibility testing of arbitrary high dimensional varieties to that of curves and points. The theorem plays an important role in many routines in computer algebra, a comprehensive survey of its applications can be found in [Dic05, Section 9.1.3]. The effective versions of the theorem are usually stated for hypersurfaces, reduction from the general case to hypersurfaces follows from a projection argument.

In order to apply the theorem to arbitrary varieties, a preprocessing projection step is applied. The following two results are from [CM06], where the effective Bertini theorem is used to obtain improved Lang-Weil bounds. While the results in [CM06] are stated for varieties over $\mathbb{F}_q$, the following two results also hold for varieties over any perfect field, with the same proofs.

Lemma 2.24 ([CM06, Proposition 6.1, Proposition 6.3])**.** *Let V be an equidimensional affine variety of dimension r and degree D. Let Λ be a $(r+1) \times n$ matrix of variables, and let Γ be a $r+1$ dimensional vector of variables. There exists a nonzero polynomial $G \in \overline{\mathbb{Q}}[\Lambda, \Gamma]$ of degree at mot $2(r+1)D^2$ such that for any values λ, γ of the variables Λ, Γ satisfying $G(\lambda, \gamma) \neq 0$, the projection π defined by $\ell_i = \sum \lambda_{ij} x_j + \gamma_j$ is a birational map between V and its image $\pi(V)$. Further, there is a polynomial g of degree D such that $\pi(V) \setminus Z(g)$ and $V \setminus \pi^{-1}(Z(g))$ are isomorphic.*

Lemma 2.25 ([CM06, Corollary 3.4])**.** *Suppose $f \in \mathbb{Z}[x_1, \ldots, x_n]$ is an absolutely irreducible polynomial of degree D. For a point $(v, w, z) \in \overline{\mathbb{Q}}^n \times \overline{\mathbb{Q}}^{n-1} \times \overline{\mathbb{Q}}^{n-1}$, let $f_{v,w,z}(x, y) := f(x + v_1, w_2 x + z_2 y + v_2, \ldots, w_n x + z_n y + v_n)$. There exists a polynomial H of degree at most $10D^5$ such that for any v, w, z satisfying $H(v, w, z) \neq 0$, the polynomial $f_{v,w,z}$ remains absolutely irreducible.*

Corollary 2.26 (Effective Bertini Theorem)**.** *Suppose V is an equidimensional affine variety of dimension r and degree D. There is a randomised algorithm that returns linear equations $\ell_1, \ldots, \ell_{r-1}$ with $\mathrm{ht}(\ell_i) \leqslant (nD)^c$ such that the following holds with high probability: $V \cap Z(\ell_1, \ldots, \ell_{r-1})$ is an equidimensional affine curve, with the same number of irreducible components as V.*

Proof. We focus on each irreducible component of V and apply a union bound at the end, therefore assume without loss of generality that V is irreducible. Fix set $B := \{1, \ldots, 10D^7\}$, a subset of $\mathbb{N}$. Pick linear polynomials $h_1, \ldots, h_{r+1}$ with coefficients from B, and let π be the projection map with coordinate functions h_i. By Lemma 2.24, $\pi(V)$ is birational with its image with probability at least $1 - 2(r+1)D^2/|B|$. The image is an irreducible hypersurface of degree D, say $Z(f)$. Now pick a point (v, w, z) from $B^n \times B^{n-1} \times B^{n-1}$. By Lemma 2.25, $f_{v,w,z}$ is absolutely irreducible with probability at least $1 - 10D^5/|B|$. Pick a linear subspace $L \subset \mathbb{A}^{r+1}$ such that $f_{v,w,z}$ is isomorphic to $f|_L$, this can be performed by elementary linear algebra. By construction, $L \cap \pi(V)$ is an irreducible curve. Since L is a random linear subspace, $(\pi(V) \cap Z(g)) \cap Z(L)$ is a finite set, where g is the polynomial guaranteed by Lemma 2.25 [7]. Now we show that $\pi^{-1}(L)$ is the required subspace.

First observe that $\pi^{-1}(L)$ has defining equations of logarithmic height $\mathrm{poly}(D, n)$, this is because the equations are obtained by inverting matrices with entries of logarithmic height $\mathrm{poly}(D, n)$. Now we know that $\pi(V) \cap Z(L)$ is an irreducible curve. If ψ is the inverse of the birational map π, then $\pi^{-1}(L) \cap V$ is exactly $\overline{\psi(\pi(V) \cap Z(L))}$, which is irreducible and dimension 1. □

3 Height bounds

In this section we recall some height bounds for certain operations over polynomial rings over the integers, as well as of certain operations over polynomial rings over certain algebraic numbers. Then, we proceed to prove some height bounds for membership problems in given ideals, which is done in Section 3.3. We then conclude the section with some lemmas on absolutely irreducible factors of bivariate polynomials, where we discuss the work of Kaltofen and its implications to our setting.

[7] Note that Lemma 2.14 does not apply since the distribution of L is different, but the same proof works

3.1 Height bounds for elementary operations

The following are some elementary height bounds on the results of basic arithmetic operations performed on polynomials. The next two statements are from [KPS01, Lemma 1.2].

Lemma 3.1. *Let* $f_1, \ldots, f_m \in R$ *be polynomials with* $\deg f_i \leqslant d$ *and* $\mathrm{ht}(f_i) \leqslant h$. *Then*

1. $\mathrm{ht}(\sum f_i) \leqslant h + \log(m)$.
2. $\mathrm{ht}(\prod f_i) \leqslant hm + md\log(n+1)$.
3. *If* $g \in \mathbb{Z}[y_1, \ldots, y_m]$ *then* $\mathrm{ht}(g(f_1, \ldots, f_m)) \leqslant \mathrm{ht}(g) + \deg g\,(h + \log(m+1) + d\log(n+1))$.

The above lemma allows us to get height bounds for the determinant of a matrix with ring elements.

Corollary 3.2. *Let* M *be a* $m \times m$ *matrix with* $M_{ij} \in R$ *such that* $\deg(M_{ij}) \leqslant d$ *and* $\mathrm{ht}(M_{ij}) \leqslant h$ *for all* i, j. *Then* $\mathrm{ht}(\det M) \leqslant m(h + \log(m) + d\log(n+1))$.

Once we have the above height bounds, the next corollary yields height bounds on the resultant and on the cofactors of the resultant identity.

Corollary 3.3. *Let* $f, g \in \mathbb{Z}[y]$ *be coprime polynomials with* $\deg f, \deg g \leqslant d$ *and* $\mathrm{ht}(f), \mathrm{ht}(g) \leqslant h$. *Then there exists* $a, b \in \mathbb{Z}[y]$ *with* $\deg a < \deg g$ *and* $\deg b < \deg f$ *such that* $\mathrm{res}_y(f, g) = af + bg$, *where*

$$\mathrm{ht}(a), \mathrm{ht}(b), \mathrm{ht}(\mathrm{res}_y(f, g)) \leqslant 2d(h + \log(2d+1)).$$

In particular, if Δ *is the discriminant of* f *then* $\mathrm{ht}(\Delta) \leqslant 2d(h + \log(d) + \log(2d+1))$.

Proof. In $\mathbb{Q}[y]$ we have the Bézout identity $sf + tg = 1$, with $\deg s < \deg g$ and $\deg t < \deg f$. We can write this as a linear system of at most $2d$ equations in at most $2d$ variables. The solution to this system is given by Cramer's rule: each coefficient of s, t has numerator given by the determinant of a matrix of size $2d \times 2d$ with entries the coefficients of f, g. The denominator is common among all the coefficients, and is similarly given by such a determinant. In fact this denominator is exactly $\mathrm{res}_y(f, g)$. Clearing common denominators from the equation gives us the claimed identity with the claimed bounds. The last statement follows from the fact that the discriminant is exactly $\mathrm{res}_y(f, f')$, and $\mathrm{ht}(f') \leqslant h + \log d$. □

We also need to bound the logarithmic height of any rational root of an integral polynomial.

Lemma 3.4. *Let* $f \in \mathbb{Z}[y]$ *be a degree* d *polynomial and suppose* a/b *is a rational root of* f *in minimal form. Then* $\mathrm{ht}(a), \mathrm{ht}(b) \leqslant \mathrm{ht}(f)\,d + 1$.

Proof. Suppose $f = \sum f_i y^i$, where $f_i \in \mathbb{Z}$. Multiplying throughout by f_d^{d-1}, and replacing $f_d y$ with a variable x gives us a monic integer polynomial $g = x^d + \sum_{i=1}^{d-1} g_i x^i$ with $\deg g = d$ and $\mathrm{ht}(g) \leqslant \mathrm{ht}(f)\,d$. Since a/b is a root of f, we have tht $c := f_d a/b$ is a root of g. Further, $c \in \mathbb{Z}$ since g is a monic integer polynomial. By [Mig83, Theorem 2] we have $\mathrm{ht}(c) \leqslant 1 + \mathrm{ht}(g) \leqslant \mathrm{ht}(f)\,d + 1$. We can obtain a/b by putting c/f_d in minimal form, therefore $\mathrm{ht}(a), \mathrm{ht}(b) \leqslant \mathrm{ht}(c), \mathrm{ht}(f_d) \leqslant \mathrm{ht}(f)\,d + 1$. □

We can also bound the heights of the quotient and the remainder upon division by a monic polynomial, as the following proposition shows.

Proposition 3.5. *Let* $f, g \in \mathbb{Z}[y]$ *be polynomials where* $d := \deg f \geqslant \deg g =: e$, $\mathrm{ht}(f), \mathrm{ht}(g) \leqslant h$, *and* g *is a monic polynomial. If* $f = g \cdot q + r$, *where* $g, r \in \mathbb{Z}[y]$ *with* $\deg r < e$, *then we have* $\mathrm{ht}(g), \mathrm{ht}(r) \leqslant (d+1)\cdot(h + 2\log(d+1))$.

Proof. Note that g, r are the unique solution to the following system of linear equations: $N \cdot \begin{pmatrix} g \\ r \end{pmatrix} = (f)$, where N is a $(d+1) \times (d+1)$ matrix of the form $N = \begin{pmatrix} M_q & P \end{pmatrix}$ and $P = \begin{pmatrix} 0 \\ I_e \end{pmatrix}$. Since q is monic, we know that N is a unipotent lower triangular matrix. Thus, $\det(N) = 1$, and by Cramer's rule we have that each coefficient of g and r is given by the determinant of a matrix in $\mathbb{Z}^{(d+1)\times(d+1)}$ with height $\leqslant h$. By Corollary 3.2 we have the desired bound. □

As discussed earlier, we will often have to work with polynomials with coefficients lying in some finite extension generated by an algebraic integer α. If α has monic minimal polynomial $q(z) \in \mathbb{Z}[z]$ with $\deg q(z) \leqslant e$ then we will represent elements of $A = \mathbb{Z}[\alpha][x_1, \ldots, x_n] \simeq R[z]/(q)$ as polynomials in $R[z]$ of degree less than e in z. Thus, elements of A inherit the logarithmic height function from $R[z]$.

After some algebraic operations, for example taking products of such polynomials, the resulting polynomial will no longer have degree less than e in the variable z, and we must perform a step of going modulo $q(z)$. The following lemma bounds the change in logarithmic height by this operation. A sharper version of this lemma can be found in [Kal95, Lemma 1].

Lemma 3.6. *Let $q(z) \in \mathbb{Z}[z]$ be a monic polynomial with $\deg q = e$. Suppose $f \in R[z]$ is a polynomial with $\deg_z f = d \geqslant e$. Lastly, suppose $\mathrm{ht}(f), \mathrm{ht}(q) \leqslant h$. Then $\mathrm{ht}(f \bmod q) \leqslant (d+1)(h + 2\log(d+1))$.*

Proof. Let $\mathcal{M}_f$ be the set of exponent vectors of f, when viewed as a polynomial in $\mathbb{Z}[z][x_1, \ldots, x_n]$. Hence, we have $f = \sum_{\alpha \in \mathcal{M}_f} f_\alpha(z) \cdot x^\alpha$. Therefore, $f \bmod q = \sum_{\alpha \in \mathcal{M}_f} (f_\alpha(z) \bmod q) \cdot x^\alpha$ and thus

$$\mathrm{ht}(f \bmod q) = \max_{\alpha \in \mathcal{M}_f} \mathrm{ht}(f_\alpha(z) \bmod q) \leqslant (d+1) \cdot (h + 2\log(d+1)).$$

Where the last inequality follows from Proposition 3.5. □

Combining Corollary 3.2 with Lemma 3.6 we are able to obtain the following height bounds for determinants with entries in rings of the form $\mathbb{Z}[\alpha]$, where α is an algebraic integer.

Corollary 3.7. *Let $q(z) \in \mathbb{Z}[z]$ be monic with $\deg q = e$, $\mathrm{ht}(q) \leqslant h$ and let $\alpha \in \overline{\mathbb{Q}}$ be a root of q. If $M \in \mathbb{Z}[\alpha]^{m \times m}$ is such that $\mathrm{ht}(M_{ij}) \leqslant h$ for all $i, j \in [m]$, we have that $\mathrm{ht}(\det M) \leqslant e \cdot m \cdot (m \cdot (e + h + \log m) + 2\log(em))$.*

Proof. By regarding $M \in \mathbb{Z}[z]^{m \times m}$, with $\deg M_{ij} < e$ and $\mathrm{ht}(M_{ij}) \leqslant h$ for all $i, j \in [m]$ Corollary 3.2 implies that $\deg \det M \leqslant (e-1) \cdot m$ and $\mathrm{ht}(\det M) \leqslant m \cdot (h + \log(m) + (e-1))$. Now, applying Lemma 3.6 to $\det M$ with quotient $q(z)$, we have that

$$\begin{aligned} \mathrm{ht}((\det M) \bmod q) &\leqslant ((e-1) \cdot m + 1) \cdot (m \cdot (h + \log(m) + (e-1)) + 2\log((e-1)m + 1)) \\ &\leqslant e \cdot m \cdot (m \cdot (e + h + \log(m)) + 2\log(em)). \end{aligned}$$

□

The following corollary obtains height bounds for basic operations over the ring A.

Corollary 3.8. *Let $q(z) \in \mathbb{Z}[z]$ be monic with $\deg q = e$, $\mathrm{ht}(q) \leqslant h$ and let $\alpha \in \overline{\mathbb{Q}}$ be a root of q. Let $A := \mathbb{Z}[\alpha][x_1, \ldots, x_n]$. Let $f_1, \ldots, f_m \in A$ be polynomials with $\deg f_i \leqslant d$ and $\mathrm{ht}(f_i) \leqslant h$. Then*

1. $\mathrm{ht}(\sum_{i=1}^m f_i) \leqslant h + \log(m)$.
2. $\mathrm{ht}(\prod_{i=1}^m f_i) \leqslant em \cdot (hm + dm\log(n+2) + 2\log(em))$.
3. *If $g \in \mathbb{Z}[\alpha][y_1, \ldots, y_m]$ with $d_g := \deg(g)$ then*

$$\mathrm{ht}(g(f_1, \ldots, f_m)) \leqslant 2ed_g \cdot [\mathrm{ht}(g) + d_g \cdot (h + \log(m+1) + \max\{e, d\} \cdot \log(n+2)) + 2\log(2ed_g)]$$

Proof. For item 1, note that the coefficients of $\sum_{i=1}^m f_i$ are still given by polynomials in $\mathbb{Z}[z]$ with degree less than e, since addition does not increase the degree. Thus, the bounds follow from Lemma 3.1 item 1.

For item 2, let $p = \prod_{i=1}^m f_i$. In this case, we have that $p = \hat{p} \bmod q$, where $\hat{p} \in R[z]$ is the product of the polynomials f_i when seen as elements of $R[z]$. By Lemma 3.1 item 2, $\mathrm{ht}(\hat{p}) \leqslant h \cdot m + md\log(n+2)$. If $\deg_z(\hat{p}) < e$, we have $\mathrm{ht}(p) = \mathrm{ht}(\hat{p})$ and we are done. Else, by Lemma 3.6 and the bound $\deg_z(\hat{p}) \leqslant em - 1$, we have

$$\mathrm{ht}(p) = \mathrm{ht}(\hat{p} \bmod q) \leqslant h \cdot em^2 + dem^2\log(n+2) + 2em\log(em).$$

For item 3, consider the representation $\hat{g} \in \mathbb{Z}[y_1, \ldots, y_m][z]$ and the representation of $\widehat{f_i} \in R[z]$. Let $p := \hat{g}(\widehat{f_1}, \ldots, \widehat{f_m}) \in R[z]$. Thus, $\deg_z(p) \leqslant (e-1) \cdot (d_g + 1)$ and by Lemma 3.1 item 3,

$$\mathrm{ht}(p) \leqslant \mathrm{ht}(g) + d_g \cdot (h + \log(m+1) + \max\{e-1, d\} \cdot \log(n+2))$$

Since $g(f_1, \ldots, f_m) = p \bmod q$, by Lemma 3.6 we have

$$\mathrm{ht}(g(f_1, \ldots, f_m)) \leqslant 2ed_g \cdot [\mathrm{ht}(g) + d_g \cdot (h + \log(m+1) + \max\{e, d\} \cdot \log(n+2)) + 2\log(2ed_g)]. \qquad \square$$

3.2 Height bounds for primitive elements

In this section we gather height bounds for representations of primitive elements. The following theorem is stated in [Koi96, Theorem 4].

Theorem 3.9 (Complexity of primitive element)**.** *There is a universal constant $c \geqslant 1$ such that the following holds. Let $\alpha_1, \alpha_2, \ldots, \alpha_m$ be m algebraic numbers which are roots of polynomials $P_i \in \mathbb{Z}[z]$ with $\deg(P_i) \leqslant d$ and $\mathrm{ht}(P_i) \leqslant h$. There is a primitive element γ for $\alpha_1, \ldots, \alpha_m$ which is a root of an irreducible polynomial $q \in \mathbb{Z}[z]$ with $\deg(q) \leqslant d^{m^c}$ and $\mathrm{ht}(q) \leqslant h \cdot d^{m^c}$. Moreover, there are non-zero integers a_i and polynomials $Q_i \in \mathbb{Z}[z]$ with $\mathrm{ht}(a_i) \leqslant h \cdot d^{m^c}$, $\deg(Q_i) < \deg(q)$ and $\mathrm{ht}(Q_i) \leqslant h \cdot d^{m^c}$ such that $\alpha_i = Q_i(\gamma)/a_i$ for every $i \in [m]$.*

We now state [Koi96, Theorem 7].

Theorem 3.10. *There is an absolute constant $c > 0$ such that if the system $f_1, \ldots, f_m \in R$ has a solution over $\overline{\mathbb{Q}}$, then there is a solution $\alpha = (\alpha_1, \ldots, \alpha_n)$ such that each α_i is a root of a polynomial $P_i \in \mathbb{Z}[z]$ satisfying $\deg(P_i) \leqslant 2^{(n \log \sigma)^c}$ and $\mathrm{ht}(P_i) \leqslant h \cdot 2^{(n \log \sigma)^c}$, where $\sigma := 2 + \sum_{i=1}^{m} \deg(f_i)$.*

3.3 Height bounds for membership in ideals

We now state some height bounds on the polynomials occurring in the Nullstellensatz over $\mathbb{Z}$ and $\mathbb{Z}[\alpha]$, where α is an algebraic integer. Results of this nature are known as effective arithmetic Nullstellensatz, and the current best known bounds can be found in [KPS01]. While the optimal bounds for polynomials with integer coefficients are easy to state, the optimal bounds for the more general case of polynomials with algebraic integers are significantly more technical. Their statements involve extending the notion of logarithmic height to algebraic number fields, which is fairly intricate. We choose to forego these optimal bounds in the interest of simplicity, as the more elementary (and weaker) bounds suffice for our purposes.

Subsection setup. For the rest of this subsection, we let $q \in \mathbb{Z}[z]$ be a monic irreducible polynomial with $\deg q = e$ and $\mathrm{ht}(q) \leqslant h$, and let $\alpha \in \overline{\mathbb{Q}}$ be a root of q. Let $A := \mathbb{Z}[\alpha][x_1, \ldots, x_n]$. By our convention, we represent polynomials in A as elements of $R[z]$ with z-degree less than e. In the statements that follow, we will consider a degree parameter $d \geqslant 3$, and polynomials $f_1, \ldots, f_m \in A$ with $\deg f_i \leqslant d$ and $\mathrm{ht}(f_i) \leqslant h$. We will denote by $I := (f_1, \ldots, f_m) \subseteq S$ the ideal generated by $f_1, \ldots, f_m$ and by $V := Z(I) \subseteq \overline{\mathbb{Q}}^n$.

Lemma 3.11. *If $V(I) = \emptyset$, then there exists $a \in \mathbb{Z}[\alpha] \setminus \{0\}$ and $g_1, \ldots, g_m \in A$ satisfying*

$$a = f_1 g_1 + \cdots + f_m g_m, \tag{1}$$

with $\deg g_i \leqslant 2d^n$ and $\mathrm{ht}(g_i), \mathrm{ht}(a) \leqslant h \cdot e^5 \cdot d^{5n^2}$.

Proof. Since $V(f_1, \ldots, f_m) = \emptyset$, the effective Nullstellensatz Theorem 2.21 shows that there are polynomials $h_1, \ldots, h_m \in S$ with with $\deg h_i \leqslant 2d^n$ such that

$$1 = \sum f_i h_i. \tag{2}$$

Let $M := M_D(f_1, \ldots, f_m)$ be the matrix of the linear system Eq. (2) (as defined in Section 2) with $D := 2d^n$. Note that the entries of M are elements of $\mathbb{Z}[\alpha]$ (hence given by elements of $\mathbb{Z}[z]$ of degree less than e). Theorem 2.21 implies that the linear system $1 = Mv$ has a solution with coefficients in $\overline{\mathbb{Q}}$, and since the coefficients of M lie in $\mathbb{Q}[\alpha]$, this implies the existence of a solution with coefficients in $\mathbb{Q}[\alpha]$. By Lemma 2.10, there exists a solution where each unknown v_j is either 0, or of the form $\frac{\det N_j}{\det N}$, where $\det N_j, \det N$ are minors of $[M|1]$. Note that the denominator is common among all the unknowns. We will take our polynomials $g_1, \ldots, g_m$ to be the polynomials in A corresponding to the vector $\det N \cdot v$ (note that this is a vector with entries in A). By the above, this choice gives us a solution to Eq. (1) with $g_i \in A$ and $\deg g_i \leqslant 2d^n$. We now need to bound $\mathrm{ht}(g_i)$.

Since $\mathrm{ht}(g_i) \leqslant \max_j\{\mathrm{ht}(\det N_j)\}$, it is enough to bound $\mathrm{ht}(\det N)$ and $\mathrm{ht}(\det N_j)$. Each square submatrix N, N_j has entries in $\mathbb{Z}[\alpha]$ and dimension upper bounded by $D' := \binom{D+d+n}{n} \leqslant d^{n^2+n}$, which is the number of rows of M. As $\mathrm{ht}(M_{ij}) \leqslant h$ for all i, j, Corollary 3.7 yields

$$\begin{aligned}\mathrm{ht}(\det N), \mathrm{ht}(\det N_j) &\leqslant eD'(D'(e + h + \log D') + 2\log(eD')) \\ &\leqslant ed^{n^2+n}\left(d^{n^2+n}\left(e + h + (n^2+n)\cdot \log d\right) + 2\log\left(ed^{n^2+n}\right)\right) \leqslant h \cdot e^5 \cdot d^{5n^2} \qquad \square\end{aligned}$$

The following easy corollary extends this result to arbitrary membership in radical ideals.

Corollary 3.12. *Let* $f \in A$ *be such that* $\deg f \leqslant d$ *and* $\operatorname{ht}(f) \leqslant h$. *If* $f \in \operatorname{rad}(I)$ *then there exist* $t \in \mathbb{N}$, $a \in \mathbb{Z}[\alpha] \setminus \{0\}$, *and* $g_1, \ldots, g_m \in A$ *satisfying*

$$af^t = f_1 g_1 + \cdots + f_m g_m,$$

with $t \leqslant 2(d+1)^{n+1}$, $\deg g_i \leqslant 2(d+1)^{n+2}$, *and* $\operatorname{ht}(g_i), \operatorname{ht}(a) \leqslant 12 \cdot h \cdot e^5 \cdot (d+1)^{6(n+1)^2}$.

Proof. Let y be a new variable, and consider the polynomials $f_1, \ldots, f_m, 1-yf \in A[y]$. As $f \in \operatorname{rad}(f_1, \ldots, f_m)$, we have $Z(f_1, \ldots, f_m, 1-yf) = \emptyset$, and thus Lemma 3.11 implies that there exist $a \in \mathbb{Z}[\alpha] \setminus \{0\}$ and $h, h_1, \ldots, h_m \in A$ such that $a = \sum f_i h_i + h(1-yf)$. Moreover, Lemma 3.11 implies that $\deg(h_i), \deg(h) \leqslant D := 2(d+1)^{n+1}$. Substituting $y = 1/f$ in the above equation and multiplying by f^t to clear denominators, where $t \leqslant D$ (due to the degree bounds), we obtain equation $af^t = \sum f_i g_i$, where we have $g_i = f^t \cdot h_i(x_1, \ldots, x_n, 1/f)$.

We now show the bounds on the heights and degrees. Let $h_i'(x_1, \ldots, x_n, y) := y^{\deg_y(h_i)} \cdot h_i(x_1, \ldots, x_n, 1/y)$. Since h_i and h_i' have the same coefficient set, we have $\operatorname{ht}(h_i') = \operatorname{ht}(h_i)$. Moreover, $\deg h_i' = \deg(h_i)$.

By Lemma 3.11 we also have that $\operatorname{ht}(h_i), \operatorname{ht}(a) \leqslant h \cdot e^5 \cdot (d+1)^{5(n+1)^2}$, and also $\deg h_i \leqslant D$. Since $g_i = f^{t-\deg_y(h_i)} \cdot h_i'(x_1, \ldots, x_n, f)$, we have that $\deg(g_i) \leqslant t \cdot \deg(f) + \deg(h_i) \leqslant D \cdot (d+1)$. Moreover, Corollary 3.8 items 2 and 3 imply the desired height bounds. □

We will use the above lemmas to show bounds on the defining equations for the image of $Z(f_1, \ldots, f_m)$ under certain projections. This is an arithmetic version of Lemma 2.11, albeit with stronger assumptions.

Lemma 3.13. *Suppose* I *is radical and* $\dim I = 1$. *Define* $\pi : \mathbb{A}^n \to \mathbb{A}^2$ *to be projection to the first two coordinates. If* $\dim \pi(V) = 1$ *and* $\pi(V)$ *is equidimensional, then* $\pi(V) = Z(g)$ *for some* $g \in \mathbb{Z}[\alpha][x_1, x_2]$. *Further, there exists* $a \in \mathbb{Z}[\alpha] \setminus \{0\}$, *and* $h_1, \ldots, h_m \in A$ *satisfying*

$$ag = \sum f_i h_i,$$

with $\deg h_i \leqslant d^{4n^2}$ *and* $\operatorname{ht}(g), \operatorname{ht}(h_i) \leqslant h \cdot e^2 \cdot (d)^{12n^3}$.

Proof. The ideal of $\pi(V)$ is exactly the elimination ideal $I \cap \overline{\mathbb{Q}}[x_1, x_2]$. By the assumption that $\pi(V)$ is equidimensional with $\dim \pi(V) = 1$, this ideal is principal. Let G be the reduced Gröbner basis of I with respect to the lexicographic order induced by the variable order $x_n \succ x_{n-1} \succ \cdots \succ x_1$. Since I is generated in A, we have $G \subset \mathbb{Q}[\alpha][x_1, \ldots, x_n]$. Thus, $I \cap \overline{\mathbb{Q}}[x_1, x_2]$ is generated by the elements of $G \cap \overline{\mathbb{Q}}[x_1, x_2]$. Since $I \cap \overline{\mathbb{Q}}[x_1, x_2]$ is principal, we must have $G \cap \overline{\mathbb{Q}}[x_1, x_2] = \{\hat{g}\}$. Since $\hat{g} \in G$ we have $\hat{g} \in \mathbb{Q}(\alpha)[x_1, x_2]$ and $\operatorname{LT}(\hat{g}) = 1$. Let $b \in \mathbb{Z}$ be the common denominator of $\hat{g}$ and $g := b \cdot \hat{g} \in \mathbb{Z}[\alpha][x_1, x_2]$. By the above, $\pi(V) = Z(g)$.

By Bézout's theorem (Theorem 2.12) we have $\deg g \leqslant d^n$. By the above paragraph, g is the unique (up to scalar multiplication) lowest degree polynomial in $I \cap \overline{\mathbb{Q}}[x_1, x_2]$. Let $B := d^{4n^2}$. Since $\dim(I) = 1$ and B is larger than $\deg g$ and also larger than the representation bound from Corollary 2.23, there are forms $h_i \in \mathbb{Q}[\alpha][x_1, \ldots, x_n]$ with $\deg(f_i h_i) \leqslant B$ such that $g = \sum f_i h_i$.

Let $M := M_B(f_1, \ldots, f_m)$ be the matrix corresponding to the linear system as described in Section 2. Let M' be the matrix obtained by dropping the rows of M that correspond to monomials in x_1, x_2 that are smaller than $\operatorname{LM}(g)$ in the monomial order. Any solution $\widehat{h_1}, \ldots, \widehat{h_m}$ of the linear system $M'v = \operatorname{LM}(g)$ is such that $\sum f_i \widehat{h_i} \in I \cap \overline{\mathbb{Q}}[x_1, x_2]$, and $\operatorname{LM}\left(\sum f_i \widehat{h_i}\right) = \operatorname{LM}(g)$. Therefore, for any such solution $\widehat{h_1}, \ldots, \widehat{h_m}$ we have $\sum f_i \widehat{h_i} = \frac{1}{\operatorname{LT}(g)} \cdot g = b^{-1} \cdot g$.

By Cramer's rule (Lemma 2.10), there is a solution $\widetilde{h_i}$ whose coefficients are of the form $\det M_j / \det N$, where M_j, N are submatrices of the augmented matrix $[M' | \operatorname{LM}(g)]$. Let $B' := \binom{B+n}{n}$, that is, B' is the number of monomials in the variables $x_1, \ldots, x_n$ of degree at most B. By our upper bound on B, we have $B' \leqslant (4d)^{4n^3}$. Note that B' is an upper bound on the number of rows of the matrix $[M' | \operatorname{LM}(g)]$. Thus, by Corollary 3.7, we have $\operatorname{ht}(\det M_j), \operatorname{ht}(\det N) \leqslant e^2 \cdot h \cdot (4d)^{9n^3}$.

Multiplying the equation $\sum f_i \widetilde{h_i} = \frac{1}{b} \cdot g$ throughout by $\det N$, we get that $\frac{\det N}{b} \cdot g = (\det N) \cdot \sum f_i \widetilde{h_i} \in \mathbb{Z}[\alpha][x_1, x_2] \setminus \{0\}$. Let $h_i := \det(N) \cdot \widetilde{h_i}$. As the coefficients of h_i are given by the minors $\det(M_j)$, we

have $\operatorname{ht}(h_i) \leqslant \max_j \operatorname{ht}(\det(M_j)) \leqslant e^2 \cdot h \cdot (4d)^{9n^3}$. Since $\frac{\det N}{b} \cdot g \in \mathbb{Z}[\alpha][x_1, x_2]$, and by the definition of g, we must have $a := \frac{\det N}{b} \in \mathbb{Z}[\alpha] \setminus \{0\}$. In particular, this implies $b \mid \det N$ and therefore we have $\operatorname{ht}(a), \operatorname{ht}(b) \leqslant \operatorname{ht}(\det N) \leqslant e^2 \cdot h \cdot (4d)^{9n^3}$. As $\operatorname{ht}\left(b^{-1} \cdot g\right) \leqslant e^2 \cdot h \cdot (4d)^{9n^3}$ and $b \in \mathbb{Z}$, we have that $\operatorname{ht}(g) \leqslant 2 \cdot e^2 \cdot h \cdot (4d)^{9n^3}$. □

We will need the following extension of the above lemma, where we drop the assumption that I is radical.

Lemma 3.14. *Suppose* $\dim I = 1$. *Define* $\pi : \mathbb{A}^n \to \mathbb{A}^2$ *to be projection to the first two coordinates. If* $\dim \pi(V) = 1$ *and* $\pi(V)$ *is equidimensional, then* $\pi(V) = Z(g)$ *for some* $g \in \mathbb{Z}[\alpha][x_1, x_2]$. *Moreover, there exist an absolute constant* $c > 0$, $t \in \mathbb{N}$, $a \in \mathbb{Z}[\alpha]$, *and* $h_1, \ldots, h_m \in A$ *such that* $ag^t = \sum f_i h_i$, *with* $t \leqslant 4d^{2n}$, $\deg(f_i h_i) \leqslant 5d^{3n}$, *and* $\operatorname{ht}(a), \operatorname{ht}(h_i) \leqslant h \cdot e^c \cdot d^{cn^c}$.

Proof. Since $\pi(V)$ is equidimensional and $\dim \pi(V) = 1$, we have that $I(\pi(V)) = (\hat{g})$ for some $\hat{g} \in \overline{\mathbb{Q}}[x_1, x_2]$. Let G be the reduced Gröbner basis of I with respect to the lexicographic order induced by the variable order $x_n \succ x_{n-1} \succ \cdots \succ x_1$. Since I is generated in A, we have $G \subset \mathbb{Q}[\alpha][x_1, \ldots, x_n]$. Thus, $J := I \cap \overline{\mathbb{Q}}[x_1, x_2]$ is generated by the elements of $G \cap \overline{\mathbb{Q}}[x_1, x_2]$. Since $\operatorname{rad}(J) = \operatorname{rad}(I) \cap \overline{\mathbb{Q}}[x_1, x_2]$, by the closure theorem [CLO97, Chapter 3.2, Theorem 3] we have that $\operatorname{rad}(J) = (\hat{g})$.

Now, let $J' := J \cap \mathbb{Q}[\alpha][x_1, x_2]$ and let $\operatorname{rad}(J')$ be the radical ideal of J' in $\mathbb{Q}[\alpha][x_1, x_2]$. Since $\dim J' = 1$, we have that $\operatorname{rad}(J') = (g')$, for some $g' \in \mathbb{Q}[\alpha][x_1, x_2]$. Since $\mathbb{Q}$ is a perfect field, by [Sta24, Tag 030U], we know that $\operatorname{rad}(J) = \operatorname{rad}(J') \otimes \overline{\mathbb{Q}}$. Thus, we have that $\operatorname{rad}(J) = (g')\overline{\mathbb{Q}}[x_1, x_2]$. Letting $g \in \mathbb{Z}[\alpha][x_1, x_2]$ be the polynomial obtained by clearing the common denominator of g', we have $\operatorname{rad}(J) = (g)$ as we wanted.

By Bézout's theorem (Theorem 2.12) we have $\deg g \leqslant d^n$. Further, by the effective nullstellensatz (Theorem 2.21) and the Rabinowitsch trick, there exists $t \in \mathbb{N}$ and $h_i' \in \mathbb{Q}[\alpha][x_1, \ldots, x_m]$ such that $g^t = \sum f_i h_i'$. Moreover, we have $t \leqslant 4d^{2n}$ and $\deg(f_i h_i') \leqslant 5d^{3n}$. We assume that t is the smallest power for which we can write g^t with these degree bounds.

Consider the linear system $M := M_D(f_1, \ldots, f_m)$ with $D = 5d^{3n}$. Let M' be the system obtained by dropping those rows of M corresponding to monomials that are smaller than $\operatorname{LM}(g)^t = \operatorname{LM}(g^t)$. Any solution to $M'v = \operatorname{LM}(g)^t$ is a monic polynomial in $\mathbb{Q}[\alpha][x_1, x_2]$ with leading coefficient 1, and leading monomial $\operatorname{LM}(g)^t$. Let $h_1', \ldots, h_m'$ be a solution to this system obtained by applying Cramer's rule (Lemma 2.10). Thus, there is a submatrix N of $\left[M' \mid \operatorname{LM}(g)^t\right]$ such that $\det N$ is the common denominator of every coefficient of h_i' for $i \in [m]$. Moreover, there are submatrices M_j of $\left[M' \mid \operatorname{LM}(g)^t\right]$ such that each coefficient $\det N \cdot h_i'$ is given by some $\det M_j$. In particular, Corollary 3.7 implies that $\operatorname{ht}(\det N \cdot h_i') \leqslant e \cdot D \cdot (2ehD + D \log D + 2\log(eD)) \leqslant 4e^2 hD^3 \leqslant h \cdot e^2 \cdot d^{10n}$.

Let $g_1 := \sum f_i \cdot (\det N \cdot h_i')$. Hence, we have $g_1 \in A$, and Corollary 3.8 implies $\operatorname{ht}(g_1) \leqslant h \cdot e^4 \cdot d^{14n}$. Let $\gamma \in \mathbb{Z}[\alpha]$ be the leading coefficient of g_1. Treating γ as an element of $\mathbb{Z}[z]$ of degree at most e, we see that γ and q are relatively prime. By Corollary 3.3, we can find an element $\delta \in \mathbb{Z}[z]$ with $\deg(\delta) < e$, and $b := \operatorname{res}_z(\gamma, q) \in \mathbb{N}^*$ such that $\gamma\delta \equiv b \bmod q$. Moreover, Corollary 3.3 implies $\operatorname{ht}(\delta), \operatorname{ht}(b) \leqslant h \cdot e^4 \cdot d^{18n}$. Replacing g_1 by δg_1, we can assume that the leading coefficient of g_1 is b, and $\operatorname{ht}(g_1) \leqslant h \cdot e^6 \cdot d^{22n}$.

Since $g_1 \in I \subset \operatorname{rad}(I) = (g)$, we have $g \mid g_1$ in $\mathbb{Q}[\alpha][x_1, x_2]$. Let us write $g_1 = gg_2$. We use this to deduce a bound on $\operatorname{ht}(g)$. Let O_L be the ring of integers of $\mathbb{Q}[\alpha]$. Thus, we have $\mathbb{Z}[\alpha] \subset O_L \subset \mathbb{Z}[\alpha]_\Delta$, where $\Delta := \operatorname{disc}_z(q)$. In the ring $(O_L)_b[x_1, x_2]$, the polynomial g_1 is monic. By Gauss's lemma Lemma 2.17 we can assume that $g, g_2 \in (O_L)_b[x_1, x_2]$. Further, after rescaling, we can assume that g', g_2' are the monic forms of g, g_2 in $(O_L)_b[x_1, x_2]$. By [Len84, pp. 64-67], there is a universal constant $c_1 > 0$ such that every coefficient of g', when written as a monic element of $\mathbb{Z}[\alpha]_{Db}$, uses rational numbers of absolute value $\leqslant 2^{h \cdot e^{c_1} \cdot d^{c_1 n^{c_1}}}$. Multiplying by Δb shows that there is a universal constant $c_2 > 0$ such that $\operatorname{ht}(g) \leqslant h \cdot e^{c_2} \cdot d^{c_2 n^{c_2}}$.

Now that we have control on $\operatorname{ht}(g)$, by Cramer's rule applied to $Mv = g^t$, there are $a \in \mathbb{Z}[\alpha]$ and $h_1, \ldots, h_m \in A$ such that $ag^t = \sum f_i h_i$, with $\deg f_i h_i \leqslant D$, and $\operatorname{ht}(a), \operatorname{ht}(h_i) \leqslant h \cdot e^c \cdot d^{cn^c}$ for some universal constant $c > 0$. This height bound comes from noticing that, by Cramer's rule, a and every coefficient of h_i is a minor of the matrix $[M \mid g^t]$, and we have upper bounds on $\operatorname{ht}(M)$ and $\operatorname{ht}(g^t)$. □

3.4 Absolutely irreducible factors of bivariate polynomials

In [Kal95], Kaltofen gives an algorithm to factor bivariate polynomials with coefficients in a field $\mathbb{F}$ over the algebraic closure $\overline{\mathbb{F}}$. The algorithm uses the usual template of factoring algorithms: given $f(x, y)$, the algorithm begins by picking a root $\alpha \in \overline{\mathbb{F}}$ of $f(x, 0)$ in $\overline{\mathbb{F}}$, which is the same as the linear factor $x - \alpha$. This factorisation is then lifted, using Newton iteration, modulo y^t for increasing powers of t. After some steps of lifting, a reconstruction step is performed and a factor f_1 of f, over $\overline{\mathbb{F}}$, is obtained.

A key observation made in [Kal95] is that the lifting and reconstruction steps only use arithmetic in the field $\mathbb{F}$, even though the factor $(x - \alpha)$ that the Newton iteration begins with has coefficients in $\overline{\mathbb{F}}$. In particular this implies that the factor f_1 has coefficients in $\mathbb{F}(\alpha)$. Since the minimal polynomial of α is a factor of $f(x, 0)$, the result can be interpreted as the fact that factors of f can be found in low complexity extensions of $\mathbb{F}$. Further, every irreducible factor of f can be obtained this way, therefore every irreducible factor of f lies in some low complexity extension of $\mathbb{F}$. While each factor lies in a low complexity extension of $\mathbb{F}$, since the number of factors of f can be large, the compositum of all these extensions might be very large. This is a common phenomenon in computational algebraic geometry, see [BM93, Section 4]. Following the principle, we will only factor as much as is required, in order to ensure that we can work with low complexity extensions.

The next lemmas are from [Kal95], and they formalise the discussion in the previous two paragraphs.

Lemma 3.15. *Given polynomial $f \in \mathbb{F}[x, y]$ that monic in x, such that $\deg f = d$, and such that the resultant $\operatorname{res}_x(f(x, 0), \partial f(x, 0)/\partial x) \neq 0$, there exists an algorithm that returns a list of absolutely irreducible factors $f_1, \ldots, f_r$ of f, with each $f_i \in \mathbb{F}(\alpha_i)[x, y]$, where α_i is a root of $f(x, 0)$.*

While the result gives an actual algorithm that computes the factors, we are just interested in the fact that every factor of f is of this form. The list of factors returned by the algorithm is not necessarily distinct, and computing the set of distinct factors is a non-trivial task depending on the model used to represent the extensions $\mathbb{F}(\alpha_i)$. We now state bounds on the logarithmic heights of the factors. The following is [Kal95, Corollary 1, Theorem 3], although we state a simplified statement specialised to our setting.

Lemma 3.16. *Suppose $f \in \mathbb{Z}[x, y]$ is a squarefree polynomial with $\deg(f) = d$ and $\operatorname{ht}(f) = h$ that is monic in x, and such that $f(x, 0)$ is squarefree. Suppose f_1 is an absolutely irreducible factor of f.*

There exists $g(x)$ an irreducible (in $\mathbb{Z}[x]$) factor of $f(x, 0)$, and polynomials $\widetilde{f}_1, f_2, h \in \mathbb{Z}[x, y, z]$ and integers Δ_1, Δ_2 with $\operatorname{ht}(\Delta_1), \operatorname{ht}(\Delta_2), \operatorname{ht}\left(\widetilde{f}_1\right) = O(hd^c)$ such that $\Delta_1\Delta_2 f = \widetilde{f}_1 f_2 + g(z)h$. Further, there exists a root α of g such that $f_1(x, y) = \widetilde{f}_1(x, y, \alpha)$.

Proof. The algorithm of Lemma 3.15 shows the existence of g, α, and also monic polynomials $\widehat{f}_1, \widehat{f}_2 \in \mathbb{Q}[x, y, z]$ such that $f = \widehat{f}_1\widehat{f}_2 + g(z)h$ for some $h \in \mathbb{Q}[x, y, z]$. By [Kal95, Corollary 1], the coefficients of $\widehat{f}_1$ are the solutions of a certain linear system, and by [Kal95, Theorem 3] the minors of this system have logarithmic height $O(h \cdot d^{c_1})$. Using Corollary 3.2 to invert the denominators, we can write the coefficients of $\widehat{f}_1$ as elements of $(1/\Delta_1)\mathbb{Z}[z]$, where $\operatorname{ht}(\Delta_1) = O(h \cdot d^{c_2})$. Further, the coefficients themselves also have logarithmic height $O(h \cdot d^{c_2})$.

No similar bounds are provided on the logarithmic height of $\widehat{f}_2$, however since $f, \widehat{f}_1, \widehat{f}_2$ are monic we can use the fact that the coefficients of $\widehat{f}_2$ are algebraic integers, in particular they are contained in $(1/\Delta_2)\mathbb{Z}[z]$ where $\Delta_2 = \operatorname{disc}_y(z)$. Writing $\widehat{f}_2$ in this form and clearing denominators from the equation $f = \widehat{f}_1\widehat{f}_2 + g(z)h$ gives us the required equation. □

Recall the notation of the number theory preliminaries. As a corollary of the factoring algorithm, [Kal95, Theorem 8] obtains a bound on primes p for which an absolutely irreducible polynomial remains absolutely irreducible after applying the map $\mathbb{Z} \to O_{\mathbb{F}}/\mathfrak{p}$ for some prime $\mathfrak{p}$ above p.

Theorem 3.17. *Let $A = \mathbb{Z}[\alpha][x_1, \ldots, x_m]$ for an algebraic integer α with minimal polynomial q satisfying $\deg q = e$ and $\operatorname{ht}(q) \leqslant h$. Suppose $f \in A$ is absolutely irreducible with $\operatorname{ht}(f) \leqslant h$ and $\deg f \leqslant d$. There exists an integer Δ with $\operatorname{ht}(\Delta) \leqslant c\left(ed^6 \log(h) + ed^6 n \log(d)\right)$ for a universal constant c, such that for any prime $p \nmid \Delta$ and any prime $\mathfrak{p}$ above p, the image of f under the map $A \to O_{\mathbb{F}}/\mathfrak{p}$ remains absolutely irreducible.*

4 Geometric Irreducibility and base change

In this section, we prove our technical base change results. We first show that certain mod p reductions preserve dimensions of zerosets Lemma 4.1. We then show that irreducibility of dimension 0 and dimension 1 varieties is also preserved under such base changes. Finally, we show that reducible zerosets of dimension 0 and 1 remain reducible under base changes, and further, for sufficiently many primes, there are irreducible components of the zeroset after base change that are definable over $\mathbb{F}_p$ itself. This last fact will be crucial when we apply the Lang-Weil bounds Theorem 2.18 in a subsequent section. We split the study of irreducibility and reducibility into two parts, based on the dimension. The dimension 0 case (Corollary 4.3 and Lemma 4.4) is relatively straightforward, and uses ideas from [Koi96] The dimension 1 case (Theorem 4.5 and Theorem 4.6) is significantly more involved.

We show Corollary 4.3, Lemma 4.4, and Theorem 4.6 for zerosets that are defined over $\mathbb{Z}$. Our proof of Theorem 4.6 invokes Lemma 4.1 and Theorem 4.5 for zerosets that are defined not just over $\mathbb{Z}$, but also over small algebraic extensions of $\mathbb{Q}$. Therefore, we prove Lemma 4.1 and Theorem 4.5 in this more general setting.

With all of the above in mind, we now set up the notation for this section. We recall the preliminaries discussed in Section 2, especially the number theory preliminaries in Section 2.6. Let $q \in \mathbb{Z}[z]$ be a monic irreducible polynomial with $\deg q = e$ and $\mathrm{ht}(q) \leqslant h$, and let α be a root of q. Let $\mathbb{F} := \mathbb{Q}[\alpha]$, and let $O_{\mathbb{F}}$ be the ring of integers of $\mathbb{F}$, we have $\mathbb{Z}[\alpha] \subset O_{\mathbb{F}}$. Recall that for each prime $\mathfrak{p}$ above p with $p \not| \mathrm{disc}_z(q)$, we have a quotient map $\mathbb{Z}[\alpha] \to O_{\mathbb{F}}/\mathfrak{p} \cong \mathbb{F}_p[z]/q_1$, where q_1 is an irreducible factor of $q \pmod p$.

We define $A := \mathbb{Z}[\alpha][x_1, \ldots, x_n]$. Suppose $f_1, \ldots, f_m \in A$ are polynomials with $\deg f_i \leqslant d$ with $d \geqslant 3$ and $\mathrm{ht}(f_i) \leqslant h$. Let $I := (f_1, \ldots, f_m)$ and $V := Z(I)$. For a prime $p \in \mathbb{N}$ such that $p \not| \mathrm{disc}_z(q)$, and for any prime $\mathfrak{p} \subset O_{\mathbb{F}}$ that lies above p corresponding to a factor q_1 of $q \pmod p$, let $I_{\mathfrak{p}}$ be the ideal of $(\mathbb{F}_p[z]/q_1)[x_1, \ldots, x_n]$ generated by the images of $f_1, \ldots, f_m$ under the map $\mathbb{Z}[\alpha] \to \mathbb{F}_p[z]/q_1$ described above. Let $V_{\mathfrak{p}}$ denote the algebraic subset of $I_{\mathfrak{p}}$ in $\overline{\mathbb{F}}_p[x_1, \ldots, x_n]$. Define $\sigma := dm + 2$.

Lemma 4.1. *There exists a universal constant* c *and an integer* Δ *with* $\mathrm{ht}(\Delta) \leqslant h \cdot e^c \cdot 2^{(n \log \sigma)^c}$ *such that for any prime* $p \not| \Delta$ *and for every* $\mathfrak{p}$ *above* p *we have* $\dim(V) = \dim(V_{\mathfrak{p}})$.

Proof. Let r be the dimension of V. We first control the primes p for which $\dim V_{\mathfrak{p}} < r$, and then control the primes p for which $\dim V_{\mathfrak{p}} > r$.

By Lemma 2.11, there is a subset of r variables such that I does not contain any polynomial supported on these r variables. Without loss of generality, suppose these variables are $x_1, \ldots, x_r$. Set $D := 4d^{3n}$, and $D_r := rD$. Set $D_b := D_r + 5d^{3n}$. Let $g := x_1^D x_2^D \cdots x_r^D$, so $\deg g = D_r$. Let $M := M_{D_b}(f_1, \ldots, f_m)$ be the linear system corresponding to membership in I as defined in Section 2. Let M' be the system obtained by dropping all the rows of M corresponding to monomials in $x_1, \ldots, x_r$ that are smaller than g in the graded lexicographic order. If $h_1, \ldots, h_m$ is a solution to $M'v = g$, then $\sum f_i h_i \in \mathbb{Q}[x_1, \ldots, x_r]$, with $\mathrm{LM}(\sum f_i h_i) = g$.

The system $M'v = g$ is unsatisfiable since I does not contain any polynomial supported on these r variables. Therefore $\mathrm{rank}\, M' < \mathrm{rank}\,[M'|g]$. Let $\delta_0 \in \mathbb{Z}[\alpha]$ be any nonzero minor of $[M'|g]$ of size $\mathrm{rank}\,[M'|g]$, the size of the submatrix corresponding to δ_0 is at most $1 + \binom{D_b + D_r + n}{n}$. By Corollary 3.2 and Lemma 3.6 we can deduce that $\mathrm{ht}(\delta_0) \leqslant h \cdot e^{c_1} \cdot 2^{(n \log \sigma)^{c_1}}$. Observe that δ_0 and q are coprime as elements of $\mathbb{Z}[z]$, since q is irreducible. By Corollary 3.3 we can deduce that $\Delta_0 := \mathrm{disc}_z(q) \cdot \mathrm{res}_z(q, \delta_0)$ satisfies $\mathrm{ht}(\Delta_0) \leqslant h \cdot e^{c_2} \cdot 2^{(n \log \sigma)^{c_2}}$.

We show that for every $p \not| \Delta_0$, and for every $\mathfrak{p}$ above p we have $\dim V_{\mathfrak{p}} \geqslant r$. Let $J_{\mathfrak{p}} := I_{\mathfrak{p}} \overline{\mathbb{F}}_p[x_1, \ldots, x_n]$. Suppose towards contradiction that $V_{\mathfrak{p}}$ has dimension less than r. By Lemma 2.11, the ideal of $V_{\mathfrak{p}}$, that is $\mathrm{rad}(J_{\mathfrak{p}})$, contains a polynomial supported on $x_1, \ldots, x_r$. By Theorem 2.12, the degree of $V_{\mathfrak{p}}$ is bounded by d^n, therefore the degree the projection of V_p to the coordinates $x_1, \ldots, x_r$ is bounded by d^n. By [Sch07, Corollary 1.9], the ideal $\mathrm{rad}(J_p)$ contains a polynomial f in the variables $x_1, \ldots, x_r$ of degree at most d^n. By Theorem 2.21 and the Rabinowitsch trick, there exist $e \leqslant 4d^{2n}$ and g_i with $\deg g_i \leqslant 5d^{3n}$ such that $f^e = \sum f_i g_i$. The degree of f^e is at most $D = 4d^{3n}$. By potentially multiplying the equation throughout by a monomial, we can assume that there is a polynomial f' with leading monomial $x_1^D x_2^D \cdot x_r^D$ such that $f' = \sum f_i g_i'$, where each g_i' has degree at most $5d^{3n} + rD$, which is exactly D_b.

Observe that the matrix of the linear system corresponding to membership in $I_{\mathfrak{p}}$ is exactly the matrix M, where we apply the map $\mathbb{Z}[\alpha] \to \mathbb{F}_p[z]/q_1$ to each coefficient. Denote this by matrix by $M_{\mathfrak{p}}$. If we

similarly drop the rows of M_p corresponding to monomials in $x_1, \ldots, x_r$ that are larger than g in the graded lexicographic order, the matrix obtained is exactly the reduction of M' under the same map, denote this matrix by M'_p. We have $\operatorname{rank} M'_p \leqslant \operatorname{rank} M'$, since any zero minor of M' remains zero after reduction. Since $p \not| \Delta_0$, in particular $p \not| \operatorname{res}_z(q, \delta_0)$. Therefore, the polynomials δ_0 and q remain relatively prime in $\mathbb{F}_p[z]$. In particular, the minor δ_0 remains nonzero under the map $\mathbb{Z}[\alpha] \to \mathbb{F}_p[z]/q_1$. The augmented matrix $\left[M'_p|g\right]$ satisfies $\operatorname{rank}\left[M'_p|g\right] = \operatorname{rank}[M_p|g] = \operatorname{rank} M' + 1$. In particular this implies that $M'_p v = g$ is unsolvable, which is in contradiction to the existence of f'. This completes the first part of the proof, that is controlling the primes p for which $\dim V_p < r$.

We now move on to the second part of the proof, that is, controlling the primes p for which $\dim V_p > r$. Similar to the proof above, since V has dimension r, for every subset of $r+1$ variables, say for example $x_1, \ldots, x_{r+1}$, the ideal I contains a polynomial f' in $x_1, \ldots, x_{r+1}$ with leading monomial $g := x_1^D x_2^D \cdots x_{r+1}^D$. Further, we can write $f' = \sum f_i g_i$ with $\deg g_i \leqslant 5d^{3n} + (r+1)D$. Set $D_b := 5d^{3n} + (r+1)D$. If we consider the linear system M_{D_b}, and drop the rows corresponding to monomials in $x_1, \ldots, x_{r+1}$ that are smaller than $x_1^D x_2^D \cdots x_{r+1}^D$ and call this new matrix M', then $M'v = g$ is now satisfiable. In other words, $\operatorname{rank} M' = \operatorname{rank}[M'|g]$. By Lemma 2.10, there is a solution to this system, where the entries of v have common denominator a minor of $[M'|g]$. Call this minor $\delta_{[r+1]}$, by Corollary 3.2 and Lemma 3.6 we can deduce that $\operatorname{ht}\left(\delta_{[r+1]}\right) \leqslant h \cdot e^{c_3} \cdot 2^{(n \log \sigma)^{c_3}}$. Set $\Delta_{[r+1]} := \operatorname{res}_z\left(q, \delta_{[r+1]}\right)$, for any $p \not| \Delta_{[r+1]}$ and $p \not| \operatorname{disc}_q(z)$, the image of this denominator is nonzero in $\mathbb{F}_p[z]/q_1$, and the ideal I_p contains a polynomial in $x_1, \ldots, x_{r+1}$.

Similarly, we can construct δ_U, Δ_U for every $U \subset [n]$ with $|U| = r+1$. If we let $\Delta_1 := \operatorname{disc}_q(z) \prod \Delta_U$, then for every prime $p \not| \Delta_1$ we must have $\dim V_p \leqslant r$. We have $\operatorname{ht}(\Delta_1) \leqslant h \cdot e^{c_4} \cdot 2^{(n \log \sigma)^{c_4}}$. This controls all primes for which $\dim V_p > r$. Defining $\Delta := \Delta_1 \cdot \Delta_2$ completes the proof. □

4.1 Dimension zero

Note that when $\dim I = 0$, we have that V is a finite set of points. In this subsection, we assume that $f_1, \ldots, f_m \in R$.

We first show that if V consists of a single point, then the logarithmic height of this point is not too large.

Lemma 4.2. *Suppose $V = \{(\alpha_1, \ldots, \alpha_n)\}$. Then $\alpha_i \in \mathbb{Q}$ and there is a universal constant $c > 0$ such that $\operatorname{ht}(\alpha_i) \leqslant h \cdot 2^{(n \log \sigma)^c}$ for all $i \in [n]$.*

Proof. Let K be the smallest Galois extension of $\mathbb{Q}$ that contains $\alpha_1, \ldots, \alpha_n$. Suppose $K \neq \mathbb{Q}$, and without loss of generality assume $\alpha_1 \notin \mathbb{Q}$. Then, there exists some element τ of the Galois group of $K/\mathbb{Q}$ that is not identity on α_1. Since τ fixes f_i, for $i \in [m]$, as $f_i \in R$, we must have $(\tau(\alpha_1), \cdots, \tau(\alpha_n)) \in V$, contradicting the fact that $V = \{(\alpha_1, \ldots, \alpha_n)\}$.

We now obtain the bound on $\operatorname{ht}(\alpha_i)$. By Theorem 3.10, there is a universal constant $c_1 > 0$ and a polynomial P_i with $\deg(P_i) \leqslant 2^{(n \log \sigma)^{c_1}}$, and $\operatorname{ht}(P_i) \leqslant h \cdot 2^{(n \log \sigma)^{c_1}}$ such that α_i is a root of P_i. By Lemma 3.4, there is a universal constant $c > 0$ such that $\operatorname{ht}(\alpha_i) \leqslant h \cdot 2^{(n \log \sigma)^c}$. □

With the above lemma at hand, we now prove that if $|V| = 1$, then for all but finitely many primes p, we must have that V_p consists of exactly one point, and this point lies in $\mathbb{F}_p^n$.

Corollary 4.3. *Suppose $V = \{(\alpha_1, \ldots, \alpha_n)\}$. There is a universal constant c and $\Delta \in \mathbb{Z}$ with $\operatorname{ht}(\Delta) \leqslant h \cdot 2^{(n \log \sigma)^c}$, such that for any prime p with $\Delta \notin (p)$, $V_p \subset \mathbb{F}_p^n$ and $|V_p| = 1$.*

Proof. By Lemma 4.2, there is a universal constant c_1 such that $\operatorname{ht}(\alpha_i) \leqslant h \cdot 2^{(n \log \sigma)^{c_1}}$, for all $i \in [n]$. Let $\alpha_i = \beta_i/\gamma_i$ with $\beta_i, \gamma_i \in \mathbb{Z}$ in reduced form, and let $g_i := \gamma_i x_i - \beta_i$. By the Nullstellensatz $g_i \in I$ for all $i \in [n]$.

By Corollary 3.12, there are: a universal constant $c_2 > 0$, positive integers a_i, e_i and polynomials $h_{ij} \in R$ such that $a_i g_i^{e_i} = \sum_j f_j h_{ij}$, where $e_i \leqslant 4(n+1)(d+1)^{n+1}$, $\deg(h_{ij}) \leqslant 8(n+1)(d+1)^{n+2}$ and lastly we have $\operatorname{ht}(a_i), \operatorname{ht}(h_{ij}) \leqslant h \cdot 2^{(n \log \sigma)^{c_1}} (d+1)^n (nd \log m)^{c_2}$.

Define $\Delta := \prod_i a_i \gamma_i$. For any prime p such that $\Delta \notin (p)$, the equation $a_i g_i^{e_i} = \sum_j f_j h_{ij}$ reduced mod p implies that $g_i \in \operatorname{rad}(I_p)$ over $\mathbb{F}_p[x_1, \ldots, x_n]$. Therefore, $V_p = \{(\beta_1 \gamma_1^{-1}, \ldots, \beta_n \gamma_n^{-1})\}$. From the bounds

on $\mathrm{ht}(a_i), \mathrm{ht}(\gamma_i)$ we have $\mathrm{ht}(\Delta) \leqslant (n+1) \cdot h \cdot 2^{(n \log \sigma)^{c_1}} (d+1)^n (nd \log m)^{c_2}$, which can be bounded by $h \cdot 2^{(n \log \sigma)^c}$ for an appropriate universal constant $c > c_2$. □

We now show that if $|V| \geqslant 2$, then the density of primes p such that the set V_p has at least two points in $\mathbb{F}_p^n$ is large enough. Our argument is an extension of the main argument of [Koi96].

Lemma 4.4. *Suppose $|V| \geqslant 2$. There exist: a universal constant c, a polynomial $G \in \mathbb{Z}[z]$ with $\deg(G) \leqslant 2^{(n \log \sigma)^c}$ and $\mathrm{ht}(G) \leqslant h \cdot 2^{(n \log \sigma)^c}$, and $\Delta \in \mathbb{Z}$ with $\mathrm{ht}(\Delta) \leqslant h \cdot 2^{(n \log \sigma)^c}$, such that for any prime p satisfying $\Delta \notin (p)$ and $G \bmod p$ has a root in $\mathbb{F}_p$, we have $|V_p \cap \mathbb{F}_p^n| \geqslant 2$.*

Proof. Let $(\alpha_1, \ldots, \alpha_n)$ and $(\beta_1, \ldots, \beta_n)$ be two points in V. By [Koi96, Lemma 2], there are a universal constant $c_1 > 0$ and polynomials $P_{ij} \in \mathbb{Z}[x_1, \ldots, x_n]$, where $i \in [2], j \in [n]$, with $\deg(P_{ij}) \leqslant 2^{(n \log \sigma)^{c_1}}$ and $\mathrm{ht}(P_{ij}) \leqslant h \cdot 2^{(n \log \sigma)^{c_1}}$ such that α_j is a root of P_{1j} and β_j is a root of P_{2j}. [8]

Let $c_2 \geqslant 1$ be the universal constant from Theorem 3.9. Applying Theorem 3.9 to the $2n$ elements $\alpha_1, \ldots, \alpha_n, \beta_1, \ldots, \beta_n$ with the polynomials P_{ij} gives us a primitive element γ, with minimal polynomial G, where $\deg(G) \leqslant 2^{(n \log \sigma)^{c_2}}$ and $\mathrm{ht}(G) \leqslant h \cdot 2^{(n \log \sigma)^{c_2}}$. Further, we also obtain $Q_{ij} \in \mathbb{Z}[z]$ with $\deg(Q_{ij}) < \deg(G)$ and $\mathrm{ht}(Q_{ij}) \leqslant h \cdot 2^{(n \log \sigma)^{c_2}}$, and $a_{ij} \in \mathbb{Z}^*$ with $\log(|a_{ij}|) \leqslant h \cdot 2^{(n \log \sigma)^{c_2}}$ such that $\alpha_j = Q_{1j}(\gamma)/a_{1j}$ and $\beta_j = Q_{2j}(\gamma)/a_{2j}$.

Let p be a prime such that $G \bmod p$ has a root x_0 in $\mathbb{F}_p$ and that $\prod_{i,j} a_{ij} \notin (p)$. Then, the points $((Q_{11}(x_0)/a_{11}, \ldots, Q_{1n}(x_0)/a_{1n})$ and $((Q_{21}(x_0)/a_{21}, \ldots, Q_{2n}(x_0)/a_{2n})$ are in V_p. To complete the proof, it suffices to ensure that these two points differ on at least one coordinate, so we can claim that V_p has at least two distinct points with coefficients in $\mathbb{F}_p$. To this end, assume without loss of generality that $\alpha_1 \neq \beta_1$, and therefore $Q_{11}/a_{11} \neq Q_{21}/a_{21}$. Consider the resultant $\mathrm{res}_z(a_{21}Q_{11} - a_{11}Q_{21}, G)$. Since G is irreducible and $\deg(a_{21}Q_{11} - a_{11}Q_{21}) \leqslant \max(\deg(Q_{11}), \deg(Q_{21})) < \deg(G)$, we can deduce $\mathrm{res}_z(Q_{11} - Q_{21}, G) \in \mathbb{Z}^*$. Further, $\mathrm{ht}(G), \mathrm{ht}(a_{21}Q_{11} - a_{11}Q_{21}) \leqslant h \cdot 2^{(n \log \sigma)^{c_2}}$. Since $\mathrm{res}_z(Q_{11} - Q_{21}, G) \in \mathbb{Z}$ is the determinant of a matrix in the coefficients of $G, Q_{11} - Q_{21}$, it has logarithmic height at most $h \cdot 2^{(n \log \sigma)^{c_3}}$.

Define $\Delta := a \cdot \mathrm{res}_z(Q_{11} - Q_{21}, G)$, so Δ has logarithmic height at most $h \cdot 2^{(n \log \sigma)^c}$. Suppose p is a prime such that $p \nmid \Delta$ and such that G has a root in $\mathbb{F}_p$. Since $p \nmid \mathrm{res}_z(Q_{11} - Q_{21}, G)$, the polynomials $Q_{11} - Q_{21}$ and G remain relatively prime in $\mathbb{F}_p[z]$. Therefore, for any root x_0 of G, we must have $Q_{11}(x_0) \neq Q_{21}(x_0)$. Combined with the arguments above, this shows that for any such p, the algebraic set V_p has at least two distinct points with coefficients in $\mathbb{F}_p$. □

4.2 Dimension one

In the following theorem, we assume once again the more general setting where $f_1, \mathtt{dots}, f_m \in A$.

Theorem 4.5. *Suppose V is an irreducible curve. There exists an integer Δ with $\mathrm{ht}(\Delta) \leqslant h \cdot e^c \cdot 2^{(n \log \sigma)^c}$ such that for any prime $p \nmid \Delta$ and for every $\mathfrak{p}$ above p the zeroset $V_{\mathfrak{p}}$ is irreducible.*

Proof. Let $D := d^n$ so $\deg(V) \leqslant D$ by Theorem 2.12. Define $B := D^2$. Let Φ be the set of all linear maps $\mathbb{A}^n \to \mathbb{A}$, where each coefficient is picked from the set $[B]$. The total number of such linear maps is B^{n+1}. For maps $\phi_i, \phi_j \in \Phi$, define $\phi_{ij} : \mathbb{A}^n \to \mathbb{A}^2$ to be the map with coordinate functions ϕ_i, ϕ_j.

Let $\overline{V}$ be the projective closure of V, and let $V^p := \overline{V} \cap \{x_0 = 0\}$ be the intersection of V with the hyperplane at infinity. Since $\deg(V) \leqslant D$, and since V is defined by polynomials in R, the zeroset V^p is finite and consists of at most D points. By [SR94, Theorem 1.15], for any ϕ_i such that the projective closure of the hyperplane $\phi_i = 0$ does not contain a point in V^p, the map $\phi_i : V \to \mathbb{A}^1$ is a finite map. Define $\Phi' \subset \Phi$ to be the set of maps that are finite, the above argument shows that $|\Phi'| \geqslant (1 - D/B)|\Phi| = (1 - 1/D)|\Phi|$.

Consider the pairs i, j such that $\phi_i \in \Phi', \phi_j \in \Phi'$. Since ϕ_i is a finite map, in particular it is surjective. Since ϕ_i is a projection of ϕ_{ij}, we can deduce that $\overline{\phi_{ij}(V)}$ has dimension 1. Further, $\overline{\phi_{ij}(V)}$ is irreducible since V is irreducible. For all such ϕ_{ij}, define g_{ij} to be the generator of $I(\overline{\phi_{ij}(V)})$. Each g_{ij} is a form of degree at most D, and is absolutely irreducible. The ideal I is prime, and satisfies $\dim I = 1$. For every i, j we can pick an invertible linear map $\mathbb{A}^n \to \mathbb{A}^n$ such that ϕ_{ij} is projection to the first two coordinates after applying the linear map. Therefore, we can apply Lemma 3.14 to deduce that $g_{ij} \in A$, and we obtain

[8] Note than while Theorem 7 in [Koi96] is stated in terms of a single root, Lemma 2 of [Koi96] holds for every root.

$h_{ijk} \in A$, and $t_{ij} \in \mathbb{N}$, and $a_{ij} \in \mathbb{Z}[\alpha]$ such that $a_{ij} g_{ij}^{t_{ij}} = \sum_k f_k h_{ijk}$. Further, $\deg h_i, t \leqslant 8d^{n^2}$, and $\mathrm{ht}(a), \mathrm{ht}(h_i), \mathrm{ht}(g_{ij}) \leqslant h \cdot \log m \cdot e^{c_1} \cdot d^{c_1 n^{c_1}}$. By Theorem 3.17, for each such ij there exists an integer Δ_{ij} with $\mathrm{ht}(\Delta_{ij}) \leqslant h \cdot e^{c_2} \cdot d^{cn^{c_2}}$ such that for any $p \nmid \Delta_{ij}$ and $p \nmid \mathrm{disc}_q(z)$, and for any $\mathfrak{p}$ above p, the image of polynomial g_{ij} under the map $\mathbb{Z}[\alpha] \to \mathbb{F}_p[z]/q_1$ remains absolutely irreducible.

Define $\Delta' := \prod_{i,j} \mathrm{res}_z(a_{ij}, q) \cdot \Delta_{ij}$, where the product is over i, j such that $\phi_i, \phi_j \in \Psi'$. Finally, define $\Delta := \Delta' \cdot \Delta'' \cdot (B+1)!$, where Δ'' is the integer obtained by applying Lemma 4.1 to V. We have $\mathrm{ht}(\Delta') \leqslant h \cdot e^{c_2} \cdot d^{cn^{c_2}} \cdot D^{(2n+2)} \leqslant h \cdot e^{c_3} \cdot 2^{(n \log \sigma)^{c_3}}$. Further $\mathrm{ht}(\Delta'') \leqslant h \cdot e^{c_4} \cdot 2^{(n \log \sigma)^{c_4}}$, and $\mathrm{ht}((B+1)!) \leqslant d^{n^3}$. Therefore, $\mathrm{ht}(\Delta) \leqslant h \cdot e^c \cdot 2^{(n \log \sigma)^c}$.

Pick a prime $p \nmid \Delta$ and $\mathfrak{p}$ above p, we will show that $V_{\mathfrak{p}}$ is irreducible, which will complete the proof. Since $p \nmid \Delta''$, $V_{\mathfrak{p}}$ has dimension 1. Further, since $p \nmid (S+1)!$, we have $p > S+1$. This ensures that the set of maps in Φ are distinct, even when considered as maps from $\overline{\mathbb{F}}_p^n \to \overline{\mathbb{F}}_p^2$.

Fix any i, j such that $\phi_i \in \Phi', \phi_j \in \Phi'$. The equation $a_{ij} g_{ij}^{t_{ij}} = \sum_k f_k h_{ijk}$ holds in $\mathbb{F}_p[z]/q_1$, and $a_{ij} \neq 0$, therefore $g_{ij}^{t_{ij}} \in I_{\mathfrak{p}}$ for every such $\mathfrak{p}$. Further, since $p \nmid \Delta_{ij}$, the polynomial g_{ij} remains absolutely irreducible. These facts together imply that $\overline{\phi_{ij}(V_{\mathfrak{p}})}$ is either empty or irreducible, here we treat ϕ_{ij} as a map from $\overline{\mathbb{F}}_p^n \to \overline{\mathbb{F}}_p^2$. The total number of such maps is $(1 - 1/D)^2 |\Phi|^2$.

Assume towards a contradiction that $V_{\mathfrak{p}}$ is reducible. Let C_1, C_2 be two components of $V_{\mathfrak{p}}$, we have $\deg(C_1), \deg(C_2) \leqslant D$. Define $\Phi_{\mathfrak{p}} \subset \Phi$ such that $\phi_i \in \Phi_{\mathfrak{p}}$ if $\phi_i : C_1 \to \overline{\mathbb{F}}_p^1$ and $\phi_i : C_2 \to \overline{\mathbb{F}}_p^1$ are both finite maps. Similar to the argument for Φ', we can deduce that $|\Phi_{\mathfrak{p}}| \geqslant (1 - 2/D)|\Phi|$ and we can deduce that for any i, j with $\phi_i \in \Phi_{\mathfrak{p}}$ we have $\dim \overline{\phi_{ij}(C_1)} = 1$ and $\dim \overline{\phi_{ij}(C_2)} = 1$. We bound the number of maps ϕ_{ij} such that $\phi_{ij}(C_1) = \phi_{ij}(C_2)$. Fix i, and pick a point $\alpha \in C_1 \setminus C_2$. Since $\phi_i : C_2 \to \overline{\mathbb{F}}_p^1$ is finite, in particular it is surjective and therefore the curve C_2 is not contained in the hyperplane defined by $\phi_i = \phi_i(\alpha)$. Therefore, $C_2 \cap \{\phi_i = \phi_i(\alpha)\}$ is a finite set consisting of at most D points, call these points $\beta_1, \ldots, \beta_D$. The set of points $\alpha - \beta_1, \ldots, \alpha - \beta_D$ is a finite set of nonzero points, therefore at most $(D/S)|\Phi|$ many ϕ_j are such that ϕ_j is zero at one of these points. For every remaining ϕ_j, the intersection of C_2 with the hyperplanes $\phi_i = \phi_i(\alpha)$ and $\phi_j = \phi_j(\alpha)$ is empty, whence $\phi_{ij}(C_1) \neq \phi_{ij}(C_2)$. Since the maps ϕ_{ij} are finite on C_1, C_2, this implies $\overline{\phi_{ij}(C_1)} \neq \overline{\phi_{ij}(C_2)}$. Therefore, the total number of pairs i, j such that $\phi_i, \phi_j \in \Phi_{\mathfrak{p}}$ and $\overline{\phi_{ij}(C_1)} \neq \overline{\phi_{ij}(C_2)}$ is at least $(1 - 2/D)(1 - 3/D)|\Phi|^2$. Recall however that for at least $(1 - 1/D)^2 |\Phi|^2$ maps, $\overline{\phi_{ij}(V_{\mathfrak{p}})}$ was either irreducible or empty, in particular $\overline{\phi_{ij}(C_1)} = \overline{\phi_{ij}(C_2)}$. This is a contradiction to our assumption that $V_{\mathfrak{p}}$ is reducible, and completes the proof. □

The following result again assumes $f_1, \ldots, f_m \in R$. The proof will invoke Lemma 4.1 and Theorem 4.5 in the general setting on a new system of polynomials that extends $f_1, \ldots, f_m$. We note that this is a limitation of our proof, and we expect the result to also hold if $f_1, \ldots, f_m \in A$.

Theorem 4.6. *Suppose $f_1, \ldots, f_m \in R$. Suppose $\dim I = 1$, that V is equidimensional, and that V has at least two irreducible components. There exists an integer Δ with $\mathrm{ht}(\Delta) \leqslant h \cdot 2^{(n \log \sigma)^c}$ and a polynomial $G \in \mathbb{Z}[z]$ with $\deg G \leqslant D^4$ and $\mathrm{ht}(G) \leqslant h \cdot 2^{(n \log \sigma)^c}$ such that for any prime $p \nmid \Delta$ and such that $G \pmod p$ has a root in $\mathbb{F}_p$, the zeroset V_p has at least two irreducible components that are $\mathbb{F}_p$-definable.*

Proof. Let k be the number of irreducible components of V, and suppose $V = V_1 \cup \cdots V_k$ is the irreducible decomposition of B. By Bézout's theorem (Theorem 2.12) we have $k \leqslant D := d^n$. As a first step, we will find find defining equations for two components of V. The components of V might not be $\mathbb{Q}$-definable, therefore the defining equations we find will potentially lie in some algebraic extension of $\mathbb{Q}$.

Pick a linear map $\phi : \mathbb{A}^n \to \mathbb{A}^2$, where the coefficients of each coordinate map are picked from the set $[D^6]$, such that $\pi|_V : V \to \phi(V)$ is finite, and such that each of the k components of V have distinct images under ϕ. A random linear map has these properties, and therefore such a map exists. Since $\dim I = 1$ we have $\dim \phi(V) = 1$. Further, since I is equidimensional, $\overline{\pi(V)}$ has exactly k components that are each of codimension 1, therefore $\overline{\phi(V)} = Z(g_1)$ for some g_1. We can apply Lemma 3.14 to deduce that $g_1 \in R$, and the existence of polynomials $h_{1k} \in R$, and $a_1, t_1 \in \mathbb{Z}$ such that $a_1 g_1^{t_1} = \sum_k f_k h_{1k}$. Further $\mathrm{ht}(g_1), \mathrm{ht}(h_{1k}) \leqslant h \cdot d^{c_1 n^{c_1}}$.

Now g_1 has exactly k absolutely irreducible factors, $g_1 = g_{11} \cdot g_{12} \cdot \cdots \cdot g_{1k}$. Each of these factors are coprime, and occur with multiplicity 1. Further, each factor g_{1i} corresponds to the irreducible component

V_i. Consider $I_1 := I + g_{11}$ and $I_2 := I + g_{12}$. We have $V_1 \subset Z(I_1)$, and $V_2 \subset Z(I_2)$. None of the components $V_2, \ldots, V_k$ vanish on g_{11}. Therefore $(V_2 \cup \cdots \cup V_k) \cap Z(g_{11})$ has dimension 0. By Bézout's inequality, and the fact that $\deg g_{11} \leqslant \deg g_1 \leqslant d^n$, we deduce that $(V_2 \cup \cdots \cup V_k) \cap Z(g_{11})$ consists of at most d^{2n} points. Similarly, $(V_1 \cup V_3 \cup \cdots \cup V_k) \cap Z(g_{12})$ consists of at most d^{2n} points.

In order to define V_1 (respectively V_2), we need to add equations to I_1 (respectively I_2) that vanish on V_1 (respectively V_2) but not on the above points. Towards this, pick a second linear map $\psi : \mathbb{A}^n \to \mathbb{A}^2$ where the coefficients of each coordinate map are picked from the set $\left[D^6\right]$, such that $\psi|_V : V \to \psi(V)$ is finite, and such that each of the k components of V have distinct images under ψ. We further require that none of the d^{2n} points of $(V_2 \cup \cdots \cup V_k) \cap Z(g_{11})$ are mapped to $\psi(V_1)$, and that none of the d^{2n} points of $(V_1 \cup V_3 \cup \cdots \cup V_k) \cap Z(g_{12})$ are mapped to $\psi(V_2)$. Such a map exists, since a random map satisfies these properties. As before, we have g_2 such that $\psi(V) = Z(g_2)$, and by Lemma 3.14 we have $g_2 \in R$, and the existence of polynomials $h_{2k} \in R$, and $a_1, t_1 \in \mathbb{Z}$ such that $a_2 g_2^{t_2} = \sum_k f_k h_{2k}$. Further $\mathrm{ht}(g_2), \mathrm{ht}(h_{2k}) \leqslant h \cdot d^{c_1 n^{c_1}}$. We can also factor g_2 as $g_2 = g_{21} \cdot g_{22} \cdot \cdots \cdot g_{2k}$ with g_{2i} corresponding to V_i.

We now set $J_1 := I_1 + g_{21} = I + g_{11} + g_{21}$ and $J_2 := I_2 + g_{22} = I + g_{12} + g_{22}$. By construction, $Z(J_1) = V_1$ and $Z(J_2) = V_2$.

The next step is to control the heights of $g_{11}, g_{12}, g_{21}, g_{22}$, in a number field that contains all these elements. To this end, we treat g_1, g_2 as polynomials in two variables each, say y_1, y_2. We can assume that the leading coefficients of g_1, g_2 in y_1 belong to $\mathbb{Z}$: if not then this can be achieved by a random linear transformation, and we could have applied this transformation to the linear maps defining g_1, g_2 themselves. If a is the leading coefficient of g_1, then we write $\widetilde{g}_1 := a^{\deg g_1 - 1} g(y_1/a, y_2 + b)$, for a random b with $\mathrm{ht}(b) \leqslant \cdot 2^{(n \log \sigma)^{c_2}}$. After this transformation, the polynomial $\widetilde{g}_1$ is monic in y_1, and $\widetilde{g}_1(y_1, 0)$ is squarefree. Further, factors of g_1, in particular g_{11}, g_{12} are in bijection with factors of $\widetilde{g}_1$.

We now focus on $\widetilde{g}_{11}$, the factor corresponding to g_{11} under this bijection. By Lemma 3.16, there is an irreducible factor $q_{11}(z)$ of $\widetilde{g}_1(z, 0)$, integers $\Delta_{11}, \Delta'_{11}$, and polynomials $u_{11}, v_{11}, w_{11} \in \mathbb{Z}[y_1, y_2, z]$ such that

$$\Delta_{11}\Delta'_{11}\widetilde{g}_1(y_1, y_2) = u_{11}(y_1, y_2, z)v_{11}(y_1, y_2, z) + w_{11}(y_1, y_2, z)q_{11}(z).$$

Further, $\deg_z(u_{11}), \deg_z(v_{11}) < e$. Finally, there is a root α_{11} of q_{11} such that $g_{11} = u_{11}(x, y, \alpha_{11})$, potentially after scaling. By Gelfond's inequality [HS13, B.7.3] we have $\mathrm{ht}(q_{11}) \leqslant h \cdot 2^{(n \log \sigma)^{c_3}}$, and we also have $\deg q_{11} \leqslant \deg g_1 = D$. We also have $\mathrm{ht}(\Delta_{11}), \mathrm{ht}(\Delta'_{11}), \mathrm{ht}(u_{11}) \leqslant h \cdot 2^{(n \log \sigma)^{c_3}}$ by Lemma 3.16.

We similarly construct $q_{ij}, \alpha_{ij}, u_{ij}, v_{ij}, w_{ij}, \Delta_{ij}, \Delta'_{ij}$. We now construct an extension of $\mathbb{Q}$ that contains every α_{ij}. To this end, we construct a primitive element of $\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22}$. Using Theorem 3.9 we can deduce the existence of a polynomial G with $\deg G \leqslant D^4$ and $\mathrm{ht}(G) \leqslant h \cdot 2^{(n \log \sigma)^{c_4}}$, and a root γ of q, along with Q_{ij}, r such that $\alpha_{ij} = Q_{ij}(\gamma)/r$. We also have $\mathrm{ht}(Q_{ij}, r) \leqslant h \cdot 2^{(n \log \sigma)^{c_4}}$. In equation

$$\Delta_{11}\Delta'_{11}\widetilde{g}_1(y_1, y_2) = u_{11}(y_1, y_2, z)v_{11}(y_1, y_2, z) + w_{11}(y_1, y_2, z)q_{11}(z)$$

we substitute $z = Q_{11}/r$, multiply throughout by r^{D^2}, and reduce the coefficients mod G to obtain

$$\Delta_{11}\Delta'_{11}r^{D^2}\widetilde{g}_1(y_1, y_2) = \widetilde{u}_{11}(y_1, y_2, z)\widetilde{v}_{11}(y_1, y_2, z) + \widetilde{w}_{11}(y_1, y_2, z)G(z).$$

We now have $\widetilde{u}(y_1, y_2, \gamma) = \widetilde{g}_{11}$, potentially after scaling. Up to scaling by a, the polynomial g_{11} is exactly $\widetilde{u}(ay_1, y_2 + b, \gamma)$. This shows that $\mathrm{ht}(g_{11}) \leqslant h \cdot 2^{(n \log \sigma)^{c_5}}$. Therefore, we have obtained the height bounds on g_{11} that we were looking for, and the same bound holds for g_{ij}.

Recall that J_1, J_2 correspond to different irreducible components of I, and that $J_1 + J_2$ has dimension 0. Pick any prime $p \nmid \mathrm{disc}_z(G)$ such that $G \pmod p$ has a linear factor. Let $\mathfrak{p}$ be a prime of the ring of integers of $\mathbb{Q}[z]/G(z)$ that corresponds to the linear factor. The quotient map $\mathbb{Z}[\gamma] \to \mathbb{Z}[\gamma]/\mathfrak{p}$ has image in $\mathbb{F}_p$. We now apply Lemma 4.1 and Theorem 4.5 to the ideals J_1, J_2, and $J_1 + J_2$. We deduce that there is an integer Δ_0 such that for any prime $p \nmid \Delta_0$, the images of J_1, J_2 remain irreducible, and the dimensions of $I, J_1, J_2, J_1 + J_2$ are one, one, one, zero respectively. We also have $Z((J_1)_\mathfrak{p}) \subset Z(I_\mathfrak{p})$. Therefore, $Z((J_1)_\mathfrak{p})$ is an $\mathbb{F}_p$-definable irreducible component of $Z(I_\mathfrak{p})$. The same holds for $Z((J_2)_\mathfrak{p})$. Since $\dim(J_1 + J_2)_\mathfrak{p} = 0$, these two components are different from each other. □

5 Interactive proofs of primality testing

In this section, we prove our main theorem: interactive protocols for testing non-primality of natural classes of ideals. As we have discussed in Section 1, we first give some protocols for certain special cases, as the protocol in the main theorem will invoke these special cases as subroutines.

The setting of this section is that of our main theorem: we assume that we have an algebraic circuit C' of size $5sm$ with integer coefficients that computes polynomials $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_m]$, and the partial derivatives $\partial_i f_j$. The existence of C' is guaranteed by Lemma 2.9 applied to the input circuit C. By Lemma 2.8 also applied to the input circuit C, we have $\mathrm{ht}(f_i) \leqslant 2^{2s}$ and $\deg f_i \leqslant 2^s$. Define $d := \max \deg f_i$. Define $\sigma := dm + 2$.

Let $I := (f_1, \ldots, f_m)$, and let $V := Z(I)$ denote the algebraic set corresponding to I. For any prime p, define I_p to be the ideal in $\mathbb{F}_p[x_1, \ldots, x_n]$ generated by the images of $f_1, \ldots, f_m$ under the quotient map $R \to \mathbb{F}_p[x_1, \ldots, x_n]$, and let V_p denote the algebraic set of I_p in $\overline{\mathbb{F}}_p[x_1, \ldots, x_n]$.

We begin with a protocol to decide if a zero dimensional ideal has at least two points in its zeroset.

Lemma 5.1. *Assume GRH. If I is promised to be zero dimensional, then there is an AM protocol that verifies $|V| \geqslant 2$.*

Proof. The proof uses Corollary 4.3, Lemma 4.4 and the Goldwasser-Sipser protocol (Lemma 2.1). Let $h := \max_i \mathrm{ht}(f_i)$, by Lemma 2.8 we have $h \leqslant 2^{2s}$.

Define $\pi_{red}(x)$ to be the set of primes $p \leqslant x$ such that V_p has at least two points in $\mathbb{F}_p$. By Corollary 4.3, if $|V| = 1$ then $\pi_{red}(x) \leqslant h \cdot 2^{(n \log \sigma)^{c_1}}$ for all x. Suppose $|V| \geqslant 2$. Let G be the polynomial guaranteed by Lemma 4.4. By the effective Chebotarev density theorem (Theorem 2.19) we have

$$\pi_G(x) \geqslant \frac{1}{\deg G}\left(\pi(x) - \log \Delta_G - c_2 x^{1/2} \log\left(\Delta_G x^{\deg(G)}\right)\right),$$

where Δ_G is the discriminant of G, and $\pi_G(x)$ is the set of primes $p \leqslant x$ such that G has a root in $\mathbb{F}_p$. By Lemma 4.4 $\deg G \leqslant 2^{(n \log \sigma)^{c_3}}$, and $\mathrm{ht}(G) \leqslant h \cdot 2^{(n \log \sigma)^{c_4}}$. Bounding the discriminant using Corollary 3.3, we have

$$\pi_G(x) \geqslant \frac{\pi(x)}{2^{(n \log \sigma)^{c_3}}} - c_5 h x^{1/2} 2^{(n \log \sigma)^{c_6}} - c_7 x^{1/2} \log x.$$

By Theorem 2.20 we have $\pi_{red}(x) \geqslant \pi_G(x) - h \cdot 2^{(n \log \sigma)^{c_8}}$. Finally, by [Dus10, Theorem 6.9] we have $\pi(x) > \frac{x}{\log_e x}$ for all x larger than 600. For $x_0 = h^{c_9} \cdot 2^{(n \log \sigma)^{c_{10}}}$ for a large enough constants c_9, c_{10}, we see that $\pi_{red}(x_0) \geqslant 2h \cdot 2^{(n \log \sigma)^{c_1}}$. Note that $\mathrm{ht}(x_0)$ is polynomial in the input.

We can now invoke the Goldwasser-Sipser protocol (Lemma 2.1) on the set $\pi_{red}(x_0)$, with $K = h \cdot 2^{(n \log \sigma)^{c_1}}$. Membership in this set is in NP: the certificate for membership are two distinct roots, which has polynomial bit complexity, and we can evaluate the input polynomials on these roots to verify them. As shown above, $|V| \geqslant 2$ if and only if $|\pi_{red}(x_0)| \geqslant 2K$ and $|V| = 1$ if and only if $|\pi_{red}(x_0)| \leqslant K$. Therefore, the protocol correctly verifies $|V| \geqslant 2$. □

The second special case is a protocol that decides if a zeroset has at least two components of highest dimension.

Lemma 5.2. *Assume GRH. If I is promised to be equidimensional and have dimension r, then there is an AM protocol that verifies that V has at least two irreducible components of dimension r.*

Proof. Let $D := d^n$. Apply the algorithm of Corollary 2.26 in order to obtain a linear subspace H. The space H is defined by linear equations $h_1, \ldots, h_{r-1}$, with $\mathrm{ht}(h_i) \leqslant (nD)^{c_1}$. Let $J := I + (h_1, \ldots, h_{r-1})$, and let $W := Z(J)$. By Corollary 2.26, with high probability we have $\dim(W) = 1$, and the number of irreducible components of dimension 1 of W is the same as the number of irreducible components of dimension r in V. For the rest of the proof, we assume that this event occurs, this will only change the soundness and completeness of our protocol by a small constant, which does not change the complexity class AM. Let J_p, W_p be defined the same way as I_p, V_p.

Let $h := \max(\max_j \mathrm{ht}(h_j), \max_i \mathrm{ht}(f_i))$, we have $h \leqslant (nD)^{c_1} + 2^{2s}$. Let $\pi_{red}(x)$ be the set of primes $(150)^2 D^8 < p < x$ such that $\left|W_p \cap \mathbb{F}_p^n\right| \geqslant 2p - 10D^4 p^{1/2}$. By an effective Lang-Weil bound Theorem 2.18,

$p \in \pi_{red}(x)$ if and only if $(150)^2 D^8 < p < x$ and W_p has at least two irreducible components of dimension 1 defined over $\mathbb{F}_p$.

Suppose W has only one irreducible component of dimension 1 (this component is then automatically defined over $\mathbb{F}_p$). By Theorem 4.5, for all x we have $\pi_{red}(x) \leqslant h \cdot 2^{(n \log \sigma)^c}$, since for all but $h \cdot 2^{(n \log \sigma)^c}$ primes W_p has at most one irreducible component defined over $\mathbb{F}_p$. On the other hand, suppose W has at least two irreducible components of dimension 1. Let G be the polynomial guaranteed by Theorem 4.6. By the effective Chebotarv density theorem Theorem 2.19 we have

$$\pi_G(x) \geqslant \frac{\pi(x)}{2^{(n \log \sigma)^{c_2}}} - c_3 h x^{1/2} 2^{(n \log \sigma)^{c_4}} - c_5 x^{1/2} \log x.$$

By Theorem 4.6, for all but $h \cdot 2^{(n \log \sigma)^{c_6}}$ primes in $\pi_G(x)$, the zeroset W_p has at least two components, therefore every such prime larger than $(150)^2 D^8$ is in $\pi_{red}(x)$. For $x_0 = h^{c_7} \cdot 2^{(n \log \sigma)^{c_8}}$ for a large enough constants c_7, c_8, we see that $\pi_{red}(x_0) \geqslant 2h \cdot 2^{(n \log \sigma)^c}$.

We can now invoke the extension of the Goldwasser-Sipser protocol (Corollary 2.2) on the set $\pi_{red}(x_0)$, with $K = h \cdot 2^{(n \log \sigma)^c}$. Membership in this set is itself in AM: note that $\left|W_p \cap \mathbb{F}_p^n\right|$ is either at most $p + 5D^4 p^{1/2} + D$ or at least $2p - 10D^4 p^{1/2} - D$. Since $p > (150)^2 D^8$ the ratio of these sizes is at least 1.9. If the protocol accepts, then with high probability W, and therefore V has at least two irreducible components of top dimension. If not, then with high probability V has only one component of top dimension. □

Remark 5.3. We give an example that shows that Theorem 4.6 is tight in some sense. The theorem, combined with the effective Chebotarev density theorem Theorem 2.19 shows that the density of primes for which $\mathbb{F}_p$ has two $\mathbb{F}_p$ definable components is inverse exponential. We give an example to show that this inverse exponential behaviour is unavoidable.

Consider $I = \left(x_1^2 - 2, x_2^2 - 3, \ldots, x_{n-1}^2 - p_{n-1}\right)$, where p_j is the j^{th} prime. The zeroset V consists of 2^{n-1} components, each of dimension 1. Similarly, for any prime $p > p_{n-1}$, the zeroset V_p also consists of 2^{n-1} components, each of dimension 1. However, these components might not be defined over $\mathbb{F}_p$ itself. For any prime p, the components will be defined over the smallest extension of $\mathbb{F}_p$ that contains square roots of $2, 3, \cdots, p_{n-1}$. The only primes p for which there exist $\mathbb{F}_p$ definable components of V_p are those where $2, 3, \ldots, p_{n-1}$ are all quadratic residues. By quadratic reciprocity and Dirichlet's theorem, the density of such primes within all primes is 2^{-n+1}.

5.1 Interactive proof for radical ideals

We now give an AM protocol for deciding non primality of radical ideals, using the above two protocols as subroutines.

Algorithm 3: AM protocol for non-primality of radical ideals

Input : A circuit C' that computes polynomials $f_1, \ldots, f_m \in R$ along with the derivatives $\partial_i f_j$, with $\deg f_i \leqslant d$ such that $I := (f_1, \ldots, f_m)$ is a radical ideal, and the dimension r of I.

Arthur : If $r = 0$, then perform the protocol in Lemma 5.1, and accept if and only if the protocol accepts. If $r > 0$, then send the empty string to Merlin.

Merlin : Compute the Jacobian matrix $\mathcal{J}_{ij} = \partial f_i / \partial x_j$. Check if there is a minor M of $\mathcal{J}$ of size at least $n - r + 1$ such that $M \neq 0, f_1 = 0, \ldots, f_m = 0$ is a satisfiable system. If such a minor exists, send the rows and columns that correspond to M to Arthur. If not, then send the empty string to Arthur.

Arthur : If Merlin sends the empty string then perform the protocol in Lemma 5.2, and accept if and only if the protocol accepts. If not, then construct an algebraic circuit for $1 - yM$, where M is the minor of $\mathcal{J}$ corresponding to the rows and columns received from Merlin. Perform the protocol in Theorem 2.4 on the inputs $f_1, \ldots, f_m, 1 - yM$. Accept if and only if this protocol accepts.

Theorem 5.4. *Assume GRH. If I is a radical ideal of dimension r then Algorithm 3 is a valid AM protocol for deciding I is not prime.*

Proof. Suppose I is not prime. Since I is radical, it has at least two irreducible components. If $\dim I = 0$, then $|V| \geqslant 2$, and the protocol in Lemma 5.1 correctly accepts with high probability, and therefore Algorithm 3 correctly accepts with high probability in the first step. If $r > 0$, then either V has at least one component of dimension at most $r-1$, or V is equidimensional and has at least two components of dimension r. If V has a component of dimension less than r, then at any point $x \in V$ that lies in this component, we have $\operatorname{rank} \mathcal{J}(x) > n - r$, therefore Merlin can always find a minor of $\mathcal{J}$ of size greater than $n-r$ that makes the polynomial system created by Arthur in round 3 satisfiable. Observe that the system has degree at most $nd+1$, and logarithmic height at most $h \cdot (dn)^c$, where $h := \max_i \operatorname{ht}(f_i)$. Further, given the circuit C' that computes $f_1, \ldots, f_m$ and $\partial_j f_i$, Arthur can easily produce a circuit C'' that also computes $1 - yM$. In this case, the protocol in Theorem 2.4 correctly accepts with high probability. If V has two components of dimension r, then the protocol in Lemma 5.2 correctly accepts with high probability.

Suppose I is a prime ideal. If $r = 0$, then $|V| = 1$, and the protocol in Lemma 5.1 correctly rejects with high probability, and therefore Algorithm 3 correctly fails to accept with high probability in the first step. If $r \geqslant 0$, then V has only one irreducible component. The Jacobian $\mathcal{J}$ has rank at most $n-r$ at every point in V, therefore no matter what choice of rows and columns Merlin picks in round 2, the system created by Arthur in round 3 is unsatisfiable, and therefore the protocol in Theorem 2.4 fails to accept with high probability. If Merlin sends the empty string, then the protocol in Lemma 5.2 fails to accept with high probability. Therefore, when I is a prime ideal, the protocol in Algorithm 3 correctly fails to accept with high probability.

The above protocol is in AM[4], the theorem now follows from the fact that AM[4] = AM. □

5.2 Equidimensional Cohen-Macaulay ideals

We now give an AM protocol for deciding non primality of equidimensional Cohen-Macaulay ideals.

Algorithm 4: AM protocol for non-primality of equidimensional CM ideals

Input : A circuit C' that computes polynomials $f_1, \ldots, f_m \in R$ along with the derivatives $\partial_i f_j$, with $\deg f_i \leqslant d$ such that $I := (f_1, \ldots, f_m)$ is an equidimensional Cohen-Macaulay ideal, and the dimension r of I.

Arthur : If $r = 0$, then perform the protocol in Lemma 5.1, and accept if the protocol accepts.
If $r > 0$, then perform the protocol in Lemma 5.2, and accept if the protocol accepts.
If the above protocols fail to accept, let Y, Z be $(n-r) \times m$ and $n \times (n-r)$ symbolic matrices (in new variables), and perform the protocol in Theorem 2.5 on the system $f_1, \ldots, f_m, \det(Y \mathcal{J} Z)$, with parameter $(n-r)(m+n)+r$. Accept if and only if this protocol accepts.

Theorem 5.5. *Assume GRH. If I is a equidimensinal Cohen-Macaulay ideal of dimension r then Algorithm 4 is a valid AM protocol for deciding I is not prime.*

Proof. Suppose I is not prime. Then either I has at least two components of dimension r, or I has exactly one irreducible component but is not radical. If I has at least two irreducible components of dimension r, then depending on r either the protocol in Lemma 5.1 or the protocol in Lemma 5.2 correctly accepts with high probability. Suppose now that I has a unique minimal prime, call this prime $\mathfrak{p}$. Since S is an affine domain of dimension n, by [Eis13, Cor 13.4] we have $\operatorname{codim}(I) = \operatorname{codim}(\mathfrak{p}) = n - r$. Let J be the ideal generated by the $(n-r) \times (n-r)$ minors of $\mathcal{J}$. Since S/I is Cohen-Macaulay and I is not radical, by [Eis13, Thm 18.15] we deduce that J has codimension 0 in S/I, equivalently every generator of J lies in $\mathfrak{p}$. The polynomial $\det(Y \mathcal{J} Z)$ lies in the ideal $J \otimes \overline{\mathbb{Q}}[Y, Z]$, therefore $\det(Y \mathcal{J} Z) \in \mathfrak{p} \otimes \overline{\mathbb{Q}}[Y, Z]$. The ideal $\mathfrak{p} \otimes \overline{\mathbb{Q}}[Y, Z]$ has dimension $r + (n-r)(m+n)$, and is the unique minimal prime of $I \otimes \overline{\mathbb{Q}}[Y, Z]$, therefore $\dim \det(Y \mathcal{J} Z) + I \otimes \overline{\mathbb{Q}}[Y, Z] = r + (n-r)(m+n)$. Therefore the protocol in Theorem 2.5 correctly accepts with high probability.

Now suppose I is prime. Since I has only one irreducible component, depending on r either the protocol in Lemma 5.1 or the protocol in Lemma 5.2 fails to accept with high probability. Further we have $I + J \subsetneq \mathfrak{p}$, equivalently there is at least one generator of J that fails to lie in $\mathfrak{p}$. Therefore, $\det(Y \mathcal{J} Z) \notin \mathfrak{p} \otimes \overline{\mathbb{Q}}[Y, Z]$, and $\dim \det(Y \mathcal{J} Z) + I \otimes \overline{\mathbb{Q}}[Y, Z] = r + (n-r)(m+n) - 1$. The protocol in Theorem 2.5 correctly fails to accept with high probability. □

5.3 Proof of main theorems

We now use the above protocols to prove our main theorems, which we restate here for convenience.

Theorem 1.2 (Interactive protocols for primality)**.** *Let* C *be an algebraic circuit of size* s *with integer constants that computes* $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$*. If* $I := (f_1, \ldots, f_m)$ *is either radical or equidimensional Cohen-Macaulay, and if the dimension of* I *is given, then the complexity of testing if* I *is prime lies in* coAM*, assuming GRH.*

Proof. Since we know the dimension of I, by Theorem 5.4 and Theorem 5.5 (depending on whether I is radical or equidimensional Cohen-Macaulay), the complexity of deciding if I is not prime lies in AM, therefore the complexity of deciding if I is prime lies in coAM. □

Theorem 1.3 (Primality testing in PH)**.** *Let* C *be an algebraic circuit of size* s *with integer constants that computes* $f_1, \ldots, f_m \in \mathbb{Z}[x_1, \ldots, x_n]$*. If the ideal* $I := (f_1, \ldots, f_m)$ *is either radical or equidimensional Cohen-Macaulay then the complexity of testing if* I *is prime lies in* $\Sigma_3^p \cap \Pi_3^p$*, assuming GRH.*

Proof. We first compute the dimension of I. By Theorem 2.5, there exists an AM protocol for checking if $\dim I \geqslant r$ for any r, and since $AM \subset \Pi_2^p$, the complexity of computing the dimension of I exactly lies in $P^{\Pi_2^p} \subset \Sigma_3^p \cap \Pi_3^p$. With the dimension at hand, we can apply Theorem 1.2 and decide whether I is prime in $P^{\Pi_2^p}$. Thus, the complexity of deciding if I is prime is in $\Sigma_3^p \cap \Pi_3^p$. □

6 Conclusion & open problems

In this work, we proved that the ideal primality testing problem is in the third level of the polynomial hierarchy for the natural classes of ideals comprised of radical ideals and equidimensional Cohen-Macaulay ideals. This significantly tightens the complexity-theoretic gap for the primality testing problem for these classes of ideals, given that the primality testing problem is already coNP-hard for such classes. Prior to our work, the best upper bounds for testing whether a radical ideal or a zero-dimensional ideal is prime was PSPACE, whereas for complete intersections the best upper bound was EXP.

- A common nice feature of the two classes of ideals that we studied is that any associated prime must be a minimal prime, and therefore we avoid the issue of having to detect *embedded* primes. The issue of embedded primes, which causes non-reduced behavior in "very small" parts of our algebraic set, is in general a very hard problem to identify (and therefore to handle). We leave it as an open question, to detect embedded primes either "via zeros" or algorithmically better than EXPSPACE.
- We studied the problems over the field of $\mathbb{C}$ (or $\mathbb{Q}$). This allows us to go modulo several primes p, leading to the application of arithmetic-geometry results over finite fields (e.g. Lang-Weil points count and Chebotarev primes count). The ideal primality testing problem over an input field $\overline{\mathbb{F}}_q$, for prime q, rules out the idea of going modulo other primes $p \neq q$. This makes it an open question to extend Koiran's result [Koi96] to input field $\overline{\mathbb{F}}_q$. Similarly, our paper leaves it as an open question, to improve ideal primality testing beyond EXPSPACE, for natural classes of ideals over an algebraically closed field of positive characteristic. (E.g. what about the dimension $\leqslant 1$ case here?)

Bibliography

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach*. Cambridge University Press, 2009. 11

[AO83] Leonard M Adleman and Andrew M Odlyzko. Irreducibility testing and factorization of polynomials. *Mathematics of Computation*, 41(164):699–709, 1983. 15

[BC04] Peter Bürgisser and Felipe Cucker. Variations by complexity theorists on three themes of Euler, Bézout, Betti, and Poincaré. *Complexity of computations and proofs*, 13:73–152, 2004. 3, 4

[BCGW93] Chanderjit Bajaj, John Canny, Thomas Garrity, and Joe Warren. Factoring rational polynomials over the complex numbers. *SIAM Journal on Computing*, 22(2):318–331, 1993. 3

[BM93] Dave Bayer and David Mumford. What can be computed in algebraic geometry? *arXiv preprint alg-geom/9304003*, 1993. 3, 6, 23

[BS83] Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoretical computer science*, 22(3):317–330, 1983. 13

[BS07] Peter Bürgisser and Peter Scheiblechner. Differential forms in computational algebraic geometry. In *Proceedings of the 2007 international symposium on Symbolic and algebraic computation*, pages 61–68, 2007. 3, 4

[BS10] Peter Bürgisser and Peter Scheiblechner. Counting irreducible components of complex algebraic varieties. *computational complexity*, 19:1–35, 2010. 3

[Bür13] Peter Bürgisser. *Completeness and reduction in algebraic complexity theory*, volume 7. Springer Science & Business Media, 2013. 13

[Chi86] AL Chistov. Algorithm of polynomial complexity for factoring polynomials and finding the components of varieties in subexponential time. *Journal of Soviet Mathematics*, 34:1838–1882, 1986. 3

[CLO97] David Cox, John Little, and Donal O'shea. *Ideals, varieties, and algorithms*, volume 3. Springer, 1997. 22

[CM06] Antonio Cafure and Guillermo Matera. Improved explicit estimates on the number of solutions of equations over a finite field. *Finite Fields and Their Applications*, 12(2):155–185, 2006. 15, 17

[DFGS91] Alicia Dickenstein, Noaï Fitchas, Marc Giusti, and Carmen Sessa. The membership problem for unmixed polynomial ideals is solvable in single exponential time. *Discrete Applied Mathematics*, 33(1-3):73–94, 1991. 3

[Dic05] Alicia Dickenstein. *Solving polynomial equations*. Springer, 2005. 17

[Dub90] Thomas W Dubé. The structure of polynomial ideals and Gröbner bases. *SIAM Journal on Computing*, 19(4):750–773, 1990. 3

[Dus10] Pierre Dusart. Estimates of some functions over primes without RH. *arXiv preprint arXiv:1002.0442*, 2010. 15, 29

[EHV92] David Eisenbud, Craig Huneke, and Wolmer Vasconcelos. Direct methods for primary decomposition. *Inventiones mathematicae*, 110(1):207–235, 1992. 3

[Eis13] David Eisenbud. *Commutative algebra: with a view toward algebraic geometry*, volume 150. Springer Science & Business Media, 2013. 3, 4, 5, 31

[Gao03] Shuhong Gao. Factoring multivariate polynomials via partial differential equations. *Mathematics of computation*, 72(242):801–822, 2003. 3

[GD65] Alexandre Grothendieck and Jean A. Dieudonné. Éléments de géométrie algébrique iv. *Publ. Math. IHES, Ibid.*, 24, 1965. 5

[Gri86] Dima Grigoriev. Factorization of polynomials over a finite field and the solution of systems of algebraic equations. *Journal of Soviet Mathematics*, 1986. 3

[GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, pages 59–68, 1986. 11

[GSS19] Zeyu Guo, Nitin Saxena, and Amit Sinhababu. Algebraic dependencies and PSPACE algorithms in approximative complexity over any field. *Theory of Computing*, 15(1):1–30, 2019. 14

[GTZ88] Patrizia Gianni, Barry Trager, and Gail Zacharias. Gröbner bases and primary decomposition of polynomial ideals. *Journal of Symbolic Computation*, 6(2-3):149–167, 1988. 3

[Hei83] Joos Heintz. Definability and fast quantifier elimination in algebraically closed fields. *Theoretical Computer Science*, 24(3):239–277, 1983. 14

[Her26] Grete Hermann. Die Frage der endlich vielen Schritte in der Theorie der Polynomideale: Unter Benutzung nachgelassener Sätze von K. Hentzelt. *Mathematische Annalen*, 95(1):736–788, 1926. 3, 8

[HS81] Joos Heintz and Malte Sieveking. Absolute primality of polynomials is decidable in random polynomial time in the number of variables. In *International Colloquium on Automata, Languages, and Programming*, pages 16–28. Springer, 1981. 3

[HS13] Marc Hindry and Joseph H Silverman. *Diophantine geometry: an introduction*, volume 201. Springer Science & Business Media, 2013. 28

[Jel05] Zbigniew Jelonek. On the effective nullstellensatz. *Inventiones mathematicae*, 162(1):1–17, 2005. 8, 16

[Kal85] Erich Kaltofen. Fast parallel absolute irreducibility testing. *Journal of Symbolic Computation*, 1(1):57–67, 1985. 3

[Kal95] E Kaltofen. Effective Noether irreducibility forms and applications. *Journal of Computer and System Sciences*, 50(2):274–295, 1995. 3, 6, 8, 19, 23

[Koi96] Pascal Koiran. Hilbert's nullstellensatz is in the polynomial hierarchy. *Journal of complexity*, 12(4):273–286, 1996. 4, 5, 9, 12, 15, 20, 24, 26, 32

[Koi97] Pascal Koiran. Randomized and deterministic algorithms for the dimension of algebraic varieties. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 36–45. IEEE, 1997. 4, 5, 12

[KPS01] Teresa Krick, Luis Miguel Pardo, and Martin Sombra. Sharp estimates for the arithmetic Nullstellensatz. *Duke Mathematical Journal*, 109(3):521 – 598, 2001. 18, 20

[Las05] Emanuel Lasker. Zur theorie der moduln und ideale. *Mathematische Annalen*, 60(1):20–116, 1905. 3

[Len84] Arjen K. Lenstra. Polynomial-time algorithms for the factorization of polynomials, 1984. 22

[LO77] Jeffrey C Lagarias and Andrew M Odlyzko. Effective versions of the Chebotarev density theorem. In *Algebraic number fields: L-functions and Galois properties (Proc. Sympos., Univ. Durham, Durham, 1975)*, volume 7, pages 409–464, 1977. 15

[Mig83] Maurice Mignotte. Some useful bounds. In *Computer Algebra: Symbolic and Algebraic Computation*, pages 259–263. Springer, 1983. 18

[Mil20] James S. Milne. Algebraic number theory (v3.08), 2020. Available at www.jmilne.org/math/. 15

[MR13] Ernst W Mayr and Stephan Ritscher. Dimension-dependent bounds for Gröbner bases of polynomial ideals. *Journal of Symbolic Computation*, 49:78–94, 2013. 8, 16

[MT17] Ernst W Mayr and Stefan Toman. Complexity of membership problems of different types of polynomial ideals. *Algorithmic and experimental methods in algebra, geometry, and number theory*, pages 481–493, 2017. 3

[Poo08] Bjorn Poonen. Rational points on varieties. *Graduate Studies in Mathematics*, 186:337, 2008. 5

[Sch07] Peter Scheiblechner. *On the Complexity of Counting Irreducible Components and Computing Betti Numbers of Algebraic Varieties*. PhD thesis, University of Paderborn, 2007. 24

[Sei74] Abraham Seidenberg. Constructions in algebra. *Transactions of the American Mathematical Society*, 197:273–313, 1974. 3, 5, 12

[Sei78] Abraham Seidenberg. Constructions in a polynomial ring over the ring of integers. *American Journal of Mathematics*, 100(4):685–703, 1978. 3

[SR94] Igor Rostislavovich Shafarevich and Miles Reid. *Basic algebraic geometry*, volume 2. Springer, 1994. 14, 26

[Sta24] The Stacks project authors. The stacks project. https://stacks.math.columbia.edu, 2024. 4, 22

[SY10] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends® in Theoretical Computer Science*, 5(3–4):207–388, 2010. 13

[WR76] Peter J. Weinberger and Linda Preiss Rothschild. Factoring polynomials over algebraic number fields. *ACM Transactions on Mathematical Software (TOMS)*, 2(4):335–350, 1976. 15